



OpenText™ Server Backup Portal 9.5

User Guide

Copyright 2025 Open Text

The only warranties for products and services of Open Text and its affiliates and licensors (“Open Text”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

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

For terms and conditions, see [Terms and Conditions](#).

Version History

Version	Date	Description
1	January 2026	Initial user guide for Portal 9.5x.

Contents

1 Get started with Server Backup Portal	7
1.1 Pages in Portal	7
1.2 Sign in	10
1.3 Set or change your password	11
1.4 Verify your account	16
1.5 Navigate to a page	19
1.6 View Portal text in another language	20
1.7 View and hide Dashboard messages	20
1.8 View records	21
1.9 Change your default page settings	24
1.10 Sign out	25
2 Add and manage computers in Portal	26
2.1 Add a Windows computer	26
2.2 Add a Windows cluster	39
2.3 Add a Linux computer	41
2.4 Add a Hyper-V environment	53
2.5 Add a vSphere environment	72
2.6 Assign computers to groups	87
2.7 View job and backup status information for offline computers	88
2.8 Undelete Hyper-V environments	89
2.9 Move computers between sites	89
2.10 Minimum Agent and plug-in permissions	90
3 Configure computers	92
3.1 Add vault settings	92
3.2 Add a description	94
3.3 Add retention types	94
3.4 Configure bandwidth throttling	96
3.5 Set the data read error handling method for a Windows computer	97
3.6 Change credentials or the network address for accessing Hyper-V	98
3.7 Start configuring computers from Dashboard notifications	99
4 Add and edit backup jobs	102
4.1 Add a Windows Local System backup job	102
4.2 Add a Windows Image backup job	108
4.3 Add a UNC file backup job	112
4.4 Add backup jobs for a Windows cluster	115

4.5 Add a SQL Server database backup job	116
4.6 Add an Exchange backup job	121
4.7 Add an Oracle database backup job	123
4.8 Add a Linux backup job	125
4.9 Add a UNIX backup job	130
4.10 Add an NFS backup job	132
4.11 Add a vSphere backup job	134
4.12 Add and schedule a Hyper-V backup job	140
4.13 Log file options	154
4.14 Encryption settings	155
4.15 Advanced backup options	156
4.16 Filter subdirectories and files in backup jobs	157
4.17 Edit a backup job	158
5 Delete jobs and computers, and delete data from vaults	160
5.1 Delete a backup job without deleting data from vaults	160
5.2 Delete a backup job and delete job data from vaults	161
5.3 Cancel a scheduled job data deletion	163
5.4 Delete a computer without deleting data from vaults	164
5.5 Delete a computer and delete computer data from vaults	165
5.6 Cancel a scheduled computer data deletion	167
5.7 Delete specific backups from vaults	168
6 Run and schedule backups, synchronizations and custom commands	170
6.1 Schedule a backup	172
6.2 Schedule a backup to run multiple times per day	177
6.3 Maximum number of restore points for a job	182
6.4 Specify whether scheduled backups retry after a failure	183
6.5 Trigger backups when events occur on Windows desktop computers	184
6.6 Disable or enable all scheduled backup jobs	186
6.7 Run an ad-hoc backup	187
6.8 Plan Full and Incremental Exchange backups	190
6.9 Synchronize a job	191
6.10 Schedule a custom command	191
7 Resolve certificate failures and potential threats	195
7.1 Resolve certificate failures	195
7.2 Manage potential ransomware threats	196
8 Restore data	200
8.1 Restore Windows data	200

8.2 Recover a Windows cluster	211
8.3 Restore databases and application data	215
8.4 Restore Linux or UNIX files and folders	232
8.5 Restore a Linux system from a BMR backup	235
8.6 Restore a Linux system without a BMR backup	239
8.7 Restore vSphere data	242
8.8 Restore Hyper-V data	254
8.9 Recover jobs and settings from an offline Hyper-V Agent	268
8.10 Search for files to restore	274
8.11 Filter subdirectories and files when restoring data	275
8.12 Restore data to a replacement computer	277
8.13 Restore data from another computer	279
8.14 Advanced restore options	280
9 Monitor computers, jobs and processes	282
9.1 Monitor recent events using the Status Feed	282
9.2 Monitor backups and computers using the Current Snapshot	285
9.3 Monitor storage usage using Site Usage charts and emailed alerts	286
9.4 View computer and job status information	287
9.5 View skipped rates and backup status histories	290
9.6 View an unconfigured computer's logs	294
9.7 View current process information for a job	295
9.8 Monitor backups using email notifications	297
9.9 View a job's process logs and safeset information	302
9.10 View, export and email backup statuses on the Monitor page	304
9.11 View a Hyper-V VM's backup history and logs	307
9.12 Determine whether an agent has been configured automatically	310
10 View and schedule reports	313
10.1 Portal report descriptions	313
10.2 Schedule the Daily Status Report	322
10.3 View a report	324
10.4 View the Backup Verification Report	327
10.5 View the Aggregated Usage Summary Report	329
10.6 Switch to another report view	331
10.7 Save a report view	332
10.8 Delete a customized report view	332
10.9 Export a report	332
10.10 Email a report	333

- 10.11 Schedule an emailed report 334
- 10.12 View a Usage Summary Report chart 337
- 11 Server Backup Support 341**
- 11.1 Contacting Support 341

1 Get started with Server Backup Portal

Server Backup Portal provides a central access point for managing server backups and restores. You can create and run backup jobs, restore data, monitor computers and processes, and generate reports— all within Portal.

1.1 Pages in Portal

Server Backup Portal includes a series of pages where you can perform tasks for managing backups and restores. Depending on your user type and site, you can navigate to some or all of the following pages:

- [Dashboard](#)
- [Computers page](#)
- [Monitor page](#)
- [Reports page](#)
- [Policies page](#)
- [Sites page](#)
- [Users page](#)

Note: Additional pages are available for Super users and Support users, who manage or view all Portal sites.

Your user type determines which pages you can view. For example, Admin users and regular users can create and run backup jobs on the Computers page, but Super users cannot access the Computers page.

1.1.1 Dashboard

The Dashboard in Portal provides notifications and summary information about computers, backups, restores and activity in your site.

Information that appears in the center of the Dashboard depends on the item you select in the Notification Center:

- If you select **What's New** in the Notification Center, messages from your service provider and notifications of newly-added computers in your site appear in the Dashboard. You can view and hide messages, and start to configure newly-added computers. See [View and hide Dashboard messages](#) and [Start configuring computers from Dashboard notifications](#).
- If you select **Status Feed** in the Notification Center, notifications of recent backups, restores and configuration changes appear in the Dashboard, with the most recent events at the top of the list. See [Monitor recent events using the Status Feed](#).

The Dashboard also includes the Current Snapshot. The Current Snapshot shows total numbers of backups and computers in your site in categories such as Backups Requiring Attention, Successful Backups and Offline Computers. You can navigate from totals in the Current Snapshot to detailed information in Portal. For example, by clicking the number of missed backups in the Current Snapshot, you can navigate to the

Monitor page to determine which computers have missed backups. See [Monitor backups and computers using the Current Snapshot](#).

A Quick Links area might appear on the Dashboard. This area only appears if links have been added in the Portal instance.

In some Portal instances, a Tools area appears on the Dashboard. This area can include links for downloading Server Backup software and obtaining assistance.

In some Portal instances, a Site Usage area appears on the Dashboard. This area shows the amount of data backed up for a site in a billing period compared to a usage checkpoint amount. See [Monitor storage usage using Site Usage charts and emailed alerts](#).

1.1.2 Computers page

On the Computers page in Portal, you can view computers and environments in your site, configure and run backup jobs and restore data.

You can view information and options available for a computer by clicking the computer row. You can then create backup jobs, edit jobs, schedule and run backups, run restores, view logs and more.

For each computer in the list, a computer type icon indicates the computer or environment type. Computer type icons include:

-  Desktop
-  Server
-  Node in a cluster
-  Virtual cluster server
-  VMware vSphere environment
-  Microsoft Hyper-V environment

The Computers page shows the site name for each computer. If a site is a parent site, a parent site icon () appears beside the site name.

An Actions list appears for Admin and regular users. Using a command in the Actions list, Admin and regular users can add computers. Admin users can also delete offline computers, enable and disable scheduled jobs for computers, assign policies to computers, and unassign policies from computers.

1.1.3 Monitor page

On the Monitor page in Portal, you can view the last backup statuses for your jobs, monitor processes and view logs.

The Monitor page shows the computer name and site name for each job. If a site is a parent site, a parent site icon () appears beside the site name.

You can navigate from information on the Monitor page to related information on the Computers page or in the Logs window. If you click an online computer name or job name, you can view the computer's record on the Computers page. If you click a job's last backup status, you can view the job's logs in the History / Logs window.

1.1.4 Reports page

The Reports page in Portal, which is only available for Admin and Support users, shows reports that can be generated. Customization options and saved report views are also available.

Available reports include the Backup Verification Report, Daily Status Report, Usage Summary Report, Backup Status Report, Activity Details Report, Backup Details Report and Custom Command Report. See [Portal report descriptions](#).

1.1.5 Policies page

The Policies page in Portal, which is only available for Admin users, allows Admin users to create and delete policies, and assign policies to computers. A policy is a collection of settings for computers and jobs.

Note: The Policies tab is not available for Admin users with E3 appliances.

Dashboard Computers Monitor Reports Policies Sites Users					Search Knowledge Base
Actions	1 2	Show 25	Records per page	Save View	All
Name	Description	Assigned	Created		
☐		All			
☐ AK Policy		No	2017-09-28 5:22:50 AM		
☐ Demo policy		Yes	2013-11-20 9:10:40 AM		
☐ dtsPolicy		No	2014-08-29 2:34:19 PM		
☐ encryptionpolicy		Yes	2015-10-16 2:11:29 AM		
☐ Compression_policy		Yes	2016-08-17 10:49:37 AM		

1.1.6 Sites page

The Sites page in Portal is available for Super users, and for Admin users who are allowed to manage child sites. The Sites page is used to manage sites and users.

1.1.7 Users page

On the Users page in Portal, Super users and Admin users can create and manage users.

Sites Users Global Settings Usage Metrics Custom Feed Items					
Create New User	Show 25	Records per page	Save View	All	
Email Address (Username)	First Name	Last Name	Status	Role	Site Name
			All	All	
admin@test.com	Admin	A	Active	Admin	Bit_child2
child1@test.com	a	a	Active	Admin	Bit_child1
ro@test.com	Read	Only	Active	Read Only	Bit_child2
super	Super		Active	Super User	Default Entire Portal Tenant
support@test.com	a	a	Disabled	Support User	Default Entire Portal Tenant
user@test.com	User	U	Active	User	Bit_child1

1.2 Sign in

Before you can manage backups and restores, you must sign in to the Portal website. If you do not know the Portal website address or do not have a Portal user name and password, please contact your service provider. If you have forgotten your password, you can reset it and choose a new password. See [Reset your password](#).

When you sign in to Portal, your user type determines which information and functionality you can access.

Note: In some Portal instances, users sign in with credentials that they also use to sign in to other systems. See [Sign in using single sign-on credentials](#).


To sign in:

1. In a web browser, go to the Portal website.

The Sign In page appears.

Note: If a cookie banner appears, you cannot sign in until you specify your cookie preferences. You can change your cookie preferences at any time by clicking the cookie preferences link at the bottom of a Portal page.

Note: If an Authentication Required dialog box appears, you must sign in using single sign-on credentials from your organization. See [Sign in using single sign-on credentials](#).



2. In the **Email Address** box, type your email address for signing in to Portal.
3. In the **Password** box, type your password.

If you have forgotten your password, you can reset it. See [Reset your password](#).

4. Click **Sign In**.

If the **Dashboard** appears, you are now signed in and can start managing computers, backups and restores.

If the **Change Password** page appears, you must change your password. See [Change your password when required at sign-in](#).

If the **Now we need to verify your account** page appears, you must enter an account verification code from a text message or automated voice call. See [Verify your account](#).

If the **Set up two-factor account verification** page appears, you can set up this extra layer of account security. See [Set up or skip account verification at sign-in](#).

1.2.1 Sign in using single sign-on credentials

In some Portal instances, users sign in with user names and passwords that are also used to sign in to other systems. These single sign-on credentials are managed and authenticated by an external identity server (e.g., Active Directory).

Admin and regular users who sign in using single sign-on credentials must change their Agent Registration passwords in Portal before they can register agents. See [Change your Agent Registration password](#).


If you sign in using single-sign on credentials, we recommend closing all web browser windows after you sign out. See [Sign out](#).

Note: If you sign in to Portal using credentials that you also use to sign in to other systems, you cannot reset your password in Portal.

To sign in using single sign-on credentials:

1. In a web browser, navigate to the Portal website.

If an Authentication Required dialog box appears, you can enter single sign-on credentials from your organization. Continue to [Step 2](#).



If a Sign In page appears, you must sign in using credentials that are managed in Portal. See [Sign in](#).

2. In the **User name** box, type your user name.

You can enter your user name in any of the following formats: username@domain.com, username, or domain\username

3. In the **Password** box, type your password.
4. Click **OK**.

Note: If the sign-in is not successful, you might need to change your single-sign on password outside of Portal.

1.3 Set or change your password

If you are a new Portal user and receive an email with a "Set My Password" button, you can set your initial Portal password. See [Set your password](#).

If you forget your password for signing in to Portal, you can reset it and choose a new password. See [Reset your password](#).

You might be required to change your password immediately after you sign in to Portal. See [Change your password when required at sign-in](#). You can also change your Portal password at any time. See [Change your password](#).

You must choose a password that meets requirements set by your organization.

If you sign in to Portal using credentials that are also used to sign in to other systems, you cannot change your password using Portal. However, Admin and regular users with single sign-on credentials can set Agent Registration passwords in Portal. See [Change your Agent Registration password](#).

Super users and Admin users can also set and change users' passwords.

1.3.1 Set your password

If you are a new Portal user and receive an email with a "Set My Password" button, you can set your initial Portal password. This email is sent to the email address that you will use for signing in to Portal.

This email is valid for 24 hours. If you do not set your password before it expires, you can request a new email for setting your password.

To set your password:

1. In the email you received for setting your Portal password, click **Set My Password**.



The Set Your Password page opens in a web browser.

2. In the **New Password** and **Confirm New Password** boxes, enter your new password.

Your password must meet the requirements shown on the page.

3. Click **Set my password**.
4. Do one of the following:
 - If a Take Me To Portal button appears, click **Take Me To Portal**. If the Set up two-factor account verification page appears, you can set up two-factor verification. See [Set up or skip account verification at sign-in](#). You can then sign in using your email address and new password. See [Sign in](#).
 - If a message states that the code is no longer valid, do the following:
 - a. Click **Forgot my password**.
 - b. In the **Enter your username** box, enter the email address that you will use to sign in to Portal. Click **Submit**.
 - c. Check your email to find the new email with a "Set My Password" button, then return to Step 1 to set your Portal password.

1.3.2 Reset your password

If you have forgotten your password, you can reset the password and choose a new one.


If your Portal account is locked because of failed login attempts, you cannot reset your password using this procedure. Instead, please contact your service provider to unlock your account.

To reset your password:

1. In a web browser, go to the Portal website.

The Sign In page appears.

Note: If an Authentication Required dialog box appears, you cannot reset your password using Portal. See [Sign in using single sign-on credentials](#).

A screenshot of a web form for signing in. It has two input fields: 'Email Address' with the placeholder 'yourname@domain.com' and 'Password' with the placeholder 'Password'. Below the fields is a blue 'Sign In' button. At the bottom of the form is a link that says 'Forgot my password?' which is highlighted with a red rectangular box.

2. On the Sign In page, click **Forgot my password**.

The **Forgot Your Password** page appears.

3. In the **Enter your username** box, enter the email address that you use to sign in to Portal.
4. Click **Submit**.
5. If the **Now we need to verify your account** page appears, you must enter an account verification code from a text message or automated voice call. See [Verify your account](#).
6. Check your email. Find the Password Reset Request email.
7. In the Password Reset Request email, click **Reset My Password**.

Note: The link in the email is valid for 24 hours. You must reset your password within 24 hours of requesting the reset link, or you cannot use the link.

The Set Your Password page opens in a web browser.

8. In the **New Password** and **Confirm New Password** boxes, enter your new password.
Your password must meet the requirements shown on the page.
9. Click **Set my password**.

A Take Me To Portal button appears on the Set Your Password page.

10. Click **Take Me To Portal**.

You can then sign in to Portal using your email address and new password. See [Sign in](#).

1.3.3 Change your password


You can change your password any time you are signed in to Portal.

If you sign in to Portal using single sign-on credentials, you cannot change your password using this procedure; you must change your password outside of Portal. However, Admin and regular users with single sign-on credentials can set Agent Registration passwords in Portal. See [Change your Agent Registration password](#).

To change your password:

1. At the top of the Portal page, click the email address that you used to sign in to Portal.

A user menu appears.



2. Click **Profile Settings**.

Your profile appears. Your profile includes a section for changing your password.

3. In the **Current Password** box, type your current password.
4. In the **New Password** and **Confirm New Password** boxes, type your new password.

Your new password must meet the requirements shown on the page.

5. Click **Update Password**.

1.3.4 Change your password when required at sign-in

You might be required to change your password when you sign in to Portal. Your password must meet the requirements set by your organization.

You can also choose to change your password any time you are signed in to Portal. See [Change your password](#).

To change your password at sign-in:

1. Sign in to Portal. See [Sign in](#).

The Change Password page appears.

Note: If the Change Password page does not appear, you do not have to change your password.

2. In the **Current Password** box, type your current password.
3. In the **New Password** and **Confirm New Password** boxes, type your new password.

Your new password must meet the requirements shown on the page.

4. Click **Change Password**.

If the **Dashboard** appears, you are now signed in and can start managing computers, backups and restores.

If the **Now we need to verify your account** page appears, you must enter an account verification code sent as a text message or automated voice call. See [Verify your account](#).

If the **Set up two-factor account verification** page appears, you can set up two-factor account verification. See [Set up or skip account verification at sign-in](#).

1.3.5 Change your Agent Registration password

If you sign in to Portal using credentials that are also used to sign in to other systems, you cannot change your sign-in password using Portal. Because single sign-on credentials are managed and authenticated by an external identity server (e.g., Active Directory), you must change your sign-in password outside of Portal.


If you are an Admin user or regular user that signs in using single sign-on credentials, you must enter an Agent Registration password before you can register agents to Portal. Changing the Agent Registration password does not change your password for signing in to Portal.

Super users and Admin users can also set and change users' Agent Registration passwords.

To change your Agent Registration password:

1. At the top of the Portal page, click the user name that you used to sign in to Portal.

A user menu appears.



2. Click **Profile Settings**.

Your profile appears. If you sign in to Portal using credentials that are also used to sign in to other systems, your profile includes a section for changing your Agent Registration password.

Change Agent Registration Password

Your new password cannot be the same as your old password, and must conform to stated policies.

Changing the password will not affect the SSD credentials but will impact agent registrations.

Current Password

New Password

Confirm New Password

Password must contain at least:

- 8 characters
- 1 capital letter(x)
- 1 lowercase letter(x)
- 1 numeric character(x)
- 1 special character(x) from the following list: !@#%&'&()*+,-=;:[]{}^_`~|/ ?

3. If the **Current Password** box appears, type your current password.

The **Current Password** box does not appear if you have not entered an Agent Registration password in Portal yet.

4. In the **New Password** and **Confirm New Password** boxes, type your password.

Your new password must meet the requirements shown on the page.

5. Click **Change**.

1.4 Verify your account

If you are prompted to verify your account when you sign in to Portal, you must enter an account verification code from a text message or automated voice call. The code is sent to a phone number that you specified when you set up two-factor account verification. To view or change the phone number, see [Set up or change two-factor account verification](#).

When two-factor verification is set up for your account, you are sometimes prompted to enter a verification code when you sign in to Portal or reset your password. Two-factor verification provides an extra layer of security for your account, and is only available in some Portal instances.

Note: If your phone number has changed and you cannot sign in to Portal because you can no longer receive account verification codes, please contact your Portal administrator. You can then be prompted to set up two-factor account verification again when you try to sign in.

To verify your account:

1. Sign in to Portal. See [Sign in](#).

If the **Verify your account** page appears, you must verify your account. If problems occur during this process, click **Return to Sign In** to return to the Sign In page.

If the **Set up two-factor account verification** page appears, you can set up two-factor account verification. In some cases, you can skip setting up two-factor account verification. See [Set up or skip account verification at sign-in](#).

If neither of these pages appear, you do not have to verify your account. You can check your profile settings to see if you can set up two-factor account verification. See [Set up or change two-factor account verification](#).

2. Do one of the following:

- To receive a verification code in a text message, select **Send me a text message**.
- To receive a verification code in an automated voice call, select **Call me**.

3. Click **Submit**.

The **Enter your code** page appears.

4. Check your phone for a text message or automated voice call with an account verification code. Enter the code.

5. Do one of the following:

- If you only want to verify your account periodically (e.g., every 30 days) when you sign in to Portal from this web browser, select **Remember me on this device**.
- To verify your account every time you sign in to Portal from this web browser, clear **Remember me on this device**.

Note: Beginning in Portal 9.10, if two-factor account verification is set up in your Portal instance, you must verify your account every time you sign in from a new web browser.

6. Click **Verify**.

If the Dashboard appears, you are now signed in and can download and install agent software and set up backups.

If a message states that the code you entered is incorrect, enter the code again, and then click **Submit**. You can also click **Resend my code** to receive a new account verification code.

1.4.1 Set up or skip account verification at sign-in

You might be prompted to set up two-factor account verification when you sign in to Portal. When two-factor account verification is set up, you are sometimes prompted to enter a verification code when you sign in. This extra layer of account security is only available in some Portal instances.

In some Portal instances, you can skip setting up two-factor account verification when prompted at sign-in if you did not set it up previously. In other Portal instances, you are required to set up two-factor account verification.

You can also set up or change two-factor account verification at any time by entering a phone number in your profile settings. See [Set up or change two-factor account verification](#).

To set up or skip two-factor account verification when prompted at sign-in:

1. Sign in to Portal. See [Sign in](#).

If the **Set up two-factor account verification** page appears, you can set up two-factor verification. If this page does not appear, check your profile settings to see if you can set up two-factor verification. See [Set up or change two-factor account verification](#).

2. Do one of the following:

- To set up two-factor verification for your account, do the following:
 - a. Specify a phone number for receiving verification codes.
 - b. Select **Send me a text message** to receive a verification code in a text message, or select **Call me** to receive a code in an automated voice call.
 - c. Click **Submit**.

The Enter your code page appears.

- d. Check your phone for a text message or automated voice call with an account verification code. Enter the code in the **Enter your code** box.
- e. If you do not want to verify your account every time you sign in to Portal from this web browser, select **Remember me on this device**.
- f. Click **Verify**.

If a You've set up two-factor account verification message appears, click **Take me to Portal**.

If a message states that the code you entered is incorrect, enter the code again, and click **Submit**. You can also click **Resend my code** to receive a new account verification code, or click **change my phone number** to enter a new phone number for receiving codes.

- If a **Skip this step** option appears at the bottom of the page and you want to skip setting up two-factor verification, click **Skip this step**.

The **Skip this step** option does not appear if two-factor account verification is required in your Portal instance or if you already set up two-factor account verification.

- To return to the Sign In page, click the **Return to Sign In** option. See [Sign in](#).

1.4.2 Set up or change two-factor account verification


If two-factor account verification is available in your Portal instance, you can set account verification up at any time by entering a phone number in your profile settings. You can also change your phone number for receiving verification codes.

With two-factor account verification, you are sometimes prompted to enter a verification code when you sign in to Portal. The code is sent in a text message or automated voice call to a phone number that you specify.

To set up or change two-factor account verification:

1. At the top of the Portal page, click the email address that you used to sign in to Portal.

A user menu appears.



2. Click **Profile Settings**.

Your profile appears. If your profile includes a **Two-Factor Account Verification** section, you can set up two-factor account verification or change your phone number for receiving verification codes.

If this section does not appear, two-factor account verification is not available in your Portal instance. Your Portal administrator can find instructions for enabling two-factor account verification in the *Portal Administration Guide*.

3. Do one of the following:

- To set up two-factor account verification, click **Setup**.

Two-Factor Account Verification

Two-factor account verification keeps your portal account and your data safe by periodically requiring account validation via phone.

Phone Number

Setup

- To change your phone number for receiving verification codes, click **Change phone**.

Two-Factor Account Verification

Two-factor account verification keeps your portal account and your data safe by periodically requiring account validation via phone.

Phone Number

###-###-3720

Change phone

The **Choose a phone number** box appears.

4. Specify a phone number for receiving verification codes.

5. Do one of the following:

- To receive a verification code in a text message, select **Send me a text message**.
- To receive a verification code in an automated voice call, select **Call me**.

6. Click **Send code**.

The **We've sent a code to your phone** box appears.

7. Check your phone for a text message or automated voice call with an account verification code. Enter the code in the box, and then click **Submit**.

If a **You've set up two-factor account verification** message appears, click **Close**.

If a message states that the code you entered is incorrect, enter the code again, and click **Submit**. You can also click **Resend my code** to receive a new account verification code, or click **change my phone number** to enter a new phone number for receiving codes.

1.5 Navigate to a page

To navigate to a page, click the page name in the navigation bar.

The navigation bar appears at the top of every Portal page. The name of the page that you are viewing is highlighted.

Dashboard Computers Monitor Reports Policies Sites Users

1.6 View Portal text in another language

By default, Portal text appears in American English (en-US). Depending on which languages are available in your Portal instance, you might also be able to view text in UK English (en-GB), French (fr-FR), German (de-DE) or Spanish (es-ES).


When you view Portal text in UK English, times appear in 24-hour format.

If you can specify cookie preferences in your Portal instance and you rejected functional cookies, Portal text does not appear in the language selected using this procedure. Instead, Portal text appears in the language specified in your browser or in American English, if Portal text is not supported in your browser language. To view Portal text in the language selected using this procedure, click the cookie preferences link at the bottom of a Portal page and accept functional cookies.


Note: If a cookie preferences link does not appear at the bottom of each Portal page, you cannot specify cookie preferences in your Portal instance.

To view Portal text in another language:

1. Do one of the following:
 - Before signing in to Portal, click the language list at the top right of the Sign In page.



- After signing in to Portal, click the language list at the top of any Portal page.



If a list of languages appears, you can view Portal text in another language.

2. Click the language for Portal text.

1.7 View and hide Dashboard messages

Your service provider can enter messages that appear in the Dashboard. You can view these messages and, in some cases, hide messages.

Notifications of newly-added computers also appear in the Dashboard. You can start to configure computers using links in these notifications. See [Start configuring computers from Dashboard notifications](#).

To view and hide messages in the Dashboard:

1. On the navigation bar, click **Dashboard**.
2. In the **Notification Center**, click **What's New**.

Messages from your service provider and notifications (e.g., notifications of newly-added computers in your site) appear in the center of the Dashboard.

3. To hide a message that has a **Hide** link, click the link.

If you are signed in as a regular user, a confirmation dialog box appears. Click **Yes** to hide the message.

If you are signed in as an Admin user, a Hide Status dialog box appears. Do one of the following:

- To hide the message from your own message list, click **Hide From My Status Feed**.
- To hide the message from the message list for all users in your site, click **Hide For All Users**.

Note: Users can still view hidden messages using the Show Hidden Items toggle. See [Step 4](#) of this procedure.

If a **Hide** link does not appear under a message, you cannot hide the message.

4. To show all messages in the What's New list, including hidden messages, click **Yes** beside **Show hidden items**.
5. To hide hidden messages in the What's New list, click **No** beside **Show hidden items**.

1.8 View records


To make it easier to find and view information in Portal, you can sort records, filter records and set the number of records on a page.

1.8.1 Set the number of records on a page

On pages with many records, such as the Computers and Monitor pages, you can set the number of records that appear.

To set the number of records on a page:

1. On a page with many records, click the **Show <number of> records per page** list at the top of the page.



2. Click a number in the list.

The page shows the selected number of records.

1.8.2 Filter records on a page

On pages with many records, such as the Computers and Monitor pages, you can filter the records that appear. You can filter records by:

- Selecting a view. A view is a set of saved or predefined criteria that records must match in order to appear on the page. See [Filter records using a view](#).
- Entering criteria that records must match. You can also save the criteria as a view. See [Enter filter criteria and save a view](#).


1.8.2.1 Filter records using a view

On pages with many records, such as the Computers and Monitor pages, you can filter the records that appear by choosing a view. A view specifies criteria that records must match in order to appear, as well as the sort order and number of records on the page.

Note: When you choose a view, the records that appear might not be the same as when the view was saved. Because filter criteria are saved rather than a list of specific records, records that match the criteria may change.

To filter records using a view:

1. On a page of records, click the views list at the top of the page.



2. In the views list, click the view that you want to apply.
3. To view all records on the page, click **Clear Filters**.

1.8.2.2 Enter filter criteria and save a view

On pages with many records, such as the Computers and Monitor pages, you can filter records by entering criteria that the records must match. For example, on the Computers page, you can enter criteria so that computers only appear if they are online and have warnings.


If you want to reuse filter criteria, you can save the criteria as a view. The view is then available for all Admin and regular users in the site. In addition to the filter criteria, the sort order and number of records on the page are saved.

The following information is not saved in a view: the specific records that appear, the records that are selected, and the record that is expanded on the page.

Note: Because filter criteria are saved in a view rather than a list of specific records, the records that match the criteria may change. When a user chooses a view, records that appear might not be the same as when the view was saved.

To enter filter criteria and save a view:

1. On a page of records, find the filter row under the column headings.



2. In the filter row, in each column where you want to apply a filter, do one of the following:
 - In the empty box, type text that records must match.
 - In the list, click the value that records must match.

Records only appear on the page if they match all specified criteria.


3. To save the filter criteria as a view, click **Save View**. In the box that appears, enter a view name.

The new view appears in the views list.

1.8.3 Change the order of records on a page

On pages with many records, such as the Computers and Monitor pages, records appear in alphabetical order by values in one column.

As shown below, an arrow appears beside the name of the column that determines the order of the records. If the arrow points up, records appear in ascending order by values in the column. If the arrow points down, records appear in descending order.



You can change the order of records on a page.

To change the order of records on a page:

1. On a page of records, click the name of the column for determining the order of records.
2. To reverse the order of records, click the name of the column again.

1.9 Change your default page settings

You can specify which items appear by default in your Status Feed, on the Computers page and on the Monitor page.


You can also temporarily change items in the Status Feed and filter records on a page. See [Monitor recent events using the Status Feed](#) and [Filter records on a page](#).

Super users and Admin users can also change a user's default page settings. Super users do not have their own default page settings, since they do not have access to the Status Feed, Computers page or Monitor page.

To change your default page settings:


1. At the top right of the Portal page, click your email address.

A user menu appears.



2. Click **Profile Settings**.

Your profile appears. Your profile includes a Default Views section for changing page settings.



3. Do one or more of the following:

- To choose items that should appear in your Status Feed, click the **Status Feed** list. Click items in the list until a check mark appears beside each item that you want in the Status Feed, and then click outside the **Status Feed** list.
- To specify which computers should appear on the Computers page, click a view in the **Computer Page** list.

A view is a set of saved or predefined criteria that records must match in order to appear on the page. See [Filter records on a page](#).


- To specify which jobs should appear on the Monitor page, click the view in the **Monitor Page** list.

4. Click **Update default views**.

1.10 Sign out

To sign out, do one of the following:

- At the top right of the Portal page, click the Sign Out button. 
- At the top of the Portal page, click the email address or user name that you used to sign in to Portal. A user menu appears. Click **Sign Out**.



If you were signed in using single sign-on credentials, we recommend closing all browser windows after you sign out.

2 Add and manage computers in Portal

Before you can manage backups and restores for a computer, the computer must be added in Server Backup Portal using the following procedures:

- To add a Windows, Linux or UNIX computer in Portal, install or update agent software on the computer. During the installation, register the agent to Portal. For more information, see [Add a Windows computer](#), [Add a Linux computer](#), or the specific Agent guide.
- To add a Windows cluster in Portal, install the Windows Agent, Cluster Support Plug-in and any other required plug-ins on each cluster node and register each agent to Portal. You can then use Portal to configure a virtual server for the cluster. See [Add a Windows cluster](#).
- To add a VMware vSphere environment in Portal, install the vSphere Recovery Agent and register the agent to Portal. For more information, see [Add a vSphere environment](#).
- To add a Hyper-V environment in Portal, install a Hyper-V Agent Management service, specify Hyper-V information and credentials, and install one or more Hyper-V Agent Host services. See [Add a Hyper-V environment](#) and the *Hyper-V Agent User Guide*.

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

You must enter the name of a Portal user or Admin user when adding a computer. The added computer appears on the Computers page when this user or an Admin user in the user's site logs in to Portal.


To make it easier for users to find and manage computers, Admin users can assign computers to groups. See [Assign computers to groups](#).

You can view information for offline configured computers. See [View job and backup status information for offline computers](#).

2.1 Add a Windows computer

To add a Microsoft Windows computer in Portal, you must install Windows agent software on the computer and register the Agent with Portal.

The agent is installed on Windows systems where you want to back up and restore data. As shown in the following diagram, you can use Portal to manage the agent and jobs, back up data to a secure vault, and restore data from the backups.



Note: You can also use the legacy Windows CentralControl to manage the agent and jobs. However, if an agent is registered to Portal, the agent's vault settings are read-only in Windows CentralControl and you must use Portal to add and edit the vault settings.

The Windows Agent can back up:

- Files and folders on the Windows system.
- System files required for recovering the operating system, including registry and boot files.
- The entire system so that, in a disaster recovery situation, it can be restored to other hardware using System Restore.
- Files and folders on UNC shares.
- Data on Windows Storage Spaces.

Note: The Agent does not back up or restore the configuration of Windows Storage Spaces. In a disaster recovery, you can configure storage spaces manually, and then restore data to the storage spaces.

For additional functionality, you can install plug-ins with the agent. See [Windows Agent Plug-ins](#).

Beginning with Windows Agent 8.90, you can schedule a backup job to run multiple times per day, as often as hourly, when the agent backs up data to a Director version 8.60 or later vault. You can schedule a backup job to run multiple times per day by creating an intra-daily schedule in Portal 8.88 or later. See [Schedule a backup to run multiple times per day](#).

If you enable ransomware detection in an Image or Local System backup job, the agent can check for potential ransomware threats when running the job. Ransomware detection is available in Image jobs beginning with Windows Agent 9.40 and Portal 9.40, and in Local System backup jobs beginning with Windows Agent 9.00 and Portal 8.90. See [Add a Windows Image backup job](#), [Add a Windows Local System backup job](#) and [Manage potential ransomware threats](#).

Beginning with Windows Agent 9.30 and Portal 9.30, Image and Local System backups can be triggered by system events on supported Windows desktop operating systems. Backups can be triggered, or start automatically, when a user logs on to the computer or when the computer starts to shut down. See [Trigger backups when events occur on Windows desktop computers](#).

2.1.1 Windows Agent Plug-ins

When installing or upgrading a Windows Agent, you can install plug-ins with additional functionality. The following table lists and describes plug-ins that can be installed with the Windows Agent.

Plug-in	Description
Cluster Support Plug-in	Backs up and restores files and folders on shared cluster disks. The plug-in also works with the SQL Server Plug-in to protect SQL Server databases on Windows clusters, and the Image Plug-in to back up cluster volumes as images. Jobs are automatically redirected to the active node after a failover. For more information, see Add a Windows cluster .
Exchange Plug-in	Backs up and restores Exchange databases. You can also restore individual mailboxes and messages using this plug-in and the Granular Restore for Microsoft Exchange and SQL application.
Exchange Plug-in (Legacy)	Legacy plug-in. Backs up and restores Exchange 2007 databases (no longer supported).
Image Plug-in	Backs up Windows volumes as images rather than backing up individual files and folders. You can restore complete volumes and specific files and folders from Image backups. You can also restore entire systems from Image backups using System Restore. Using Image Plug-in version 7.5 or later, you can create application-consistent SQL Server database backups and restore database files. For more information, see Image Plug-in. <i>Note:</i> You must use Portal to manage Image Plug-in backups and restores. This plug-in is not supported in the legacy Windows CentralControl.
Oracle Plug-in	Backs up and restores Oracle databases.
SQL Server Plug-in	Backs up and restores SQL Server databases. The plug-in also works with the Cluster Support Plug-in to protect Microsoft SQL Server databases on Windows clusters. You can also use the SQL Server Plug-in to back up and restore SharePoint databases. You can restore individual items from a SQL Server or SharePoint database backup using the Granular Restore for Microsoft Exchange and SQL application.

2.1.1.1 Image Plug-in

To back up Windows volumes as images, install the Image Plug-in with the 64-bit Windows Agent. Unlike the Windows Agent, which enumerates and backs up individual files and folders during a backup, the Image Plug-in sequentially backs up all blocks on a volume.

Because backups with the Image Plug-in require significantly less processing than backups with the Windows Agent, the time required for a backup can be significantly reduced. To back up an entire server,

create an Image job instead of a Local System job. We recommend Image backup jobs over Local System backup jobs when backing up larger number of files on slow disks.

After the first “seed” backup of a volume, in which all data from the volume is sent to the vault, the Image Plug-in uses Changed Block Tracking to determine which blocks have changed. In subsequent Image backups, the plug-in only reads and backs up changed blocks to the vault.

When creating an Image backup job, you can select specific volumes to back up, or create a Bare Metal Restore (BMR) job that backs up all volumes, partitions, and data required for restoring a system to new hardware. You can also back up data on Windows Storage Spaces. See [Add a Windows Image backup job](#).

Note: The Image Plug-in is not supported with volumes created from Microsoft Storage Spaces Direct (S2D) storage pools.

Note: The Image Plug-in does not back up or restore the configuration of Windows Storage Spaces. In a disaster recovery, you can configure storage spaces manually, and then restore volumes to the storage spaces.

Note: The Image Plug-in is not supported on servers where Windows Offloaded Data Transfer (ODX) is enabled.

You can restore entire volumes and specific files and folders from Image backups. See [Restore Windows volumes from an Image backup](#) and [Restore files and folders from an Image backup](#). You can also use the System Restore application to restore systems from Image Plug-in BMR backups to new hardware. For more information, see the *System Restore User Guide*.

Beginning in version 7.5, the Image Plug-in can back up volumes with SQL Server database files. This option creates application-consistent database backups, so that separate SQL Server Plug-in jobs are not required. You can then mount these safesets, and restore database files from the backups. See [Add a Windows Image backup job](#) and [Restore SQL Server databases](#).

You can use the Image Plug-in only on supported 64-bit Windows operating systems with the NTFS file system. The Image Plug-in is not supported with ReFS, FAT or FAT32 file systems (except for volumes that are required to start the system). To back up a system with the ReFS, FAT or FAT32 file system, use a Windows Agent Local System job. The Image Plug-in supports both UEFI and BIOS, and MBR and GPT disks. For a complete list of supported platforms, see the Windows Agent release notes.

2.1.2 Install the Windows Agent and plug-ins

Beginning in version 9.20, the Windows Agent is only available as a 64-bit application; there is no 32-bit version of the agent. For supported platforms and system requirements, see the Windows Agent release notes.

Beginning with Windows Agent 8.90a and Portal 8.89, backups on Windows servers can be configured automatically based on job templates. If you install the Windows agent with a default encryption password and register the agent to a Portal child site where agent auto-configuration is enabled, a backup job and schedule can be created automatically. For more information, see [Determine whether an agent has been configured automatically](#).

Note: Agent auto-configuration must be enabled in the child site when the agent first registers to Portal. If you enable auto-configuration after an agent is registered to Portal, the agent will not be configured automatically. To set up agent auto-configuration, see the *Portal Administration Guide* or Server Backup online help.

You can automate the deployment of the Windows agent across your organization using Active Directory Group Policy. For more information, see the *Agent for Microsoft Windows: Automating Agent Deployment* guide.

To install the Windows Agent and plug-ins:

1. Double-click the Windows Agent installation kit.

The language selection dialog box appears.

2. In the language list, click the language for agent messages, and then click **OK**.

The installation wizard starts.

3. On the Welcome page, click **Next**.

4. On the Support Information and Release Notes page, click **Next**.

5. On the License Agreement page, read the license agreement. Click **I accept the terms in the license agreement**, and then click **Next**.

6. On the Setup Type page, do one of the following:

- To install the Agent only and use default settings, click **Typical**, and then click **Next**. Go to [Step 13](#).
- To install plug-ins and choose settings for the Agent, click **Custom**, and then click **Next**.

Note: If you want the agent to be auto-configured, and are registering the agent to a Portal child site where an Image job template is selected, click **Custom** and select the Image Plug-in in [Step 10](#).

7. On the Logon Credentials for Agent Services page, specify an account for running Agent services:

Note: The account must be in the Administrators group and have the “Log on as a service” right.

- To run Agent services using the local system account, select **Use ‘Local System’ Account**.

A local system account is required for restoring files and folders from Image backups.

A local system account cannot be used to back up UNC files and folders.

- To automatically create an account for running Agent services, select **Create account automatically**.

- To run Agent services using a custom account, select **Use custom account**. In the **Username** and **Password** boxes, enter the custom account username and password.

8. Click **Next**.

9. On the Destination Folder page, do one of the following:

- To install the Agent in the default location, click **Next**.
- To install the Agent 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**.

The Custom Setup page lists each Windows Agent component and plug-in that can be installed with the Agent that you are installing. For more information, see [Windows Agent Plug-ins](#).

The following hard disk icon appears for each component that will be installed: 

The following X icon appears for each component that will not be installed: 

Note: The “Backup Agent” is the Windows Agent and is always selected and installed.

10. On the Custom Setup page, do the following:

- For each component that you want to install, click the button to the left of the component name, and then click **This feature will be installed on local hard drive**.
- For each component that you do not want to install, click the button to the left of the component name, and then click **This feature will not be available**.

Note: If you want the agent to be auto-configured and are registering the agent to a Portal child site where an Image job template is selected, select the Image Plug-in.

11. Click **Next**.

12. On the Data Encryption Method page, do one of the following:

- For best agent performance and to encrypt data using the optimized AES 256 encryption method that is integrated in the agent, click **Encrypt data using the integrated encryption method**, and then click **Next**.
- To encrypt data using an external AES 256 encryption library that is provided with the agent, click **Encrypt data using the external encryption library**, and then click **Next**. Some organizations require the external encryption library for audit purposes.

The data encryption method is used for data at rest.

IMPORTANT: The agent is only supported with the external encryption library that is provided with the agent. It has not been tested with other encryption libraries.

Note: You cannot change the data encryption method when you modify or repair the agent. You can only change the data encryption method when you install or upgrade the agent.

13. On the Register Agent with Portal page, specify the following information:

- In the **Network Address** box, type the host name or IPV4 address of the Portal for managing the agent.

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

- In the **Port** box, type the port number for communicating with the Portal. The default port is 8086.
- In the **Username** box, type the name of the Portal user for the agent. Typically, the user name is an email address. The user must be an Admin user or regular user.

If you want the agent to be auto-configured, the user must be in a child site where agent auto-configuration is enabled.

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

14. Click **Next**.

15. On the Default Encryption Password page, do one of the following:

- If you do not want the agent to be auto-configured, select **No** and then click **Next**.
- To auto-configure the agent based on a job template, select **Yes**. In the **Password** and **Confirm Password** boxes, enter the data encryption password for the auto-configured backup job. Click **Next**.

If you install Windows agent 8.90a or later with a default encryption password and register the agent to a Portal child site where agent auto-configuration is enabled, a backup job and schedule can be created automatically. For more information, see [Determine whether an agent has been configured automatically](#).

16. On the Ready to Install the Program page, click **Install**.

The Installing Agent page appears while the Agent is being installed.

17. On the InstallShield Wizard Completed page, click **Finish**.

The Windows computer appears on the Computers page for the specified user, and for other Admin users in the user's site. If the agent is registered to a site where agent auto-configuration is enabled, wait for the agent to be configured. If the agent is not waiting for auto-configuration, you can add a backup job. See [Add the first backup job for a Windows computer](#).

2.1.2.1 Upgrade the Windows agent and plug-ins

You can upgrade a Windows agent by manually running the agent installation kit. For supported upgrade paths and system requirements, see the Windows agent release notes.

IMPORTANT: Windows Agent 8.70 and later versions can be upgraded automatically. When automatic agent upgrades are set up in Portal, the agent downloads the installer and upgrades itself automatically. Agents can only be upgraded automatically on computers where Windows Agent version 8.70 or later is installed. Windows agents must be manually upgraded to version 8.70.

Note: Agents with the Cluster Plug-in cannot be upgraded automatically. You must upgrade these agents by running the installation kit, to ensure that all nodes in a cluster have the same agent version. See [Upgrade Windows agents in a cluster](#).

Note: Support ended for SharePoint Plug-in, Exchange MAPI and SQL Server VDI jobs as of Agent version 7.50. You must delete jobs with these legacy types from a pre-version 7.50 agent before you can upgrade the agent. If you upgrade a pre-version 7.50 agent with a legacy plug-in, the plug-in will be removed. To restore from these legacy job types on the vault, install a version 7.34 or earlier agent with the appropriate plug-in, and use the *Restore from another computer* procedure.

Note: Support ended for the Agent Assistant as of Agent version 7.50. If you upgrade a pre-version 7.50 agent where the Agent Assistant is installed, the Agent Assistant will be removed from the system.

To upgrade the Windows agent and plug-ins:

1. Double-click the Windows Agent installation kit.
A message box asks if you want to continue the Agent upgrade.
2. Click **Yes**.
3. On the Data Encryption Method page, do one of the following:
 - For best agent performance and to encrypt data using the optimized AES 256 encryption method that is integrated in the agent, click **Encrypt data using the integrated encryption method**, and then click **Next**.
 - To encrypt data using an external AES 256 encryption library that is provided with the agent, click **Encrypt data using the external encryption library**, and then click **Next**. Some organizations require the external encryption library for audit purposes.

Note: The data encryption method is used for data at rest.

4. On the Portal registration page, do one of the following:
 - If the page states that the agent is already registered with Portal and you want to keep the same Portal registration information, select **Keep my current registration**, and then click **Next**.
 - If the page states that the agent is already registered with Portal and you want to change the Portal registration information, select **Change Registration**, and then click **Next**. On the Register Agent page, specify the Portal host name or IPV4 address, port number, username and password. Click **Next**.
 - If the page states that you can register the agent with Portal, specify the Portal host name or IPV4 address, port number, username and password. Click **Next**.

We recommend specifying the Portal host name when you register an agent to Portal. If the Portal IP address changes in the future, DNS can handle the change and you will not have to manually register the agent to Portal again.
5. On the Resuming the Installshield Wizard page, click **Next**.
6. On the InstallShield Wizard Completed page, click **Finish**.

2.1.2.2 Upgrade Windows agents in a cluster

The same Windows Agent version should be installed on all nodes in a Windows cluster. Use the following upgrade procedure in a Windows cluster to avoid problems with mixed Agent versions.

Note: A Windows agent with the Cluster plug-in cannot be upgraded automatically. You must upgrade agents with the Cluster plug-in by running the installation kit.

To upgrade Windows agents in a cluster:

1. Ensure that no backups or restores are running.
2. Upgrade the agent on the active node in the cluster. See [Upgrade the Windows agent and plug-ins](#).
3. Upgrade the agent on each passive node in the cluster. See [Upgrade the Windows agent and plug-ins](#).

2.1.2.3 Modify a Windows agent installation

You can modify a Windows agent installation to change the credentials for running agent services, the plug-ins that are installed, or the Portal registration.

You cannot change the data encryption method (integrated encryption method or external encryption library) when you modify an agent. You can only change the data encryption method when you uninstall or upgrade the agent.

To change the language of an agent, uninstall the agent program files, and then reinstall the agent.

To modify a Windows agent installation:

1. Double-click the Windows Agent installation kit.
2. On the Welcome page, click **Next**.
3. On the Program Maintenance page, click **Modify**, and then click **Next**.
4. On the Logon Credentials for Agent Services page, do one of the following:
 - To continue using the same credentials for running Agent services, select **Leave unchanged**.
 - To run Agent services using the local system account, select **Use 'Local System' Account**.
A local system account is required for restoring files and folders from Image backups.
A local system account cannot be used to back up UNC files and folders.
 - To automatically create an account for running Agent services, select **Create account automatically**.
 - To run Agent services using a custom account, select **Use custom account**. In the **Username** and **Password** boxes, enter the custom account username and password.

Note: The account must be in the Administrators group and have the "Log on as a service" right.
5. Click **Next**.

6. On the Custom Setup page, do the following:
 - For each component that you want to install, click the button to the left of the component name, and then click **This feature will be installed on local hard drive**.
 - For each component that you do not want to install, click the button to the left of the component name, and then click **This feature will not be available**.
7. Click **Next**.
8. On the Portal registration page, do one of the following:
 - If the page states that the Agent is already registered with Portal and you want to keep the same Portal registration information, select **Keep my current registration**, and then click **Next**.
 - If the page states that the Agent is already registered with Portal and you want to change the Portal registration information, select **Change Registration**, and then click **Next**. On the Register Agent page, specify the Portal host name or IPV4 address, port number, username and password. Click **Next**.
 - If the page states that you can register the Agent with Portal, specify the Portal host name or IPV4 address, port number, username and password. Click **Next**.
9. On the Ready to Modify the Program page, click **Install**.
10. On the InstallShield Wizard Completed page, click **Finish**.

2.1.2.4 Install or upgrade the Windows agent and plug-ins in silent mode

You can install or upgrade the Windows agent and plug-ins by running the installation in silent mode.

Note: You can download agent installation kits from some Portal instances. However, you cannot install or upgrade the Windows agent in silent mode if the installation kit was downloaded from Portal.

Beginning with Windows Agent 8.90a and Portal 8.89, backups on Windows servers can be configured automatically based on job templates. If you install the Windows agent with a default encryption password and register the agent to a Portal child site where agent auto-configuration is enabled, a backup job and schedule can be created automatically. See [Determine whether an agent has been configured automatically](#).

Note: Agent auto-configuration must be enabled in the child site when the agent first registers to Portal. If you enable auto-configuration after an agent is registered to Portal, the agent will not be configured automatically. To set up agent auto-configuration, see the *Portal Administration Guide* or Server Backup online help.

IMPORTANT: Windows Agent 8.70 and later versions can be upgraded automatically. When a new installer is available in Portal, the agent downloads the installer and upgrades itself automatically. Agents can only be upgraded automatically on computers where Agent version 8.70 or later is installed. Windows agents must be manually upgraded to version 8.70. Agents with the Cluster plug-in cannot be upgraded automatically. You must upgrade these agents manually, to ensure that all nodes in a cluster have the same agent version.

Note: Support ended for legacy Exchange MAPI, SQL Server VDI and SharePoint Plug-in jobs as of Agent version 7.50. Before upgrading an agent, you must delete jobs with these legacy types, or the upgrade will fail. If you upgrade an agent with a legacy plug-in to version 7.50 or later, the plug-in will be removed.

To install or upgrade the Windows agent and any plug-ins in silent mode, run the following command in the directory where the installation kit is located:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /v"/qn [parameters] [featureParameters]"  
[/l"language"]
```

Where:

- `Agent-Windows-x64-x-xx-xxxx.exe` is the name of the Windows Agent installation kit. `x-xx-xxxx` represents the Agent version number.
- `parameters` are optional parameters for running the installation kit in silent mode. For a list of available parameters, see [Windows agent installation parameters](#).
- `featureParameters` are optional parameters for installing plug-ins and features in silent mode. See [Windows Agent feature parameters](#).
- `/l"language"` is an optional parameter that specifies the language for the Agent. Available *language* values are:
 - 1033 – English (United States). This is the default value.
 - 1036 – French (Standard)
 - 1031 – German
 - 1034 – Spanish

For example, to install the French version of the agent, include the following parameter: `/L"1036"`

Windows agent installation parameters

Parameter	Description	Default Value
ACCOUNTTYPE	Possible values are LocalSystem, AutoCreate, and Custom.	LocalSystem
SERVICEACCOUNTNAME	If ACCOUNTTYPE is Custom, this field is required.	
SERVICEACCOUNTPASSWORD	If ACCOUNTTYPE is Custom, this field is required.	
REGISTERWITHWEBCC	Turns on/off registration of the agent with Portal.	False
AMPNWADDRESS	Host name or IPV4 address of the Portal for managing the agent. If REGISTERWITHWEBCC is True, this field is required.	
AMPPASSWORD	Password of the specified Portal user. If REGISTERWITHWEBCC is True, this field is required.	
AMPPORT	Port number for communicating with Portal.	8086

Parameter	Description	Default Value
AMPUSERNAME	Portal user for the agent. The user must be an Admin user or regular user. If REGISTERWITHWEBCC is True, this field is required.	
DEFAULTJOBENCRYPTIONKEY	Data encryption password for automatically-configured backup jobs. Beginning with Windows Agent 8.90a and Portal 8.89, if you install the Windows agent with a default encryption password and register the agent to a child site where agent auto-configuration is enabled, a backup job and schedule can be created automatically. <i>Note:</i> This parameter is applicable for new installations, but not for upgrades. <i>Note:</i> Agent auto-configuration must be enabled in the child site when the agent first registers to Portal. If you enable auto-configuration after an agent is registered to Portal, the agent will not be configured automatically.	
EXTRACTMSI	Turns on/off extraction of the Microsoft Installer (MSI) package.	False
INTEGRATEDENCRYPTION	Specifies whether to use the optimized AES 256 data encryption method that is integrated with the agent, or use the external AES 256 encryption library that is provided with the agent. Available values are: - On – the agent uses the internal, optimized AES 256 data encryption method - Off – the agent uses the external encryption library The data encryption method is used for data at rest. You cannot change the data encryption method when you modify or repair the agent. You can only change the data encryption method when you install or upgrade the agent.	On
KEEPAMPREGISTRATION	Set this property to True to retain the previous Portal registration.	True
MSIPATH	If EXTRACTMSI is True, this property denotes the location of the extracted MSI and MST files.	C:\
SILENTINSTALLDIR	Specifies an installation folder for the Agent. The installation folder must be enclosed in double quotation marks if there are spaces in the name or path.	

Windows Agent feature parameters

Feature Parameter	Description	Default Value
FEATURECLUSTER={On Off}	Turns on/off installation of the Cluster Plug-in.	Off

FEATUREEXCHANGE= {On Off}	Turns on/off installation of the Exchange Plug-in (Legacy).	Off
FEATUREEXCHANGE2010= {On Off}	Turns on/off installation of the Exchange Plug-in.	Off
FEATUREORACLE={On Off}	Turns on/off installation of the Oracle Plug-in.	Off
FEATURESQL={On Off}	Turns on/off installation of the SQL Server Plug-in.	Off
FEATUREVOLUMEIMAGE= {On Off}	Turns on/off installation of the Image Plug-in. <i>Note:</i> After the Image Plug-in is installed silently, the machine must be restarted before the Plug-in can use Changed Block Tracking (CBT) to identify data that has changed since a previous backup. Without CBT, the Agent reads all data when backing up a volume.	Off

For example, to install the Windows Agent in a different directory, run the following command:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /v"/qn SILENTINSTALLDIR=\"C:\Program Files\Acme Software\" "
```

Note: In each example shown, x-xx-xxxx represents the Agent version number.

To install the French version of the Agent, run the following command:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /v"/qn" /l"1036"
```

where 1036 indicates that the French version of the Agent is installed.

To install the Windows Agent, register the agent to Portal, specify a default data encryption password, and install the Image Plug-in, run a command similar to this:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /v"/qn REGISTERWITHWEBCC=True  
AMPNWADDRESS=123.456.com AMPUSERNAME=user@test.com AMPPASSWORD=password  
DEFAULTJOBENCRYPTIONKEY=encryptionpassword FEATUREVOLUMEIMAGE=On"
```

To install the Windows Agent and SQL Server Plug-in:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /v" FEATURESQL=On /qn"
```

2.1.3 Windows Agent default ports

The following table shows default ports that must be open for Windows Agent to communicate with other systems:

Port	Communication	Protocol
Outbound: 8086, 8087	To Portal	TCP
Outbound: 443	To Portal (for automatic agent upgrades)	TCP
Outbound: 2546	To vault	TCP
Outbound: 8031	To Windows CentralControl	TCP
Inbound: 2548	From Windows CentralControl	TCP

2.1.4 Uninstall the Windows agent and plug-ins

To uninstall the Windows agent and plug-ins:

1. Double-click the Windows Agent installation kit.
2. On the Welcome page, click **Next**.
3. On the Program Maintenance page, click **Remove**, and then click **Next**.
4. On the Uninstallation Type page, click **Total Install**, and then click **Next**.
5. On the Remove the program page, click **Remove**.
6. When the uninstallation is finished, click **Finish**.

2.1.4.1 Uninstall the Windows agent and plug-ins in silent mode

To uninstall the Windows agent and any plug-ins in silent mode and remove all of its configuration files, run the following command in the directory where the installation kit is located:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /x /v"/qn TOTALUNINSTALL=True"
```

To uninstall the Windows agent and any plug-ins in silent mode but leave its configuration files, run the following command in the directory where the installation kit is located:

```
Agent-Windows-x64-x-xx-xxxx.exe /s /x /v"/qn TOTALUNINSTALL=False"
```

x-xx-xxxx represents the Agent version number.

2.2 Add a Windows cluster

To add a Windows failover cluster in Portal, install the Windows Agent and Cluster Support Plug-in on each node in the cluster. You can also install the Image Plug-in to back up Windows volumes as images, and install the SQL Server Plug-in to back up SQL Server databases.

When installing the Windows Agent and plug-ins on each cluster node, register the agent to Portal using the same user name and password. You can then sign in to Portal using these credentials and do the following:

- Register a virtual server for the cluster core and for each cluster role (e.g., file server, SQL Server) that you want to protect.
- Add the same vault setting for each virtual server.
- Create and run backup jobs on each virtual server. When a backup job runs on a virtual server, the job is automatically directed to the active cluster node and will not reseed after a failover. You can also create backup jobs on the cluster nodes. Jobs on a cluster node will not fail over when the cluster fails over. See [Add backup jobs for a Windows cluster](#).

Note: Agents with the Cluster Plug-in cannot be upgraded automatically. You must upgrade these agents by running the installation kit, to ensure that all nodes in a cluster have the same agent version.

To add a Windows cluster:

- On each node in the Windows cluster, install the Windows Agent and the following plug-ins:
 - Cluster Support Plug-in
 - Image Plug-in (recommended)
 - SQL Server Plug-in (required for point-in-time database protection in a SQL Server cluster)

See [Install the Windows Agent and plug-ins](#).

IMPORTANT: During the installation, register each agent to the same Portal instance using the same credentials.

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

- Sign in to Portal using the credentials that you used in Step 1.
- In Portal, on the navigation bar, click **Computers**.

The Computers page shows the registered cluster nodes.

SQL1	8.00.6800	Windo...	Unconfigured	Online	vault
SQL2	8.00.6800	Windo...	Unconfigured	Online	vault

- Find the active cluster node, and expand its view by clicking its row. Click **Configure Manually**.
- Click the **Cluster Virtual Servers** tab.

The tab lists the cluster core and each cluster resource (e.g., file server, SQL Server).

SQL1	8.00.6800	Windo...	Unconfigured	Online	vault																		
Jobs Vault Settings Cluster Virtual Servers Advanced Completing the registration of any Virtual Server on this agent will require that you manually restart the 'BUAgent' service. <table> <thead> <tr> <th>Name</th><th>Domain</th><th>IP Address</th><th>Status</th><th>Drive Location</th><th>Action</th></tr> </thead> <tbody> <tr> <td>CLUSTER</td><td></td><td>192.168.2.165</td><td>Unregistered</td><td>Q</td><td>Register</td></tr> <tr> <td>SQLSERV</td><td></td><td>192.168.2.8</td><td>Unregistered</td><td>S</td><td>Register</td></tr> </tbody> </table>						Name	Domain	IP Address	Status	Drive Location	Action	CLUSTER		192.168.2.165	Unregistered	Q	Register	SQLSERV		192.168.2.8	Unregistered	S	Register
Name	Domain	IP Address	Status	Drive Location	Action																		
CLUSTER		192.168.2.165	Unregistered	Q	Register																		
SQLSERV		192.168.2.8	Unregistered	S	Register																		

- Click **Register** for the cluster core and for each role that you want to protect.

A virtual server appears on the Computers page for the registered cluster core and each registered role. Initially, each virtual server is Offline.

- On each cluster node, restart the BUAgent service.

On the Computers page in Portal, each virtual server changes to Online.

CLUSTER	8.00.6800	Windo...	Unconfigured	Online	vault
SQL1	8.00.6800	Windo...	Unconfigured	Online	vault
SQL2	8.00.6800	Windo...	Unconfigured	Online	vault
SQLSERV	8.00.6800	Windo...	Unconfigured	Online	vault

8. In Portal, for each cluster node and virtual server, add the same vault setting. Each cluster node and virtual server must be registered to the same vault using the same credentials. See [Add vault settings](#).

You can then create and run jobs on the virtual server, and the jobs will run after a failover. You can also create and run jobs on each cluster node. See [Add backup jobs for a Windows cluster](#).

2.2.1 Protect a SQL Server cluster

To protect a SQL Server cluster, you must install the Windows Agent with the Cluster Support Plug-in and SQL Server Plug-in on each node in the cluster. In Portal, you can then register a virtual server for the SQL Server role in Portal and create and run backup jobs on the virtual server. Backup jobs on a virtual server are automatically directed to the active cluster node and will not reseed after a failover.

To fully protect a SQL Server cluster, you must back up:

- the quorum disk
- each physical node in the cluster
- cluster volumes
- the SQL Server databases to provide point-in-time database recovery.


When a cluster is fully protected, you can recover the cluster if components are lost, are corrupted or fail.

For detailed information, see Windows cluster information in the Portal online help or the Windows Agent guide.

2.3 Add a Linux computer

To add a Linux computer in Portal, you must install Linux Agent software on the computer and register the Agent with Portal.

The Agent is installed on Linux systems where you want to back up and restore data. As shown in the following diagram, you can use Portal to manage the Agent and jobs, back up data to a secure vault, and restore data from the backups.



The Linux Agent can back up:

- Files and folders on a Linux system.
- System files required for recovering the operating system, including boot files.
- Files and folders that are saved on mounted NFS shares

Beginning with Linux Agent 8.90, when the agent is backing up data to a Director version 8.60 or later vault, you can schedule the backup to run multiple times per day, as often as hourly. To schedule a backup job to run multiple times per day, create an intra-daily schedule using Portal 8.88 or later. See [Schedule a backup to run multiple times per day](#).

Beginning with version 8.83 of the 64-bit Linux Agent, the Linux Agent can create Bare Metal Restore (BMR) backups for restoring entire Linux systems. The Linux Agent can create Bare Metal Restore (BMR) backups for restoring entire Linux systems. A Linux BMR backup includes an .iso file for starting the destination system and running the restore, and a backup in the vault that includes all data on the system by default. See [Restore a Linux system from a BMR backup](#).

An Oracle Plug-in, which backs up and restores Oracle databases, can be installed with the Linux Agent. There is a separate installation kit for the Oracle Plug-in for Linux.

2.3.1 Install the Linux Agent

Beginning in version 9.20, the Linux Agent is only available as a 64-bit application; there is no 32-bit version of the agent. For supported platforms and system requirements, see the Linux Agent release notes.

Beginning with version 8.83 of the 64-bit Linux Agent, the Linux Agent can create Bare Metal Restore (BMR) backups for restoring entire Linux systems. The Linux Agent can create Bare Metal Restore (BMR) backups for restoring entire Linux systems. A Linux BMR backup includes an .iso file for starting the destination system and running the restore, and a backup in the vault that includes all data on the system by default. If you want to enable Linux BMR backups, the Relax-and-Recover tool must be installed on the system. See [Install and verify Relax-and-Recover for Linux BMR backups](#). You can enable Linux BMR backups when you install the Agent, or after the Agent is installed. See [Register the Linux Agent with Relax-and-Recover and enable BMR backups](#).

Note: The Linux Agent installation process configures Relax-and-Recover for use with the Agent. If Relax-and-Recover is installed on the server for another use, you can install a second copy of the tool in a different location to avoid overwriting your settings. When installing the Linux Agent, enter the Relax-and-Recover location for the Agent to use.

The Linux Agent installation kit is provided as a tar.gz file. Only unzip this file on the machine where it will be installed. Unzipping the file on another type of machine can cause unpredictable results.

To install the Linux Agent, you must have root privileges on the target system.

The installation program will check whether there is enough disk space for the installation. If the available disk space is insufficient, the installation directory will roll back to its original state.

To install the Linux Agent:

1. If you want to enable support for BMR backups, ensure that the correct version of the Relax-and-Recover tool is installed on the Linux system. See [Install and verify Relax-and-Recover for Linux BMR backups](#).
2. Download the Linux Agent tar.gz installation package on the machine where you are installing the Agent.
3. Run the following command to extract files from the installation package:

```
tar -zxvf packageName.tar.gz
```

Where *packageName* is the name of the Agent installation kit.

4. Run the following command to change to the Agent installation kit directory:

```
cd packageName
```

5. Run the following command to start the installation:

```
./install.sh
```

For available command options, see [Install or upgrade the Linux Agent in silent mode](#).

6. Press **Enter** to read the software license agreement. If you accept the agreement, enter **Y**.

```
Do you accept the terms and conditions of the license agreement?
If yes, enter 'y' to accept the license agreement. If no, enter 'n' to cancel th
e installation: y
user accepted license agreement.

                                Installing Backup Agent

Installation directory? [/opt/BUAgent] _
```

7. At the Installation directory prompt, do one of the following:
 - To accept the default installation directory (/opt/BUAgent), press **Enter** .
 - Specify an installation directory and press **Enter**.

The directory, disk space required and available disk space are shown.

```
Directory          : /opt/BUAgent
Disk Space Required : 139 MB (estimated)
Available          : 45397 MB

Preparing for installation ...
/opt/BUAgent doesn't exist. Create it? ([Y]/n) _
```

8. Enter **Y** to create the BUAgent directory.
9. At the Select language prompt, enter the language for Agent messages. The default language is English [en-US].

```
Specify the language that should be used by default for e-mail
notifications. The Agent knows the following languages:

    de-DE    German (Germany)
    en-US    English (US)
    es-ES    Spanish (Spain)
    fr-FR    French (France)

Your default language has been detected as en_US.UTF-8 [English (US)].

Type in a supported language from the list above or press ENTER to use this
language.

Select language: [en-US] _
```

You are then prompted to choose the data encryption method.

By default, the Agent encrypts data-at-rest using an encryption method that is integrated in the Agent. For audit purposes, some organizations require the Agent to use the external encryption library that is provided with the Agent. Using the external encryption library can degrade Agent performance.

IMPORTANT: The Agent is only supported with the external encryption library that is provided with the Agent. It has not been tested with other encryption libraries.

```
By default, the Agent encrypts data using an encryption method that is integrated
in the Agent. For audit purposes, some organizations require the Agent to use an
external encryption library that is provided. Using the external encryption library
can degrade Agent performance.

Please select one of the following:
[A] Encrypt data using the Integrated encryption method. Select this encryption method
    for the best Agent performance.
[B] Encrypt data using the External encryption library. Select this encryption method
    if it is required for audit purposes.

Note: To change the encryption method that is used, you must reinstall the Agent.
Select option (A|B) (default A)
selecting A
```

10. Do one of the following:

- To use the integrated encryption method, enter **A**. This is the default value.
- To use the external encryption library that is provided with the Agent, enter **B**.

11. At the Bare Metal Restore (BMR) backup support prompt, do one of the following:

Note: BMR backup support is only available with Linux Agent version 8.83 or later.

- To enable BMR backups, enter **Y**. When prompted, enter the path to the Relax-and-Recover tool.

By default, the tool is installed in `/usr/sbin/rear`. The Relax-and-Recover tool must already be installed on the Linux server. See [Install and verify Relax-and-Recover for Linux BMR backups](#).

- If you do not want to enable BMR backups, enter **N**.

12. When prompted to register to Portal, enter **Y**.

13. At the Portal address prompt, enter the Portal host name or IPV4 address.

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

14. At the Portal connection port prompt, enter the Portal connection port. The default value is 8086.

15. At the Portal username prompt, enter the Portal username for registering the Agent.

16. At the Portal password prompt, enter the password for the Portal user specified in the previous step.

The installation proceeds. When complete, a message appears and the Agent starts.

The installation log (`Install.log`) is located in the installation directory.

2.3.1.1 Install and verify Relax-and-Recover for Linux BMR backups

Beginning with version 8.83 of the 64-bit Linux Agent, the Linux Agent can create Bare Metal Restore (BMR) backups for restoring entire Linux systems.

The Linux Agent can create Bare Metal Restore (BMR) backups for restoring entire Linux systems.

For BMR backup support with Linux Agent 8.90 or later, Relax-and-Recover (rear) version 2.7 or 2.6 must be installed on the Linux system. This section describes how to install this open-source tool and verify the installation. For more information, see the Relax-and-Recover website: <http://relax-and-recover.org/>

Note: Relax-and-Recover version 2.5 was required for BMR backups with Linux Agent 8.83. Before upgrading the Linux Agent from version 8.83 to version 8.90 or later, we recommend uninstalling Relax-and-Recover version 2.5 and performing a fresh install of Relax-and-Recover version 2.6, as described in this procedure.

If Relax-and-Recover is installed on a Linux system, you can enable BMR backups when you install the Agent on the system. See [Install the Linux Agent](#). You can also enable BMR backups by running a script after the Agent is installed. Running this script is also required if you reinstall Relax-and-Recover or change its installation path after enabling BMR backups with the Linux Agent. See [Register the Linux Agent with Relax-and-Recover and enable BMR backups](#).

When you enable Linux Agent BMR backups, Relax-and-Recover is configured for use with the Agent. If Relax-and-Recover is installed on a server for another use, you can install a second copy of the tool in a different location to avoid overwriting existing settings. When installing the Linux Agent, you can enter the Relax-and-Recover location for the Agent to use.

If you uninstall Relax-and-Recover after enabling BMR backups for a Linux Agent, non-BMR backups will continue to work.

To install and verify Relax-and-Recover for Linux BMR backups:

1. Ensure that the following Relax-and-Recover requirements are installed on the Linux system:

- bash
- mkisofs or genisoimage
- mingetty

The Relax-and-Recover website also lists nfs-utils and cifs-utils as requirements. These packages are not required for use with the Linux Agent.

Note: Some Linux distributions may have additional requirements (e.g., binutils and isolinux for Ubuntu 18.04). Any missing packages will be identified in [Step 6](#) of this procedure.

2. If Relax-and-Recover version 2.5 was installed on the system for use with Linux Agent 8.83, back up any files that were customized for Relax-and-Recover, and then uninstall Relax-and-Recover version 2.5.
3. Download and install Relax-and-Recover as described on the Relax-and-Recover website: <http://relax-and-recover.org/documentation/installation>

4. Change to the folder where Relax-and-Recover (rear) is installed (/usr/sbin/rear, by default). Check the installed rear version by running the following command:

```
rear -V
```

If a Relax-and-Recover version earlier than 2.6 is installed, go to the Relax-and-Recover downloads page and download the stable release of Relax-and-Recover version 2.7 or 2.6. Install Relax-and-Recover version 2.7 or 2.6, and verify that it is installed by running the `rear -V` command again.

5. If you backed up customized files for Relax-and-Recover in [Step 2](#) of this procedure, restore the customized files.

6. Verify that the installation was successful by running the following command:

```
rear -D -v mkrescue
```

If the installation was successful, a rescue .iso file is created in /var/lib/rear/output.

If a rescue .iso file was not created, check the log to determine whether a dependency is missing. By default, the log is located in /var/log/rear. For Relax-and-Recover support, go to <http://relax-and-recover.org/support/>.

2.3.2 Install or upgrade the Linux Agent in silent mode

To install or upgrade the Linux Agent in silent mode, run the following command in the directory where the installation kit is located:

```
install.sh [options]
```

Where *options* are optional parameters for running the installation kit in silent mode. For a list of available parameters, see [Linux Agent installation parameters](#).

Linux Agent installation parameters

Parameter	Description
-shutdown -s	Force the Agent to shut down, if running.
-force -F	Force the installation; skip the initial free space check.
-defaults -D	Use the default values for installation.
-force-defaults	Force the installation using the defaults (assumes -s and -F).
-web-registration=off -W-	Turns off Portal registration.
-web-registration= file -W=file	Attempts to register to Portal with the values found in the <i>file</i> . See Linux Agent registration options .
-quiet -Q	Quiet install; does not echo output to the screen. If user interaction is required in quiet mode, the install will fail unless -force-defaults is specified.
-log=NAME -L=NAME	Writes the installation log to the specified file <i>NAME</i> .

Parameter	Description
<code>-lang=NAME -l=NAME</code>	Selects <i>NAME</i> as the language. Must begin with an ISO language code. May optionally be followed by a dash or underscore and an ISO country code (e.g., fr, fr-FR, and fr_FR are acceptable). Character set markers (e.g., UTF-8) are ignored. Languages that cannot be matched will report an error and the language will be defaulted to en-US [English (US)]. If not specified, the language will be guessed from your system value of "en_US.UTF-8".
<code>-backup=DIR -B=DIR</code>	Backs up the current installation of the Agent to the specified directory.
<code>-verify -V</code>	Verifies the integrity of the installation kit.
<code>-enable-bmr=Y -rear-path=[path]</code>	Turns on support for Bare Metal Restore (BMR) backup jobs. <i>path</i> is the location of the Relax-and-Recover tool for the Agent to use (e.g., /user/sbin/rear) to create an iso file for restoring the system. The Relax-and-Recover tool (https://relax-and-recover.org/) must be installed on the Linux system before you install the Agent. See Install and verify Relax-and-Recover for Linux BMR backups . <i>Note:</i> When you install the Linux Agent, it configures the Relax-and-Recover tool for use with the Linux Agent. If you use the Relax-and-Recover tool for another purpose, you can avoid overwriting your Relax-and-Recover tool settings by installing a second copy of the tool in a different location.
<code>-enable-bmr=N</code>	Turns off support for Bare Metal Restore (BMR) backup jobs. <i>Note:</i> If you do not specify the <code>-enable-bmr=Y -rear-path=[path]</code> parameter, <code>-enable-bmr=N</code> is the default value.
<code>-help</code>	Shows install.sh command options.

Linux Agent registration options

For the `-web-registration=FILE` command, you can create a separate file to supply the following values as responses:

```
wccAddress=ADDRESS_OF_AMP_SERVER
wccPort=PORT_OF_AMP_SERVER # Defaults to 8086
wccLogin=PortalUserName
wccPassword=PortalPassword
```

Use the values provided by your administrator in these lines for address, port, and login name/password.

Note: This command only applies during installation. It works with the install.sh script, but not the register script.

2.3.3 Register the Linux Agent with Relax-and-Recover and enable BMR backups

Before you can enable BMR backups with the Linux Agent, the Relax-and-Recover tool must be installed on the Linux system. See [Install and verify Relax-and-Recover for Linux BMR backups](#). You can then enable Linux BMR backups when you install the Agent. See [Install the Linux Agent](#).

After the Linux Agent is installed, you can enable Linux BMR backups using this procedure. You must also follow this procedure if you reinstall or change the installation path of the Relax-and-Recover tool after enabling Linux BMR backups.

You can also disable Linux BMR backups on Agents where they are enabled. See [Disable BMR backups](#).

To register the Linux Agent with Relax-and-Recover and enable BMR backups:

1. In the Agent installation directory (/opt/BUAgent, by default), run the following command:

```
./bmrregister
```

2. At the Enable Bare Metal Restore (BMR) prompt, enter **Y**.
3. When prompted, enter the path to the Relax-and-Recover tool.

By default, the tool is installed in /usr/sbin/rear. The Relax-and-Recover tool must already be installed on the Linux server. See [Install and verify Relax-and-Recover for Linux BMR backups](#).

2.3.3.1 Disable BMR backups

You can disable Linux BMR backups on a Linux Agent where they are enabled.

To enable Linux BMR backups, see [Register the Linux Agent with Relax-and-Recover and enable BMR backups](#).

To disable Linux BMR backups:

1. In the Agent installation directory (/opt/BUAgent, by default), run the following command:

```
./bmrregister
```

2. At the Enable Bare Metal Restore (BMR) prompt, enter **N**.

2.3.4 Upgrade the Linux Agent

You can upgrade a Linux Agent by manually running the Agent installation kit. Before you upgrade the Agent, ensure that your system meets the requirements for the new Agent version as described in the Linux Agent release notes.

During the upgrade, specify the installation directory of the Linux Agent that is currently installed. Otherwise, the upgrade will proceed as if it is a new installation.

Relax-and-Recover version 2.5 was required for BMR backups with Linux Agent 8.83. Before upgrading the Linux Agent from version 8.83 to version 8.90 or later, we recommend uninstalling Relax-and-Recover version 2.5 and performing a fresh install of Relax-and-Recover version 2.7 or 2.6. See [Install and verify Relax-and-Recover for Linux BMR backups](#).

Note: When you enable Linux BMR backups, Relax-and-Recover is configured for use with the Agent. If you use Relax-and-Recover for another purpose, you can avoid overwriting your settings by installing a second copy of the tool in a different location. When upgrading the Agent, you will be prompted to enter the tool location that the Agent will use.

Note: After upgrading the Agent, we recommend running each of the Agent's backup jobs. This allows the Agent to upload new configuration information to the vault.

To upgrade the Linux Agent:

1. Download the Linux Agent tar.gz installation kit on the machine where you are installing the Agent.
2. Run the following command to extract files from the installation package:

```
tar -zxvf packageName.tar.gz
```

Where *packageName* is the name of the Agent installation kit.

3. Run the following command to change to the Agent installation kit directory:

```
cd packageName
```

4. Run the following command to start the upgrade:

```
./install.sh
```

5. Press Enter to read the software license agreement. If you accept the agreement, enter **Y**.
6. If a message states that VVAgent is running, enter **Y** to stop the Agent.
7. At the Installation directory prompt, enter the Agent installation directory. The default Agent installation directory is /opt/BUAgent.

IMPORTANT: Specify the installation directory of the Linux Agent that is currently installed. Otherwise, the upgrade will proceed as if it is a new installation.

8. At the Select language prompt, enter the language for Agent messages. The default language is English [en-US].
9. At the Select encryption option prompt, do one of the following:
 - To use the integrated encryption method, enter **A**. This is the default value.
 - To use the external encryption library that is provided with the Agent, enter **B**.

By default, the Agent encrypts data-at-rest using an encryption method that is integrated in the Agent. For audit purposes, some organizations require the Agent to use the external encryption library that is provided with the Agent. Using the external encryption library can degrade Agent performance.

IMPORTANT: The Agent is only supported with the external encryption library that is provided with the Agent. It has not been tested with other encryption libraries.

10. At the Bare Metal Restore (BMR) backup support prompt, do one of the following:
 - To enable support for BMR backups, enter **Y**. When prompted, enter the path to the Relax-and-Recover tool. The default installation directory is /usr/sbin/rear.

The Relax-and-Recover tool must already be installed on the Linux server. See [Install and verify](#)

[Relax-and-Recover for Linux BMR backups.](#)

- If you do not want support for BMR backups, enter **N**.
11. If a message states that you are already registered to a Portal, and asks whether you want to register as a new computer, do one of the following:
 - To change the Portal registration, enter **Y** and then enter the new Portal information.
 - To keep the same Portal registration, enter **N**.

The upgrade proceeds. When complete, a message appears, and the Agent starts.

2.3.5 Change the Portal registration for a Linux Agent

When you install a Linux Agent, you can register the Agent to Portal. You can also change the Portal registration at any time.

The Agent is restarted when you change the Portal registration.

To change the Portal registration for a Linux Agent:

1. In the directory where the Agent is installed, run the following command:

```
./register
```

2. If you are prompted to register as a new computer, enter **Y**.
3. At the Register to a Web-based Agent Console server prompt, enter **Y**.
4. At the Portal address prompt, enter the Portal host name or IPV4 address.

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

5. At the Portal connection port prompt, enter the Portal connection port. The default value is 8086.
6. At the Portal username prompt, enter the Portal username for registering the Agent.
7. At the Portal password prompt, enter the password for the Portal user specified in the previous step.

The Agent is restarted and the Portal registration is changed.

2.3.6 Uninstall the Linux Agent

To uninstall the Linux Agent:

1. In the directory where the Agent is installed, run the following command:

```
./uninstall.sh
```

The default Agent installation directory is `/opt/BUAgent`.

2. If a message states that VVAgent is running, enter **Y** to stop the Agent.
3. At the confirmation prompt, enter **Y**.

2.3.7 Install the Oracle Plug-in for Linux

To protect Oracle databases, you can install the Oracle Plug-in with the Linux Agent. For supported platforms and database versions, see the Oracle-Plug-in for Linux release notes.

You can determine which version of Oracle is installed by querying `BANNER` from `V$VERSION` or `VERSION` from `V$INSTANCE`:

```
SELECT banner
FROM v$version

SELECT version
FROM v$instance
```

The Oracle Plug-in can *only* find the TNS name list (`tnsnames.ora`) in the global location `/etc/oratab`. This may be a copy or symbolic link to the `tnsnames.ora` that was used to start the listener.

The Oracle Plug-in installation kit is provided as a `tar.gz` file. You must install the Oracle Plug-in on the system that has the Oracle database server. The Linux Agent must be installed before the Plug-in.

To install the Oracle Plug-in for Linux, you must have root privileges on the target system.

To install the Oracle Plug-in for Linux:

1. Download the Oracle Plug-in for Linux `tar.gz` installation package on the machine where you are installing the Plug-in.
2. Run the following command to extract files from the installation package:

```
tar -zxvf packageName.tar.gz
```

Where *packageName* is the name of the Oracle Plug-in installation kit.

3. Run the following command to change to the Oracle Plug-in installation kit directory:

```
cd packageName
```

4. Run the following command to start the installation:

```
./install.sh
```

5. Follow the installation instructions on the screen.

2.3.8 Uninstall the Oracle Plug-in for Linux

Uninstall the Oracle Plug-in as a **root** user.

To uninstall the Oracle Plug-in, run the uninstall script:

```
# ./uninstall-oracle.sh
```

This script is in the install kit directory (typically `/tmp/Oracle-Plugin-Linux<version>`).

After you run the uninstall script, use the systemctl command to stop and start the Agent.

2.4 Add a Hyper-V environment

To add a Microsoft Hyper-V environment in Portal, you must:

- Install a Hyper-V Agent Management service.
- Specify Hyper-V environment information and credentials so that the Agent can authenticate with the environment that you want to protect.

IMPORTANT: Beginning with Hyper-V Agent 8.84, you must provide Hyper-V environment information and credentials for the Management service before you can install Hyper-V Agent Host services.

- Install one or more Hyper-V Agent Host services.

You can then back up and restore virtual machines (VMs) without installing Agent software on individual VMs.

Note: Beginning in Hyper-V Agent version 9.12, the Agent can also back up and restore VMs in Microsoft Azure Stack HCI clusters.

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

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

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

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

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

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

2.4.1 Hyper-V Agent components


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

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

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

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

Components for protecting a Hyper-V environment



2.4.2 Prepare for a Hyper-V Agent deployment

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

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

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

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

2.4.2.1 Portal for managing a Hyper-V Agent

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

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

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

2.4.2.2 Vaults for Hyper-V backups

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

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

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

2.4.2.3 Recommended deployment for protecting a Hyper-V standalone host

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

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

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

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

Alternate deployment for protecting a Hyper-V standalone host

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

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

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

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

2.4.2.4 Recommended deployment for protecting a Hyper-V cluster

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

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

- The Hyper-V Agent Management service automatically distributes the backup processing load across the hosts.

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

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

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

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

Alternate deployment for protecting a Hyper-V cluster

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

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

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

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

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

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

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

2.4.2.5 Hyper-V Rapid VM Restore requirements

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

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

Component	Rapid VM Restore requirement
Hyper-V Agent	Hyper-V Agent version 9.00 or later. To perform a Rapid VM Restore in a Hyper-V cluster, the Hyper-V Agent Management server must be joined to the same domain as the cluster. Beginning in version 9.10, the Hyper-V Agent can restore a VM using Rapid VM Restore even if checkpoints were disabled on the VM when it was backed up. With Hyper-V Agent version 9.00, checkpoints had to be enabled on the protected VM when it was backed up.
Portal	Portal version 8.89 or later.
Director Vault	A Director version 8.50 or later vault. The vault must be installed locally (i.e., not on a cloud server or in a remote datacenter). The Rapid VM Restore feature must be enabled on the vault. This feature is enabled by default on Satellite vaults. If you have a local Base vault, you can enable the Rapid VM Restore feature by running a script. See Enable the Rapid VM Restore feature on a vault. See the Server Backup help or <i>Hyper-V Agent User Guide</i>.

Enable the Rapid VM Restore feature on a vault

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

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

To enable the Rapid VM Restore feature on a vault:

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

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

2.4.2.6 Hyper-V Agent ports

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

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

2.4.2.7 Best practices in a protected Hyper-V environment

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

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

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

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

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

Best practices in Hyper-V on Windows Server 2012 R2

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

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

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

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

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

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

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

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

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

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

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

2.4.3 Install the Hyper-V Agent Management service

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

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

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

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

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

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

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

To install the Hyper-V Agent Management service:

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

The installation wizard starts.

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

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

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

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

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

By default, the Management service communicates with Host services using port 5444. Ensure that this inbound port, or the custom communication port specified, is open.
 9. On the Ready to Install the Program page, click **Install**.
 10. On the Installation Completed page, click **Finish**.

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

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

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

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

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

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

The following table lists and describes command parameters.

Parameter	Description
<code>/L<localeID></code>	Optional. Specifies the language for installation log messages. Available <i>localeID</i> values are: • 1033 – English (United States). This is the default value. • 1036 – French (Standard) • 1031 – German • 1034 – Spanish

Parameter	Description
<code>\"<i>logFileName</i>"</code>	Optional. Specifies the path and name of the installation log file. If the <code>logFileName</code> includes spaces, enclose the value in double quotation marks. Example: <code>"C:\Logs\My Log.txt"</code> 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).
<code>UIREG_NETADDRESS=<i>webUIAddress</i></code>	Specifies the host name or IP address of the Portal for managing the Hyper-V Agent. Example: <code>UIREG_NETADDRESS=192.0.2.233</code> Specifying the host name is recommended. This will allow DNS to handle IP address changes.
<code>UIREG_PORT=<i>webUIportNumber</i></code>	Optional. Specifies the port number used to communicate with Portal. Example: <code>UIREG_PORT=8086</code> If you do not specify a <code>webUIportNumber</code>, port 8086 is used for communicating with Portal.
<code>UIREG_USERNAME=<i>webUIUser</i></code>	Specifies the name of the Portal user associated with the Hyper-V Agent. Example: <code>UIREG_USERNAME=user@site.com</code>
<code>UIREG_PASSWORD=<i>webUIUserPassword</i></code>	Specifies the password of the specified Portal user. Example: <code>UIREG_PASSWORD=password1234</code>
<code>COORDINATOR_PORT=<i>portNumber</i></code>	Optional. Specifies the port used to communicate with Hyper-V Agent Host services. Example: <code>COORDINATOR_PORT=5444</code> If you do not specify a port, port 5444 is used for communicating with Hyper-V Agent Host services.
<code>INSTALLDIR=<i>"installFolder"</i></code>	Optional. Specifies the installation folder for the Management service, if you do not want to install the Management service in the default location. The installation folder must be enclosed in double quotation marks if there are spaces in the path. Example: <code>INSTALLDIR="c:\Program Files\Management Service"</code> If you do not specify an installation folder, the Management service is installed in the default location.

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

```
Hyper-V_Agent_Management-x-xx-xxxx.exe /quiet /S /L1033 /V"/qn /L*v
"C:\logs\1.log\" UIREG_NETADDRESS=192.0.2.233 UIREG_USERNAME=user@site.com
```

```
UIREG_PASSWORD=password1234 UIREG_PORT=8086 INSTALLDIR=\"C:\Program  
Files\Management Service\\"
```

2.4.4 Configure a new Hyper-V Agent

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

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

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

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

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

To configure a new Hyper-V Agent:

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

The Computers page shows registered computers.

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

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

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

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

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

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

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

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

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

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

For a standalone host, the user can be a local or domain user with administrative rights.
 - In the **Password** box, type the password for the specified user.
5. To validate the credentials, click **Verify Information**. If the credentials are valid, a Success message appears. Click **Okay**.
 6. Click **Continue**.
 7. In the Vault Configuration section, click **Configure Vault**.

You can also add vault connections after the initial configuration. See [Add vault settings](#).
 8. On the Vault Settings tab, click **Add Vault**.
 9. In the Vault Settings dialog box, do one of the following:
 - In the **Vault Name** box, enter a name for the vault connection. In the **Address** box, enter the vault host name or IPV4 address. In the **Account**, **Username**, and **Password** boxes, enter an account and credentials for backing up data to and restoring data from the vault.

Specifying the host name of the vault is recommended. This will allow DNS to handle IP address changes.
 - If a policy with a vault profile is assigned to the computer, click the **Vault Profile** list. In the list, click the vault profile that you want to add for the computer. Vault information and credentials are then populated in the Vault Settings dialog box.
 10. (Optional) Change one or more of the following Advanced Settings for the vault connection:
 - **Agent Host Name**. Name of the computer on the vault. For a Hyper-V environment, by default, the name is the fully qualified domain name of the cluster or standalone host.
 - **Port Number**. Port used to connect to the vault.
 - **Attempt to Reconnect Every**. Specifies the number of seconds after which the agent should try to connect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 30 to 1800 seconds.

- **Abort Reconnect Retries After.** Enter the number of minutes after which the agent should stop trying to reconnect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 60 to 720 minutes. If the Agent cannot connect to the vault successfully in the specified time, the backup or restore fails.

11. Click **Save**.

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

2.4.5 Install the Hyper-V Agent Host service

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

Before you install a Host service, be sure that:

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

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

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

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

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

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

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

To install the Hyper-V Agent Host service:

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

The installation wizard starts.

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

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

If an error message states that the Hyper-V Agent Management service is unreachable, check that the Hyper-V environment network address and credentials have been specified for the Management service in Portal, and that the host where you are installing the Host service is associated with the specified Hyper-V environment. See [Configure a new Hyper-V Agent](#).
8. In the **Port** box, enter the port number for communicating with the Hyper-V Agent Management service.

By default, the Management service communicates with Host services using port 5444. However, a custom port might have been specified during the Management service installation.
9. Click **Next**.
10. Click **Install**.
11. On the Wizard Completed page, click **Finish**.

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

Before you install a Host service, be sure that:

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

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

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

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

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

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

The following table lists and describes command parameters.

Parameter	Description
/L<localeID>	Optional. Specifies the language for installation log messages. Available <i>localeID</i> values are: • 1033 – English (United States). This is the default value. • 1036 – French (Standard) • 1031 – German • 1034 – Spanish
\"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).
HOST=managementServiceAddress	Specifies the fully qualified domain name (FQDN) or IP address of the Hyper-V Agent Management service that assigns work to the Host service. Example: HOST=192.0.2.234 Specifying the FQDN is recommended. This will allow DNS to handle IP address changes.
PORT=portNumber	Optional. Specifies the port number for communicating with the Hyper-V Agent Management service. Example: UIREG_PORT=5444 If you do not specify a port number, port 5444 is used.
INSTALLDIR=\"installFolder\"	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: INSTALLDIR=\"c:\Program Files\Host Service\" If you do not specify an installation folder, the Host service is installed in the default location.

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

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

2.4.6 Upgrade the Hyper-V Agent

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

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

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

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

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

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

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

2.4.6.1 Upgrade the Hyper-V Agent Management service

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

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

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

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

To upgrade the Hyper-V Agent Management service:

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

Upgrade the Hyper-V Agent Management service in silent mode

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

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

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

The following table lists and describes command parameters.

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

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

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

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

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

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

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

1. Back up Hyper-V Agent logs in the <ManagementServiceInstallFolder>\Data folder.
This folder includes logs from both the Hyper-V Agent Management and Host services. Host services upload logs to the Management service after a process ends.
2. Uninstall the Management service that is installed on a Hyper-V host in a Hyper-V cluster.
3. Install the Hyper-V Agent Management service on a VM or server that has local network access to the protected Hyper-V environment. See [Prepare for a Hyper-V Agent deployment](#).
4. Recover jobs and settings from the offline Hyper-V Agent which you uninstalled in [Step 2](#). See [Recover jobs and settings from an offline Hyper-V Agent](#). Be sure to enter all required passwords, including Hyper-V environment, vault, and encryption passwords.

2.4.6.2 Upgrade the Hyper-V Agent Host service

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

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

To upgrade the Hyper-V Agent Host service:

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

Upgrade the Hyper-V Agent Host service in silent mode

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

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

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

The following table lists and describes command parameters.

Parameter	Description


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

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

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


2.5 Add a vSphere environment

To protect a vSphere environment, you can install and configure a vSphere Recovery Agent (VRA). As shown in the following diagram, a single VRA can back up virtual machines (VMs) and templates across all hosts managed by a vCenter Server.



Beginning in version 8.87, a VRA can also back up virtual machines (VMs) and templates on an ESXi host that is not managed by vCenter Server.

Note: A separate VRA is required for each ESXi host that is not managed by vCenter Server.



The VRA must be installed on a Windows physical or virtual machine with local network access to the vCenter or ESXi host that you want to protect. You can use Portal to configure and manage the VRA, back up VMs and templates to a secure vault, and restore data.

To minimize backup time and required vault space, the VRA only reads and backs up disk blocks that are being used on each VM. However, if a disk is encrypted using Bitlocker, the VRA must read all sectors of the disk. The VRA can back up VMs with encrypted disks, but the process might take longer than for non-encrypted disks.

To improve the performance of delta backups, the VRA can use Changed Block Tracking (CBT): a VMware feature that tracks changed disk sectors.

The VRA can back up and restore:

- VMs with VMDKs that are as large as 10 TB.
- VMs that reside partly or completely on vSAN storage. The VRA can back up and restore VMs on vSAN storage as long as the minimum number of nodes required for the vSAN cluster are up.
- VMs in vSAN stretched clusters.

The following options are available in vSphere backup jobs:

- Guest file system quiescing. Beginning with VRA 9.20 and Portal 9.30, you can specify whether to quiesce the file system of each VM before backing it up. For more information, see [Add a vSphere backup job](#).
- Application-consistent backups. Beginning in version 8.82, while protecting the entire file system of a Windows VM, the VRA can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on the VM. Application-consistent backups minimize the amount of work needed to restore applications from backups. For more information, see [Application-consistent backups on vSphere VMs](#).

- Ransomware threat detection. Beginning in version 9.10, the VRA can check for potential ransomware threats on VMs when running the backup job. If the VRA detects a potential threat on a VM, the VM backup is identified as a potential threat throughout Portal so you can investigate and resolve the threat. See [Manage potential ransomware threats](#).
- Backup verification. Beginning in version 9.00, the VRA can check whether each Windows VM can be restored from the backup. You can view the verification results in the Backup Verification report in Portal 9.00 or later or in Verification logs in Portal 9.30 or later. For more information, see [Backup verification for vSphere VMs](#) and [View the Backup Verification Report](#).

You can restore entire VMs using the VRA, and restore specific files, folders and database items from Windows VMs. See [Restore vSphere data](#). Beginning with VRA 8.80, you can restore a VM within minutes using the Rapid VM Restore feature. In a vCenter, you can restore a VM using Rapid VM Restore and then migrate it to another datastore to restore it permanently. On an ESXi host that is not managed by vCenter Server, you can restore a VM temporarily using Rapid VM Restore. For more information, see [Restore a vSphere VM within minutes using Rapid VM Restore](#).

2.5.1 Prepare for a vSphere Recovery Agent installation

Before installing a vSphere Recovery Agent (VRA), you must do the following:

- Obtain a Portal account for managing the agent. See [Portal for managing a vSphere Recovery Agent](#).
- Determine the destination vaults for vSphere backups. See [Vaults for vSphere Recovery Agent backups](#).
- Determine where to install the agent. See [Recommended vSphere Recovery Agent deployment](#).

You should also check requirements for VRA features that you want to use. See [Requirements for specific vSphere Recovery Agent features](#).

For best practices in a protected VMware vSphere environment, see [vSphere Recovery Agent limitations and best practices](#).

2.5.1.1 Portal for managing a vSphere Recovery Agent

You must manage a vSphere Recovery Agent using Portal. You cannot manage a vSphere Recovery Agent using the legacy Windows CentralControl interface.

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

2.5.1.2 Vaults for vSphere Recovery Agent backups

To provide fast, local vault access for backups and restores, back up vSphere data to a Satellite vault. A local vault is also required for restoring VMs within minutes using the Rapid VM Restore feature or verifying VM backups. See [vSphere Rapid VM Restore and backup verification requirements](#).

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

If you choose not to use a Satellite vault, consider using a standalone vault to seed and restore large backups.

For supported vault versions, see the vSphere Recovery Agent release notes.

2.5.1.3 Recommended vSphere Recovery Agent deployment

The vSphere Recovery Agent must be installed on a Windows physical or virtual machine that has network access to the vCenter or ESXi host that you want to protect. For best performance, install the vSphere Recovery Agent on a machine in the same subnet as the vCenter or ESXi host.

To distribute the workload, up to five vSphere Recovery Agents (VRAs) can protect VMs in a single vCenter.

In a vSAN stretched cluster, each VM has a preferred site. Ideally, have one local VRA in each site that backs up preferred VMs for that site. If a VM is moved to a different site (e.g., because of maintenance or failures), back up performance may be degraded but acceptable.

A separate VRA is required for each ESXi host that is not managed by vCenter Server. A VRA can only protect VMs on multiple ESXi hosts if the hosts are in the same vCenter.

For system requirements and supported platforms, see the vSphere Recovery Agent release notes.

We recommend using firewalls or other mechanisms to isolate VRAs and vSphere environments from the Internet.

2.5.1.4 Requirements for specific vSphere Recovery Agent features

To use specific VRA features, check the following requirements:

- To quiesce the guest file system on a VM before backing it up, see [Requirement for quiescing guest file systems](#).
- To perform application-consistent backups, see [Application-consistent backup requirements](#).
- To restore VMs within minutes or verify whether Windows VMs can be restored from backups, see [vSphere Rapid VM Restore and backup verification requirements](#).
- To check for potential ransomware threats on Windows VMs, see [Ransomware threat detection requirements](#).

Requirement for quiescing guest file systems

Beginning with VRA 9.20 and Portal 9.30, you can specify whether to quiesce the file system of each VM in a backup job before backing it up. Quiescing the file system on a VM brings the data into a consistent state that is suitable for backups.

To quiesce the guest file system on a VM, VMware Tools version 11 or later must be installed on the VM.

Application-consistent backup requirements

Beginning in version 8.82, the VRA can perform application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows VMs in vSphere environments, VMware Tools

version 11 or later must be installed on the VMs.

As part of an application-consistent backup, the VRA can truncate SQL Server, Exchange and SharePoint transaction logs on VMs on ESXi 8.0, 7.0, 6.7 and 6.5 hosts.

Application-consistent backups are supported on VMs with hardware version 8 or later.

Note: Application-consistent backups are not supported on Linux VMs.

vSphere Rapid VM Restore and backup verification requirements

Beginning with VRA 8.80 and Portal 8.84, you can restore a virtual machine (VM) to a vSphere environment within minutes using Rapid VM Restore. See [Restore a vSphere VM within minutes using Rapid VM Restore](#).

Beginning with VRA 9.00 and Portal 9.00, the VRA can verify whether Windows VMs can be restored from vSphere backups. See [Backup verification for vSphere VMs](#).

The following table lists and describes requirements for Rapid VM Restores and backup verification. If VRA and Vault requirements are not met, backup verification settings do not appear for a VRA and Rapid VM Restore does not appear as a restore option in Portal. If vSphere environment requirements are not met, you can start a Rapid VM Restore but it will not finish successfully.

Note: Because the VRA uses automated Rapid VM Restore processes to verify VM backups, these features share some requirements.

Component	Rapid VM Restore requirement	Backup verification requirement
VRA	vSphere Recovery Agent installed on a supported Windows Server platform. Windows File and Storage Services with the iSCSI Target Server feature must be installed on the server. If you install the iSCSI Target Server feature after installing VRA, you must stop and restart the VRA services (BUAgent and VVAgent) before you can perform backup verifications.	
Vault	A Director version 8.50 or later vault that is installed locally (i.e., not on a cloud server or in a remote datacenter). The Rapid VM Restore feature must be enabled on the vault. This feature is enabled by default on Satellite vaults. If you have a local Base vault, you can enable the Rapid VM Restore feature by running a script. See Enable the Rapid VM Restore feature on a vault . See the Server Backup help or <i>vSphere Recovery Agent User Guide</i> .	
vSphere environment		

Component	Rapid VM Restore requirement	Backup verification requirement
ESXi hosts	Each ESXi host must have the Software iSCSI Adapter installed and bound to a network port group that the VRA can reach. To migrate VMs restored using Rapid VM Restore to permanent storage, each ESXi host must have access to two datastores: one for writing changes while the VM runs using Rapid VM Restore, and one for permanent storage. Each datastore must have enough space for the restored VM. <i>Note:</i> On an ESXi host that is not managed by vCenter Server, Rapid VM Restore can be used to verify that VMs were backed up correctly, but cannot be used to restore VMs permanently. An ESXi server that is not part of a vCenter does not have the capabilities required to migrate VMs to permanent storage.	The ESXi host for running backup verifications must have the Software iSCSI Adapter installed and bound to a network port group that the VRA can reach. The ESXi host must be able to accommodate the expected load. During backup verification, the VRA starts each VM using an automated Rapid VM Restore process. One VM in each backup job is verified at a time and the original memory settings are used for each VM. If, for example, backup verification runs for five backup jobs at the same time and each VM uses 256 GB of RAM, backup verification could use up to 1268 GB of RAM on the host. <i>Note:</i> The ESXi host for running backup verifications is selected on the vSphere Settings tab for a VRA. See Configure a vSphere Recovery Agent.
License	To migrate VMs restored using Rapid VM Restore to permanent storage, your VMware license must support storage migration.	
Datastores	We recommend using supported storage from the VMware Hardware Compatibility Guide: https://www.vmware.com/resources/compatibility/search.php	
	When you restore a VM using Rapid VM Restore, you must choose a datastore for writing changes while the VM runs using Rapid VM Restore. This datastore can be local, iSCSI or vSAN storage, but cannot be NFS storage. When you migrate a VM to permanent storage, the destination datastore can be local, iSCSI, vSAN or NFS storage.	When you enter backup verification settings, you must choose a datastore for verifying VMs. This datastore can be local, iSCSI or vSAN storage, but cannot be NFS storage.
VM		Backup verification is supported with Windows VMs. Backup verification is not supported with non-Windows operating systems (e.g., Linux). VMware Tools version 11 or later must be installed on the VM.

Enable the Rapid VM Restore feature on a vault

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

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

To enable the Rapid VM Restore feature on a vault:

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

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

Ransomware threat detection requirements

Beginning with VRA 9.10 and Portal 9.10, the VRA can check for potential ransomware threats on Windows VMs when running a backup job. VMware Tools must be installed on the VMs. We recommend using the latest version of VMware Tools available.

The VRA can only check for ransomware threats on VMs that are running. The VRA cannot check for ransomware threats on VM templates.

2.5.1.5 vSphere Recovery Agent ports

The following table shows ports that must be open for the vSphere Recovery Agent to communicate with other systems:

Port	Communication	Protocol
Outbound: 8086, 8087	To Portal	TCP
Outbound: 2546	To vault	TCP
Outbound: 443	To vCenter	TCP
Outbound: 902	To ESXi	TCP/UDP
Inbound: 3260	iSCSI connections (for Rapid VM Restores and backup verification)	TCP

2.5.1.6 vSphere Recovery Agent limitations and best practices

The VRA can back up and restore VMs with VMDKs that are as large as 10 TB in size. Avoid using VMDKs that are larger than 10 TB.

The VRA skips physical Raw Device Mapping (pRDM), shared disks and independent disks when backing up VMs, because VMware does not allow them to be included in snapshots for VM-level backups. To back up data on these disks, you must install an Agent within the VM. During backup, the VRA skips disks with these features with a warning message. If a VM contains one or more disks that can be protected, the VM will still be backed up.

The VRA can back up and restore VMs that have volumes on Windows Storage Spaces. However, the VRA does not support file and folder restores of volumes from Windows Storage Spaces.

2.5.2 Install the vSphere Recovery Agent

The vSphere Recovery Agent (VRA) is a Windows application. You can install the VRA on a Windows physical or virtual machine that has local network access to the vCenter or ESXi host that you want to protect.

After installing VRA, you can configure vSphere environment, vault and other settings for the Agent. See [Configure a vSphere Recovery Agent](#).

To upgrade a VRA, see [Upgrade the vSphere Recovery Agent](#).

You cannot modify a VRA installation. To change the Portal registration for a VRA, you must uninstall the VRA, reinstall it with the new Portal registration, and then re-register the VRA with the vault.

Note: We recommend using firewalls or other mechanisms to isolate VRAs and vSphere environments from the Internet.

To protect a VMware vSphere environment, you must install the vSphere Recovery Agent (VRA) on a Windows physical or virtual machine that has local network access to the vCenter or ESXi host that you want to protect.

You cannot install VRA on a machine where the Windows Agent is installed.

Do not install VRA on an Active Directory domain controller.

Ensure that power management is disabled on the machine where you install VRA.

To install the vSphere Recovery Agent:

1. On a physical or virtual machine with a supported Windows platform, double-click the VRA installation kit.
2. On the Terms of Service page, read the license agreement. Click **I agree to the license terms and conditions**, and then click **Install**.
3. On the Welcome page, click **Next**.
4. On the Destination Folder page, do one of the following:
 - To install the VRA in the default location, click **Next**. Beginning in version 9.22, the default location is C:\Program Files\OpenText Server Backup\vSphere Recovery Agent.
 - To install the VRA in another location, click **Change**. In the Change 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 Agent with Portal page, specify the following information:
 - In the **Network Address** box, type the host name or IPV4 address of the Portal for managing the VRA.

Note: We recommend specifying the Portal host name when you register an agent to Portal. If the Portal IP address changes in the future, DNS can handle the change and you will not have to manually register the agent to Portal again.
 - In the **Port** box, type the port number for communicating with the Portal.

- In the **Username** box, type the name of the Portal user for managing the VRA.
After the VRA is installed, the VRA 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. When the installation has finished, click **Finish**.
 8. Click **Close**.

2.5.3 Install the vSphere Recovery Agent in silent mode

To install the vSphere Recovery Agent in silent mode, run the following command with administrative rights in the directory where the installation kit is located:

```
installKitName /install /quiet [AGENTDIR="installPath"]  
PORTAL_ADDRESS=PortalAddress [PORTAL_PORT=portNumber] PORTAL_USER=PortalUser  
PORTAL_PASSWORD=PortalPassword
```

Where *installKitName* is the name of the vSphere Recovery Agent installation kit.

The following table lists and describes command parameters:

Parameter	Description
AGENTDIR=" <i>installPath</i> "	Optional. Specifies the installation location for the Agent. If you do not include this parameter, the default installation location is used. Beginning in version 9.22, the default location is C:\Program Files\OpenText Server Backup\vSphere Recovery Agent).
PORTAL_ADDRESS= <i>PortalAddress</i>	Specifies the host name or IPV4 address of the Portal for managing the Agent. Example: PORTAL_ADDRESS=portal.site.com Specifying the host name is recommended. This will allow DNS to handle IP address changes.
PORTAL_PORT= <i>portNumber</i>	Optional. Specifies the port number for communicating with Portal. If you do not include this parameter, the default value (8086) is used.
PORTAL_USER= <i>PortalUser</i>	Specifies the name of the Portal user associated with the Agent. Example: PORTAL_USER=user@site.com
PORTAL_PASSWORD= <i>PortalPassword</i>	Specifies the password of the Portal user. Example: PORTAL_PASSWORD=password1234

2.5.4 Upgrade the vSphere Recovery Agent

You can upgrade a vSphere Recovery Agent (VRA) by manually running the Agent installation kit. For supported upgrade paths and system requirements, see the VRA release notes.

Note: When you first run an existing VRA backup job after upgrading from version 8.80 or earlier to version 8.82 or later, the backup could take longer than a normal delta backup. When the VRA first backs up a VM after an upgrade, the VRA reads all of the VM's data.

Beginning in version 8.82, the VRA can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows virtual machines (VMs). See [Application-consistent backups on vSphere VMs](#). When you upgrade a VRA from version 8.80 or earlier to version 8.82 or later, the application-consistent backup setting is not enabled in existing backup jobs. To enable application-consistency in a backup job, edit the job.

To upgrade the vSphere Recovery Agent:

1. On the machine where a previous VRA version is installed, double-click the VRA installation kit.
2. On the Terms of Service page, read the license agreement. Click **I agree to the license terms and conditions**, and then click **Install**.
3. On the confirmation page, click **Yes**.
4. On the Welcome page, click **Next**.
5. When the upgrade is complete, click **Finish**.
6. Click **Close**.

2.5.5 Upgrade the vSphere Recovery Agent in silent mode

To upgrade the vSphere Recovery Agent in silent mode, run the following command with administrative rights in the directory where the installation kit is located:

```
installKitName /install /quiet
```

2.5.6 Uninstall the vSphere Recovery Agent

Note: To change the Portal registration for a VRA, you must uninstall the VRA, reinstall it with the new Portal registration, and then re-register the VRA with the vault. See [Configure a vSphere Recovery Agent](#). You cannot modify a VRA installation.

To uninstall a vSphere Recovery Agent, do one of the following:

- Double-click the VRA installer. In the Modify Setup box, click **Uninstall**. When the VRA has been uninstalled, click **Close**.
- In the Control Panel, uninstall the vSphere Recovery Agent.

2.5.7 Uninstall the vSphere Recovery Agent in silent mode

To uninstall a vSphere Recovery Agent in silent mode, run the following command with administrative rights in the directory where the installation kit is located:

```
installKitName /uninstall /quiet
```

2.5.8 Configure a vSphere Recovery Agent

After a vSphere Recovery Agent (VRA) is installed and registered with Portal, you must configure the agent by doing the following:

- Provide information and credentials for the vCenter or ESXi host that you want to protect. The specified account should have administrative rights to the vSphere environment.
- Change the CBT setting. Changed Block Tracking (CBT) is a VMware feature that tracks changed disk sectors and improves the performance of VM backups. By default, the vSphere Agent enables Changed Block Tracking (CBT) for VMs.
- Add a vault connection. A vault connection provides vault information and credentials so that the agent can back up data to and restore data from the vault.

Beginning in version 9.00, you can also enter backup verification settings for a VRA. When backup verification settings are entered and backup verification is enabled for a vSphere backup job, the VRA verifies whether each Windows VM can be restored from the backup. See [Backup verification for vSphere VMs](#).

To change these settings after the initial configuration, see [Change vCenter or ESXi host information for a vSphere Recovery Agent](#), [Change the CBT Setting for a vSphere Recovery Agent](#), [Enter backup verification settings for a vSphere Recovery Agent](#) and [Add vault settings](#).

You can also:

- Add a description for the agent. The description appears for the vSphere 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 [Monitor backups using email notifications](#).
- Specify the amount of bandwidth consumed by backups. See [Configure bandwidth throttling](#).


To configure the vSphere Recovery Agent:

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

The Computers page shows registered computers.

2. Find the unconfigured vSphere Recovery Agent, and expand its view by clicking its row.

If the agent has not been configured, the Configure Automatically and Configure Manually boxes appear.



3. If an **Assign the computer to a site** list appears, choose a site for the agent.

The site list appears if you are signed in as an Admin user in a parent site that has child sites. The list includes the parent site if it has a vault profile, and all child sites in the parent site. If the parent site name is in the list, it appears in bold followed by the word "Parent" in brackets.

4. To add a vault connection for the agent, do one of the following:

- Choose a vault from the **Choose a vault** list, and then click **Configure Automatically**. If the vault connection is added successfully, a message appears. Click **Go to Agent**.

If the vault connection is not added successfully, you can add the vault connection manually.

- Click **Configure Manually**. On the Vault Settings tab, click **Add Vault**. In the Vault Settings dialog box, do the following:
 - In the **Vault Name** box, enter a name for the vault connection.
 - In the **Address** box, enter the vault host name or IPV4 address.

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

 - In the **Account**, **Username**, and **Password** boxes, enter an account name and credentials for backing up data to and restoring data from the vault.

Click **Save**.

5. On the vSphere Settings tab, do the following:

- In the **Host** box, type the host name or IPV4 address of the vCenter or ESXi host that you want to protect. Specifying the host name 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 vCenter or ESXi host. The domain is not required if you specify the domain in the **Username** box.
- In the **Username** box, type the account that is used to authenticate with the vCenter or ESXi host. You can type the account as *username*, *domain\username*, or *username@domain*.

The user must have administrator permissions.

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

Note: If the password for the specified user changes, change it for the VRA as soon as possible.

6. Click **Verify and Save**. If the credentials are valid, a Success message appears. Click **Okay**.
7. Do one of the following:
 - To enable CBT for VMs that do not have it enabled, select Enable Change Block Tracking (CBT) for Virtual Machines during backup.
 - To stop the VRA from enabling CBT for VMs, clear Enable Change Block Tracking (CBT) for Virtual Machines during backup.
8. To enter backup verification settings, do the following:
 - a. Select **Verify backups upon completion**.
 - b. In the **Temporary Datastore** list, select a datastore for running VMs during backup verification.
 - c. In the **Destination Host** list, select a host for running VMs during backup verification.

Note: Backup verification settings only appear if Portal and VRA requirements are met. See [vSphere Rapid VM Restore and backup verification requirements](#).

9. Click **Save**. A Success message appears. Click **Okay**.

The VRA is now ready for creating backup jobs. See [Add a vSphere backup job](#).

2.5.8.1 Change vCenter or ESXi host information for a vSphere Recovery Agent

Use the following procedure to change vCenter or ESXi host environment information for a vSphere Recovery Agent, including the host name or address and account and password for authenticating with the vSphere environment.

If you change the password for the account used to authenticate with a vSphere environment, change it as soon as possible for the VRA.

To change vCenter or ESXi host information for a vSphere Recovery Agent:

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

The Computers page shows registered computers.
2. Find the vSphere Recovery Agent, and expand its view by clicking its row.
3. On the vSphere Settings tab, do the following:
 - In the **Host** box, enter the host name or IP address of the vCenter or ESXi host that you want to protect. Specifying the host name 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 vCenter or ESXi host. The domain is not required if you specify the domain in the **Username** box.

- In the **Username** box, type the account that is used to authenticate with the vCenter or ESXi host. You can type the account as *username*, *domain\username*, or *username@domain*.
The user must have administrator permissions for the vCenter or ESXi host.
 - In the **Password** box, type the password for the specified user.
4. Click **Save**. A Success message appears. Click **Okay**.

2.5.8.2 Change the CBT Setting for a vSphere Recovery Agent

Changed Block Tracking (CBT) is a VMware feature that tracks changed disk sectors and improves the performance of VM backups. By default, the vSphere Agent enables Changed Block Tracking (CBT) for VMs.

However, because CBT requires some virtual disk processing overhead, you can stop the agent from enabling CBT for VMs. This does not disable CBT for VMs that already have it enabled through the agent or another mechanism. It only stops the agent from enabling CBT in the future for VMs that do not already have it enabled.

To change the CBT setting for a vSphere Recovery Agent:

1. In Portal, on the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the vSphere Recovery Agent, and expand its view by clicking its row.
3. On the vSphere Settings tab, do one of the following:
 - To enable CBT for VMs that do not have it enabled, select **Enable Change Block Tracking (CBT) for Virtual Machines during backup**.
 - To stop the VRA from enabling CBT for VMs, clear **Enable Change Block Tracking (CBT) for Virtual Machines during backup**.
4. Click **Save**.

2.5.8.3 Enter backup verification settings for a vSphere Recovery Agent

Beginning in version 9.00, you can enter backup verification settings for a VRA. When backup verification settings are entered and backup verification is enabled for a vSphere backup job, the VRA verifies whether each Windows VM in the job can be restored from the backup. See [Backup verification for vSphere VMs](#).

Note: Backup verification does not run for vSphere backups that are started by intra-daily schedules.

To enter backup verification settings for a vSphere Recovery Agent:

1. In Portal, on the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the vSphere Recovery Agent, and expand its view by clicking its row.
3. On the vSphere Settings tab, select **Verify backups upon completion**.

Note: Backup verification settings only appear if Portal and VRA requirements are met. See [vSphere Rapid VM Restore and backup verification requirements](#).

4. In the **Temporary Datastore** list, select a datastore for running VMs during backup verification.
5. In the **Destination Host** list, select a host for running VMs during backup verification.
6. Click **Save**. A Success message appears. Click **Okay**.

2.5.8.4 Change the Portal registration for a vSphere Recovery Agent

You cannot change the Portal registration of a VRA by running the installation kit. To change the Portal address or user information for a vSphere Recovery Agent, you must uninstall the VRA, reinstall it with the new Portal registration, and then re-register the VRA with the vault.

To change the Portal registration for a vSphere Recovery Agent:


1. On the machine where the VRA is installed, back up the log files in the folder where the agent is installed.
2. Uninstall the VRA.
3. Reinstall the VRA. When prompted to register the agent with Portal, enter the new Portal registration information. See [Install the vSphere Recovery Agent](#).

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


The Computers page shows registered computers.

5. Find the VRA that you installed, and expand its view by clicking its row.

The Configure Automatically and Configure Manually boxes appear.



6. Click **Configure Manually**.
7. On the Vault Settings tab, click **Re-register**.



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

- In the **Vault Profile** list, select the vault with backups from the original VRA. Vault information and credentials are then populated in the dialog box.
- In the **Vault Name** box, enter a name for the vault. In the **Address** box, enter the host name or IPV4 address of the vault with backups from the original VRA. 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.

9. Click **Load Computers**.

10. In the list of computers, click the name of the original VRA. Click **Save**.

11. In the Confirmation message box, click **Yes**.

12. On the vSphere Settings tab, type the username and password for authenticating with the vCenter or ESXi host.

13. Click **Save**. A Success message appears. Click **Okay**.

14. On the Jobs tab, do the following for each backup job:

- a. In the **Select Action** menu for the job, click **Edit Job**.
- b. In the Edit Job dialog box, re-enter the encryption password for the job in the **Password** and **Confirm Password** boxes.

IMPORTANT: To avoid reseeding the job, you must enter the encryption password that was used when the original VRA ran the backup job.

c. Save the job.

d. In the **Select Action** menu for the job, click **Synchronize**.

15. On the Advanced tab, if a Notifications tab appears and you can edit SMTP settings, enter and save SMTP credentials. Click **Save**.

2.6 Assign computers to groups

Admin users can assign computers to groups, making it easier for users to find and manage computers. After a group is created, users can filter records on the Computers page so that computers only appear if they belong to a specific group. See [Filter records on a page](#).


Admin users can also remove computers from groups. Removing a computer from a group does not delete the computer; it only removes its group membership.

To assign computers to groups:

1. On the navigation bar, click **Computers**. The Computers page shows registered computers.

2. Select the check box for each computer that you want to assign to groups.
3. In the **Actions** list, click **Manage Group Membership**.

The Group Membership Management dialog box appears. The check box is selected for any group to which all selected computers belong.



4. Do one or more of the following:
 - To remove the selected computers from a group, clear the selected group check box.
Note: If you remove the last computer from a group, the group is deleted and will no longer be available for adding computers.
 - To add the selected computers to a group, select the group check box.
 - To create a group and add the selected computers to the new group, enter a group name in the **Create new group** box, and then click **Create new group**. 
5. Click **Save**.


2.7 View job and backup status information for offline computers

You can view the last known jobs and backup status information for offline, configured computers. However, you cannot perform other actions or view logs for offline computers.

To view job and backup status information for an offline computer:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the offline computer, and expand its view by clicking its row.
The Jobs tab shows the last known backups jobs and backup status for the computer.

Note: You cannot expand an unconfigured offline computer.



Name	Job Type	Description	Last Backup Status	Date	Action
BMRjob	Local System		Completed	yesterday at 9:19 PM	
CloudServerBackup	Local System	This backup protects your entire C drive. It will be backed up to the cloud, per your retention schedule.	Completed	yesterday at 7:32 PM	
Job1	Local System		Completed with warnings	yesterday at 10:45 PM	
testjob	Local System		Completed	yesterday at 11:45 PM	


2.8 Undelete Hyper-V environments

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

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

To undelete a Hyper-V environment:

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



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

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

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

2.9 Move computers between sites

In some Portal instances, Admin users can move computers from a parent site to a child site, from a child site to the parent site, or between child sites. This can be useful for partners who manage computers for their customers.

Admin users can move both online and offline computers between sites. Computers that are scheduled for deletion or have been deleted (i.e., deleted Hyper-V environments) cannot be moved.

To move computers between sites:

1. On the navigation bar, click **Computers**. The Computers page shows registered computers.
2. Select the check box for each computer that you want to move to a site.
3. In the **Actions** list, click **Move Computer(s)**.
4. In the Move Computers box, select the site where you want to move the selected computers, and then click **Move**. If the parent site name is in the list, it appears in bold followed by the word "Parent" in brackets.

The computers are moved to the new site. If the new site name does not appear for the computers immediately, refresh the page.

2.10 Minimum Agent and plug-in permissions

Agent software requires sufficient permissions to back up and restore files. The following table lists the minimum permissions required for specific Agents and plug-ins.

Product	Required Permissions
Windows Agent	The account for running Agent services must: • belong to the Backup Operators group • have the "Log on as a service" right • have the "Replace a process level token" right To back up files and folders locally, the account must have read/write access to files and folders on the system. To back up files and folders in UNC locations, the account must have read/write permissions to the UNC locations, including security streams. Security streams might not work in some places unless the account is an Admin equivalent. If you are using Encrypting File System (EFS), additional permissions are required. After installing the Agent, you must change local security settings or the default domain policy. The service account must have the "Act as part of the operating system" right and the "Log on as a service" right. If the service account does not have the correct permissions, the service is denied access. ACLs for all subsequent files might not be backed up and error messages might appear in the log. <i>Note:</i> When you install, modify, repair or upgrade an Agent, the Agent installation kit sets or resets permissions on the Agent folder and all child items to full access for the Administrators and Backup Operators groups. Using the Modify option, the user can install Agent services under a local system account or another account that is created manually or automatically. For non-local system accounts, the created account is modified to be part of the Administrators group. If a user requires access to Agent services, the user should be included in the Administrators or Backup Operators group.

Exchange Plug-in	The account specified during the Windows Agent and Plug-in installation must belong to the following groups: • Exchange Organization Administrators • Group Policy Owners • Schema Admins • Enterprise Admins • Domain Admins
SQL Plug-in	In addition to permissions required for the Windows Agent, the account specified during the Agent and SQL Server Plug-in installation must have the public server role and sysadmin role in order to perform full SQL Server backups and transaction log backups.
Linux Agent	The Linux Agent requires read permissions to back up a file and write permissions to restore a file. To back up files that belong to the root account, root permissions are required.

3 Configure computers

After a computer is added in Server Backup Portal, you can configure settings for the computer. Settings include vault settings, retention types and bandwidth throttling settings.

Many of these settings can be configured using policies. A policy is a collection of settings for computers and jobs that Admin users create and assign to computers.

You can specify settings for a computer on the Computers page in Portal. You can also start configuring a computer by clicking a link in the Dashboard notification for a new computer.

Beginning with Portal 8.90 and Windows Agent 8.90a, backups on Windows servers can be configured automatically based on job templates. To be auto-configured, the Windows Agent must be installed with a default encryption password and registered to a Portal child site where agent auto-configuration is enabled. See [Install the Windows Agent and plug-ins](#).

3.1 Add vault settings

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

Note: If an agent is registered to Portal, the agent's vault settings are read-only in Windows CentralControl and you must use Portal to add and edit the vault settings.

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

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

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

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

When an E3 appliance reports a new IP address, the IP address is updated in Portal vault settings for computers that are registered to the appliance, and in the E3 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 E3 vault profile, the name is updated automatically in vault settings for computers that are registered to the appliance.


To add vault settings:

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

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

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

The Vault Settings dialog box appears.



4. Do one of the following:

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

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

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

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

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

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


6. Click **Save**.

3.2 Add a description

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

To add a description:

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



The screenshot shows a web interface with a navigation bar at the top containing 'Jobs', 'Vault Settings', and 'Advanced'. Below this, there are several tabs: 'Options', 'Retention Types', 'Notifications', 'Performance', and 'Agent Log Files'. The 'Options' tab is currently selected. Under the 'Options' tab, there is a label 'Agent Description:' followed by a large, empty text input box.

5. Click **Save**.

3.3 Add retention types

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

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

You cannot add, change or delete retention types for intra-daily schedules. For intra-daily schedules, you must choose one of two intra-daily retention types that are available beginning in Portal 8.88. See [Schedule a backup to run multiple times per day](#).

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

To add a retention type:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer for which you want to add a retention type, and click the row to expand its view.


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

3. On the Advanced tab, click the **Retention Types** tab.

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

4. Click **Create Retention Type**.

The Retention Type dialog box appears.



5. Complete the following fields:

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

Keep Archives For	<i>Note:</i> If data archiving is disabled in your Portal instance, this value does not appear. 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. 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.

6. Click **Save**.

3.4 Configure bandwidth throttling

Bandwidth throttling settings specify the amount of bandwidth consumed by an agent for backups and, in the case of most Agents, restores. 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.

Bandwidth throttling values are set at the computer (agent) level for most agents, and apply to both backups and restores. If three jobs are running at the same time on the same computer, each job gets 1/3 of the specified maximum bandwidth.

Bandwidth settings include:

- Maximum bandwidth (upper limit), in megabits per second, to be consumed by the Agent for backups and restores. For example, if three jobs are running at the same time on the same computer, each job gets 1/3 of the specified maximum bandwidth.
- Period of time during the day that throttling is in effect. Only one time window can be specified. Outside the window, no throttling takes place.
- Days of the week that throttling is in effect.

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

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

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

To configure bandwidth throttling:

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

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

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

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

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

The screenshot shows the 'Advanced' tab with the 'Performance' sub-tab selected. On the left, there is a slider for 'Execution Priority' ranging from 'Low' to 'High', with 'Normal' in the middle. On the right, under 'Days to Limit Bandwidth', there is a calendar grid with 'Mo' through 'Su'. Below this, 'Limit Bandwidth Between the Hours' is set from '12:00 AM' to '12:00 AM' (Agent time zone). The 'Max Bandwidth' is set to '151 Mbps'. A blue informational box at the bottom states: 'A Maximum Bandwidth value of 1.5 Mbps is recommended.'

4. Click **Save**.

3.5 Set the data read error handling method for a Windows computer

For Windows computers with Agent version 8.70 or later, you can specify how the Agent handles data read VSS errors during backups.


When the agent encounters a data read VSS error that could result in missing data during a backup, the agent can log the error and stop the backup or log the error and continue the backup.

To set the data read error handling method for a Windows computer:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer for which you want to set the data read error handling method, and click the row to expand its view.

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

3. On the Advanced tab, click the **Options** tab.

The screenshot shows the 'Advanced' tab selected in the top navigation bar, with sub-tabs 'Options', 'Retention Types', 'Notifications', 'Performance', and 'Agent Log Files'. The 'Options' sub-tab is active. Below the sub-tabs, there is a text input field labeled 'Agent Description:'. Below that, a section titled 'Data Read Error Handling' contains two radio button options: 'Log error and continue backup' (which is selected) and 'Log error and stop backup'.

Note: Data Read Error Handling options only appear for Windows computers with Agent version 8.70 or later.

4. Specify how the Agent handles data read VSS errors that could result in missing data during backups:
 - To log the error and continue the backup, click **Log error and continue backup**.
 - To log the error and stop the backup, click **Log error and stop backup**.
5. Click **Save**.

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

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

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

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

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

or *username@domain*.

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

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

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


3.7 Start configuring computers from Dashboard notifications

When a computer is first added in Portal, a notification appears on the Dashboard page. You can start to configure a computer by clicking a link in its notification.

To start configuring computers from Dashboard notifications:

1. On the navigation bar, click **Dashboard**.
2. In the Notification Center, click **What's New**.

Messages from your service provider and notifications of newly added computers in your site appear in the center of the Dashboard.




3. In the notification for the computer that you want to configure, click the **Configure Now** link.

If you are directed to a "queued for automatic configuration" message for the computer on the Computers page, wait for the agent to be configured. Portal will attempt to configure backups on the computer in the next three minutes. For more information, see [Add the first backup job for a Windows computer](#).

If you are directed to the computer on the Computers page, you can manually configure and add jobs for the computer or environment.

If a message states that you need to add a vault profile before you can configure backups, click the following link: **add a vault profile**.



If an Automatically Configure an Agent dialog box appears for a Windows or Linux computer, do the following:


- a. Enter an encryption password in the **Create a password** and **Confirm your password** boxes.

IMPORTANT: Your encryption password is required for restoring your data, so be sure to store it somewhere safe. If you forget the password, you will not be able to restore your data. The password is not maintained anywhere else and cannot be recovered.

Click here for agent settings.' and a 'Close' button." data-bbox="233 267 794 500"/>

- b. In the **Password hint** box, enter a hint to help you remember the encryption password.
- c. If the **Assign the computer to a site** list appears, choose a site for the computer.

The site list appears if you are signed in as an Admin user in a parent site that has child sites. The list includes the parent site if it has a vault profile and all child sites in the parent site. If the parent site name is in the list, it appears in bold followed by the word "Parent" in brackets.




- d. If more than one vault is available, choose the vault for receiving backups from the **Choose a vault** list.

If no vault is available, you must add a vault profile before you can automatically create a backup job.

- e. Click **Configure automatically**.

Backups are then configured for the computer. You can view the resulting "CloudServerBackup" backup job on the Computers page. See [View computer and job status information](#).

If an Automatically Configure an Agent dialog box appears for a vSphere environment, choose a vault from the **Choose a vault** list. If the **Assign the computer to a site** list appears, you can also choose a site for the computer. Click **Configure automatically**.



If you want to manually specify settings and create a backup job for the computer, click the following link: **Click here**

4 Add and edit backup jobs

Before you can back up data, you must create a backup job. A backup job specifies which data to back up on a system, specifies where to save the data, and provides other backup settings.

You can create backup jobs that protect:

- Windows systems, files and clusters. See [Add a Windows Local System backup job](#), [Add the first backup job for a Windows computer](#), [Add a Windows Image backup job](#), [Add a UNC file backup job](#) and [Add backup jobs for a Windows cluster](#).
- Databases and application data. See [Add a SQL Server database backup job](#), [Add an Exchange backup job](#) and [Add an Oracle database backup job](#).
- Linux and UNIX files and folders. See [Add a Linux backup job](#), [Add a UNIX backup job](#) and [Add an NFS backup job](#).
- Virtual machines. See [Add a vSphere backup job](#) and [Add and schedule a Hyper-V backup job](#).

After creating a backup job, you can run the job manually and schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#). You can also edit settings in existing backup jobs. See [Edit a backup job](#).

You can only create backup jobs for computers and environments that are online. Portal cannot communicate with computers that are offline. Computers might be offline if their Agent software is not running or has been uninstalled, if they have been shut down, or if they have no internet access.

4.1 Add a Windows Local System backup job

After a Windows computer is added in Portal, you can create a backup job for the computer. The backup job specifies which drives, folders and files to back up, and the vault for saving the data.

Note: Beginning with Portal 8.89 and Windows Agent 8.90a, backups can be configured automatically on Windows servers based on job templates. When agent auto-configuration is set up in a Portal instance, you do not have to manually create a backup job and schedule for each Windows server. For more information, see [Install the Windows Agent and plug-ins](#).

In a Windows backup job, you can select:

- Specific folders and files to back up
- The Bare Metal Restore (BMR) option, to back up volumes that are needed to boot up the system after a system recovery. In a disaster recovery situation, you can use the System Restore application to restore systems from BMR backups.

Note: You can also create BMR backup jobs using the Image Plug-in. When you run an Image Plug-in BMR job, the Plug-in backs up required volumes as images, instead of enumerating and backing up individual files and folders on the volumes. See [Add a Windows Image backup job](#).

Note: The Agent can back up encrypted volumes (e.g., encrypted by BitLocker) if they are unlocked at the time of backup. When unlocked volumes are restored, they are not encrypted and must be re-encrypted, if desired.

- The Entire Server option, available with Windows Agent 8.72 or later and Portal 8.87 or later. When this option is selected, the job backs up all non-removable volumes on the system, including volumes that are added after the job is created.

Note: To back up an entire server, create an Image job instead of a Local System job. The Image Plug-in sequentially backs up all blocks on a volume instead of backing up specific files and folders, which requires less processing than a traditional Windows backup job and can significantly reduce the backup time. See [Add a Windows Image backup job](#).

When the Entire Server option is selected in a Local System job, you do not need to select files and folders to include. However, you can exclude specific files and folders from the job. Beginning in Portal 8.89, the Bare Metal Restore option is automatically selected when you select the Entire Server option. This ensures that you can restore a protected server using the System Restore application, if required.

- The System State option, to back up files required for recovering the state of the operating system. System state backups typically include registry and boot files, the COM+ Class Registration Database, Windows system files and performance counters.

IMPORTANT: Instead of system state backups, we recommend using Image backups with the Bare Metal Restore option on platforms where the Image Plug-in is supported. See [Add a Windows Image backup job](#).

Note: Do not include antivirus product installation directories or resource folders in backup jobs.

Note: Some files are filtered out automatically from a backup job. For example, files specified by the FilesNotToBackup and FilesNotToSnapshot registry keys might not be backed up, and the job folder is not backed up.

Beginning with Windows Agent 9.00 and Portal 8.90, you can enable ransomware threat detection when you add a Local System backup job. When this option is enabled, the agent checks for potential ransomware threats when running the backup job. If the Windows agent detects a potential threat, the resulting backup is identified as a potential threat throughout Portal so you can investigate and resolve the threat. See [Manage potential ransomware threats](#).

Note: The agent does not check for potential ransomware threats in a seed backup or the first backup when threat detection is enabled in a job.

To back up the data, you can run the backup job manually and schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add a Windows Local System backup job:

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

The Computers page shows registered computers.

2. Find a Windows computer, and expand its view by clicking the computer row.

If a backup job has not been created for the computer, the system can attempt to create a backup job automatically. See [Add the first backup job for a Windows computer](#).

3. Click the **Jobs** tab.


If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. See [Add vault settings](#).

4. In the **Select Job Task** menu, click **Create Local System Job**.

5. 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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.
A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

6. To change log file settings or other backup options, click **Advanced Backup Options**. In the Advanced Backup Options dialog box, specify options and then click **Okay**. For more information, see [Log file options](#) and [Advanced backup options](#).



7. In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the drives, folders and files that you want to include and exclude in the backup job:
 - To back up all volumes on the system, including non-removable volumes that are added after the backup job is created, select the **Entire Server** option.

IMPORTANT: To back up an entire server, create an Image job instead of a Local System job. The Image Plug-in sequentially backs up all blocks on a volume instead of backing up specific

files and folders, which requires less processing than a traditional Windows backup job and can significantly reduce the backup time. See [Add a Windows Image backup job](#).

Note: The Entire Server option is only available with Windows Agent 8.72 or later and Portal 8.87 or later.

Note: Beginning in Portal 8.89, the Bare Metal Restore option is automatically selected when you select the Entire Server option.

- To back up system files so that you can restore the system to its state at the time of the backup, select **System State**, and then click **Include**.

Note: Instead of system state backups, we recommend using Image backups with the Bare Metal Restore option on platforms where the Image Plug-in is supported.

- To back up volumes that are needed to boot up the system after a system recovery, select **Bare Metal Restore**, and then click **Include**.

Bare Metal Restore (BMR) backups can be restored to new hardware using the System Restore application.

- To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).

Note: When the Entire Server option is selected in a Local System job, you do not need to select files and folders to include. All files will be backed up unless you exclude specific files and folders from the job.

- To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
- To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 

8. Do one of the following for a Local System job with Windows Agent 9.00 or later and Portal 8.90 or later:

- To check for potential ransomware threats while running the job, select the **Enable Threat Detection** check box.
- To disable ransomware threat detection, clear the **Enable Threat Detection** check box.

IMPORTANT: If you disable threat detection for a job where it was enabled, any potential threat warnings for backups in the job will be cleared. Only disable threat detection for a job once all potential threats have been addressed. See [Manage potential ransomware threats](#).

9. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.1.1 Add the first backup job for a Windows computer

Beginning with Portal 8.89 and Windows Agent 8.90a, backups can be configured automatically on Windows servers based on job templates. When agent auto-configuration is set up in a Portal site, you do not have to manually select a vault account and create a backup job and schedule for each Windows server. For more information, see [Install the Windows Agent and plug-ins](#).

In Portal sites where agent auto-configuration is not available, Portal can create the first backup job for a Windows computer when you click a "Configure Automatically" button. The resulting job cannot be customized using a job template. For a computer where the Windows Agent is installed with the Image Plug-in, Portal automatically creates an Image BMR backup job that protects all volumes on the computer. For a computer where the Windows Agent is installed without the Image Plug-in, Portal automatically creates a job that backs up the C drive. Automatically-created jobs are scheduled to run every night. A valid vault profile must be available before a backup job can be created automatically.

Note: Some files are filtered out automatically from a backup job. For example, files specified by the FilesNotToBackup and FilesNotToSnapshot registry keys might not be backed up, and the job folder is not backed up.

After a job is created, you can change the job settings, if desired. For example, you can specify different folders to back up or change the schedule for running the job.

To add the first backup job for a Windows computer:

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

The Computers page shows registered computers.

2. Find a Windows computer, and expand its view by clicking the computer row.

If a Configure Manually button appears, a backup job has not been created for the computer. Other messages or buttons might also appear, as described in the following step.

3. Do one of the following:

- If a "queued for automatic configuration" message appears for the computer, wait for the agent to be configured. Portal will attempt to configure backups on the server based on job templates in the next three minutes. See [Determine whether an agent has been configured](#)

[automatically.](#)

- To create a backup job manually, click **Configure Manually**. See [Add a Windows Local System backup job](#).
- If a Configure Automatically button appears, Portal can automatically create a backup job for the computer.

Do the following:

- In the **Create a password** and **Confirm your password** boxes, enter an encryption password.

Important: Your encryption password is required for restoring your data, so be sure to store it somewhere safe. If you forget the password, you will not be able to restore your data. The password is not maintained anywhere else and cannot be recovered.

- In the **Password hint** box, enter a hint to help you remember the encryption password.
- If the **Assign the computer to a site** list appears, choose a site for the computer.

The site list appears if you are signed in as an Admin user in a parent site that has child sites, and the computer is currently in the parent site. The list includes the parent site if it has a vault profile, and all child sites. If the parent site name is in the list, it appears in bold followed by the word "Parent" in brackets.

- If more than one vault is available, choose a vault from the **Choose a vault** list.
- Click **Configure automatically**.

If the configuration succeeds, a backup job appears for the computer.

If the automatic job creation fails, do the following:

- Click **Configure Manually**.
- On the Vault Settings tab, click **Add Vault**.

- iii. In the Vault Settings dialog box, enter vault information and credentials.
- iv. Create a backup job manually. See [Add a Windows Local System backup job](#).

4.2 Add a Windows Image backup job

After a Windows computer with the Image Plug-in is added and configured in Portal, you can create an Image backup job. In an Image backup job, the Image Plug-in sequentially backs up all blocks on a volume instead of backing up specific files and folders. Because this process can require less processing than a traditional Windows backup job, the backup time can be significantly reduced. We recommend Image backup jobs over Local System backup jobs when backing up larger number of files on slow disks.

In an Image backup job, you can select the following options:

- Specific volumes to back up.

Note: Encrypted volumes (e.g., encrypted by BitLocker) can be backed up if they are unlocked at the time of backup. When unlocked volumes are restored, they are not encrypted and must be re-encrypted, if desired.

Note: The Image Plug-in is not supported with volumes created from Microsoft Storage Spaces Direct (S2D) storage pools.

Note: The Image Plug-in is not supported on servers where Windows Offloaded Data Transfer (ODX) is enabled.

- Bare Metal Restore (BMR). This option backs up volumes that are needed to boot up the system after a system recovery. A BMR backup includes the volume where the operating system is installed, and the EFI system partition (ESP) on a UEFI-based system or the volume with the master boot record (MBR) on a BIOS-based system. In a disaster recovery situation, you can use the System Restore application to restore systems from BMR backups.

Note: BMR backup jobs can also be created using the Windows Agent without the Image Plug-in. Regardless of how a BMR backup was created, you can restore the backup using the System Restore application.

- Volumes with SQL Server database files. This option creates application-consistent SQL Server database backups, so that separate SQL Server Plug-in jobs are not required. Image Plug-in version 7.5 or later is required for this functionality.
- The Entire Server option, available with Windows Agent 8.72 or later and Portal 8.87 or later. When this option is selected, the job backs up all non-removable volumes on the system, including volumes that are added after the job is created.

Beginning in Portal 8.89, the Bare Metal Restore option is automatically selected when you select the Entire Server option. This ensures that you can restore a protected server using the System Restore application, if required.

Beginning with Windows Agent 9.40 with the Image Plug-in and Portal 9.40, you can enable ransomware threat detection in an Image backup job. When this option is enabled, the agent checks for potential

ransomware threats when running the Image backup job. If the agent detects a potential threat, the resulting backup is identified as a potential threat throughout Portal so you can investigate and resolve the threat. See [Manage potential ransomware threats](#).

After creating an Image backup job, you can run the job manually, and schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

In a seed backup, the Image Plug-in processes data for every block on a volume— even blocks that are empty. The amount of data reported in the backup log for a seed backup could be larger than the amount of data actually on the volumes.

Note: In some cases, a machine must be restarted before Changed Block Tracking (CBT) can identify data that has changed since a previous backup. For example, a machine must be restarted after the Image Plug-in is installed silently, a new volume is created, a new disk is added, or a disk is converted from basic to dynamic. Without CBT, the Agent reads all data on a volume before backing up the changed blocks.

Note: The Image Plug-in does not back up or restore the configuration of Windows Storage Spaces. In a disaster recovery, you can configure storage spaces manually, and then restore volumes to the storage spaces.

Note: Some files are filtered out automatically from a backup job. For example, files specified by the FilesNotToBackup and FilesNotToSnapshot registry keys might not be backed up, and the job folder is not backed up.

To add a Windows Image backup job:

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

The Computers page shows registered computers.

2. Find a Windows computer with the Image Plug-in, and click the computer row to expand its view.

If a backup job has not been created for the computer, the system can attempt to create a backup job automatically. See [Add the first backup job for a Windows computer](#).

3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. You must add a vault connection before you can create a job. See [Add vault settings](#).

4. In the **Select Job Task** menu, click **Create Image Job (recommended)**.

5. 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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.

- In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).

Note: For Image Plug-in jobs, the selected logging level does not affect the content of the logs.

- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

6. In the **Select Volumes for Backup** box, do one of the following until the **Backup Set** box shows the volumes that you want to back up:

- To back up all volumes on the system, including non-removable volumes that are added after the backup job is created, select the **Entire Server** option.

Note: The Entire Server option is only available with Windows Agent 8.72 or later and Portal 8.87 or later.

Note: Beginning in Portal 8.89, the Bare Metal Restore option is automatically selected when you select the Entire Server option.

- To back up specific volumes, select the check box for each volume that you want to back up, and then click **Include**.
- To back up volumes that are needed to boot up the system after a system recovery, select the **Bare Metal Restore** check box, and then click **Include**.

In addition to restoring systems from Bare Metal Restore (BMR) backups using the System Restore application, you can restore specific volumes, files, and folders from BMR backup jobs created using the Image Plug-in.

- To back up volumes with SQL Server database files, and create application-consistent SQL Server database backups, click **Application Aware Backup**, select the **SQL Volumes Protected** check box, and then click **Include**. The SQL Server Credentials dialog box appears. Enter the following credentials for connecting to SQL Server and then click **Okay**.
 - To connect to SQL Server using a Windows administrator account, select **Windows authentication**.
 - To connect to SQL Server using a SQL Server administrator account, select **SQL Server authentication**.
 - In the User Name box, type the user name for connecting to the instance.
 - In the Password box, type the password of the specified user.
 - If you selected Windows authentication, in the Domain box, type the domain of the specified account.

If you selected Windows authentication, to ensure that the Image Plug-in can protect SQL transaction logs, User Account Control (UAC) must be disabled on the system or the user must have explicit write permission to the folder where the Agent is installed and inherited permission to its subfolders.

7. Do one of the following for an Image backup job with Windows Agent 9.40 or later and Portal 9.40 or later:

- To check for potential ransomware threats while running the job, select the **Enable Threat Detection** check box.
- To disable ransomware threat detection, clear the **Enable Threat Detection** check box.

IMPORTANT: If you disable threat detection for a job where it was enabled, any potential threat warnings for backups in the job will be cleared. Only disable threat detection for a job once all potential threats have been addressed. See [Manage potential ransomware threats](#).

8. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. You can now create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.3 Add a UNC file backup job

After adding a Windows computer in Portal, you can create a backup job that protects files and folders on UNC shares. The backup job specifies which folders and files to back up and where to save the data. You must also provide credentials for accessing the UNC share.

Note: The Agent cannot back up files and folders in a DFS namespace in a UNC job. Instead, create a separate UNC job for each server share without using the DFS namespace.

To back up the data, you can run the backup job manually or schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add a UNC file backup job:

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

The Computers page shows registered computers.

2. Find a Windows computer, and expand its view by clicking the computer row.

If a backup job has not been created for the computer, the system can attempt to create a backup job automatically. However, this job only backs up local files. See [Add the first backup job for a Windows computer](#).

3. Click the **Jobs** tab.


If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. See [Add vault settings](#).

4. In the **Select Job Task** list, click **Create UNC Files Job**.

5. In the Connect to UNC Share dialog box, specify the following information:

- In the **Path to Network UNC Share** box, type the name of the UNC share where you want to back up files (e.g., \\server\share).
- In the **User Name** box, type the name of a user who has access to the UNC share.
- In the **Password** box, type the password of the specified user.

- In the **Domain** box, type the domain of the specified user account.

A screenshot of a Windows-style dialog box titled "Connect to UNC Share". The dialog has a blue header bar with a question mark icon and a close button (X). Below the header, there are four text input fields: "Path to Network UNC Share:" (with a hint "(e.g. \\192.168.1.1\SomeShare)"), "User Name:", "Password:", and "Domain:". At the bottom right, there are two buttons: "Okay" (green) and "Cancel" (grey).

6. Click **Okay**.

The system validates the UNC path and credentials. If the UNC path or credentials are not valid, a message appears. You must reenter information in the dialog box and click **Okay** again.

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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.
A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
- In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also

enter a password hint in the **Password Hint** box.

8. In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the folders and files that you want to include or exclude:
 - To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
 - To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
 - To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 
9. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.4 Add backup jobs for a Windows cluster

After the Windows Agent and required plug-ins are installed on cluster nodes and added in Portal as described in [Add a Windows cluster](#), you can add backup jobs to protect the Windows failover cluster.

To fully protect a Windows cluster, you must back up:

- the quorum disk
- each physical node in the cluster
- cluster volumes

In a SQL Server cluster, you must also back up the SQL Server databases to provide point-in-time database recovery.

When a backup job runs on a virtual server, the job is automatically directed to the active cluster node. However, if failover occurs when a backup is in progress, the backup will fail and must be run again.

To add backup jobs for a Windows cluster:

1. In Portal, add the backup jobs shown in the following table:

Job	Computer where job is created	Cluster component protected	Job description
A	Virtual server for the cluster core	Quorum disk	Image or local system job that backs up the quorum disk. See Add a Windows Image backup job or Add a Windows Local System backup job .
B (one job for each cluster node)	Each node in the cluster	Physical nodes in the cluster	On each node in the cluster, a Bare Metal Restore (BMR) backup job created using the Image Plug-in or Windows Agent. See Add a Windows Image backup job or Add a Windows Local System backup job .
C (one job for each cluster role)	Virtual server for each cluster role	Cluster disks	On the virtual server for each cluster role (e.g., file server or SQL Server), an Image or local system job that backs up cluster disks for the role. See Add a Windows Image backup job or Add a Windows Local System backup job .
D (for SQL Server clusters only)	Virtual server for the SQL Server role	SQL Server databases	SQL Server Plug-in job that backs up all SQL Server databases. See Add a SQL Server Plug-in backup job .

2. Schedule the backup jobs to run in the order shown in Step 1.

4.5 Add a SQL Server database backup job

You can protect SQL Server databases using two types of backup jobs:

- SQL Server Plug-in backup jobs. Using the SQL Server Plug-in, you can back up one or more databases in a SQL Server instance. You can also back up SharePoint databases. See [Add a SQL Server Plug-in backup job](#).
- Image Plug-in backup jobs. Using Image Plug-in version 7.5 or later, you can create application-consistent backups for databases in one or more SQL Server instances on a server. See [Add a Windows Image backup job](#).

You can back up SQL Server databases in AlwaysOn Availability Groups using either the SQL Server Plug-in or the Image Plug-in. See [Protect SQL Server databases in AlwaysOn Availability Groups](#).

4.5.1 Add a SQL Server Plug-in backup job

After a Windows computer with the SQL Server Plug-in is added and configured in Portal, you can create a backup job for one or more databases in a SQL Server instance. The backup job specifies which database or databases to back up, and where to save the backup data. A SQL Server Plug-in job cannot include databases from multiple SQL Server instances.

You can also back up a SharePoint database with a SQL Server Plug-in job.

When you create a SQL Server database backup job, you must specify Windows administrator or SQL Server administrator credentials with the SQL Server sysadmin role for the instance where you are backing up databases.

To back up the data, you can run the backup job manually or schedule the job to run. When scheduling or running a job, you can specify whether to back up the database, the transaction logs, or both. See [Run and schedule backups, synchronizations and custom commands](#).

From the backup, you can restore an entire database. You can also restore specific items from the database. See [Restore items from a SQL Server or SharePoint database](#).

To add a SQL Server database backup job:

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

The Computers page shows registered computers.


2. Find a Windows computer with the SQL Server Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. For information about adding a vault connection, see the Portal help.

4. In the **Select Job Task** menu, click **Create SQL Server Job**.

5. In the Connect to SQL Server dialog box, specify the following information:

- In the **Instance** list, select the SQL Server instance where you want to back up databases.
- To connect to the instance using a Windows administrator account, select **Windows authentication**
- To connect to the instance using a SQL Server administrator account, select **SQL Server authentication**.
- In the **User Name** box, type the user name for connecting to the instance.
- In the **Password** box, type the password of the specified user.
- If you selected Windows authentication, in the **Domain** box, type the domain of the specified account.



6. Click **Connect**.

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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.
A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
- In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also

enter a password hint in the **Password Hint** box.

8. In the **Select Databases for Backup** box, do one or more of the following to add databases to the backup job:

- To add specific databases to the backup job, select the check box for each database, and then click **Include**. The included databases appear in the **Backup Set** box.
- To back up all databases in the selected SQL Server instance, select the check box for the instance, and then click **Include**. The included instances appear in the **Backup Set** box.

Note: When the job runs, newly-added databases in the selected instance are automatically backed up.

- To back up databases with names that match a filter when the job runs, select the check box for the SQL Server instance, and then click **Include**. An inclusion record with an asterisk (*) appears in the **Backup Set** box.

In the **Database Filter** box, enter the names of databases to include. Separate multiple names with commas, and use asterisks (*) and question marks (?) as wildcard characters. For example, to back up databases with names that end with “Management” or include the word “database” followed by a single character, enter the following filter: *management, database?

Note: Filters are applied when the backup job runs. New databases that match the specified filters are automatically backed up when the job runs.

Note: Filters are not case-sensitive.

Note: If filter fields do not appear, you must upgrade the Agent on the computer to a version that supports database filtering.

9. To exclude databases from the backup job, do one or more of the following in the **Select Databases for Backup** box:

- To exclude specific databases from the backup job, select the check box for each database, and then click **Exclude**. The excluded databases appear in the **Backup Set** box.
- To exclude databases with names that match a filter when the backup job runs, select the check box for the SQL Server instance, and then click **Exclude**. A record with an asterisk (*) appears in the **Backup Set** box.

In the **Database Filter** box, enter the names of databases to exclude. Separate multiple names with commas, and use asterisks (*) and question marks (?) as wildcard characters. For example, to exclude databases if their names begin with “M”, enter the following filter: m*

Note: Filters are applied when the backup job runs. New databases that match the specified filters are automatically excluded when the backup job runs.

Note: Filters are not case-sensitive.

Note: If filter fields do not appear, you must upgrade the Agent on the computer to a version that supports database filtering.

10. To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the record. 
11. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. You can now create a schedule for running the backup. See [Run and schedule backups, synchronizations and custom commands](#).

Click **Cancel** if you do not want to create a schedule at this time.

4.5.2 Protect SQL Server databases in AlwaysOn Availability Groups

You can protect SQL Server databases in AlwaysOn Availability Groups using the Windows Agent and SQL Server Plug-in, or the Windows Agent and Image Plug-in version 7.5 or later.

If you back up a database in a secondary replica, a copy-only backup of the database is performed. Copy-only backups do not affect the sequence of conventional SQL Server backups. Microsoft only supports copy-only backups of secondary databases (see <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on-secondary-replicas-always-on-availability-groups>).

Note: If a backup job includes secondary databases and databases that are not in a secondary replica, a copy-only backup will be performed for all databases in the job. Do not include a secondary database in the same job as a standalone database.

To protect SQL Server databases in AlwaysOn Availability Groups, do one of the following:

- Install the Windows Agent and plug-in on the server where the primary replica is hosted.

If you use the SQL Server Plug-in, you can run a full backup of the primary databases, followed by full or transaction log backups. If the primary replica becomes a secondary replica after a failover, the Agent automatically runs copy-only database backups instead of full backups. Transaction log backups remain the same.

If you use Image Plug-in version 7.5 or later, you can run application-consistent image backups of the volumes with database files. If the primary replica becomes a secondary replica after a failover, the Agent automatically runs copy-only backups.

- Install the Windows Agent and plug-in on a server where a secondary replica is hosted. This backup strategy offloads backup processing to a non-primary server.

If you use the SQL Server Plug-in, you can run a copy-only backup of the secondary database, followed by copy-only or transaction log backups. If the secondary replica becomes the primary replica after a failover, the Agent automatically runs full backups instead of copy-only backups. Transaction log backups remain the same.

If you use Image Plug-in version 7.5 or later, you can run application-consistent image backups of the volumes with secondary database files. The Agent automatically runs copy-only backups of secondary database files. If the secondary replica becomes the primary replica after a failover, the Agent automatically runs full backups instead of copy-only backups.

Note: If the availability mode of the secondary replica is asynchronous-commit, transaction logs on the secondary database could lag behind the primary replica database. If the secondary database is being backed up, data loss could occur.

- Install the Windows Agent and plug-in on the primary replica server and on secondary replica servers. This strategy ensures that backups continue even if one of the replicas is down.

If you use the SQL Server Plug-in, you can run a full backup on the primary replica, followed by full or transaction log backups. You can also run copy-only backups on the secondary replicas, followed by copy-only or transaction log backups.

If you use Image Plug-in version 7.5 or later, you can run application-consistent image backups on both the primary replica server and the secondary replica server. The Agent automatically runs copy-only backups of secondary databases.

If a SQL database in an AlwaysOn Availability Group is hosted on a SQL Server Failover Cluster Instance, install the Agent, SQL Server Plug-in and Cluster Plug-in on each physical node, and configure jobs on the virtual node. Full backups will run if the database is a primary database, and copy-only backups will run if the database is a secondary database.

Note: Only GPT disks are supported for Image backups (including application-consistent Image backups of volumes with database files) in a cluster.

For information about restoring SQL Server databases in AlwaysOn Availability Groups, see [Restore databases in AlwaysOn Availability Groups](#).

4.6 Add an Exchange backup job

After a Windows computer with the Exchange Plug-in is added in Portal, you can create a backup job for one or more Microsoft Exchange databases. The backup job specifies which databases to back up, and where to save the backup data.

When running or scheduling an Exchange backup job, you can specify whether to run a Full or Incremental backup and whether to validate the Exchange data. See [Run and schedule backups, synchronizations and custom commands](#) and [Plan Full and Incremental Exchange backups](#).

After an Exchange backup job runs successfully, transaction logs for databases in the job are truncated so that the logs only contain changes that occurred after the backup.

When an Exchange server has multiple databases, you can put the databases into separate jobs and run the jobs simultaneously. Do not create parallel jobs for the same database or conflicts could prevent the jobs from completing successfully. Conflicts could also occur if you create backups using third-party applications or Agents on other Database Availability Group members.

When an Exchange backup job runs, databases in the job that are mounted or healthy are backed up. Other databases are skipped. If a database is skipped when a job runs but is mounted or healthy for the following run, the database backup does not reseed during the following run. However, if a database is skipped in two or more consecutive runs, the database backup reseeds during the next backup when the database is mounted or healthy. If no databases in a backup job are mounted or healthy when the job runs, the backup fails.

To add an Exchange backup job:

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

The Computers page shows registered computers.

2. Find a Windows computer with the Exchange Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. For more information, see the Portal online help or Windows Agent guide. See [Add vault settings](#).

4. In the **Select Job Task** menu, click **Create Exchange Server Job**.
5. 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, optionally type a description for the backup job.
 - In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.

- In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

6. In the **Include in Backup** box, do one of the following:
 - To add specific Exchange databases to the backup job, select the check box for each database, and then click **Include**.
 - To back up Exchange databases that match a filter when the job runs, select the check box for the server, and then click **Include**. An inclusion record appears in the **Backup Set** box.

In the **Filter** box, enter the names of databases to back up. Separate multiple names with commas, and use asterisks (*) and question marks (?) as wildcard characters. For example, to back up databases with names that end with “Management” or include the word “database” followed by a single character, enter the following filter: *management, database?

Filters in a backup job are applied when the job runs. New databases that match the filters are automatically backed up when the job runs.

7. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. You can now create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.7 Add an Oracle database backup job

After a Windows or Linux computer with the Oracle Plug-in is added in Portal, you can create a backup job for one or more Oracle databases. The backup job specifies which databases to back up, and where to save the backup data. You must also specify credentials for the Agent to use to connect to the Oracle server.

The Oracle Plug-in performs what Oracle Corporation deems an “inconsistent” whole database backup, requiring that the database be run in ARCHIVELOG mode. During a live backup, any changes to the database will be written to archived logs. The database administrator should ensure that the database is in ARCHIVELOG mode.

The Oracle Plug-in backs up redo and archive logs that are created while the database backup job is running. For example, if an Oracle database backup job runs from 22:00 to 01:00 each day, the plug-in backs up redo and archive logs that are created between 22:00 and 01:00. To back up logs that are created after the Oracle database backup job runs, we recommend running a Local System or Image job at another time each day. Using the Local System or Image job, you will be able to recover the database to a point in time that is later than the time when the Oracle database backup job ran.

To ensure that archived log files do not take up too much disk space on your system, the Oracle Plug-in can delete archived redo logs after a successful backup. This functionality is available with the Oracle Plug-in for Windows or Linux Agent version 8.60 or later. If you specify that archived logs should be deleted after a backup, ensure that the logs are backed up using a Local System or Image job.

To back up the data, you can run the backup job manually, or schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add an Oracle database backup job:

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


The Computers page shows registered computers.

2. Find a computer with the Oracle Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the Jobs tab.

4. In the **Select Job Task** menu, click **Create Oracle Job**.
5. In the Connect to Oracle Server dialog box, specify the following information:
 - In the **Database Service Name** box, type the service name of the database that you want to back up.
 - In the **User Name** box, type the name of a user who has sysdba privileges.

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

A screenshot of a 'Connect To Oracle Server' dialog box. The dialog has a blue title bar with a question mark icon and a close button. Inside, there are three text input fields labeled 'Database Service Name:', 'User Name:', and 'Password:'. At the bottom, there are two buttons: a green 'Connect' button and a grey 'Cancel' button.

6. Click **Connect**.
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, optionally type a description for the backup job.
 - In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
 - In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
 - For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
 - In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also

enter a password hint in the **Password Hint** box.

Create New Job

Name:

Description:

Destinations:

Log File Options:

Encryption Settings:

Encryption Password:

Confirm Password:

Password Hint:

You must remember your encryption password. Your data cannot be restored without your password.

Select Databases for Backup

☒ Database

Archived Logs

☒ Do not delete archived logs

☐ Delete archived logs older than: Days

Create job **Cancel**

8. In the **Select Databases for Backup** box, select the database to back up.
9. Do one of the following:
 - To leave Oracle archived redo logs on the system, click **Do not delete archived logs**.
 - To delete Oracle archived redo logs after a successful backup, click **Delete archived logs older than [...] days**. Enter the number of days after which archived logs can be deleted.
10. Click **Save**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.8 Add a Linux backup job

After a Linux system is added in Portal, you can create backup jobs for the system.

You can create backup jobs for files and folders that are saved locally on the computer. The backup job specifies which folders and files to back up, and where to save the data. You can also create a backup job for files and folders that are saved on mounted NFS shares. See [Add an NFS backup job](#).

Beginning with version 8.83 of the 64-bit Linux Agent, you can create Bare Metal Restore (BMR) backup jobs that can be used to restore entire Linux systems. You can also create Bare Metal Restore (BMR) backup jobs that can be used to restore entire Linux systems. A Linux BMR backup includes an .iso file for starting the

destination system and running a restore, and a backup in the vault that includes all required system volumes and files.

IMPORTANT: We recommend creating only one BMR backup job for each Linux system. If you create and run multiple BMR backup jobs, the resulting .iso file might not be usable.

Note: By default, a Linux BMR job backs up all data on the system. You can exclude folders from the BMR backup. However, if you exclude any of the following required folders, the exclusion will be ignored when the backup job runs: /bin; /boot; /etc; /lib; /lib64; /root; /usr/bin; /usr/lib; /usr/lib64; /usr/share; /usr/sbin

Note: On a server where Oracle or another database is running, we recommend that you shut down database services when running a BMR job. Alternatively, on a server where Oracle is running, you can exclude database directories from the BMR job and set up a separate Oracle Plug-in job for the database. Otherwise, database data might be inconsistent after it is restored.

A symbolic link (also called a symlink or soft link) consists of a special type of file that serves as a reference to another file or directory. During a backup, a symbolic link gets backed up with the timestamp of the link. Restoring a symbolic link sets its modification date and time to the date and time of the restore (rather than the date and time of the symbolic link when it was backed up).

To back up the data, you can run the backup job manually or schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add a Linux backup job:

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

The Computers page shows registered computers.

2. Find a Linux system, and expand its view by clicking the computer row.

If a backup job has not been created for a Linux computer, the system can attempt to create a backup job automatically. See [Add the first backup job for a Linux server](#).

3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the Jobs tab. See [Add vault settings](#).


4. In the **Select Job Task** menu, click **Create Local System Job**.

5. 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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.

- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.



- To change log file or other backup options, click **Advanced Backup Options**. In the Advanced Backup Options dialog box, select options and then click **Okay**. For more information, see [Log file options](#) and [Advanced backup options](#).



- In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the folders and files that you want to include in and exclude from the backup:
 - To create an ISO file that contains the volumes and files that are needed to boot up the system, and to back up all system data, select **Bare Metal Restore**, and then click **Include**.

By default, a Linux BMR job backs up all data on the system. You can exclude folders from the backup but, if you exclude a required folder, the exclusion will be ignored when the backup job runs and a message will appear in the log file.

When a Linux BMR job runs, it creates a boot ISO file named Bare_Metal_Restore_Image.iso in the root directory (/). The file is overwritten every time a BMR job runs. We recommend reserving a minimum of 1 GB of space in the root file system for the .iso file when you first run a BMR backup job.

Note: The **Bare Metal Restore** option is only available for Linux Agent version 8.83 or later on a 64-bit system.

- To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#)
- To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
- To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 

8. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.8.1 Add the first backup job for a Linux server

Portal can automatically create a backup job for a Linux computer that does not have a backup job. An automatically-created job backs up everything from the root, and is scheduled to run every night.

After a job is automatically created, you can change the job settings, if desired. For example, you can specify different directories to back up or change the schedule for running the job.

A valid vault profile must be available before Portal can automatically create a backup job.

After a job is created, you can change the job settings, if desired. For example, you can specify different folders to back up or change the schedule for running the job.

To add the first backup job for a Linux server:

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

The Computers page shows registered computers.

2. Find a Linux computer, and expand its view by clicking the computer row.

If a backup job has not been created for the computer, the Configure Manually box appears. If a backup job has not been created for the computer and at least one vault profile is available, the Configure Automatically box also appears.

3. Do one of the following:

- To create a backup job manually, click **Configure Manually**. See [Add a Linux backup job](#).
- To automatically create a backup job for the computer, do the following:

- a. In the **Create a password** and **Confirm your password** boxes, enter an encryption password.

IMPORTANT: Your encryption password is required for restoring your data, so be sure to store it somewhere safe. If you forget the password, you will not be able to restore your data. The password is not maintained anywhere else and cannot be recovered.

- b. In the **Password hint** box, enter a hint to help you remember the encryption password.
- c. If the **Assign the computer to a site** list appears, choose a site for the computer.

The site list appears if you are signed in as an Admin user in a parent site that has child sites, and the computer is currently in the parent site. The list includes the parent site if it has a vault profile, and all child sites. If the parent site name is in the list, it appears in bold followed by the word "Parent" in brackets.

- d. If more than one vault is available, choose a vault from the **Choose a vault** list.

- e. Click **Configure automatically**.

If the configuration succeeds, a backup job appears for the computer.

If the automatic job configuration fails, do the following:

- i. Click **Configure Manually**.
- ii. On the Vault Settings tab, click **Add Vault**.
- iii. In the Vault Settings dialog box, enter vault information and credentials.
- iv. Create a backup job manually. See [Add a Linux backup job](#).

4.9 Add a UNIX backup job

After a UNIX system is added in Portal, you can create a backup job for files and folders that are saved locally on the computer. The backup job specifies which folders and files to back up, and where to save the data. You can also create a backup job for files and folders that are saved on mounted NFS shares. See [Add an NFS backup job](#).

To back up the data, you can run the backup job manually, or schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add a UNIX backup job:

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

The Computers page shows registered computers.

2. Find a UNIX system, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab.

4. In the **Select Job Task** menu, click **Create Local System Job**.

5. 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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.

- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also

enter a password hint in the **Password Hint** box.

6. To change log file or other backup options, click **Advanced Backup Options**. In the Advanced Backup Options dialog box, select options and then click **Okay**. For more information, see [Log file options](#) and [Advanced backup options](#).

7. In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the folders and files that you want to include in and exclude from the backup:
- To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).

- To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
- To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 

8. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.10 Add an NFS backup job

After a Linux or UNIX system is added in Portal, you can create a backup job for files and folders that are saved on mounted NFS shares. The backup job specifies which folders and files to back up, and where to save the data.

NFS servers must share their exports in order to make them available to client systems. If you want to perform a mount-point backup or restore, the NFS server must be available, and it must provide sufficient privileges to your client system. Also, the NFS must be mounted on your client system at the time of the backup or restore.

Note: If you restore an NFS backup, and the NFS mount does not exist, the restore will proceed as if it were a local restore. It will put the data on the local disk (with a similar path that is local) without using a mount-point (NFS) path. It will not indicate a "failure".

NFS does not export extended attributes from remote file systems. On Linux NFSv3 clients, remote file system ACLs will be presented as standard Linux ACLs if possible. NFSv4 clients will present remote file system ACLs as native NFSv4 ACLs, but the Agent will protect them as extended attributes.

To back up the data, you can run the backup job manually, or schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add an NFS backup job:

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

The Computers page shows registered computers.

2. Find a Linux or UNIX system, and expand its view by clicking the computer row.

In some Portal instances, if a backup job has not been created for a Linux computer, the system can attempt to create a backup job automatically.

3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab.

4. In the **Select Job Task** menu, click **Create NFS Files Job**.
5. 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, optionally type a description for the backup job.
 - In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
 - In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
 - For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
 - In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.
6. In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the folders and files that you want to include in and exclude from the backup:
 - To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
 - To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
 - To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 
7. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups, synchronizations and custom commands](#).

4.11 Add a vSphere backup job

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

You must add vault settings and vSphere environment information before you can add a backup job. See [Configure a vSphere Recovery Agent](#).

You can also enable or disable the following options in a vSphere backup job:

- Guest file system quiescing. Beginning with VRA 9.20 and Portal 9.30, you can specify whether to quiesce the file system of each VM before backing it up. Quiescing the file system on a VM brings the data into a consistent state that is suitable for backups.

Trying to quiesce a guest file system that cannot be quiesced can take significant time and resources and cause the VM to become unresponsive. When backing up VMs that cannot be quiesced, turning off guest file system quiescing can save backup time and system resources.

- Application-consistent backups. Beginning in version 8.82, while protecting the entire file system of a Windows VM, the VRA can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on the VM. Application-consistent backups minimize the amount of work needed to restore applications from backups. You can also specify whether application transaction logs should be truncated during application-consistent backups. For more information, see [Application-consistent backups on vSphere VMs](#).

If you do not enable guest file system quiescing or application-consistency in a backup job, the backups are crash-consistent. A crash-consistent backup includes data on disk at the time of the backup and does not include data that is still in memory.

Note: Beginning with Portal 9.30 and VRA 9.20, the application-consistent option can only be enabled in a backup job if the guest file system quiescing option is enabled.

- Ransomware threat detection. Beginning in version 9.10, the VRA can check for potential ransomware threats on Windows VMs when running the backup job. If the VRA detects a potential threat on a VM, the VM backup is identified as a potential threat throughout Portal so you can investigate and resolve the threat. See [Manage potential ransomware threats](#).

Note: The VRA does not check for potential ransomware threats in a seed backup or the first backup when threat detection is enabled in a job.

- Backup verification. Beginning in version 9.00, the VRA can back up VMs in the job and then check whether each Windows VM can be restored from the backup. See [Backup verification for vSphere VMs](#). Backup verification settings must also be entered for the VRA. See [Enter backup verification settings for a vSphere Recovery Agent](#).

Note: Backup verification does not run for vSphere backups that are started by intra-daily schedules.

For requirements for these vSphere backup options, see [Requirements for specific vSphere Recovery Agent features](#).

To back up the data, you can run the backup job manually or schedule the backup job to run. See [Run and schedule backups, synchronizations and custom commands](#).

To add a vSphere backup job:

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


The Computers page shows registered computers and environments.

2. Click the vSphere environment row. 

If a message states that the Agent needs to be configured, you must add vault settings and vSphere environment information before adding a backup job. See [Configure a vSphere Recovery Agent](#).

If the vSphere environment does not have vault settings, the Configure Manually box appears. To add vault settings manually, click **Configure Manually**, and add a vault on the Vault Settings tab. See [Add vault settings](#).

If the vSphere environment does not have vault settings and at least one vault profile is available, the Configure Automatically box appears. To add vault settings, choose a vault from the **Choose a vault** list. If the **Assign the computer to a site** list appears, you can also choose a child site for the computer. Click **Configure Automatically**.



3. Click the **Jobs** tab.
4. In the **Select Job Task** menu, click **Create VMware vSphere Job**.

If the Connect to vSphere dialog box appears, specify the following information in the dialog box:

- In the **User Name** box, type the Windows domain account user name used to authenticate the VRA with the vCenter or ESXi host.
- In the **Password** box, type the password for the specified user.
- In the **Domain** box, type the domain of the specified user account. The domain is optional if you specified the domain in the **User Name** box (e.g., *domain\username*).

Note: vSphere environment settings entered in this dialog box are populated on the Agent's **vSphere Settings** tab.

5. 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, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it is assigned to the user, or if the user added it on the computer's Vault Settings tab.

- In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

6. In the Include in Backup box, do one or more of the following until the **Backup Set** box shows the VMs that you want to include and exclude in the backup job:

- To add specific VMs to the backup job, select the check box for each VM, and then click **Include**.
- To exclude specific VMs from the backup job, select the check box for each VM, and then click **Exclude**.
- To add VMs to the backup job by name, select the **Virtual Machines** check box, and then click **Include**. In the **Filter** field, enter names of VMs to include. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to include VMs in a backup if their names end with "x64" or start with "SQL", enter the following filter: *x64, SQL*

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To exclude VMs from the backup job by name, select the **Virtual Machines** check box, and then click **Exclude**. In the **Filter** field, enter names of VMs to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to exclude VMs from a backup if their names start with "test" or end with "x32", enter the following filter: test*, *x32
- To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the record. 

7. Specify whether you want the VRA to quiesce the file system of each VM before backing it up by doing one of the following:

- To quiesce the guest file system before backing up a VM, select the **Quiesce guest file system** check box.

- To back up each VM without trying to quiesce the guest file system, clear the **Quiesce guest file system** check box.
8. To perform application-consistent backups of SQL Server, Exchange, SharePoint, and Active Directory on Windows VMs in the backup job, while protecting the entire file system of each VM, do the following:

- a. Select the **Enable Application Consistent Backups** check box.

Note: You can only select this check box if the **Quiesce guest file system** check box is selected.


- b. Do one of the following:

- To preserve application transaction logs on VMs in the job, clear the **Truncate Database Transaction Logs** check box.
- To truncate application transaction logs on VMs in the job, select the **Truncate Database Transaction Logs** check box and enter credentials for connecting to VMs in the job.

To enter credentials for multiple VMs in the job, enter a username and password in the **Guest VM Credentials** area.

To enter credentials for a specific VM in the job, click the arrow at the right side of the VM name in the Backup Set area, and enter a username and password in the **Guest VM Credentials** area for the VM.

You can enter a username as *username* or *domain\username*. The specified user or users must have admin access to VMs in the backup job, but do not need admin rights for applications on the VMs.



Note: If you enter credentials for a specific VM in the job, the VRA will not attempt to connect to the VM using the credentials specified for multiple VMs in the job.

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

- To perform an application-consistent backup of a domain controller with Active Directory, enable the **Truncate Database Transaction logs** option, and enter domain admin credentials for the VM using the *domain\username* format.

Note: There are no logs to truncate when performing application-consistent backups of domain controllers with Active Directory. However, credentials with domain admin privileges are required for application-consistent backups of domain controllers. If the log truncation option is enabled, you can enter the required credentials.

9. Specify whether you want the VRA to check for potential ransomware threats by doing one of the following:

- To back up VMs without checking for potential ransomware threats, clear the **Enable Threat Detection** check box.

IMPORTANT: If you disable threat detection for a job where it was enabled, any potential threat flags for backups in the job will be cleared. Only disable threat detection for a job once all potential threats have been addressed. See [Manage potential ransomware threats](#).

- To back up VMs and check for potential ransomware threats on the VMs, select the **Enable Threat Detection** check box. If you did not enter credentials for truncating application transaction logs in [Step 7](#), enter credentials for connecting to VMs in the job.

To enter credentials for multiple VMs in the job, enter a username and password in the **Guest VM Credentials** area.

To enter credentials for a specific VM in the job, click the arrow at the right side of the VM name in the Backup Set area, and enter a username and password in the **Guest VM Credentials** area for the VM.

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

Note: The same credentials are used for truncating transaction logs in application-consistent backups and checking for potential ransomware threats.

Note: If you enter credentials for a specific VM in the job, the Agent will not attempt to connect to the VM using the credentials specified for multiple VMs in the job.

10. Specify whether you want the VRA to check whether VMs can be restored by doing one of the following:

- To back up VMs without checking whether they can be restored, clear the **Verify this backup upon completion** check box.
- To back up VMs and check whether Windows VMs can be restored from the backup, select the **Verify this backup upon completion** check box.

Note: You can only enable backup verification if the selected vault supports this feature and backup verification settings are entered for the VRA. See [vSphere Rapid VM Restore and backup verification requirements](#) and [Enter backup verification settings for a vSphere Recovery Agent](#).

11. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. To create a schedule for running the backup, see [Run and schedule backups, synchronizations and custom commands](#). If you do not want to create a schedule at this time, click **Cancel**.

4.11.1 Application-consistent backups on vSphere VMs

While protecting the entire file system of a Windows VM, the VRA can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on the VM.

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

In an application-consistent backup, pending application transactions are written to disk before the data is backed up. This minimizes the amount of work required to restore the application.

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

Application-consistent backups are only supported on Windows VMs. If Linux VMs are included in backup jobs where the application-consistent backup setting is enabled, warning messages for the Linux VMs may appear in the backup logs.

To create an application-consistent backup on a VM, VMware Tools version 11 or later must be installed on the VM.

Note: The VRA cannot create application-consistent backups on encrypted VMs. VMware does not support application-quiesced snapshots for encrypted VMs.

Note: The VRA cannot back up or restore an application database on a physical Raw Device Mapping (pRDM), shared or independent disk. VMware does not allow these disk types to be included in snapshots for VM-level backups. To back up an application on a pRDM, shared or independent disk, install the Windows Agent and SQL Server or Exchange Plug-in on the VM.

Log truncation in application-consistent backups

When performing application-consistent backups, the vSphere Recovery Agent can truncate SQL Server, Exchange and SharePoint transaction logs on VMs. This prevents the transaction logs from taking up a significant amount of disk space and reducing system performance. There are no logs to truncate when performing application-consistent backups of domain controllers with Active Directory.

Note: The vSphere Recovery Agent can truncate transaction logs for the default SQL Server instance and for all Exchange Server databases. The VRA cannot truncate logs for named SQL Server instances.

To truncate transaction logs on a VM after an application-consistent backup, you must enable log truncation in the backup job and provide credentials that have admin access to the VM. The specified user does not need admin rights to applications on the VM; it only needs admin access to the VM.

You can provide guest VM credentials with admin access to multiple VMs in a backup job and/or provide credentials for specific VMs. If you provide credentials for a specific VM, the guest VM credentials for multiple VMs will never be used to connect to that VM.

Logs cannot be truncated if an application-consistent backup could not be performed for some reason (e.g., VMware tools not installed on the guest VM).

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

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

4.11.2 Backup verification for vSphere VMs

Beginning in version 9.00, the vSphere Recovery Agent (VRA) can check whether each Windows VM in a backup can be restored. You can view the verification results in the Backup Verification report in Portal 9.00 or later or in Verification logs in Portal 9.30 or later. See [View the Backup Verification Report](#) and [View a job's process logs and safeset information](#).

When backup verification settings are entered for a VRA and backup verification is enabled for a vSphere backup job, the VRA backs up VMs in the job and then checks whether each Windows VM can be restored from the backup. Using automated Rapid VM Restore processes, the VRA attempts to start each VM from the backup and takes a screenshot of the login screen for each Windows VM that can be restored.

Note: Backup verification does not run for vSphere backups that are started by intra-daily schedules.

VMs in a backup job are verified sequentially, one at a time. The verification process for each VM can take up to 10 minutes. If the VM has not started after 10 minutes, the process times out and the VRA tries to verify the next VM in the backup job.

If VMs in a backup job are being verified and you start the backup job again, verification is canceled for VMs that have not yet been verified. If VMs in a backup job have not been verified recently, the job might be scheduled too frequently to allow backup verification to complete.

Only one Rapid VM Restore process can run for a VM in a backup job at the same time, regardless of whether Rapid VM Restores are started by a user or by a backup verification process. If the VRA tries to verify a VM backup at the same time you are restoring the VM using Rapid VM Restore, the verification process could fail. Similarly, if a verification process starts for a VM backup while the previous VM backup is being verified, the new backup verification could fail.

For more information, see [vSphere Rapid VM Restore and backup verification requirements](#), [Enter backup verification settings for a vSphere Recovery Agent](#) and [Add a vSphere backup job](#).

4.12 Add and schedule a Hyper-V backup job

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

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

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

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

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

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

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

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

4.12.1 Best practices for backing up Hyper-V VMs

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

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

occur. In a reseed, all data for a VM is sent to the vault even though the VM was previously backed up.

The following list describes situations when backups reseed:

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

If a VM is moved to a different backup job with the same encryption password, the VM backup does not reseed.

- On Windows Server 2012 R2, snapshot (AVHD) files reseed after storage migration. Hard disk VHD(x) files do not reseed after storage migration.

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

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

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

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

4.12.2 Best practices for seeding Hyper-V VM backups

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

- Seed backups locally
- Ideally, use a Satellite vault to provide fast, local vault access. If you do not use a Satellite vault, consider using a temporary vault to seed Hyper-V backups locally. The data can then be imported into an offsite vault.
- Seed VM backups in separate jobs
- Seeding backups, particularly for large VMs, can take a significant amount of time. Because deferring is not available for scheduled Hyper-V backup jobs, it is best not to seed VMs in a scheduled job. If you add a VM to an existing scheduled job, the job could take a long time, and potentially cause the backup to overlap the next scheduled backup.

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

Note: Normally, it is best to include multiple VMs in a single backup job.

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

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

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.

4.12.3 Application-consistent backups on Hyper-V VMs

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

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

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

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

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

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

- The guest VM credentials are incorrect.
- The VM is not in a running state (e.g., is powered off, paused or migrating).
- The VM is running on a Hyper-V host with an operating system where application-consistent backups are not supported.
- The VM is running on a Hyper-V host where the Host service is not installed.
- The VM has a guest operating system where application-consistent backups are not supported.

- The VM configuration version is earlier than 6.2.
- The Backup (volume shadow copy) integration service is not enabled for the VM.
- A VSS writer on the host cannot connect to the VM.
- A VSS writer on the VM is not available or is in a bad state.

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

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

Log truncation in application-consistent backups

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

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

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

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

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

4.12.4 Add a Hyper-V backup job

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

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

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

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

To add a Hyper-V backup job:

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

The Computers page shows registered computers and environments.

2. Click the Hyper-V environment row. 

3. Do one of the following:

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

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

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


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

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

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

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

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



Virtual Machine	VM ID on Vault	Job	Backup Status	Date	Actions
☐			All		
☐ ProtectedTest/VM	HY_E6H473ABE9RLWVIBANQDZE		Not Protected		
☐ Protected/VM	HY_J4DK676TVI8UJB632J4B46L4		Not Protected		
☐ supported/VM	HY_AVTU7GVVQVFNHVRN478TBJAQ	hypervjob	Completed	today at 3:33 PM 8.23 MB	Select Action
☐ VM	HY_HRP26L7N8BETJ2VH4Q2B70Y	hypervjob	Completed	today at 4:30 PM 8.21 MB	Select Action
☐ VM1	HY_GWU2BVXNFA0N2HTZSH4H4DU		Not Protected		
☐ VM2	HY_WTLBWH2KPB5Q7Y8R4WV122B1		Not Protected		

- c. Do one of the following:

- Select the check box for each VM that you want to include in the backup job.
- Select the check box at the top left of the list to select all VMs on the page.

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

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

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

4. In the Create New Job dialog box, specify the following information:

- In the **Name** box, type a name for the backup job.
- In the **Description** box, type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.
- In the **Encryption Password** and **Confirm Password** boxes, enter a data encryption password. You can also enter a password hint in the **Password Hint** box.

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

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

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

Note: Application Consistent backups are available for Hyper-V Agent version 8.84 or later in Portal version 8.87 or later.

- To perform crash-consistent backups of VMs in the backup job, turn off the **Enable Application Consistent Backups** toggle.

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

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

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

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

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

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

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

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

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

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

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

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

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




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

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

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

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

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



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

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

4.12.5 Edit a Hyper-V backup job

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

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

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

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

To edit a Hyper-V backup job:

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

The Computers page shows registered computers and environments.

2. Click the Hyper-V environment row. 

3. Do one of the following:

- On the **Jobs** tab, in the **Jobs** column, click the name of the job that you want to edit.
- On the **Jobs** tab, find the job that you want to edit. In its **Select Action** menu, click **Edit Job**.
- On the **Virtual Machines** tab, in the **Jobs** column, click the name of the job that you want to edit.
- On the **Virtual Machines** tab, click a VM that belongs to the job you want to edit. In the **Select Action** menu, click **Edit Job**.

4. If required, change the job description, encryption password or password hint at the left side of the Edit Job dialog box.

5. To change credentials for application-consistent backups on VMs in the job, do one or both of the following:

- To change credentials for VMs in the job, enter an admin user's username and password in the Guest VM Credentials area.
- To change credentials for a specific VM in the job, click the arrow at the right side of the VM name in the **Protected by this job** box, and enter an admin user's username and password for the VM.

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

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

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

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

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

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

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



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

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

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

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

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

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

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

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

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

If a job is scheduled to start at exactly the same time by multiple schedules, the job only runs once at the scheduled time. If the jobs have different retention types, the retention type of the schedule that is highest in the list is applied to the resulting safeset. For example, in the following screenshot, the job is scheduled

to run at 11 PM on the last day of the month with the Monthly retention type, and every night at 11 PM with the Daily retention type. On the last day of each month, the job runs only once at 11 PM. Because the schedule with the Monthly retention type is higher in the list than the schedule with the Daily retention type, the Monthly retention type is applied to the safeset.

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

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

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

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

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

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

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

The Configure Job Schedule dialog box opens.

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


- To run the backup on specific days each week, click **Days of Week** in the **Schedule View** list. Select the days when you want to run the job. Use the **At** field to specify the time when you want to run the job each day. Click **Okay**.

The screenshot shows the 'Configure Job Schedule' dialog box. The 'Schedule View' dropdown is set to 'Days of Week'. Below it, a row of buttons represents the days of the week: Su, Mo, Tu, We, Th, Fr, Sa. The buttons for Tu, We, Th, and Fr have a green checkmark, indicating they are selected. The 'At' field is set to '0:00 AM' and has a clock icon next to it. The text '(Agent time zone)' is displayed next to the clock icon. At the bottom, there are 'Okay' and 'Cancel' buttons.

- 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. The 'Schedule View' dropdown is set to 'Days of Month'. Below it is a calendar grid showing dates from 1 to 31. The dates 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, and 31 are all selected, indicated by a light blue background. The 'At' field is set to '12:00 AM' and has a clock icon next to it. The text '(Agent time zone)' is displayed next to the clock icon. At the bottom, there are 'Okay' and 'Cancel' buttons.


- 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 'Configure Job Schedule' dialog box features a blue title bar with a question mark icon and a close button. Below the title bar, there is a 'Schedule View' dropdown menu currently set to 'Custom'. Underneath is a 'Custom Cycle' text input field. A light blue informational box contains the following text: 'Format: 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)', and 'Keywords: * for every time, Last for the last day of any month'. At the bottom of the dialog are 'Okay' and 'Cancel' buttons.

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.



The 'Create New Job' dialog box has a blue title bar with a question mark icon and a close button. It is divided into two main sections. The left section contains fields for 'Name' (HyperVBackup), 'Description', 'Destination' (vault), 'Encryption Password' (masked), 'Confirm Password' (masked), and 'Password Hint' (A). A blue warning box at the bottom left states: 'You must remember your encryption password. Your data cannot be restored without your password.' The right section, titled 'Virtual Machines', contains a 'Schedule' sub-section with an 'Add Schedule' button. Below this is a table with columns: Retention, Schedule, Enable, and Priority. The table contains two rows: one with 'Monthly' retention and '11:00 PM Last' schedule, and another with 'Daily' retention and '11:00 PM Su,Mo,Tu,We,Th,Fr' schedule. Both 'Enable' checkboxes are checked. The 'Priority' column has up and down arrow icons for each row, with blue arrows pointing to them in the image. At the bottom right are 'Create job' and 'Cancel' buttons.

9. Click **Save**.

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

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

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

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

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

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

A grid lists available computers.

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

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

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

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

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

4.13 Log file options

When you create or edit a backup job, you can specify the level of detail for job logging. Select one of the following job logging levels from the list:

- **Files** — Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up.
- **Directory** — Provides less detail than the Files logging level. Provides information about folders that are backed up.
- **Summary** — Provides high-level information, including the vault and Agent version, and backup sizes.
- **Minimal** — Provides high-level information, including the vault and Agent version.

Changing the logging level only affects log files that are created from that point and after. It does not affect previously-created log files.

Note: For Image Plug-in jobs, the selected logging level does not affect the content of the logs.

For Windows, Linux and UNIX backup jobs, the following log file options are also available:

- **Create log file.** If this check box is selected, the system generates log files for each job. Log files can contain the start-connect-completion and disconnect times, file names (i.e., the names of the files that were copied during backup), and any processing errors.
- **Automatically purge expired log files.** If this check box is selected, the log file associated with a backup is automatically deleted when the backup has been deleted from the vault. Backups are typically deleted from the vault according to retention types. See [Add retention types](#).
- **Keep the last <number of> log files.** Specifies the number of log files to keep for a backup job. When the specified number is reached, the oldest log file for a backup job will be deleted to make space for the newest one.

Note: You must choose either the **Automatically purge expired log files** option or the **Keep the last <number of> log files** option. When a backup job runs, log files are removed according to the specified option. Log files are not removed when a backup job is synchronized.

4.14 Encryption settings

Encryption settings specify the encryption type for backup data at rest on the vault. AES 256 bit encryption is the only encryption type available for new backup jobs.

If an existing job uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES), you can continue to encrypt the job using that type. However, if you change the encryption type for an existing job, you cannot change the encryption type back to the original type. Only AES 256 bit will be available.

If you change encryption options for an existing job, it will force a new full backup (i.e., a reseed). The next backup will take longer than previous delta backups, and the amount of data stored on the vault will increase temporarily, depending on your retention settings.

Encryption password

You must enter a password for the encrypted backup data. The password is case-sensitive. To recover the data, you must provide the encryption password that was entered when the files were backed up.

You can also enter a password hint. When restoring data, you can view the password hint to remind you of the encryption password for this job. The password hint can include lowercase characters (a-z), uppercase characters (A-Z), international characters (Á-ÿ), numbers (0-9), spaces, and the following special characters: ! @ # \$ % ^ & * () _ - + = [] { } | ' " : ; , < . > ? ~ `

IMPORTANT: The encryption password is required for restoring the data, so be sure to store it somewhere safe. If you forget your password, you will not be able to restore your data. The password is not maintained anywhere else and cannot be recovered.

4.15 Advanced backup options

When you create or edit a Windows, Linux or UNIX backup job, some of the following options are available in the Advanced Backup Options dialog box.

Back up files opened for write

If the **Backup files opened for write** option is selected, files are backed up if they are open for writing or shared reading during the backup. Files that are open for exclusive writes cannot be backed up.

When this option is selected, inconsistencies in the backup can occur if an open file is modified during the backup process.

Suppress archive bit processing

Note: This option is only available for Windows backup jobs.

In some operating systems, an archive attribute is placed in a file when the file is created or modified. The archive attribute indicates that the file needs to be backed up.

If the **Suppress archive bit processing** option is selected, the Agent does not clear the archive attribute when it backs up a file. If you use other programs that rely on the archive attribute, make sure that the **Suppress archive bit processing** option is not selected.

If the **Suppress archive bit processing** option is not selected, the Agent clears the archive attribute when it backs up a file.

Back up a single instance of all selected hard linked files

Note: This option is only available for Linux and UNIX backup jobs.

A hard link is a reference, or pointer, to data on a storage volume. More than one hard link can be associated with the same data. Hard-linked files cannot cross disk boundaries and only exist on the same disk.

If the **Back up a single instance of all selected hard linked files** option is selected, only one copy of the data is backed up, along with all hard links. When the data is restored, both the data (with a new inode) and the hard links are restored. When this option is selected, a pre-scan process is required. The pre-scan reads through the file system, gets each inode and stores it in a map. The larger the file system, the more memory this map requires and the more time it takes to process. However, the resulting backup size is smaller.

If the **Back up a single instance of all selected hard linked files** option is not selected, the data is backed up separately for each hard link. When the data is restored, the hard-link relationship will not be re-established. Each file will be restored individually and applications depending on hard links may not be automatically restored. When this option is not selected, the backup is faster but the total backup size is larger.

4.16 Filter subdirectories and files in backup jobs

When you include and exclude folders in a Windows, UNC, Linux, UNIX or NFS backup job, the folder's subdirectories and files are also included or excluded by default.

If you only want to back up some subdirectories or files in a folder, you can add filters to the inclusion record. For example, you could add a filter so that files in a folder are only backed up if they have the .doc or .docx extension.

If you only want to exclude some subdirectories or files in a folder from a backup job, you can add filters to the exclusion record. For example, you could add a filter so that files in a folder are only excluded from the backup if they have the .exe extension.

If a policy is assigned to a computer, you can add filters from the policy to a folder inclusion or exclusion record.

Filters in a backup job are applied when the job runs. New subdirectories and files that match the filters are automatically backed up or excluded when the job runs.

To filter subdirectories and files in a backup job:

1. When creating or editing a Windows, UNC, Linux, UNIX or NFS backup job, view the **Backup Set** box.

	Folders Filter	Files Filter	Recursive	
+ C:		*.*	☒	
+ Documents an...	e.g., a*, b*	*. *	☒	
+ ProgramData	e.g., a*, b*	*. *	☒	

2. If editable fields do not appear for a folder inclusion or exclusion record where you want to filter subdirectories and files, click the **Edit** button in the folder row.

3. In the **Backup Set** box, for each included folder where you want to include specific subdirectories or files, do one or more of the following:

- To include specific subdirectories in the backup job, in the **Folders Filter** field, enter the names of subdirectories to include. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only include subdirectories in a backup if their names end with “-current” or start with “2015”, enter the following filter: *-current, 2015*

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To include specific files in the backup job, in the **Files Filter** field, enter the names of files to include in the backup. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only include files in a backup if they have the .doc or .docx

extension, enter the following filter: *.doc, *.docx

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To back up the specified folder but not its subfolders, clear the **Recursive** check box.
- If a policy is assigned to the computer, to apply filters from the policy to the folder inclusion record, click the **Apply Policy Filters** button. 

4. In the **Backup Set** box, for each excluded folder where you want to exclude specific subdirectories or files, do one or more of the following:

- To exclude specific subdirectories from the backup job, in the **Folders Filter** field, enter the names of subdirectories to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude subdirectories from a backup if their names end with “-old” or start with “2001”, enter the following filter: *-old, 2001*

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To exclude specific files from the backup job, in the **Files Filter** field, enter the names of files to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude files from a backup if they have the .exe or .dll extension, enter the following filter: *.exe, *.dll

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To apply the file filter to the specified folder but not to its subfolders, clear the **Recursive** check box.

If a file filter is not applied to the exclusion, clearing the **Recursive** check box will have no effect. The specified folder and all of its content will be excluded.

- If a policy is assigned to the computer, to apply filters from the policy to the folder exclusion record, click the **Apply Policy Filters** button. 

5. Click **Create Job** or **Save**.

4.17 Edit a backup job

You can edit existing backup jobs to change the following settings:

- Items to back up
- Log file options
- Encryption settings. If you change the encryption method or password in a backup job, the job will reseed the next time it runs.
- Job description

For a SQL Server Plug-in backup job, you can also change the SQL Server instance where you want to back up databases and credentials for connecting to the instance.

Depending on backup job type (e.g., Windows or vSphere), other options might appear in the Edit Job dialog box. For information about options for a specific backup job type, see [Add a Windows Local System backup job](#), [Add a Windows Image backup job](#), [Add a vSphere backup job](#) or [Add and schedule a Hyper-V backup job](#).

Note: You cannot change the name or destination of a job.

To edit a backup job:

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

The Computers page shows registered computers.

2. Find the computer or protected environment with the job that you want to edit, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. Do one of the following:

- In the **Name** column, click the name of the job that you want to edit.
- In the **Select Action** menu of the job that you want to edit, click **Edit Job**.

The Edit Job dialog box shows the current job settings.

5. For a SQL Server Plug-in backup job, to change the SQL Server instance or credentials for connecting to the instance, click **Change Instance / Credentials**. In the Connect to SQL Server dialog box, select the SQL Server instance where you want to back up databases and specify credentials for connecting to the instance. Click **Connect**.
6. Do one or more of the following:
 - In the **Description** box, type a description for the backup job.
 - In the **Log File Options** list, select the level of detail for job logging.

Note: For Image Plug-in jobs, the selected logging level does not affect the content of the logs.
 - In the **Encryption Settings** list, select the encryption method for the backup data. In most jobs, the encryption method is AES 256 bit. See [Encryption settings](#). In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.
 - In the box that shows items for backup, select items to back up.
7. If other options are available, enable or disable the options as desired. For information about options for a specific backup job type, see [Add a Windows Local System backup job](#), [Add a Windows Image backup job](#), [Add a vSphere backup job](#) or [Add and schedule a Hyper-V backup job](#).

IMPORTANT: If you disable threat detection for a backup job where it was enabled, any potential threat flags for backups in the job will be cleared. Only disable threat detection for a job once all potential threats have been addressed.

8. Click **Save**.

5 Delete jobs and computers, and delete data from vaults

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

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

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

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

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

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

Note: If you delete a Hyper-V environment using this method, backup data is deleted for all VMs in the backup job. To delete one VM from a Hyper-V backup job, see [Delete a Hyper-V VM backup from Portal and vaults](#).

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

- Delete specific backups from vaults. This option is available beginning in Portal 8.90. See [Delete specific backups from vaults](#).

Backup deletion requests are submitted to vaults immediately; there is no waiting period before the data deletion request is sent to vaults. Because backup deletion requests are submitted immediately, backup deletion requests cannot be canceled.

Note: Backup deletion requests are not supported with the Hyper-V Agent.

5.1 Delete a backup job without deleting data from vaults

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

For most agents, if a job is deleted from Portal in this way, the data can still be restored using the *Restore from Another Computer* procedure. See [Restore data from another computer](#).

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

To delete a backup job without deleting data from vaults:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the online computer with the job that you want to delete, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. In the **Select Action** menu of the job that you want to delete, click **Delete Job**.
5. If you are signed in as an Admin user in a Portal instance where the data deletion feature is enabled, a Delete Job dialog box appears.

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

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

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

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

7. Click **Confirm Deletion**.

5.2 Delete a backup job and delete job data from vaults

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

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

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

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

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

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

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


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

The Computers page shows registered computers.

2. Find the computer with the job that you want to delete, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. In the **Select Action** menu of the job that you want to delete, click **Delete Job**.

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

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



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

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

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

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

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

7. Click **Confirm Deletion**.

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

8. Click **Close**.

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



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

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

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

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

5.3 Cancel a scheduled job data deletion

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

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


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

To cancel a scheduled job data deletion:

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

The Computers page shows registered computers.

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



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

5. Click **Yes**.

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

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



5.4 Delete a computer without deleting data from vaults

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

If a computer is deleted from Portal in this way, the data can still be restored using the *Restore from Another Computer* procedure. See [Restore data from another computer](#).

Note: Hyper-V VMs cannot be restored using the *Restore from Another Computer* procedure.

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

To delete a computer without deleting data from vaults:

1. When signed in as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Select the check box for each computer that you want to delete.
3. In the **Actions** list, click **Delete Selected Computer(s)**.
4. If the data deletion feature is enabled in your Portal instance, a Delete Computer(s) dialog box appears.

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

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

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

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

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

5.5 Delete a computer and delete computer data from vaults

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

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

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

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

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

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

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

Note: If you delete a Hyper-V environment using this method, backup data is deleted for all VMs in the backup job. To delete one VM from a Hyper-V backup job, see [Delete a Hyper-V VM backup from Portal and vaults](#).

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

To delete a computer and delete computer data from vaults:


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

The Computers page shows registered computers.

2. Select the check box for each computer that you want to delete.
3. In the **Actions** list, click **Delete Selected Computer(s)**.

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

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



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

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

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

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

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

6. Click **Confirm Deletion**.

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


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

7. Click **Close**.

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

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

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



5.6 Cancel a scheduled computer data deletion

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

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

To cancel a scheduled computer data deletion:

1. When signed in as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Select the check box for each computer for which you want to cancel the scheduled data deletion.
The Status column shows *Scheduled for deletion* for each computer that is scheduled for deletion.
3. In the Actions list, click **Cancel Deletion of Selected Computers**.

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

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

4. Click **Yes**.

A Success dialog box appears.

5. Click **Okay**.

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

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

5.7 Delete specific backups from vaults

In a Portal instance where the data deletion feature is enabled, Admin users can request that specific backups (also known as safesets) be deleted from all vaults. When selecting backups to delete, Admin users can view information about each backup, including its date, retention settings, size, and whether it has a potential ransomware threat.

Backup deletion requests are submitted to vaults immediately and the data is automatically deleted from associated vaults. Because backup deletion requests are submitted immediately, backup deletion requests cannot be canceled.

When a backup deletion request is submitted, an email notification is sent to Admin users for the site and to Super users. A notification also appears in the Status Feed.

If a backup deletion request fails, an email notification is sent to a vault administrator whose email address is specified in Portal. The vault administrator can then manually delete the backup or backups from vaults.

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

To delete specific backups from vaults:

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

The Computers page shows registered computers.

2. Find the computer with the backups 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 with backups that you want to delete, click **Delete Backup**.

If the Delete backup option does not appear or a message states that the job is registered to a vault that does not support backup deletion, you cannot submit a request to automatically delete backups from vaults.

A Delete Backup dialog box appears. The dialog box shows information about each backup, including its retention settings, size, and whether it has a potential ransomware threat.

5. Select the check box for each backup that you want to delete, and then click **Delete**.

Backups that cannot be deleted (e.g., because a deletion request is scheduled for the job or computer) cannot be selected.

You cannot delete all available backups for a job. Instead, delete the entire job. See [Delete a backup job and delete job data from vaults](#).

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

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

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

7. Click **Confirm Deletion**.

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

A dialog box states that the backup data will be deleted from vaults.

8. Click **Close**.

6 Run and schedule backups, synchronizations and custom commands

After a backup job is created, you can run it manually (ad hoc) at any time and schedule it to run on specific days of the week or month. See [Run an ad-hoc backup](#) and [Schedule a backup](#).

To help you meet your recovery point objectives (RPOs), when Windows Agent 8.90 or later, Linux Agent 8.90 or later, AIX Agent 9.00 or vSphere Recovery Agent 9.11 is backing up data to a Director version 8.60 or later vault, you can schedule a backup job to run multiple times per day, as often as hourly. See [Schedule a backup to run multiple times per day](#).

Beginning with Portal 9.30 and Windows Agent 9.30, Image and Local System backups can also be triggered by system events on supported Windows desktop operating systems. See [Trigger backups when events occur on Windows desktop computers](#).

When running or scheduling a backup, you can specify the following settings:

- 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.
- Deferring. You can use deferring to prevent large backups from running at peak network times. When deferring is enabled, the backup job does not back up any new data after the specified amount of time and commits the safeset to the vault, even if some data in the job is not backed up. Changes to data that was previously backed up will be backed up, regardless of the specified amount of time.

When the job runs again, the agent checks for changes in data that was previously backed up, backs up those changes, and then backs up remaining data.

If a backup job is deferred while an item (e.g., file, database, volume, vSphere VMDK, or Hyper-V VM) is being backed up, the backup for that item is incomplete and data from the item cannot be restored. However, you can restore items that were completely backed up in the job before the job was deferred.

The following table describes deferral behavior for specific backup types:

Backup type	Behavior
System State	If the System State option is selected in a Windows backup job, the job cannot be deferred. If the agent tries to defer a job where System State data is being backed up, the backup fails.

Image Plug-in	In an Image Plug-in job, once a volume has been completely backed up, a backup for that volume cannot be deferred. When an Image Plug-in job runs, the agent backs up changes in any volumes that have been completely backed up and then starts to back up any remaining volumes. The backup can be deferred for any volume that was not previously backed up completely. If an Image Plug-in backup job is deferred while a volume is being backed up, the backup for the volume is incomplete and data from the volume cannot be restored. However, you can restore volumes, and files and folders from volumes in the job that were completely backed up.
SQL Server Plug-in	If a SQL Server Plug-in backup job is deferred while a database is being backed up, the backup for that database is incomplete and the database cannot be restored. However, you can restore databases that were completely backed up in the job before the job was deferred. If the backup is deferred while a SharePoint database is being backed up, the backup is incomplete and you cannot restore items from the database.
Exchange Plug-in	Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.
To SSI files	Backups to SSI files on disk cannot be deferred.
Hyper-V Agent	Hyper-V Agent backups can only be deferred when they are run manually (ad hoc). You cannot defer scheduled Hyper-V backups.

- For a SQL Server Plug-in backup job, you can specify whether to back up the database, the transaction logs, or both. Frequent transaction log backups are recommended for databases with a high level of activity.

Note: After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), be sure that only one tool is being used for truncating logs.

Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

- For an application-aware Image Plug-in job, you can specify whether to truncate SQL Server database transaction logs.

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

- For an Exchange Plug-in backup job, you can specify whether to:
 - Run a Full or Incremental backup. When the Full backup type is selected, the database files, checkpoint file and transaction logs are backed up. When the Incremental backup type is selected, the database files, checkpoint file and transaction logs are backed up in the first “seed” backup, but only the checkpoint file and transaction logs are backed up in subsequent runs. For more information, see [Plan Full and Incremental Exchange backups](#).

- Validate Exchange data during the backup. When this option is selected, a utility checks the Exchange data during the backup. If data corruption is detected, the backup fails and the corruption is reported.

For computers with Windows or Linux Agent version 8.60 or later, or environments with vSphere Recovery Agent version 8.80 or later, you can specify whether scheduled backups should automatically retry if they do not run successfully. You can also specify how many times a scheduled backup should retry after a failed attempt, and specify the amount of time between retries. See [Specify whether scheduled backups retry after a failure](#).

When you schedule a job to run, you can also set the compression level for the data. The compression level optimizes the amount of data stored vs. the backup speed. The default compression level is usually the optimal setting.

When a backup job first runs, all data selected in the job is backed up to the vault. This initial backup is called a “seed” backup. In subsequent backups, only data that has changed is backed up to the vault, unless a reseed is required (e.g., after a job’s encryption password has changed). In a reseed, all data selected in a backup job is sent to the vault again, even though it has already been backed up.

After a backup runs, you can view logs to check whether the backup completed successfully. See [View a job’s process logs and safeset information](#).

In some cases, you must synchronize a backup job before you run it or restore data from the job. When you synchronize a job, the agent checks which safesets for the job are online and available for restore. See [Synchronize a job](#).

You can also schedule custom commands to run on Windows and Linux computers. Custom commands are scripts that are saved on a computer where an agent is installed and are scheduled to run through Portal. For example, you could schedule a custom command that shuts down services on the computer, runs a backup, and then restarts the services. See [Schedule a custom command](#).

6.1 Schedule a backup

After creating a backup job, you can add one or more schedules for running the job on specific days of the week or month. You can create complex schedules for a 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.


Note: When Windows Agent 8.90 or later, Linux Agent 8.90 or later, AIX Agent 9.00 or vSphere Recovery Agent 9.11 is backing up data to a Director version 8.60 or later vault, you can also schedule a backup job to run multiple times per day, as often as hourly. See [Schedule a backup to run multiple times per day](#).

Note: Beginning with Portal 9.30 and Windows Agent 9.30, Image and Local System backups can also be triggered by system events on supported Windows desktop operating systems. See [Trigger backups when events occur on Windows desktop computers](#).

When scheduling multiple SQL Server database jobs in the same instance, it is good practice to schedule the jobs so that their running times do not overlap. Simultaneous backups are supported, but are not recommended.

If a job is scheduled to start at exactly the same time by multiple schedules, the job only runs once at the scheduled time, and the retention type of the schedule that is higher in the schedule list is applied to the resulting safeset. For example, in the following screenshot, a job is scheduled to run at 12 AM on Saturdays by two schedules. On Saturdays, the job runs only once at 12 AM. Because the schedule with the Weekly retention type is higher in the list than the schedule with the Daily retention type, the Weekly retention type is applied to the resulting 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. In particular, avoid overlapping schedules for SQL Server database jobs in the same instance. Simultaneous backups in the same SQL Server instance are supported but are not recommended.



When you schedule a backup, the View/Add Schedule dialog box shows the maximum number of restore points that could result from the job's current schedules and retention types. You can then change your schedules, if required. See [Maximum number of restore points for a job](#).

To schedule a Hyper-V job, see [Add or edit a schedule for a Hyper-V backup job](#).

To schedule a backup job to run at a specific time on specific days of the week or month:

- Do one of the following:
 - On the navigation bar, click **Computers**. Find the computer with the backup job that you want to schedule, and click the row to expand its view. On the Jobs tab, find the job that you want to schedule. In its **Select Action** menu, click **View/Add Schedule**.
 - Create a new backup job. The View/Add Schedule dialog box appears when you save the job.
- In the View/Add Schedule dialog box, click **Add Schedule**.
A new row appears in the dialog box.
- In the new schedule row, in the **Retention** 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. See [Add retention types](#).

The 24-Hours and 48-Hours retention types are only available for intra-daily schedules. See [Schedule a backup to run multiple times per day](#).

4. If the schedule is for a SQL Server Plug-in database backup job, do one of the following in the **Backup Type** list:

- To back up each database from the point in time when the backup starts, click **Full**.
- To back up each database and its transaction logs from the point in time when the backup starts, click **Full with transaction logs**.
- To back up the database transaction logs only from the point in time when the backup starts, click **Transaction logs only**. When **Transaction Logs only** is selected, the entire database and its transaction logs will be backed up when the job first runs. In subsequent backups, only the transaction logs will be backed up.

After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), be sure that only one tool is being used for truncating logs.

Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

6. If the schedule is for an Image Plug-in job that backs up volumes with SQL Server database files, do one of the following in the **SQL Application Settings** list:

- To truncate database transaction logs after the backup, select **Truncate transaction logs**.
- To run the backup without truncating logs, clear **Truncate transaction logs**.

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

7. If the schedule is for an Exchange database backup job, do the following:

- In the **Backup Type** list, do one of the following:
 - To only back up transaction logs and the checkpoint file after the first “seed” backup, click **Incremental**.
 - To back up the database files, checkpoint file and transaction logs, click **Full**.

For more information, see [Plan Full and Incremental Exchange backups](#).

- To validate Exchange data during the backup, select **Validate Exchange database**.

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

The Configure Job Schedule dialog box opens.

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

- To run the backup on specific days each week, select **Days of Week** in the **Schedule View** list. Select the days when you want to run the job. Then use the **At** field to specify the time when


you want to run the job.

The screenshot shows the 'Configure Job Schedule' dialog box. The 'Schedule View' dropdown is set to 'Days of Week'. Below it, a row of buttons represents the days of the week: Su, Mo, Tu, We, Th, Fr, Sa. The buttons for Mo, Tu, We, Th, and Fr are highlighted with a blue background and a green checkmark. The 'At' field is set to '0:00 AM' and has a clock icon next to it, with '(Agent time zone)' in parentheses. At the bottom, there are 'Okay' and 'Cancel' buttons.

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

The screenshot shows the 'Configure Job Schedule' dialog box. The 'Schedule View' dropdown is set to 'Days of Month'. Below it is a calendar grid showing days 1 through 31, with a 'Last' button. The 'At' field is set to '12:00 AM' and has a clock icon next to it, with '(Agent time zone)' in parentheses. At the bottom, there are 'Okay' and 'Cancel' buttons.

- To create a custom schedule, select **Custom** in the **Schedule View** list. In the Custom Cycle dialog box, enter a custom schedule. Be sure to follow the format and notation as described.



Note: If **Intra-daily** appears in the **Schedule View** list, you can also schedule the backup to run multiple times each day. See [Schedule a backup to run multiple times per day](#).

10. Click **Okay**.

The new schedule appears in the Schedule box.

11. In the **Compression** list, click a compression level for the backup data. Compression levels optimize the amount of data stored vs. the backup speed.
12. Do one of the following:
 - To allow the backup job to run without a time limit, click **None** in the Deferring list.
 - To specify a maximum amount of time that the backup job can run, click **Minutes** or **Hours** in the **Deferring** list. 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 amount of time specified.

Note: Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.

13. To run the job on the specified schedule, select the **Enable** check box near the end of the row.
14. If there is more than one schedule row, you can use the **Priority** arrows to change the order of the schedule rows. Schedules higher in the list have a 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.

15. Check the number of restore points that could result from the job's schedules and retention policies. If you want to increase or decrease the number of restore points, change the schedules or retention types.

The maximum number of restore points appears below the schedules in the View/Add Schedule dialog box. For more information, see [Maximum number of restore points for a job](#).

16. If an Automatic Retry for Scheduled Backups section appears in the View / Add Schedule dialog box, you can specify whether scheduled backups should retry after a failed backup. See [Specify whether scheduled backups retry after a failure](#).
17. If a Windows Event Backup Triggers section is available in the View / Add Schedule dialog box, you might be able to create a Windows event backup trigger. When a backup job has a trigger, the job runs automatically when a user logs on to the computer or the computer starts to shut down. See [Trigger backups when events occur on Windows desktop computers](#).
18. Click **Save**.

6.2 Schedule a backup to run multiple times per day

Beginning with Windows Agent 8.90, Linux Agent 8.90, AIX Agent 9.00 and vSphere Recovery Agent (VRA) 9.11, when an agent is backing up data to a Director version 8.60 or later vault, you can schedule the backup job to run multiple times per day by creating an intra-daily schedule. You can create an intra-daily schedule for a Windows, Linux or AIX backup job using Portal 8.88 or later. You can create an intra-daily schedule for vSphere backup jobs beginning in Portal 9.20.

Note: Backup verification does not run for vSphere backups that are started by intra-daily schedules.

Note: To schedule a backup job to run on specific days of the week or month, see [Schedule a backup](#).

Note: Beginning with Portal 9.30 and Windows Agent 9.30, Image and Local System backups can also be triggered by system events on supported Windows desktop operating systems. See [Trigger backups when events occur on Windows desktop computers](#).

Each backup job can have one intra-daily schedule. If the job has other schedules, the intra-daily schedule has the lowest priority and is at the bottom of the schedule list. If a job is scheduled to start at exactly the same time by an intra-daily schedule and another schedule, the job only runs once and the retention type of the other schedule (e.g., daily or monthly) is applied to the resulting safeset.

When you create an intra-daily schedule for a backup job, you can choose one of two retention types:

- **24-Hours.** With this retention type, each backup is kept for at least 24 hours and at least one backup with this retention type is stored online.
- **48-Hours.** With this retention type, each backup is kept for at least 48 hours and at least one backup with this retention type is stored online.


Other retention types are not available for intra-daily schedules. You cannot add, change or delete retention types for intra-daily schedules.

When you schedule a backup, the View/Add Schedule dialog box shows the maximum number of restore points that could result from the job's current schedules and retention types. You can then change your schedules, if required. See [Maximum number of restore points for a job](#).

To reduce schedule overloads, backups that are scheduled by intra-daily schedules are skipped in some cases. See [Skipped backups](#).

To schedule a backup job to run multiple times per day:

1. Do one of the following:
 - On the navigation bar, click **Computers**. Find the computer with the backup job that you want to schedule, and click the row to expand its view. On the Jobs tab, find the job that you want to schedule. In its **Select Action** menu, click **View/Add Schedule**.
 - Create a new backup job. The View/Add Schedule dialog box appears when you save the job.
2. In the View/Add Schedule dialog box, click **Add Schedule**.
A new row appears in the dialog box.
3. In the new schedule row, click the arrow in the **Schedule** box.
IMPORTANT: To create an intra-daily schedule, you must select **Intra-daily** in the Schedule box before selecting a retention type.
4. In the Configure Job Schedule dialog box, do the following:
 - a. In the **Schedule View** list, select **Intra-daily**.



- b. In the **Every x hours** list, click the frequency for running the job. You can schedule the job to run every 1, 2, 3, 4, 6, 8 or 12 hours.
- c. In the **at y minutes past the hour** box, type the number of minutes after the hour when you want to run the job. For example, enter 15 to run the job at 15 minutes past each

hour when the job runs.

d. In the Active time period area, do one of the following:

- To run the job at the specified frequency for the full 24 hour period, click **Run backups all day**.
- To run the job according to the intra-daily schedule for only part of each 24-hour day period, click **Only run between**. Click the first clock icon and specify the start of the time period for running backups at the specified frequency. Click the second clock icon and specify the end of the time period for running backups at the specified frequency.

e. Click **Okay**.

If the job has other schedules, the intra-daily schedule appears at the bottom of the schedule list and has the lowest priority. The priority of the intra-daily backup schedule cannot be changed.

5. In the **Retention** list, click one of the following retention types:

- **24-Hours**. With this retention type, each backup is kept for at least 24 hours and at least one backup with this retention type is stored online.
- **48-Hours**. With this retention type, each backup is kept for at least 48 hours and at least one backup with this retention type is stored online.

Other retention types are not available for intra-daily schedules.

6. If the schedule is for a SQL Server Plug-in database backup job, do one of the following in the **Backup Type** list:

- To back up each database from the point in time when the backup starts, click **Full**.
- To back up each database and its transaction logs from the point in time when the backup starts, click **Full with transaction logs**.
- To back up the database transaction logs only from the point in time when the backup starts, click **Transaction logs only**. When **Transaction Logs only** is selected, the entire database and its transaction logs will be backed up when the job first runs. In subsequent backups, only the transaction logs will be backed up.

After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), be sure that only one tool is being used for truncating logs.

Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

7. If the schedule is for an Image Plug-in job that backs up volumes with SQL Server database files, do one of the following in the **SQL Application Settings** list:

- To truncate database transaction logs after the backup, select **Truncate transaction logs**.
 - To run the backup without truncating logs, clear **Truncate transaction logs**.
- Note:* If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.
8. If the schedule is for an Exchange database backup job, do the following:
 - In the **Backup Type** list, do one of the following:
 - To only back up transaction logs and the checkpoint file after the first “seed” backup, click **Incremental**.
 - To back up the database files, checkpoint file and transaction logs, click **Full**.

For more information, see [Plan Full and Incremental Exchange backups](#).
 - To validate Exchange data during the backup, select **Validate Exchange database**.
 9. In the **Schedule** box, click the arrow.
The Configure Job Schedule dialog box opens.
 10. Click **Okay**.
The new schedule appears in the Schedule box.
 11. In the **Compression** list, click a compression level for the backup data. Compression levels optimize the amount of data stored vs. the backup speed.
 12. Do one of the following:
 - To allow the backup job to run without a time limit, click **None** in the Deferring list.
 - To specify a maximum amount of time that the backup job can run, click **Minutes** or **Hours** in the **Deferring** list. 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 amount of time specified. For deferral behavior for specific backup types, see [Run and schedule backups, synchronizations and custom commands](#).

Note: Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.
 13. To run the job on the specified schedule, select the **Enable** check box near the end of the row.
 14. Check the number of restore points that could result from the job's schedules and retention policies. If you want to increase or decrease the number of restore points, change the schedules or retention types.

The maximum number of restore points appears below the schedules in the View/Add Schedule dialog box. For more information, see [Maximum number of restore points for a job](#).

15. In the Automatic Retry for Scheduled Backups section at the bottom of the View / Add Schedule dialog box, specify whether scheduled backups should retry after a failed backup. See [Specify whether scheduled backups retry after a failure](#).
16. If a Windows Event Backup Triggers section is available in the View / Add Schedule dialog box, you might be able to create a Windows event backup trigger. When a backup job has a trigger, the job runs automatically when a user logs on to the computer or the computer starts to shut down. See [Trigger backups when events occur on Windows desktop computers](#).
17. Click **Save**.

6.2.1 Skipped backups

Beginning with Windows Agent 8.90, Linux Agent 8.90, AIX Agent 9.00 and vSphere Recovery Agent (VRA) 9.11, when an agent is backing up data to a Director version 8.60 or later vault, you can schedule the backup job to run multiple times per day by creating an intra-daily schedule. See [Schedule a backup to run multiple times per day](#).

To reduce schedule overloads when a backup job runs multiple times per day, backups are skipped when:

- An agent starts a backup that is scheduled by an intra-daily schedule, and a backup is already running for the job.

Note: Windows Agent 8.72 and vSphere Recovery Agent 8.87 or later also skip a backup if it is scheduled to run multiple times per day by a custom schedule and a backup is already running for the job.

- An agent contacts a Director version 8.60 or later vault to start a backup that is scheduled by an intra-daily schedule, and the vault is busy with high-priority maintenance for the job data.

Backups are not skipped if they are scheduled to run daily or less often, or are ad hoc (not scheduled). In these cases, if a backup is already running for the job, the new backup is queued and starts when the current backup is finished. If the vault is busy with high-priority maintenance for the job data, the new backup is delayed for five minutes. After this delay, the backup starts and interrupts any maintenance that is running for the job data.

If email notifications are configured centrally in a Portal instance, Admin users can receive an email when a backup is skipped. See [Set up email notifications for backups on multiple computers](#). When the last backup status reported for a job was "Skipped", this Last Backup Status appears for the job on the Computers page and Monitor page. See [View computer and job status information](#) and [View, export and email backup statuses on the Monitor page](#). The Daily Status report also shows skipped backups. See [Schedule the Daily Status Report](#).

In some Portal instances, users can also see skipped rates and 48-hour backup status histories for jobs. See [View skipped rates and backup status histories](#).

Best practices: Reducing the number of skipped backups

If you notice that some backups are skipped frequently, you can make changes to the backup job, backup schedule, or servers to ensure reliable backups. For example, you could:

- Reduce the frequency of the scheduled backups.
- Reduce the size of the job. For a vSphere job, you can reduce the number of VMs in the job.
- On a Windows server, change from a Local System job to an Image job.
- In a vSphere environment, distribute your VMs across multiple datastores instead of using a single datastore.
- Add system resources (e.g., RAM, CPU, Storage IO) on the server where the agent is running. While the resources on a server might be sufficient for backing up and restoring data periodically, the resources might not be sufficient to run backups multiple times per day.
- Add system resources to the vault server.

6.3 Maximum number of restore points for a job

Beginning in Portal version 8.88, when you schedule a backup job, the View/Add Schedule dialog box shows the maximum number of restore points that could result from the job's current schedules and retention types. The maximum number of restore points, or backups in the vault, is updated when you add or change a schedule row so you can understand the impact of your schedule changes and make additional changes, if required.

Note: Beginning in Portal 9.30, the "Maximum number of restore points" for a job is named "Maximum number of scheduled restore points". This value does not include ad hoc backups or backups that are started by Windows event backup triggers.

For example, if you schedule a backup to run daily and select the default Monthly retention type (which specifies that each backup is kept for 365 days), the maximum number of restore points shown in the View/Add Schedule dialog box is 365. If 365 restore points would use too much vault storage, you can reduce the frequency of the backups or change the retention type. For example, you could change the retention type to the default Daily retention type, which specifies that each backup is kept for 30 days.

The maximum number of restore points includes backups created from Intra-daily, Days of Week and Days of Month schedules. The maximum number of restore points does not include restore points created using:

- Custom schedules for the job.
- Retention types that are no longer used. If a schedule was deleted or the retention for a job was changed, additional backups might remain in the vault.

For example, if a job was scheduled to run daily using the default Daily retention type, but you delete that schedule and create a new schedule using another retention type, backups from the original daily schedule plus backups from the new schedule will be saved in the vault. However, backups from the original daily schedule would not be included in the Maximum number of restore points shown in the View/Add Schedule dialog box.

Note: The Maximum number of restore points is not provided for Hyper-V Agent schedules.

6.4 Specify whether scheduled backups retry after a failure

You can specify whether scheduled backups automatically retry if they do not run successfully. These settings are only available for computers with Windows or Linux Agent version 8.60 or later, and environments with vSphere Recovery Agent version 8.80 or later.

You can also specify how many times a scheduled backup should retry after a failed attempt, and specify the amount of time between retries.

Note: Automatic retry settings only apply to scheduled backups. A backup will not retry automatically after a failed ad-hoc backup attempt.

To specify whether scheduled backups retry after a failure:

1. Do one of the following:
 - On the navigation bar, click **Computers**. Find the computer for specifying automatic retry settings, and click the row to expand its view. On the **Jobs** tab, in the **Select Action** menu for a job, click **View/Add Schedule**.
 - Create a new backup job. The View/Add Schedule dialog box appears when you save the job.
2. In the Automatic Retry for Scheduled Backups section, do one of the following:

Note: Automatic retry settings are only available for computers with Windows or Linux Agent version 8.60 or later, and environments with vSphere Recovery Agent version 8.80 or later.

- To specify that scheduled backups should not retry after failed backup attempts, clear the **Retry failed job** check box.
- To specify that scheduled backups should retry after failed backup attempts, select the **Retry failed job** check box. In the **Number of retries** box, enter the number of times the backup should try again. In the **Wait before each retry attempt for [] minutes** box, enter the number

of minutes that the agent should wait before the next backup attempt.

View / Add Schedule

SERVER.job

Retention: Daily Schedule: 7:45 PM Su,Mo,Tu,We,Th,Fr,Sa Compression: Smaller Deferring (0 for none): 0 None Enable: ☒ Priority:

Maximum number of restore points (excluding custom schedules): 30

Automatic Retry for Scheduled Backups

☐ Retry failed backup

Number of retries: 1 times

Wait before each retry attempt for: 1 minutes

3. Click **Save**.

6.5 Trigger backups when events occur on Windows desktop computers

Beginning with Windows Agent 9.30 and Portal 9.30, Image and Local System backups can be triggered by system events on computers with supported Windows desktop operating systems.

On each computer, only one Image or Local System backup job can have a Windows event backup trigger. The backup trigger can be a:

- Log On trigger, where a backup starts automatically when a user logs on to the computer.
- or
- Shut Down trigger, where a backup starts automatically when the computer starts to shut down or restart.

When you create a Shut Down trigger for a backup job on a computer, the agent automatically disables Fast Startup on the computer. If you re-enable Fast Startup on the computer, the agent will disable it. If you remove the trigger, the agent will not re-enable Fast Startup on the computer.

When a computer has a Shut Down trigger, a message box appears on the computer when it starts to shut down or restart. The message box states that a backup is in progress and the system will shut down or restart after the backup is finished. The user then has the option to cancel the backup and shut down or restart the computer immediately. If the user does not cancel the backup, a "Shutting down" or "Restarting" message remains on the screen until the backup finishes and the computer shuts down or restarts. In rare cases (e.g., a seed backup), the computer might shut down before the backup is finished.

Note: If a user is connected to the computer using a remote connection, the user will not see a message box at shut down and will not be able to cancel the backup.

For each backup trigger, you can specify whether there should be at least 12 hours or 24 hours between triggered backups and specify a retention type for the resulting backups.

Triggered backups run in addition to scheduled and ad hoc backups for a computer. However:

- If a backup is triggered when a scheduled or ad hoc backup is running for the backup job, the triggered backup does not start. If the trigger is a Shut Down trigger, the computer shuts down or restarts after the scheduled or ad hoc backup finishes.
- If a scheduled or ad hoc backup starts while a triggered backup is running for the backup job, the incoming backup is queued or, if started by an intra-daily schedule, skipped.
- If a backup is triggered when a computer shuts down, and a scheduled or ad hoc backup for another backup job on the computer is already running, the computer might shut down before the scheduled or ad hoc backup is finished.

Backups cannot be triggered on computers with Windows Server operating systems.

To trigger a backup when an event occurs on a Windows desktop computer:

1. Do one of the following:
 - On the navigation bar, click **Computers**. Find the Windows desktop computer with the Image or Local System backup job that you want to trigger, and click the row to expand its view. On the Jobs tab, find the job that you want to trigger. In its **Select Action** menu, click **View/Add Schedule**.
 - Create a new Image or Local System backup job that you want to trigger. The View/Add Schedule dialog box appears when you save the job.

IMPORTANT: Do not try to create a trigger on a computer that already has a backup trigger for a job; you will not be able to save your changes. Only one backup job on a computer can have a backup trigger.

2. In the View/Add Schedule dialog box, click **Windows Event Backup Triggers**.

If a Windows Event Backup Triggers section does not appear in the dialog box or trigger options are disabled, you cannot create a trigger for the backup job. This can occur if the Windows Agent version does not support triggers or if the computer has a Windows Server operating system.

3. In the Windows Event Backup Triggers area, select the **Run backup on Windows Event** option.
4. In the **Event** list, do one of the following:
 - To run the backup job when a user logs on to the computer, click **Log On**.
 - To run the backup job when the computer starts to shut down, click **Shut Down**.

Note: When you create a shut down trigger for a backup job on a computer, the agent disables Fast Startup on the computer. If you re-enable Fast Startup on the computer, the agent will

disable it. If you remove the trigger from the backup job, the agent will not re-enable Fast Startup on the computer.

5. In the **Retention** 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. See [Add retention types](#).

The 24-Hours and 48-Hours retention types are only available for intra-daily schedules. See [Schedule a backup to run multiple times per day](#).

6. In the **Between triggered backups, wait at least x hours** list, do one of the following:
 - To wait at least 24 hours after a triggered backup on the computer before triggering another backup, click **24**.
 - To wait at least 12 hours after a triggered backup on the computer before triggering another backup, click **12**.

The number of hours applies to any triggered backups on the computer. If you remove a trigger from one backup job and add a trigger to another backup job on the computer, the next triggered backup will not start until after the specified number of hours.

7. Click **Save**.

If an error message states that another backup job on the computer has a Windows Event trigger, click **OK**. You cannot save the new trigger or any other changes that you made in the View/Add Schedule dialog box. Before you can create a trigger for this backup job, you must remove the trigger from the other backup job on the computer.

6.6 Disable or enable all scheduled backup jobs

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

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

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

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

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


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

To enable or disable all schedules:

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

A grid lists available computers.

2. Select the check box to the left of each computer or protected environment for which you want to enable or disable all schedules.



3. In the **Actions** list, do one of the following:

- To enable all schedules for the selected computers, click **Enable Scheduled Jobs**.
- To disable all schedules for the selected computers, click **Disable Scheduled Jobs**.

6.7 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. When running an ad-hoc backup, you can back up the data to a vault or to SSI files (safeset image) on disk.

Note: The AIX operating system has a default file size limit of 1 GB. If you do not change the default file size limit, backing up a dataset over 1 GB to Directory On Disk may fail. To solve this problem, change the OS file size limit using commands such as the following:


```
ulimit -f unlimited  
  
ulimit -f -H unlimited  
  
chuser fsize=-1 fsize_hard=-1 root
```

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.



The screenshot shows a 'Run Job' dialog box with the following fields and options:

- Job: Job
- Agent: WINDOWS
- Destination: vault (selected from a dropdown)
- Retention Scheme: Daily (selected from a dropdown)
- ☒ Quick File Scanning
- ☐ Use Deferring
- Backup time window: 8 (selected from a dropdown), Hours (selected from a dropdown)
- Start Backup (green button)
- Cancel (grey button)

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

To back up the data to SSI (safeset image) files on disk, select **Directory on Disk** from the **Destination** list. Click the **Browse** button. In the Select Folder dialog box, choose the location where you want to save the SSI files, and click **Okay**.

SSI files are full backups saved to disk instead of to a vault. Saving backup files on physical media and transporting them to a remote vault for importing can be quicker than backing up data directly to a vault in a remote datacenter.

Note: Backups to SSI files on disk cannot be deferred.

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. If you are backing up a SQL Server database using the SQL Server Plug-in, do one of the following:
 - To back up the database, click **Full**. To also back up the database's transaction logs, select **Include transaction logs**.
 - To back up transaction logs only, click **Transaction Log**. When **Transaction Log** is selected, the database and its transaction logs will be backed up when the backup first runs. In subsequent backups, only the transaction logs will be backed up.

After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), be sure that only one tool is being used for truncating logs.

Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

8. If you are backing up volumes with SQL Server database files using the Image Plug-in, do one of the following:

- To truncate database transaction logs after the backup, select **Truncate transaction logs**.
- To run the backup without truncating logs, clear **Truncate transaction logs**.

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

9. If you are backing up an Exchange database, do the following:

- In the **Backup Type** list, do one of the following:
 - To only back up transaction logs and the checkpoint file after the first “seed” backup, click **Incremental**.
 - To back up the database files, checkpoint file and transaction logs, click **Full**.

For more information, see [Plan Full and Incremental Exchange backups](#).

- To validate Exchange data during the backup, select **Validate Exchange database**.

10. If the **Quick File Scanning** option is available, and you want to enable it, select the **Quick File Scanning** check box.

11. To enable Quick File Scanning, select the **Quick File Scanning** check box.

Quick File Scanning (QFS) reduces the amount of data read during the backup process. Any file streams that have not changed since the last backup are skipped. Without QFS, files are read in their entirety. Note that changes in delta-file format might cause QFS to be temporarily disabled during the first backup following an upgrade. This could cause this first backup to take longer than usual.

Note: QFS is not used for vSphere backups. Instead, the Agent uses CBT (Changed Block Tracking) to identify disk sectors that have changed.

12. 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.

Note: The **Use Deferring** check box is not available if you are backing up data to SSI (safeset image) files on disk.

Note: Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.

13. 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](#).

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

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

6.8 Plan Full and Incremental Exchange backups

When you run or schedule an Exchange database backup job, you can specify whether to run a Full or Incremental backup. In a Full backup, the database files, checkpoint file and transaction logs are backed up. In an Incremental backup, only the transaction logs and checkpoint file are backed up after the first “seed” backup.

An Incremental backup takes less time to run than a Full backup. However, the time required to recover an Exchange database increases with the number of consecutive Incremental backups. To reduce the amount of time required for a recovery, we recommend performing a Full backup periodically. For example, you could schedule an Exchange backup job to run frequently with the Incremental backup type and periodically (e.g., once per week) with the Full backup type. As shown in the following table, the appropriate backup schedule can also depend on the Exchange database size and traffic.

Exchange database description	Sample backup schedule
Low traffic – approximately 250 users, 4 GB of data, 250 MB of daily data traffic	Full backup every second Saturday night; Incremental backup on other nights
Medium traffic – approximately 1000 users, 16 GB of data, 1 GB of daily data traffic	Full backup every Saturday night; Incremental backup on other nights
High traffic – approximately 4000 users, 64 GB of data, 4 GB of daily data traffic	Full backup every Wednesday and Saturday night; Incremental backup on other nights
High traffic – approximately 4000 users, 64 GB of data, 4 GB of daily data traffic, insufficient bandwidth for large backups during the week	Full backup every Saturday night (which could be deferred to Sunday, if required); Incremental backup on other nights

You should always perform a Full backup after database repair, defragmentation or recovery. These processes significantly change Exchange databases.

Exchange maintenance can affect how much data is transferred during a Full backup. If you run daily maintenance on your Exchange server, the database will change considerably each day. When performing a Full backup, these changes are incorporated into the safeset and will result in longer Full backup times.

When scheduling backup jobs, consider the maintenance window. Backup jobs have priority over mailbox database maintenance. If a backup job runs at the same time as maintenance processes, maintenance will

be put on hold until the backup is finished. A maintenance window usually provides enough time for maintenance processes to finish after an Incremental backup, but a Full backup could prevent maintenance processes from running.

6.9 Synchronize a job

When a backup job is synchronized, the agent checks which safesets for the job are online and available for restore.

A job is synchronized automatically when you restore data from the job. You can also synchronize a job manually at any time. A manual synchronization is recommended or required in the following cases:

- Before running backup jobs on re-registered computers. You must also enter the encryption passwords for the computer's existing backup jobs. See [Restore data to a replacement computer](#).
- Before running existing backup jobs on computers that were restored using the System Restore application.
- Before restoring data from jobs that are backed up to a Satellite vault and replicated to the cloud or another vault.
- To rebuild a delta (.dta) file for a job. If an error message in a log file says that the delta mapping file is corrupt, delete the delta (.dta) file from the job folder on the protected computer and then synchronize the job to rebuild the delta file.

To synchronize a job:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer with the job that you want to synchronize. Expand its view by clicking its row.
3. Click the **Jobs** tab.
4. Find the job that you want to synchronize, and click **Synchronize** in its **Select Action** menu.

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](#).

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

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

6.10 Schedule a custom command

You can schedule custom commands to run on Windows and Linux computers. Custom commands are scripts that are saved on a computer where an Agent is installed, and are scheduled to run through Portal. For example, you could schedule a custom command that shuts down services on the Agent computer, runs a backup, and then restarts the services.

Scripts can be .bat, .exe, .cmd and executable .sh files, and must be saved in the following directory on the protected computer: *<AgentInstallationDirectory>\ScheduleScripts*


To schedule a custom command:

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

The Computers page shows registered computers.

2. Find the computer with the custom command, or script, that you want to schedule. Expand its view by clicking its row.
3. Click the **Jobs** tab.
4. In the Select Job Task menu, click Create New Custom Command.

The Create / Edit Custom Commands dialog box appears. The **Custom Command File** list shows scripts in the *<AgentInstallationDirectory>\ScheduleScripts* directory that have not been scheduled.



5. In the **Custom Command File** list, click the script that you want to schedule.

Command files in the list include .bat, .exe, .cmd and executable .sh files that have not been scheduled and are saved on the protected computer in the following directory:

<AgentInstallationDirectory>\ScheduleScripts

A default schedule for the command appears in the Schedule area.

6. Do one or more of the following:
 - To add a schedule, click **Add Schedule**, and then do one of the following:
 - To run the backup on specific days each week, click **Weekly** 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, and then click **Confirm**.
 - 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, and then click **Confirm**.
 - To create a custom schedule, click **Custom** in the **Schedule View** list. In the **Custom Cycle** box, enter a custom schedule using the syntax described in the dialog box. Click **Confirm**.
 - To edit a schedule, click the schedule's Edit Schedule button.  Change the **Schedule View**, **Run on** and **At** values, and then click **Confirm**.

- To disable a schedule, click the schedule's Edit Schedule button. Clear the **Enable** option, and then click **Confirm**.
 - To enable a schedule, click the schedule's Edit Schedule button. Select the **Enable** option, and then click **Confirm**.
 - To delete a schedule, click the schedule's Delete Schedule button.
7. To change the priority order of multiple schedules, click the **Priority** arrows to change the order of the schedule rows. Schedules higher in the list have a higher priority than schedules lower in the list.
- If a command is scheduled to run at the same time by multiple schedules, the job only runs once at the scheduled time.
8. Click **Save**.


6.10.1 Edit a custom command schedule

You can edit existing schedules for custom commands.

To edit a custom command schedule:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer with the custom command schedule that you want to edit. Expand its view by clicking its row.
3. Click the **Jobs** tab.
4. Do one of the following:
 - In the **Name** column, click the name of the custom command with the schedule that you want to edit.
 - In the **Select Action** menu of the custom command with the schedule that you want to edit, click **Edit Custom Command**.

The Create / Edit Custom Commands dialog box shows existing schedules for the custom command.



5. In the row of the schedule that you want to edit, click the Edit Schedule button.  Change the **Schedule View**, **Run on**, or **At** values, and then click **Confirm**.
6. To change the priority order of multiple schedules, click the **Priority** arrows to change the order of the schedule rows. Schedules higher in the list have a higher priority than schedules lower in the list.

If a command is scheduled to run at the same time by multiple schedules, the job only runs once at the scheduled time.

7. Click **Save**.

6.10.2 Delete a custom command

To delete a custom command schedule:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer with the custom command schedule that you want to delete. Expand its view by clicking its row.
3. Click the **Jobs** tab.
4. Do one of the following:
 - To delete all schedules or the last remaining schedule for a custom command, click **Delete Custom Command** in the command's **Select Action** menu.
 - To delete one schedule, click the name of the custom command in the **Name** column, or click **Edit Custom Command** in the command's **Select Action** menu. The Create / Edit Custom Commands dialog box shows existing schedules for the custom command. In the row of the schedule that you want to delete, click the Delete Schedule button.  Click **Save**.

Note: To delete the last remaining schedule for a custom command, you must click **Delete Custom Command** in the command's **Select Action** menu.

7 Resolve certificate failures and potential threats

Beginning in Portal 8.88, Portal indicates when agents report certificate failures. If an agent reports a certificate failure, you must investigate and resolve the certificate failure before backups and restores can continue. See [Resolve certificate failures](#).

If ransomware detection is enabled in an Image or Local System backup job, the agent checks for potential ransomware threats when running the job. If an agent detects a potential ransomware threat, you must investigate and resolve the potential threat. See [Manage potential ransomware threats](#). Ransomware detection is available in Image jobs beginning with Windows Agent 9.40 and Portal 9.40, and in Local System backup jobs beginning with Windows Agent 9.00 and Portal 8.90.

7.1 Resolve certificate failures

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

A certificate failure can occur when:

- An agent tries to connect to a Director version 8.60 or later vault where the certificate pinning security feature is enabled. Beginning with Windows Agent 8.90, Linux Agent 8.90, vSphere Recovery Agent 8.87 and Hyper-V Agent 9.00, when an agent tries to connect to the vault (e.g., to run a backup or restore), it checks whether the public key of the vault's TLS certificate is the same as when the agent previously connected to the vault. If the public key of the vault certificate is different, the agent reports a certificate failure and will not connect to the vault.
- A vSphere Recovery Agent (VRA) tries to connect to the vCenter Server or ESXi host that it protects. Beginning in version 8.87, when the VRA tries to connect to a vSphere environment, it checks whether the public key of the vSphere environment certificate is the same as when the VRA previously connected to the vSphere environment. If the public key of the vSphere environment certificate is different, the VRA reports a certificate failure and will not connect to the vCenter or ESXi host.

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

If the certificate change was expected, follow the steps below to re-pin the certificate. When you re-pin a certificate, the agent securely records the new public key of the certificate. The same procedure is used to re-pin both vault and vSphere environment certificates so that backups and restores can continue.

To resolve certificate failures:

1. On the navigation bar, click **Computers**. The Computers page shows registered computers.
2. Select the check box for each computer with a certificate failure that you want to resolve.
Note: Only select computers that have the Certificate failure status, or the Re-pin certificate action will not be available.
3. In the **Actions** list, click **Re-pin certificate**.
4. In the confirmation dialog box, click **Yes**.
5. In the Success message box, click **Okay**.

7.2 Manage potential ransomware threats

When threat detection is enabled in a backup job, the agent checks for potential ransomware threats when running the backup job. You can enable ransomware threat detection in:

- Windows Image backup jobs. Beginning with Windows Agent 9.40 with the Image Plug-in and Portal 9.40, you can enable threat detection when you create or edit an Image backup job. When this option is enabled, the agent checks for potential ransomware threats when running the backup job. See [Add a Windows Image backup job](#).
- Windows Local System backup jobs. Beginning with Windows Agent 9.00 and Portal 8.90, you can enable threat detection when you create or edit a Local System backup job. When this option is enabled, the agent checks for potential ransomware threats when running the backup job. See [Add a Windows Local System backup job](#).
- vSphere backup jobs. Beginning with vSphere Recovery Agent (VRA) 9.10 and Portal 9.10, you can enable threat detection when you create or edit a vSphere backup job. When this option is enabled, the VRA checks for potential ransomware threats on Windows VMs when running the backup job. See [Add a vSphere backup job](#).

Note: The agent does not check for potential ransomware threats in a seed backup or the first backup when threat detection is enabled in a job.

If an agent detects a potential ransomware threat, the job or backup is flagged in Portal. Potential threats are flagged:

- In the Current Snapshot on the Dashboard. See [Monitor backups and computers using the Current Snapshot](#).
- On the Computers and Monitor pages. See [View computer and job status information](#) and [View, export and email backup statuses on the Monitor page](#).
- In the Daily Status report. See [Daily Status Report](#) and [Schedule the Daily Status Report](#).
- In email notifications to Admin users, if email notifications are configured centrally in a Portal instance. See [Set up email notifications for potential ransomware threats](#).

- When you restore data or delete specific backups from a Local System backup job. See [Restore Windows files and folders](#) and [Delete specific backups from vaults](#).
- When you restore data or delete specific backups from an Image backup job. See and [Delete specific backups from vaults](#).
- When you restore data or delete specific backups from a vSphere backup job. See [Restore vSphere data](#) and [Delete specific backups from vaults](#).

If a Windows server has a potential threat, the server is not scanned for threats again until the potential threat warning is cleared for the job.

If a VM has a potential threat, the VRA does not scan the VM again during backups until the potential threat warning is cleared for the job. If a VM has a potential threat but is missing from the vSphere environment during the next backup, the backup will still have a potential threat flag until an Admin user clears the potential threat warning.

When a potential threat is detected on a Windows server or VM, you can sign in to the server or VM in your environment and investigate whether it is infected with ransomware. An Admin user in Portal can then manage the threat:

- If the server or VM is not infected or the ransomware threat has been addressed, an Admin user can clear the potential threat warning from the job.
- If the server or VM is infected with ransomware, an Admin user can restore from a backup (also known as safeset) created before the attack. Backups with potential threats are identified in the Restore dialog box so you can choose a backup with no potential threat. After the restore, backups with potential ransomware threats remain in the vault and available for restore. To remove these backups (safesets), delete them from the vault and synchronize the job. See [Delete specific backups from vaults](#) and [Synchronize a job](#). An Admin user can then clear the potential threat flag from the job.

For an overview of the recommended process in VMware vSphere environments, see [Best practices: Manage ransomware threats on vSphere VMs](#).

To manage a potential ransomware threat:

1. When signed in to Portal as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Find the computer or environment with the potential threat, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. Find the job with the potential threat, and click **Manage Potential Threat** in its **Select Action** menu.

Note: The Manage Potential Threat option does not appear for a job that is restored from another computer. To manage a potential threat for a job, you must find the job on the original computer,

if it exists, or re-register a new computer to the vault as the original computer. See [Restore data to a replacement computer](#).

5. In the Manage Potential Threat box, do one of the following:

- To restore from a backup before the potential ransomware threat was detected, select **Recover** and then click **Continue**.
 - If you are restoring from a Windows Local System backup job, a calendar with a list of backups appears in the Restore dialog box. "Potential Threat" appears beside each backup where a potential ransomware threat was detected. Select the backup (also known as safeset) from which you want to restore files, select restore options, and then click **Run Restore**. See [Restore Windows files and folders](#).
 - If you are restoring from an Image backup job, restore options appear in the Image Restore dialog box. You can restore entire data volumes or individual files and folders. See [Restore Windows volumes from an Image backup](#) and [Restore files and folders from an Image backup](#).
 - If you are restoring from a vSphere backup job, restore options appear in the vSphere Restore dialog box. You can restore entire VMs, restore a VM within minutes, or restore files and folders to a VM. See [Restore vSphere data](#).

Note: After a restore, backups with potential ransomware threats remain in the vault and available for restore. To remove these backups, delete them from the vault and synchronize the job. See [Delete specific backups from vaults](#) and [Synchronize a job](#). An Admin user can clear the potential threat flag from the job.

- If you investigated or addressed the potential threat and are sure that the server or VM is not affected by ransomware, select **Clear Potential Threat Warning** and then click **Continue**. In the warning dialog box, click **Continue** to remove the potential threat flag from the job and all of its backups (safesets).

Note: Clearing potential threat warnings will clear all existing threat warnings from the job and its backups (safesets). However, warning information will still be available in the log files.

7.2.1 Best practices: Manage ransomware threats on vSphere VMs

If a potential ransomware threat is detected during a vSphere backup, you can view the backup log to see which VM or VMs might have a potential threat. You can then sign in to each VM that has a potential threat to investigate whether it is infected with ransomware.

If VMs in the backup job are not infected with ransomware, clear the potential threat warning from the backup job. See [Manage potential ransomware threats](#).

If one or more VMs in the backup job are infected with ransomware, we recommend the following:

1. Delete each infected VM from your vSphere environment.
2. Restore each infected VM from a backup that was created before the VM was infected with ransomware. During a restore, the Restore dialog box shows which safesets and VMs have potential ransomware threats. See [Restore vSphere VMs](#).

If you deleted the infected VM from the vSphere environment before restoring it (as recommended in [Step 1](#)), the restored VM replaces the deleted VM and you do not need to add the restored VM to a backup job. If you did not delete the infected VM from the vSphere environment, the restored VM will have a new name and will not be included in a backup job unless you add it.

3. Delete the backup (safeset) with one or more infected VMs from the vault so VMs with potential threats cannot be restored. When you delete a safeset, the Delete Backup dialog box shows which safesets have potential ransomware threats. See [Delete specific backups from vaults](#).
4. Clear the potential threat warning from the backup job. See [Manage potential ransomware threats](#).

If you do not clear the potential threat warning from a job, VMs in the backup job will not be scanned for ransomware in subsequent backups but will still be flagged as having a potential ransomware threat.

5. Synchronize the backup job with the vault. See [Synchronize a job](#).

8 Restore data

This section describes how to restore data on Windows, Linux and UNIX servers, and in vSphere and Hyper-V environments.

You can restore data to the computer where you backed up the data or to a different computer.

To restore some or all of a computer's backed up data to another computer without replacing the original computer, see [Restore data from another computer](#).

To register a new computer with the vault as if it were the old computer (i.e., re-register), see [Restore data to a replacement computer](#). Re-registering a computer can be useful if you are replacing a system and want to migrate all data to a new computer (e.g., at the end of a lease or if hardware is failing).

8.1 Restore Windows data

After backing up Windows servers, you can:

- [Restore Windows files and folders](#)
- [Restore files from multiple UNC jobs](#)
- [Restore Windows volumes from an Image backup](#)
- [Restore files and folders from an Image backup](#)
- [Recover a Windows cluster](#)

You can also use the System Restore application to restore an entire system from a Bare Metal Restore (BMR) backup. A BMR backup includes the operating system, applications, system state and data. For more information, see [Add a Windows Local System backup job](#) and the *System Restore Guide*.

Note: Although a BMR backup includes a computer's system state, you can only restore the system state from a BMR backup when you restore the entire computer using the System Restore application.


8.1.1 Restore Windows files and folders

After backing up data from a Windows computer, you can restore files and folders from the backup.


To restore Windows files and folders:

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

The Restore dialog box appears. If the job does not have a potential ransomware threat, the most recent safeset for the job appears in the Safeset box.



If a potential ransomware threat was detected when running the job, a calendar with a list of backups appears. "Potential Threat" appears beside each backup where a potential ransomware threat was detected.



5. To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:

- To restore data from an older safeset, if a calendar with a list of backups does not already appear, click the calendar button. In the calendar, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
- To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button. In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

6. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
7. Select a Restore Destination option.
 - To restore files and folders to the location where they were backed up, select **Restore files to their original location**.
 - To restore files and folders to a different location, select **Restore files to an alternate location**. Click the folder button.  In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.
8. Select an Overwrite Option. This option specifies how to restore a file, folder or symbolic link to a location where there is a file, folder or symbolic link with the same name.
 - To overwrite the existing item with the restored item, select **Overwrite existing items**.

Note: If you try to restore multiple files with the same name to an alternate location and select **Overwrite existing items**, only the last file restored will remain. Other files with the same name will be overwritten.

IMPORTANT: Using Agent version 8.70 or later, if you select **Overwrite existing items** and restore a file that has the same name as a folder in the restore location, the file will overwrite the folder. The folder and all of its contents will be removed.
 - To skip restoring the item that has the same name as an item in the destination location, select **Do not overwrite existing items**.
 - To add a numeric extension (e.g., .0001) to the **restored** item name, select **Rename incoming items**. For example, if you restore a file named “filename.txt” to a location where there is a file with the same name, an extension is added to the **restored** file name (e.g., “filename.txt.0001”).
 - To add a numeric extension (e.g., .0001) to the **existing** item name, select **Rename existing items**. For example, if you restore a file named “filename.txt” to a location where there is a file with the same name, an extension is added to the **existing** file name (e.g., “filename.txt.0001”). The name of the restored file is “filename.txt”.
9. To change the locked file, data streams, log detail level or bandwidth options, click **Advanced Restore Options**. Specify settings in the Advanced Restore Options dialog box, and click **Okay**. See [Advanced restore options](#).

10. In the **Select Items** box, do one or more of the following until the **Restore Set** box shows the folders and files that you want to restore:
 - Select the check box for each folder and file that you want to restore, and then click **Include**. The **Restore Set** box shows the included folders and files. If you include a folder, all of the folder's subdirectories and files are restored by default. If you do not want to restore all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).
 - To exclude one or more folders or files from the restore, select the check box for each folder or file, and then click **Exclude**. The **Restore Set** box shows the excluded folders and files. If you exclude a folder, all of the folder's subdirectories and files are excluded from the restore by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).
 - To search for files to restore or exclude from the restore, click the **Search** button.  In the **Search for files** box, enter search criteria and select files. See [Search for files to restore](#). Click **Include Selected** or **Exclude Selected**. The **Restore Set** box shows the included or excluded files.
 - To remove an inclusion or exclusion record from the **Restore Set** box, click the Delete button beside the folder or file record. 
11. Click **Run Restore**.

If you are restoring from a safeset that has a potential threat, a warning message appears. To cancel the restore, click **Cancel**. To continue restoring the volume, click **Continue**.

The Process Details dialog box shows the restore progress and indicates when the restore is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).
12. To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.1.1.1 Restore NTFS hard links, symbolic links, mount points and junctions

When you restore files to their original locations and overwrite existing files, NTFS hard links, symbolic links, mount points and junctions are preserved. If you restore files to alternate locations or do not overwrite existing files, the links break.

Note: Remote hard links and mount points (e.g., UNC paths) are not supported.

When you restore junctions to their original locations, all link functionality is preserved. If you restore to an alternate location, the junction will revert to an empty directory. To recover to an alternate location, the junction must be explicitly selected for backup and will duplicate the contents of its target directory without preserving junction functionality.

Note: Remote/alternate junctions are not supported.

8.1.1.2 Restore a domain controller

You can restore a domain controller if the system was fully backed up using system state and system volume backups, or a Bare Metal Restore (BMR) backup. You can also restore a domain controller from a Bare Metal Restore (BMR) backup using the System Restore application.

Note: A Windows Agent or BMR backup is not sufficient for an authoritative restore of Active Directory objects. For an authoritative restore, a System State backup with the Windows Agent is required. For more information, see [Best-Practices-Backing-up-Domain-Controllers-or-Active-Directory](#).

When you have more than one domain controller, you must decide whether to perform an authoritative or non-authoritative restore before restarting the machine. For more information, see documentation from Microsoft.

8.1.2 Restore files from multiple UNC jobs

When a Windows computer has more than one UNC backup job, you can search for and restore files from multiple UNC jobs at the same time. This functionality is available for computers where Windows Agent version 8.0 or later is installed.

You can restore files from UNC jobs to a local folder or to a UNC share. When you restore files to a UNC share, files are only restored from UNC jobs with credentials that have access to the share. If required, you can change the credentials in a UNC backup job until you have restored the files.

To restore files from multiple UNC jobs:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find a Windows computer with multiple UNC jobs, and expand its view by clicking the computer row.
3. In the **Select Job Task** menu, click **Restore from multiple UNC jobs**.

The Restore from multiple UNC jobs dialog box appears.

Restore from multiple UNC jobs

Search pattern (ex. *.doc):

Date: 11/11/2015

Search

Your inclusions will be preserved across multiple searches.

Search Results

Filename	Directory	Size	Job	Modified	Backup Date

Confirm options and files to restore

Cancel

4. In the **Search Text** field, enter some or all of the name of a file that you want to restore. Use asterisks (*) and question marks (?) as wildcard characters. For example, to find all files with the .pdf extension, enter the following: *.pdf

5. Do one of the following:


- To search for files in the most recent safeset for each UNC job, leave the **Date** box blank.
- To search for files in safesets with a specific date, click the calendar button.  In the calendar that appears, click the date.

If a job does not have a safeset for the specified date, the system searches for files in the safeset with the date that is closest to and before the specified date.

If a job includes multiple safesets for the specified date, the system searches for files in the last safeset on the date.

6. Click **Search**.

Files found in the safesets are listed in the lower part of the dialog box.



Restore from multiple UNC jobs

Search pattern (ex. *.doc):

Date: 

Search

Searches the safeset on or closest before the date specified. Date is in agent's time zone. Most recent if left blank.

 Your inclusions will be preserved across multiple searches.

Search Results

Filename	Directory	Size	Job	Modified	Backup Date	
Test.txt	0172.16.144.63\Test2\	0.00 Bytes	Test_UNC_Search	08/25/2015	10/15/2015	
Test1.txt	010.S1.3.52\saumari\	5.00 Bytes	unc	01/23/2015	08/25/2015	
Test2 (1).txt	010.S1.3.52\saumari\	5.00 Bytes	unc	01/23/2015	08/25/2015	
Test2.txt	010.S1.3.52\saumari\	5.00 Bytes	unc	01/23/2015	08/25/2015	


Confirm options and files to restore

Cancel

7. Do one of the following:

- To restore specific files, click the **Include for restore** button for each file. 
- To restore all files in the list, click the **Include All** button at the top of the file list. 

After a file is selected for restore, the file is highlighted in the list.




Search Results

Filename	Directory	Size	Job	Modified	Backup Date	
Test.txt	0172.16.144.63\Test2\	0.00 Bytes	Test	08/25/2015	10/15/2015	
Test1.txt	010.S1.3.52\saumari\	5.00 Bytes	unc	01/23/2015	08/25/2015	
Test2 (1).txt	010.S1.3.52\saumari\	5.00 Bytes	unc	01/23/2015	08/25/2015	
Test2.txt	010.S1.3.52\saumari\	5.00 Bytes	unc	01/23/2015	08/25/2015	

Files selected for restore

8. To search for more files to restore, repeat Steps 4 to 7.
9. To view all files that are selected for restore, click **Confirm options and files to restore**.

The Restore from multiple UNC jobs dialog box shows files that are selected for restore.



10. To select more files to restore, click **Include more files**. The Restore from multiple UNC jobs dialog box returns. Repeat Steps 4 to 9.
11. To restore the selected files, do the following:

- a. In the **Target Directory** box, do one of the following:

- To select a local folder as the restore destination, click the **Browse** button.  In the Select Folder dialog box, choose the folder and then click **Okay**.
- To specify a local folder or UNC share as the restore destination, type the name of the folder or UNC share (e.g., \\server\\share) where you want to restore files.

If the destination is a UNC share, files will only be restored from jobs with credentials that have access to the share. If required, you can change credentials in a UNC backup job until you have restored the files.

If the restore destination folder does not exist, it will be created during the restore.

- b. In the **Encryption Password** box, enter the encryption password.

If you are restoring files from UNC jobs with different encryption passwords, files will only be restored from jobs with the password specified in this box.

- c. In the **Log Level** list, click the level of detail for logging. See [Advanced restore options](#).
- d. To restrict the amount of bandwidth used, select the **Enable Bandwidth Throttling** check box. See [Advanced restore options](#).
- e. Click **Restore**.

The selected files are restored. If you restore a file with the same name as another file in the same location, a numeric extension (e.g., .0001) is added to the restored file name (e.g., filename.txt.0001).

8.1.3 Restore Windows volumes from an Image backup

After backing up volumes on a Windows computer using the Image Plug-in, you can restore volumes from the backup to selected live volumes (target volumes).

IMPORTANT: When you restore a volume from a backup, any data on the target volume will be lost.

A target volume must meet the following requirements:

- The target volume must be as large as or larger than the volume that was backed up.
- The Windows operating system cannot be installed on the target volume.
- The Windows Agent cannot be installed on the target volume.

To restore Windows volumes from an Image backup:

IMPORTANT: Before restoring volumes from an Image backup, stop any services on the system that are using the target volume (e.g., SQL Server or Exchange services). Restart the services after the restore is completed.


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

A grid lists available computers.


2. Find the Windows computer with the Image Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the Image backup job with volumes that you want to restore, and click **Restore** in the job's **Select Action** menu.

You can restore volumes from regular or Bare Metal Restore backups created using the Image Plug-in.

5. In the Image Restore dialog box, select **Entire Data Volume(s)**, and then click **Next - Configure Source**.



The Restore dialog box shows volumes that can be restored from the backup. If a potential ransomware threat was not detected in the job, the most recent safeset for the job appears in the **Safeset** box.



If a potential ransomware threat was detected when running the job, a calendar with a list of backups appears. "Potential Threat" appears beside each backup where a potential ransomware threat was detected.

6. To restore data from an older safeset or from SSI (safeset image) files, do one of the following:
 - To restore data from an older safeset, if a calendar with a list of backups does not already appear, click the calendar button.  In the calendar, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
 - To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

7. In the **Encryption Password** box, enter the encryption password. To view the password hint, click the **Hint** button. 
8. To change the log detail level or bandwidth options, click **Advanced Restore Options**. Specify settings in the Advanced Restore Options dialog box, and click **Okay**. See [Advanced restore](#)

[options](#).

9. In the **Volume to Restore** column, select each volume that you want to restore.
10. For each selected volume, do the following to choose the live volume where it will be restored:

IMPORTANT: Data on the selected volume will be lost when the backed-up volume is restored.

- a. Click the folder button. 

The Select Volume dialog box lists all live volumes on the computer. If you cannot restore the selected volume to a specific live volume (e.g., because the live volume is too small, contains the Windows operating system, or contains Windows Agent software), the volume is unavailable and cannot be clicked.

- b. Click the volume where you want to restore the backed up volume.
- c. Click **Okay**.

11. Click **Run Restore**.

If you are restoring from a safeset that has a potential threat, a warning message appears. To cancel the restore, click **Cancel**. To continue restoring the volume, click **Continue**.

The Process Details dialog box shows the restore progress and indicates when the restore is completed. The target volume goes offline until the backed-up volume is restored.

Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).

12. To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.1.4 Restore files and folders from an Image backup

After backing up volumes from a Windows server using the Image Plug-in, you can restore individual files and folders from the backup.

To restore files and folders from an Image backup, you mount volumes from a safeset as drives on the computer where you want to restore files and folders. You can then browse the drives using Windows Explorer, and copy files and folders that you want to restore.

Note: To restore files and folders from Image backups, agent services must be running using the local system account. If the local system account did not have full permission to the files and folders during backup, you might not have permission to access files and folders during the restore. In this case, you can either:


- Restore the whole volume and grant appropriate permissions.
- Restore on a computer in the same domain using a domain account that has appropriate permissions.

Note: When SQL Server databases are backed up using Image Plug-in version 7.5 or later, you can also restore database files, tables and objects from the application-consistent backups.

To restore files and folders from an Image backup:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the Windows computer with the Image Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the Image backup job with files and folders to restore, and click **Restore** in the job's **Select Action** menu.
You can restore files and folders from any Image Plug-in job.
5. In the Image Restore dialog box, select **Individual File or Folders**, and then click **Next - Configure Source**.

The Restore dialog box shows volumes that you can mount as drives. If a potential ransomware threat was not detected in the job, the most recent safeset for the job appears in the **Safeset** box.



If a potential ransomware threat was detected when running the job, a calendar with a list of backups appears. "Potential Threat" appears beside each backup where a potential ransomware threat was detected.


6. To restore data from an older safeset, if a calendar with a list of backups does not already appear, click the calendar button.  In the calendar, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
7. In the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will be unmounted automatically. The **Idle Time** can range from 2 to 180 minutes.
8. In the **Encryption Password** box, enter the encryption password. To view the password hint, click the **Hint** button. 

9. In the **Volume to Restore** column, select each backed-up volume from which you want to restore files and folders.
10. In the **Volume Mount Drive Letter** column, choose the drive letter for mounting each volume.
11. Click **Mount Volumes**.

If you are restoring from a safeset that has a potential threat, a warning message appears. To cancel the restore, click **Cancel**. To continue restoring the volume, click **Continue**.

If a confirmation message appears, read the message, and then click **Continue**.

The Process Details dialog box appears. When each volume is mounted, the drive letter is shown at the right side of the dialog box.



12. Use Windows Explorer to navigate to the drive or drives and copy files and folders that you want to restore.
13. When you are done restoring files and folders, click **Stop** to remove the mounted drives.

If you do not click **Stop**, the drive will be unmounted automatically after the number of minutes of inactivity specified in the Idle Time box. See Step 7.

8.2 Recover a Windows cluster

When a Windows cluster is protected as described in [Add backup jobs for a Windows cluster](#), you can recover the cluster if components are lost, are corrupted or fail. The following table indicates how to recover a cluster after encountering specific issues.

Issue	Recovery Process	Jobs Used
Cluster disk data loss, corruption or failure	Restore volumes on the cluster disk. If the cluster disk failed or was corrupted, clean partition and volume formatting from the disk before restoring the data. See Recover volumes in a Windows cluster .	On the virtual server for each cluster role (e.g., file server or SQL Server role), an Image or local system job that backs up cluster disks for the role. See Job C in Add backup jobs for a Windows cluster .

Cluster quorum corruption, checkpoint loss, failure or rollback required	Create a new quorum disk. See Recover the quorum disk in a Windows cluster .	On the virtual server for the cluster core, an Image or local system job that backs up the quorum disk. See Job A in Add backup jobs for a Windows cluster .
Cluster node corruption or failure	Recover the cluster node using the System Restore application. See Recover a node in a Windows cluster .	On the cluster node, a Bare Metal Restore (BMR) backup job created using the Image Plug-in or Windows Agent. See Job B in Add backup jobs for a Windows cluster .
Complete cluster failure	Recover all components of the cluster. See Recover an entire Windows cluster .	Jobs B, A and C in Add backup jobs for a Windows cluster . In addition, for a SQL Server cluster, a SQL Server Plug-in job is required for point-in-time database recovery. The job is created on the virtual server for the SQL Server role. See Job D in Add backup jobs for a Windows cluster .


8.2.1 Recover volumes in a Windows cluster

If data has been lost from a cluster disk, or a cluster disk has become corrupted or failed, you can recover the cluster volumes.

Cluster volumes must be backed up using an Image or local system job on the virtual server for a cluster role (e.g., file server or SQL Server). See Job C in [Add backup jobs for a Windows cluster](#).

To recover volumes in a Windows cluster:

- If you are recovering volumes to a disk that became corrupted or failed, do the following:
 - Remove the disk from the cluster.
 - Clean partition and volume formatting from the disk using a tool such as diskpart.
 - Add the disk back to the cluster.
- Using the Failover Cluster Manager on any cluster node, stop the cluster resources. **Do not** stop the shared disk resource.



3. Using Portal, run a “Restore from another computer” on the cluster node where the disk is mounted. Restore the cluster volume or volumes from an Image or local system backup job on the virtual server for the SQL Server role (Job C in [Add backup jobs for a Windows cluster](#)). Restore volumes to their original locations.
4. Using the Failover Cluster Manager on any cluster node, start the SQL Server and SQL Server Agent cluster resources.
5. Using SQL Management Studio, ensure that the SQL Server database is running and operational.


8.2.2 Recover the quorum disk in a Windows cluster

If the quorum disk in a Windows cluster is corrupted or fails, or if a rollback is required, you can create a new quorum disk and recover any required data.

Quorum disks are protected using an Image or local system job on the virtual server for the cluster core. See Job A in [Add backup jobs for a Windows cluster](#).

To recover the quorum disk in a Windows cluster:

1. Connect a new disk to the cluster.
2. On one cluster node only, bring the disk online and initialize it. Partition the disk and assign the same drive letter that was previously assigned to the quorum disk.
3. In the Failover Cluster Manager on any cluster node, click the cluster name. Under **Cluster Core Resources**, right-click the quorum disk, click **More Actions** and then click **Repair**. Choose the newly formatted disk and click **OK**. Wait until the quorum disk is successfully repaired.



4. Bring the quorum disk online.
5. Ensure that the quorum disk is online and that there are no errors in the cluster logs.
6. If required, restore quorum data from the Image or local system job on the virtual server for the cluster core (Job A in [Add backup jobs for a Windows cluster](#)).

8.2.3 Recover a node in a Windows cluster

If a node in a Windows cluster is corrupted or fails, you can recover the node.

Each cluster node must be backed up using an Image or Windows Agent Bare Metal Restore (BMR) job on the node. See Job B in [Add backup jobs for a Windows cluster](#).

To recover a node in a Windows cluster:

1. On a replacement machine that has similar hardware to the original cluster node, launch the System Restore application.

For detailed procedures and information, see the *System Restore User Guide*.

2. Restore system volumes from a BMR backup of the original cluster node (Job B in [Add backup jobs for a Windows cluster](#)).
3. If required at the end of the restore, repair drivers on the system.
4. Reboot and then log in to the replacement node.
5. Configure network adaptors on the replacement node with the same network settings as the original node (i.e., same IP addresses and DNS entries).
6. If required, re-connect all cluster disks.
7. Open the Failover Cluster Manager and ensure that the restored node is up and running.
8. Fail over the SQL Server and SQL Server Agent cluster resources to the restored node to ensure that the restored node is functioning correctly.

8.2.4 Recover an entire Windows cluster

If all components of a Windows cluster fail and the cluster is protected as described in [Add backup jobs for a Windows cluster](#), you can recover the cluster.

To recover an entire Windows cluster:

1. Recover one cluster node by doing the following:
 - a. On a replacement machine that has similar hardware to one of the original cluster nodes, launch the System Restore application.
For detailed procedures and information, see the *System Restore User Guide*.
 - b. Restore system volumes from a BMR backup of the original cluster node (Job B in [Add backup jobs for a Windows cluster](#)).
 - c. If required at the end of the restore, repair drivers on the system.
 - d. Reboot and then log in to the replacement node.
 - e. Configure network adaptors on the replacement node with the same network settings as the original node (i.e., same IP addresses and DNS entries).
2. On the restored cluster node, recreate the cluster disks. Format the disks and assign the original drive letters.

3. On the restored cluster node, stop the Cluster service. For a SQL Server cluster, also stop the SQL Server service.
4. If required, restore quorum data to its original location. Using Portal, run a “Restore from another computer” on the restored cluster node. Restore the quorum disk from an Image or local system backup job on the virtual server for the cluster core (Job A in [Add backup jobs for a Windows cluster](#)).
5. Restore cluster volumes to their original locations. Using Portal, run a “Restore from another computer” on the restored cluster node for each cluster role. Restore cluster volumes from an Image or local system backup job on the virtual server for each cluster role (Job C in [Add backup jobs for a Windows cluster](#)).
6. Start the cluster service.
7. Using the Failover Cluster Manager, connect to the cluster.
8. Start the cluster roles.
9. Repair the cluster disks and assign the original drive letters. Bring the cluster disks online.
For more information about repairing cluster disks and bringing them online, see documentation from Microsoft.
10. For a SQL Server cluster, use Portal to run a “Restore from another computer” on the restored cluster node. Restore SQL Server databases from a SQL Server Plug-in job on the virtual server for the SQL Server role (Job D in [Add backup jobs for a Windows cluster](#)). Restore the databases to their original locations.
11. For each remaining cluster node, do the following:
 - a. On a replacement machine that has similar hardware to the original cluster node, launch the System Restore application.
For detailed procedures and information, see the *System Restore User Guide*.
 - b. Restore system volumes from the BMR backup of the original cluster node (Job B in [Add backup jobs for a Windows cluster](#)).
 - c. If required at the end of the restore, repair drivers on the system.
 - d. Reboot the machine.
 - e. Log in to the machine.
 - f. Configure network adaptors with the same network settings as the original node (i.e., same IP addresses and DNS entries).
 - g. If required, reconnect all cluster disks.

8.3 Restore databases and application data

After backing up data from a specific application, you can restore data. You can:

- [Restore Exchange databases](#)
- [Restore Exchange mailboxes, messages and other objects](#)
- [Restore SQL Server databases](#)
- [Restore items from a SQL Server or SharePoint database](#)
- [Restore Oracle databases](#)

8.3.1 Restore Exchange databases

You can restore a Microsoft Exchange database to its original location or to an alternate Exchange database (e.g., a recovery database). To overwrite an existing database, the database must be unmounted and marked for overwrite.

When restoring an Exchange database, you can specify whether or not to replay transaction logs into the database and mount the database in Exchange. If this option is selected, the logs are rolled forward if they are in the original directory and no log files are missing or corrupt. If this option is not selected, transaction logs are restored to the system but are not replayed into the restored database. The Administrator must review the restored files and manually mount the database.

The process of restoring the database files to the system is recorded in the job logs. The process of replaying transaction logs into the database is recorded in the Windows Event Viewer.

For more information about Exchange restore strategies, see [Exchange database restores](#) and documentation from Microsoft.

To restore Exchange databases to flat files, see [Restore Exchange databases to flat files](#).

To restore an Exchange database:

1. If you are restoring the database to its original location, delete all files (e.g., .edb, .log and .chk) in the folder where the original database files are located.

Note: Do not delete the folder that has a GUID in its name (e.g., 523E7980-EA56-440-9ACB-AF0AE2CF1F0212...).

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

A grid lists available computers.

3. Find the computer with the Exchange database you want to restore, and expand its view by clicking the computer row.

4. Click the **Jobs** tab.

5. Find the job whose Exchange database you want to restore, and click **Restore** in the **Select Action** menu for the job.

6. In the Choose What You Want to Restore dialog box, select **Exchange Databases**, and click **Okay**.

The Restore dialog box shows the most recent safeset for the job.

7. To restore data from an older safeset or from SSI (safeset image) files, do one of the following:

- To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
- To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

8. In the **Items to Restore** box, select the check box for each database that you want to restore.
9. Select a Restore Destination option:
 - To restore data to the location where it was backed up, select **Restore files to their original location**.
 - To restore to an alternate Exchange database, select **Restore to an alternate Exchange database**. In the **Current Chosen Destination** box, click **Browse**. In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.
10. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
11. To change the log detail level, bandwidth throttling setting or hard recovery option, click **Advanced Restore Options**. In the dialog box, do one or more of the following:
 - In the **Log Level Detail** list, select the level of detail for job logging. See [Advanced restore options](#).
 - Select or clear the **Use all available bandwidth** option. See [Advanced restore options](#).
 - Select or clear the **Start Hard Recovery** option.

If the **Start Hard Recovery** option is selected, transaction logs are replayed after the database is restored and the restored database is mounted by Exchange. Logs are also rolled forward if they are in the original directory and no log files are missing or corrupt.

If the **Start Hard Recovery** option is not selected, transaction logs are restored to the system but are not replayed after the database is restored. The restored database is not mounted by Exchange. The Administrator must review the restored Exchange files and manually mount the database.

Click **Okay**.

12. Click **Run Restore**.

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

To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.3.1.1 Exchange database restores

You can only restore databases to the active copy in a Database Availability Group. If you restore to the replica copy, you will not be able to mount the database or make it active. You will need to copy the restored files to the active copy node in order to successfully mount (and take precedence over the other copies). You will also need to update each copy through the Exchange Management Console. For more information, see documentation from Microsoft.

If you are restoring to a new location and you want to mount the database, you must first create a recovery database through the Exchange Management Shell. For more information, see documentation from Microsoft.

To restore an Exchange database to an alternate location on a non-Exchange server, you must clear the **Start Hard Recovery** option.

If a database's transaction log files are missing or damaged, an incremental backup after the recovery will not succeed. Perform a full backup before you attempt another incremental backup.

8.3.1.2 Restore Exchange databases to flat files

On a computer where the Exchange Plug-in is installed, you can restore an Exchange database to flat files. The Eseutil utility can then be used to bring the data into a database.

To restore an Exchange database to flat files:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with the Exchange database that you want to restore, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job with the Exchange database that you want to restore, and click **Restore** in the **Select Action** menu for the job.
5. In the Choose What You Want to Restore dialog box, select **Restore to folder**, and click **Okay**.
The Restore dialog box shows the most recent safeset for the job.
6. To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:

- To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
- To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

7. In the **Files to Restore** box, select the check box for each database that you want to restore.
8. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the Hint button. 
9. Under Restore Destination, enter a path for the destination, or click the folder button.  In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.
10. To change the log detail level, bandwidth options or hard recovery option, click **Advanced Restore Options**. In the dialog box, do one or more of the following:
 - In the **Log Level Detail** list, select the level of detail for job logging. See [Advanced restore options](#).
 - Select or clear the **Use all available bandwidth** option. See [Advanced restore options](#).
 - Select or clear the **Start Hard Recovery** option.

If the Start Hard Recovery option is selected, transaction logs are replayed after backup data is restored. The restored database is prepared for use by Exchange, and logs are rolled forward if they are in the original directory and no log files are missing or corrupt. These processes are recorded in the Windows Event viewer.

If the Start Hard Recovery option is not selected, the database will not be available to Exchange after a restore. The Administrator must review the restored Exchange files and database and manually mount the database. For more information, see documentation from Microsoft.

Click **Okay**.

11. Click **Run Restore**.

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

To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.3.2 Restore Exchange mailboxes, messages and other objects

If a Microsoft Exchange database is backed up using the Exchange Plug-in, you can restore mailboxes, messages, and other objects from the backup.

To restore items from a Microsoft Exchange backup, you must first use Portal to expose the Exchange safeset as a shared resource. You can then use the Granular Restore for Microsoft Exchange and SQL application to find and restore mailboxes, messages and other Exchange objects.

For more information, or to obtain the Granular Restore for Microsoft Exchange and SQL application, contact your service provider.

To restore Exchange mailboxes, messages and other objects:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with Exchange objects that you want to restore, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job with Exchange objects (e.g., messages) that you want to restore, and click **Restore** in the **Select Action** menu for the job.

The Choose What You Want to Restore dialog box appears.

5. Select **Mailboxes, messages and other Exchange objects**, and click **Okay**.
6. In the Restore dialog box, choose a safeset from which to restore.
7. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
8. In the **Idle Time** box, enter the number of minutes of inactivity after which the share should automatically stop. The value can range from 2 to 180.
9. Select or clear the **Use all available bandwidth** option.
10. Click **Share**.

The Process Details dialog box shows the status of the share process. When the share is available, the share path appears at the right side of the dialog box.

11. Click the **Copy Path to Clipboard** button.  The path to the safeset share is now available for you to paste into the Granular Restore application.

12. Launch the Granular Restore for Microsoft Exchange and SQL application. Paste the path to the Exchange safeset share into the Granular Restore application, and then select and restore your data. For more information, see the *Granular Restore for Microsoft Exchange and SQL User Guide*.

13. When you no longer need the safeset share, click **Stop**.

When you click **Stop** or the share idle time is reached, the **Current Process** information indicates that the share is no longer available.

8.3.3 Restore SQL Server databases

After backing up SQL Server databases using the SQL Server Plug-in, you can restore databases directly to a SQL Server instance, or restore databases to flat files. See [Restore databases directly to SQL Server](#), [Restore SQL Server databases to files](#) or [Restore a SQL Server master database](#).

After backing up SQL Server databases using the Image Plug-in, you can restore database files from the safeset. See [Restore SQL Server database files or objects from Image backups](#).

When restoring a SQL Server database in an Always On Availability Group, you must always restore the database to the primary replica. See [Restore databases in AlwaysOn Availability Groups](#).

8.3.3.1 Restore databases directly to SQL Server

After backing up SQL Server databases using the SQL Server Plug-in, you can restore databases directly to a SQL Server instance.

Note: If a database was backed up using the Image Plug-in, you must restore it using the Image Plug-in.

If transaction logs have been backed up using an alternative method (e.g., native SQL Server backup), you can restore a database in the restoring state so that you can apply transaction logs to the database after the restore.

You must specify a Windows or SQL Server administrator account for connecting to SQL Server during a restore.

After restoring a SQL Server 2016 database that is stretched to Microsoft Azure, you must run a stored procedure ([sys.sp_rda_reauthorize_db](#)) to reconnect the local restored database to the remote Azure data. See “Restore the connection between the SQL Server database and the remote Azure database” on the Microsoft Developer Network website: <https://msdn.microsoft.com/en-us/library/mt733205.aspx#reconnect>

To restore a database directly to SQL Server:

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

A grid lists available computers.

2. Find the computer with the SQL Server database backup that you want to restore, and expand its view by clicking the row for the computer.
3. Click the **Jobs** tab.

4. Find the job with the database that you want to restore. In the job's **Select Action** menu, click **Restore**.
5. In the Choose how to restore dialog box, select **Restore database to a SQL Server instance**.
6. In the **Instance** list, click the SQL Server instance where you want to restore the database.
7. Do one of the following:
 - To connect to the instance using a Windows administrator account, select **Windows authentication**. Enter the user name, password, and domain in the appropriate fields.
 - To connect to the instance using a SQL Server administrator account, select **SQL Server authentication**. Enter the user name and password in the appropriate fields.
8. Click **Continue**.

The SQL Server Restore dialog box shows the most recent safeset for the job.

9. To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:
 - To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
 - To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

10. In the **Database Selection** box, select the check box for each database that you want to restore.
11. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
12. Do one of the following:
 - To restore one or more databases with their original names, select **Original Database Names**.
 - To restore one database with a new name, select **Alternate Database Name**. In the field that appears, enter the new name for the restored database.

You can only restore one database if **Alternate Database Name** is selected.

13. Do one of the following:

- To overwrite the existing database if you restore a database with the same name as the existing database, select **Overwrite existing databases**.
- To fail the restore if a database with the same name already exists, clear **Overwrite existing databases**.

If **Overwrite existing databases** is not selected, and you are restoring multiple databases, the restore fails for all databases if even one database has the same name as an existing database.

14. To restore the database in restoring state, select **Restore using No Recovery option**.

If this option is selected, and transaction logs have been backed up using an alternative method (e.g., native SQL Server backup), you can apply transaction logs to the database after it has been restored.

15. To specify an alternate location for database files, select **Alternate Path**. Click the folder button. In the Select Folder dialog box, select the alternate file location, and click **Okay**.

Note: The alternate file location is only used if the original location for database files is not available.

16. To change the log detail level or bandwidth throttling setting, click **Advanced Restore Options**. In the dialog box, do one or more of the following:

- In the **Log Level Detail** list, select the level of detail for job logging. See [Advanced restore options](#).
- Select or clear the **Use all available bandwidth** option. See [Advanced restore options](#).

17. Click **Run Restore**.

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

To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.3.3.2 Restore SQL Server databases to files

After backing up SQL Server databases using the SQL Server Plug-in, you can restore the databases to SQL database files (.mdf and .ldf files). You can then use SQL Server tools to attach the files to a SQL Server instance.

To restore a SQL Server database to files:

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

A grid lists available computers.

2. Find the computer with the SQL Server database backup that you want to restore, and expand its view by clicking the row for the computer.
3. Click the **Jobs** tab.
4. Find the job with the database you want to restore, and click **Restore** in the **Select Action** menu for the job.
5. In the Choose how to restore dialog box, select **Restore to folder**.
6. Click **Continue**.

The SQL Server Restore dialog box shows the most recent safeset for the job.

7. To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:
 - To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
 - To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

8. In the **Database Selection** box, select the check box for each database that you want to restore.
9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
10. Under **Restore Destination**, enter a path for the destination, or click the folder button.  In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.
11. To change the log detail level or bandwidth throttling setting, click **Advanced Restore Options**. In the dialog box, do one or more of the following:
 - In the **Log Level Detail** list, select the level of detail for job logging. See [Advanced restore options](#).
 - Select or clear the **Use all available bandwidth** option. See [Advanced restore options](#).
12. Click **Run Restore**.

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

To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.3.3.3 Restore a SQL Server master database

In a SQL Server instance, a master database contains information about all other databases and configuration information for the instance. The master database must be restored by itself while the instance is offline.

Other system databases (i.e., msdb and model) can be restored while the SQL Server instance is online, and do not require a special restore procedure.

To restore a SQL Server master database:

1. Stop the SQL Server services on the machine where you want to restore the master database.
2. Manually copy the existing master database files (master.mdf and mastlog.ldf) to another location so they will not be overwritten.

These copies can be used to roll back the database if a problem occurs with the restored master database.

3. Restore the master database files (master.mdf and mastlog.ldf) from a SQL Server Plug-in backup or application-aware Image Plug-in backup. See [Restore SQL Server databases to files](#) or [Restore SQL Server database files or objects from Image backups](#).
4. Manually copy the master database files to the original location of the master database files (with replace).
5. Restart all SQL Server services.

8.3.3.4 Restore SQL Server database files or objects from Image backups

After backing up SQL Server databases using Image Plug-in version 7.5 or later, you can restore database files, tables and objects from the application-consistent backups.

To restore database files from Image Plug-in backups, you must mount a safeset as a drive on the computer where you want to restore database files. You can then restore files using Windows Explorer or tables and other objects using the Granular Restore for Microsoft Exchange and SQL application.

To restore SQL Server database files or objects from an Image Plug-in backup:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with the SQL Server database backup created using the Image Plug-in, and expand its view by clicking the row for the computer.
3. Click the **Jobs** tab.
4. Find the job with the database you want to restore, and click **Restore** in the **Select Action** menu for the job.


The Choose how to restore dialog box appears.

5. Select **Individual File or Folders**.
6. Click **Continue**.

The Restore dialog box shows the most recent safeset for the job.

7. To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
8. In the **Idle Time** text box, enter the number of minutes of inactivity after which the share should automatically stop. This value can range from 2 to 180.
9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
10. In the **Volume to Restore** column, select each backed-up volume from which you want to restore files and folders.
11. In the **Volume Mount Drive Letter** column, choose the drive letter for mounting each volume.
12. Click **Mount Volumes**.
13. If a confirmation message appears, read the message, and then click **Continue**.

The Process Details dialog box appears. When each volume is mounted, the drive letter is shown at the right side of the dialog box.



14. Do one of the following:
 - To restore SQL Server databases, use Windows Explorer to browse to the location of the database files and copy the database files (.mdf and .ldf) that you want to restore. If a *databaseName_mod* folder exists, copy the *databaseName_mod* folder as well. You can then attach the database in Microsoft SQL Server.
 - To restore SQL Server database tables or objects, launch the Granular Restore for Microsoft Exchange and SQL application on a SQL Server system. In the Granular Restore application,

choose **File > Open**, drill down into the volume to choose the .mdf file, then select and restore your data. For more information, see documentation for the Granular Restore for Microsoft Exchange and SQL application.

15. When you are done restoring database files, click **Stop** to remove the mounted drives.

If you do not click **Stop**, the drive will be unmounted automatically after the number of minutes of inactivity specified in the **Idle Time** box. See Step 8.

8.3.3.5 Restore databases in AlwaysOn Availability Groups

You must always restore a SQL Server database to the primary replica in an AlwaysOn Availability Group. If a Windows Agent and plug-in are not installed on the primary replica server, you must fail over to a server where the Agent and plug-in are installed before restoring the database.

After restoring a database to the primary replica and adding the database back into the AlwaysOn Availability Group, it will be replicated to the secondary replicas. To reduce the amount of replication traffic after a restore, you can run a “Restore from another computer” on any secondary replica server where the Windows Agent and plug-in are installed.

For information about backing up databases in AlwaysOn Availability Groups, see [Protect SQL Server databases in AlwaysOn Availability Groups](#).

To restore a primary database in an AlwaysOn Availability Group:

1. If the Agent and plug-in are not installed on the primary replica server, fail over to the secondary database instance where the Agent is installed.
The formerly secondary replica where you backed up the database becomes the primary replica.
2. Remove the primary database from the AlwaysOn Availability Group.
3. Delete the database from all secondary replicas.
4. Restore the primary database to the original database name using the Overwrite Existing Databases option.
5. Add the restored primary database to the AlwaysOn Availability Group using the Full Synchronization option.

After restoring a SQL Server database to the primary replica, to reduce the amount of required replication traffic, you can restore the database to secondary replica servers.

To restore a secondary database in an AlwaysOn Availability Group:

1. If you did not delete the database from all secondary replicas when restoring the primary database (see Step 3 in the previous procedure), remove the secondary database from the AlwaysOn Availability Group.
2. On a secondary replica server where the Agent and plug-in are installed, restore the database by running a Restore From Another Computer using the No Recovery option.

3. Add the restored secondary database to the AlwaysOn Availability Group using the Join option.

8.3.4 Restore items from a SQL Server or SharePoint database

If a Microsoft SharePoint database is backed up using the SQL Server Plug-in, you can restore items such as site collections, websites, lists and documents from the backup.

If a Microsoft SQL Server database is backed up using the SQL Server Plug-in or Image Plug-in, you can restore specific tables and objects from the backup.

To restore items from a database backup, you must first use Portal to expose the safeset as a shared resource. You can then use Granular Restore for Microsoft Exchange and SQL to find and restore items from the backup. For more information, or to obtain the Granular Restore application, contact your service provider.

To restore items from a SQL Server or SharePoint database:

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

A grid lists available computers.

2. Find the computer with the safeset with SharePoint or SQL Server data that you want to restore, and expand its view by clicking the computer row.

3. Click the **Jobs** tab.

4. Find the job with the SharePoint data that you want to restore, and click **Restore** in the **Select Action** menu for the job.

The Choose how to restore dialog box appears.

5. Select **Restore items to a SharePoint or SQL Server database**, and click **Continue**.

The SQL Server Restore dialog box shows the most recent safeset for the job.

6. To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.

7. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 

8. In the **Idle Time** box, enter the number of minutes of inactivity after which the share should automatically stop. The value can range from 2 to 180 minutes.

9. Select or clear the **Use all available bandwidth** option.

10. Click **Share**.

The Process Details dialog box shows the status of the share process. When the share is available, the share path appears at the right side of the dialog box.

11. Click the Copy Path to Clipboard button.  The path is now available for you to paste into the Granular Restore application.
12. Launch the Granular Restore for Microsoft Exchange and SQL application on a SQL Server system.
13. Paste the path for the SQL safeset share into Granular Restore for Microsoft Exchange and SQL.
14. Select and restore your data. For more information, see documentation for Granular Restore for Microsoft Exchange and SQL.
15. When you no longer need to share the safeset, click **Stop**.

When you click Stop or the share idle time is reached, the Process Details dialog box indicates that the share is no longer available.

8.3.5 Restore Oracle databases

After backing up an Oracle database using the Oracle Plug-in, you can restore the database.

You might also need to recover the entire system, by performing a “bare metal restore” (installing the OS, applications, and then the full database (plus any transaction logs) onto a new system).

If there is an Oracle backup and a full-system backup:

1. Restore the system (putting back the contents of ORACLE_HOME – specifically the database installation). If you like, you can exclude the data files and archive logs that are backed up by the plug-in.
2. Restore the Oracle backup, and then copy the required components to the appropriate directories. Follow the standard user-managed Oracle recovery procedure from the Oracle backup and recovery guide (available from Oracle) that is appropriate for the operating system.

An Oracle restore process is performed by a Database Administrator. Briefly, the steps are:

- Shut down the database.
- Restore the files.
- If necessary, reset the control information for the database.
- Start and recover the database.
- Re-open the database for use.


Note: The Plug-in does not do table-level restores.

To restore an Oracle database:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with the Oracle database that you want to restore, and expand its view by clicking the row for the computer.

3. Click the **Jobs** tab.
4. Find the job with the database that you want to restore, and click **Restore** in the **Select Action** menu for the job.

The Restore dialog box shows the most recent safeset for the job.



5. To restore the database from an older safeset, or from SSI (safeset image) files, do one of the following:
 - To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
 - To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.
 - SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.
6. In the **Files to Restore** box, select the items that you want to restore.
7. Select a Restore Destination option.
 - To restore files and folders to the location where they were backed up, select **Restore files to their original location**.

- To restore files and folders to a different location, select **Restore files to an alternate location**. Click the folder button.  In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.
8. Select an Overwrite Option. This option specifies how to restore an item (e.g., a file) to a location where there is an item with the same name.
- To overwrite the existing item with the restored item, select **Overwrite existing items**.
If you try to restore multiple files with the same name to an alternate location and select **Overwrite existing items**, only the last file restored will remain. Other files with the same name will be overwritten.
IMPORTANT: Using Agent version 8.70 or later, if you select **Overwrite existing items** and restore a file that has the same name as a folder in the restore location, the file will overwrite the folder. The folder and all of its contents will be removed.
 - To skip restoring the item that has the same name as an item in the destination location, select **Do not overwrite existing items**.
 - To add a numeric extension (e.g., .0001) to the *restored* item name, select **Rename incoming items**. For example, if you restore a file named “filename.txt” to a location where there is a file with the same name, an extension is added to the **restored** file name (e.g., “filename.txt.0001”).
 - To add a numeric extension (e.g., .0001) to the *existing* item name, select **Rename existing items**. For example, if you restore a file named “filename.txt” to a location where there is a file with the same name, an extension is added to the *existing* file name (e.g., “filename.txt.0001”). The name of the restored file is “filename.txt”.
9. To change the log detail level or bandwidth settings, click **Advanced Restore Options**. Specify settings in the Advanced Restore Options dialog box, and click **Okay**. See [Advanced restore options](#).
10. Click **Run Restore**.
- The Process Details dialog box shows the restore progress and indicates when the restore is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).
11. To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.
- Note:* For a full disaster recovery (in which the full database instance is restored), be careful when you recover the database because the plug-in does not back up TEMPORARY tablespaces.

In Solaris and Linux, start the database recovery with an explicit PFILE or SPFILE reference:

```
SQL> STARTUP PFILE='path-to-pfile\initSIDNAME.ora'
```

It might be necessary to take the temporary tablespace files offline:

```
SQL> ALTER DATABASE DATAFILE 'path-to-datafile' OFFLINE
```

Restore the database as usual, but when you open it after recovery, use this command:

```
SQL> ALTER DATABASE OPEN NORESETLOGS
```

TEMPORARY tablespaces should be dropped, the data files for the temporary tablespaces should be removed, and the TEMPORARY tablespaces should be recreated (this can include the default TEMP tablespace).

At this point, you can close the database normally and restart it (with RESETLOGS, for example).

8.4 Restore Linux or UNIX files and folders

After backing up data from a Linux or UNIX computer, you can restore files and folders from the backup.


Beginning with version 8.83 of the 64-bit Linux Agent, you can also restore Linux systems from BMR backups. See [Restore a Linux system from a BMR backup](#).

You can also restore Linux systems from BMR backups. See [Restore a Linux system from a BMR backup](#).

To restore Linux or UNIX files and folders:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the Linux or UNIX computer with data that you want to restore, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job with data that you want to restore, and click **Restore** in the job's **Select Action** menu.

The Restore dialog box shows the most recent safeset for the job.



5. To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:

- To restore data from an older safeset, click the calendar button. In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
- To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

6. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 

7. Select a Restore Destination option.

- To restore files and folders to the location where they were backed up, select **Restore files to their original location**.
- To restore files and folders to a different location, select **Restore files to an alternate location**. Click the folder button.  In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.

8. Select an Overwrite Option. This option specifies how to restore a file, folder or symbolic link to a location where there is a file, folder or symbolic link with the same name.

- To overwrite the existing item with the restored item, select **Overwrite existing items**.

Note: If you try to restore multiple files with the same name to an alternate location and select **Overwrite existing items**, only the last file restored will remain. Other files with the same name will be overwritten.

IMPORTANT: Using Agent version 8.70 or later, if you select **Overwrite existing items** and restore a file that has the same name as a folder in the restore location, the file will overwrite the folder. The folder and all of its contents will be removed.

- To skip restoring the item that has the same name as an item in the destination location, select **Do not overwrite existing items**.
- To add a numeric extension (e.g., .0001) to the *restored* item name, select **Rename incoming items**. For example, if you restore a file named "filename.txt" to a location where there is a file with the same name, an extension is added to the **restored** file name (e.g.,

"filename.txt.0001").

- To add a numeric extension (e.g., .0001) to the *existing* item name, select **Rename existing items**. For example, if you restore a file named "filename.txt" to a location where there is a file with the same name, an extension is added to the *existing* file name (e.g., "filename.txt.0001"). The name of the restored file is "filename.txt".
9. To change locked file, data streams, log detail level or bandwidth options, click **Advanced Restore Options**. Specify settings in the Advanced Restore Options dialog box, and click **Okay**. See [Advanced restore options](#).
 10. In the **Select Items** box, do one or more of the following until the **Restore Set** box shows the folders and files that you want to restore:
 - Select the check box for each folder and file that you want to restore, and then click **Include**. The **Restore Set** box shows the included folders and files. If you include a folder, all of the folder's subdirectories and files are restored by default. If you do not want to restore all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).
 - To exclude one or more folders or files from the restore, select the check box for each folder or file, and then click **Exclude**. The **RestoreSet** box shows the excluded folders and files. If you exclude a folder, all of the folder's subdirectories and files are excluded from the restore by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).
 - To search for files to restore or exclude from the restore, click the **Search** button.  In the **Search for files** box, enter search criteria and select files. See [Search for files to restore](#). Click **Include Selected** or **Exclude Selected**. The **RestoreSet** box shows the included or excluded files.
 - To remove an inclusion or exclusion record from the **Restore Set** box, click the Delete button  beside the folder or file record.
 11. Click **Run Restore**.

The Process Details dialog box shows the restore progress and indicates when the restore is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).
 12. To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

8.4.1 Restore ACLs

You can back up and restore Access Control Lists (ACLs). The following behaviors can occur when you restore ACLs on a Linux server.

ACLs control the access of users or groups to particular files. Similar to regular file permissions (e.g., owner, group, world), ACLs are tracked by the ID of the user/group. ACLs provide access-control granularity beyond regular file permissions, and unlike regular permissions, they are not always enabled.

ACL implementations might differ by variety of Linux, and by the type of file system. Not all ACL implementations are "portable" (i.e., ACLs on one OS/file system may be incompatible with ACLs on another OS/file system). In addition, you might need to enable ACL support on a partition before you can configure it.

If you attempt to restore ACLs to an incompatible system (e.g., a file system that does not support ACLs), the ACLs will not be restored. An error message will appear in the backup log.

If you restore to a compatible system (e.g., the original system, or a different system with the same variety of Linux), ACLs will also be restored.

Since ACLs are associated with user and group IDs, you will observe the following on a compatible system:

- If the group, user names, and IDs on the restored system match those of the original system, the ACLs will be associated with the same user name as on original system.
- If the group, user names, and IDs on the restored system do not match those on the original system, the ACLs will be associated with a different user or group name compared to the original system.
- If the group or user name ID does not exist on the restored system, the ACLs will be associated with the user ID or group ID respectively. Therefore, browsing ACLs on these files will show user/group IDs as opposed to user/group names.

8.5 Restore a Linux system from a BMR backup

Beginning with version 8.83 of the 64-bit Linux Agent, you can restore entire Linux servers from Bare Metal Restore (BMR) backups. A BMR backup includes:

- An .iso file for starting the destination system and running the restore. The .iso file is created on the source system during a BMR backup and is backed up to the vault
- A backup in the vault of all folders and files that are required for the system. By default, a Linux BMR backup includes all folders and files from the root (/), although some files can be excluded.

When restoring a Linux server from a BMR backup, the destination machine must have:

- At least 4 GB of RAM.
- The same boot type (BIOS or UEFI) as the source system, and compatible hardware.
- Hard drives that are the same size or larger than drives on the source system.
- A connection to the network, so that it can communicate with the vault.


Note: Restores are not supported to systems with different types of firmware.

To restore a Linux system from a BMR backup:

1. Do one of the following:
 - If the source system is still available, copy the `/Bare_Metal_Restore_Image.iso` file from the root directory (/) of the source system to another machine.


- Use the *Restore from another computer* procedure to restore the /Bare_Metal_Restore_Image.iso file from the Linux BMR backup to another machine. See [Restore data from another computer](#).
2. Create a bootable USB device, CD or DVD from the Bare_Metal_Restore_Image.iso file and mount it on the destination system.
 3. Boot the destination system from the bootable file.
 4. Do one of the following:
 - If the Relax-and-Recover screen appears, select **Recover sourceSystemName** and then press Enter. This screen appears if the protected system is BIOS-based.

Do not select the Automatic Recover *sourceSystemName* option or the system will not start successfully.




- If the following screen appears, select **Relax-and-Recover (no Secure Boot)** and then press

Enter. This screen appears if the protected system is UEFI-based.



5. At the login prompt, log in as root.

Note: If a login prompt does not appear initially, press Enter.



6. (Optional) To ensure that the connection to the vault is active, ping the vault IP address. If there are network connection issues, change the network settings.
7. Enter the following command:

```
./bmragent
```
8. On the Vault Login screen, enter information for connecting to the vault where the BMR backup is saved.

In the Address field, enter the vault IP address or fully qualified domain name (FQDN). In the Port field, enter the port number for connecting to the vault (2546, by default). In the Account, Username and Password fields, enter an account and credentials used for the backup.

```
Vault Login:
=====
Address      :
Port Number  : 2546
Account      :
Username     :
Password     :

'ESC' : quit

-----
Enter information about the vault where the BMR backup is saved.
Address: IP address or FQDN
Default port: 2546

=====
```

9. On the Protected Servers screen, press the up and down arrow keys to select the protected server to restore, and then press Enter.
10. On the Job List screen, press the up and down arrow keys to select the BMR job to restore, and then press Enter.
11. On the Safeset List screen, press the up and down arrow keys to select the backup to restore, and then press Enter.
12. On the Start Restore screen, enter the encryption password for the BMR backup job.
On the CONFIRM line, press the right arrow key to choose yes, and then press Enter.
13. If the destination system is larger than the protected system, a *Confirm the recreated disk layout or go back one step* prompt appears. Press Enter to select the default option.
If an error occurs at this stage, please go to <http://relax-and-recover.org/support/> for Relax-and-Recover support.
14. If the destination system has a larger disk than the protected system, a *Confirm restored config files are OK or adapt them as needed* prompt appears. Press Enter, and then press Enter again to select the default options.

The system restore begins. The restore time depends on the size of the backup.

If the restore takes a long time, the screen might go blank. To refresh it, press Enter.

When the restore is finished, a *Completed the bare metal restore* message appears, followed by the RESCUE prompt.

```
Finished recovering your system. You can explore it under '/mnt/local'.
Exiting rear recover (PID 772) and its descendant processes ...
Running exit tasks
Completed the bare metal restore.
RESCUE localhost:/opt/BUAgent # _
```

15. Run the following command to view the restore log:

```
./xlogcat jobName/RSTyyyyymmdd-hhmmss.XLOG | tail -n 25
```

Where *jobName* is the name of the BMR job from which you restored the system, and *yyyyymmdd-hhmmss* is the date and time of the restore.

Review the restore log. Check that the restore completed with no errors and that the restored system size is correct

16. Restart the system.

Depending on the source system platform and configuration, the system could restart once or twice automatically.

17. Log in to the restored system with credentials from the protected system, and verify that the system is working.

8.6 Restore a Linux system without a BMR backup

Beginning with version 8.83 of the 64-bit Linux Agent, you can restore entire 64-bit Linux systems from BMR backups. See [Restore a Linux system from a BMR backup](#).

You can restore entire Linux systems from BMR backups. See [Restore a Linux system from a BMR backup](#).

If a Linux system is not protected by a BMR backup, you must recover the system using techniques described in this section. This section also describes the minimum resources required to rebuild a file system to its state at the last system backup.

The basic recovery procedure is:

1. Install the minimal operating system, including networking.
2. Install and configure the Agent.
3. Restore the backed up system state, programs, and data using the Agent.
4. Perform M-restore maintenance.
5. Verify the restore.

Before performing a recovery, ensure that your hardware configuration is sufficient for the programs, data, and system state of the protected system.

8.6.1 Hardware requirements

It is crucial for local storage on the system to be sufficient for a full restore of programs, system state, and data. Otherwise, the restore will fail, and your system may be left in an indeterminate state.

If any configuration files for your operating system depend on specific identifiers of installed hardware (such as the MAC address of a network card), ensure that this information is noted, as the values may be different than when the system was backed up using the Agent.

Note: When performing a complete system restore (DR), you need to ensure there is ample disk space for the creation of large recovery logs from our Agent and other possible logging or auditing from the operating system. Using file level logging on a system containing a large file system can generate a large log, which can potentially fill up the available or allocated disk space. If the logs are on the same partition as the root file system, this may prevent the OS from booting.

8.6.2 Software requirements

Ensure that the appropriate installation media is available. The minimum system software includes:

- Installation media identical to that installed on the original system.
- Any necessary OS patches to install the Agent, as described in the installation instructions for the Agent on the OS.
- Agent Installation media identical to that installed on the original system.

8.6.3 Recovery steps

This section describes the steps to perform a system recovery.

Install the minimal operating system

Follow the instructions in your operating system manual and installation media to install a minimal operating system.

- When prompted to partition your drive(s), ensure that the partitions are large enough to restore to; they should be at least as large as the original partitions.
- If restoring over the network, TCP/IP network services must be installed and configured appropriately, and there must be a connection between the system and the backup vault.
- If restoring from a directory on disk, there must be sufficient disk space to handle all the restored data.

Install and configure the Agent

1. Install the Agent for your operating system.
2. Configure the Agent. Re-register the Agent to the vault where the data was backed up.
3. Synchronize the job to ensure that local copies of job catalogs are created.

Restore the backed up system

1. Start a restore.
2. Select the files you wish to restore. The Agent will restore most files to their original locations and protect against many known restore problems (for file systems mounted in their default locations), but some files may cause unpredictable results if restored. These files vary and may generally be restored to alternative locations without problems.
3. Ensure that the files are not being restored to a file system that is mounted read-only.

Note: The Agent will prevent recovery of files to critical locations, but not all critical locations are necessarily detected.

When the recovery procedure is complete, the process of verifying the integrity of the restore can commence.

Perform post-recovery maintenance

If any modifications to the configuration of the restored system are required after restore, these should be performed now. Known post-restore maintenance steps are noted below.

Verify the recovery

Once the restore procedure is complete, determine if the recovery is complete and correct. The listing and testing of the jobs should be performed as part of the systems recovery planning. The specific jobs to be performed for verification depend on the application environment and the system's importance.

Once the system is restored, the integrity of the recovery must be verified. The test can be as simple as placing a duplicate file in a different directory structure and testing for any differences within the file. Then, confirm that the file can be opened using a known application and that you are able to send e-mail to a known address. It can also be as complex as completing an SQL query on a known database set.

Whatever the test, both the list and the test itself must be planned and executed during normal system operation.

8.6.4 Recovery problems

Should any of the recovery jobs fail, consider these questions:

- Was the system restored using the same version of OS?
- What possible differences were there in the hardware or software settings that could have affected the recovery?
- Were any errors reported in the error log file?
- Were all the necessary drivers installed?
- Were the applicable OS patches added?
- Was there sufficient disk space to handle all of the restored data?

8.7 Restore vSphere data

When VMs are protected in a vSphere environment, you can:

- [Restore vSphere VMs](#)
- [Restore a vSphere VM within minutes using Rapid VM Restore](#)
vSphere Recovery Agent 8.80 or later is required for Rapid VM Restores.
- [Restore files, folders and database items using a vSphere Recovery Agent](#)

8.7.1 Restore vSphere VMs


Before you restore a vSphere VM, the vSphere Recovery Agent (VRA) checks whether sufficient storage space is available. If there is not enough space, the restore fails and a message appears in the log file.

If you restore a VM or template to a vSphere environment and the original VM is present, the VM will be restored as a clone of the original with the following name: <VMname>-vra-restored-<Date>. This name will appear for the clone in both the vCenter environment and the datastore. The VM will be restored as a clone whether the original VM is powered on, off, or suspended. The original VM name will not change and its data will not be overwritten. Beginning with VRA 8.87, the restored VM is assigned a new MAC address. An IP address conflict will not occur when the original and newly-restored VMs are powered on.

After you restore a VM from a crash-consistent backup, the VM may perform a disk check when it first starts.


To restore vSphere VMs:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the vSphere environment with the VM that you want to restore, and expand its view by clicking the row.
3. Click the **Jobs** tab.
4. Find the backup job with the VM that you want to restore, and click **Restore** in the job's **Select Action** menu.
5. In the Choose What You Want to Restore dialog box, select **Virtual Machines**.




6. Click **Continue**.

The Restore dialog box appears. If a potential ransomware threat was not detected in the job, the most recent safeset for the job appears in the Safeset box.



If a potential ransomware threat was detected when running the job, a calendar with a list of safesets appears. "Potential Threat" appears beside each safeset where a potential ransomware threat was detected.

Safeset



Note: If you are restoring data as described in [Restore data to a replacement computer](#) or [Restore data from another computer](#), "Potential Threat" does not appear for any safesets even if a potential threat was detected during a backup in the original vSphere environment.

7. To restore data from an older safeset, or from SSI (safeset image) files on disk, do one of the following:

- To restore from an older safeset, if a calendar with a list of backups does not already appear, click the **Browse Safesets** button. In the calendar, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset from which you want to restore.
- To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button. In the Select Folder dialog box, select the directory


where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

8. In the **Items to Restore** box, select the check box for each VM that you want to restore.

If a VM has a potential ransomware threat, "Potential Threat" appears beside the VM name.



9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
10. In the **Destination Datastore** list, click the datastore for the restored VMs.
11. Select one of the following options for restoring VMs to the selected datastore:
 - **Restore all selected Virtual Machines to the selected datastore only**
 - **Restore to the selected datastore only when a Virtual Machine's original datastore is not available.** If the backed-up VM contains multiple VMDKs that resided on two or more datastores, and one or more of the datastores is unavailable, the entire VM will be restored to the selected datastore.
12. In the **Destination Host** list, click the host where you want to register the VMs.

The list only shows hosts that have access to the selected datastore. If only one ESXi host is available, it is populated as the Destination host when you select a datastore.
13. If the VRA is protecting a vCenter Server, select one of the following options for registering restored VMs with the selected host:
 - **Register all selected Virtual Machines with the selected host only**
 - **Register with the selected host only when a Virtual Machine's original host is not available**

Note: If the VRA is protecting a single ESXi host that is not managed by vCenter Server, registration options do not appear in the Restore dialog box.
14. To power on the VMs after they are restored, select **Power VMs on after restoring**.
15. In the **Log Level Detail** list, click the logging level. See [Advanced restore options](#).
16. To use all available bandwidth for the restore, select **Use all available bandwidth**.

To ensure the best possible performance for your restore, we recommend selecting **Use all available bandwidth**.

17. Click **Run Restore**.

If you are restoring from a safeset that has a potential threat, a warning message appears. To cancel the restore, click **Cancel**. To continue restoring the volume, click **Continue**.

8.7.2 Restore a vSphere VM within minutes using Rapid VM Restore

Using Rapid VM Restore, you can restore a virtual machine (VM) to your vCenter or ESXi host within minutes.

In a vCenter, you can restore a VM using Rapid VM Restore and then migrate it to a second datastore to restore it permanently. This can be useful in a disaster recovery situation, where critical servers must be restored and available to users and applications as soon as possible. You can also restore a VM temporarily, to quickly verify that the VM backup can be restored.

On an ESXi host that is not managed by vCenter Server, you can restore a VM temporarily using Rapid VM Restore. Restoring a VM temporarily can be useful as a test, to quickly verify that a VM backup can be restored.

When you first restore a vSphere VM using Rapid VM Restore, disks from the selected VM backup are mounted as storage devices (virtual RDMs) on a VM for immediate access. While the VM runs, changes are written to a temporary datastore. At this stage, the VM requires a running Rapid VM Restore process, requires connections to the VRA and vault, and is intended for temporary use. The longer a VM runs using Rapid VM Restore, the more its performance will degrade and the more vault and VRA resources it will use.

After you migrate a restored VM to permanent storage in a vCenter, the VM does not require a running Rapid VM Restore process and is independent from the VRA and vault. We recommend migrating a VM to permanent storage as soon as possible after it is restored using Rapid VM Restore. See [Migrate a vSphere VM restored using Rapid VM Restore to permanent storage](#). If the network connection to the VRA, vault or ESXi host is interrupted before a VM is migrated to permanent storage, VM data could be lost.

IMPORTANT: If the VRA is protecting a single ESXi host that is not managed by vCenter Server, you cannot restore a VM permanently using Rapid VM Restore. An ESXi server that is not part of a vCenter does not have the capabilities required to migrate VMs to permanent storage.

Notes:


- Before a VM is restored using Rapid VM Restore, the VRA checks that sufficient storage space is available. If there is insufficient space, the restore fails and a message appears in the log file.
- If you restore a template using Rapid VM Restore, it is restored as a running virtual machine and not as a template.
- After you restore a VM from a crash-consistent backup, the VM may perform a disk check when it first starts.

- We highly recommend backing up virtual machines (VMs) that are restored using Rapid VM Restore. See [Best practice: Back up vSphere VMs restored using Rapid VM Restore](#).
- Rapid VM Restore is available with vSphere Recovery Agent (VRA) version 8.80 or later. For complete requirements, see [vSphere Rapid VM Restore and backup verification requirements](#).

To restore a vSphere VM within minutes using Rapid VM Restore:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the vSphere environment with the VM that you want to restore, and expand its view by clicking the row.
3. Click the **Jobs** tab.
4. Find the backup job with the VM that you want to restore, and click **Restore** in the job's **Select Action** menu.
5. In the Choose What You Want to Restore dialog box, select **Virtual Machine using Rapid VM Restore**.

If the **Virtual Machine using Rapid VM Restore** option does not appear, this restore method is not available. This could occur with a VRA version earlier than 8.80, if backups are not available in a local vault that supports Rapid VM Restores, or if other requirements are not met. For complete requirements, see [vSphere Rapid VM Restore and backup verification requirements](#).




6. Click **Continue**.

The Restore dialog box appears. If a potential ransomware threat was not detected in the job, the most recent safeset for the job appears in the Safeset box.

If a potential ransomware threat was detected when running the job, a calendar with a list of safesets appears. "Potential Threat" appears beside each safeset where a potential ransomware threat was detected.


Note: If you are restoring data as described in [Restore data to a replacement computer](#) or [Restore data from another computer](#), "Potential Threat" does not appear for any safesets even if a potential threat was detected during a backup in the original vSphere environment.



- To restore from an older safeset, if a calendar with a list of backups does not already appear, click the **Browse Safesets** button.  In the calendar, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset from which you want to restore.

- In the **VM to Restore** list, select the VM that you want to restore.

If a potential ransomware threat was detected on a VM, "Potential Threat" appears beside the VM name.



- In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button .
- In the **Log Level Detail** list, select the level of detail for job logging. For more information, see [Log file options](#).
- In the Restore Settings box, do the following:

- In the **Restored VM Name** box, type a name for the restored VM.

If you specify the name of a VM that already exists in the vSphere environment (e.g., the VM that was backed up), the restored VM will have the following name: *VMname-rvmr-yyyy-Mon-dd--hh-mm-ss*, where *yyyy-Mon-dd--hh-mm-ss* is the date and time when the VM was restored (e.g., VM-rvmr-2019-Nov-27--06-14-09).

- In the **Datastore** list, select a datastore for writing changes while the VM is restored using Rapid VM Restore (i.e., while disks from the selected backup are mounted as storage devices).

If the VRA is protecting a vCenter and you later want to migrate the VM to permanent storage, do not choose the datastore that you want to use as permanent storage.

- In the **Destination Host** list, select a host for running the restored VM.

If only one ESXi host is available, it is populated as the Destination host when you select a datastore.

If the VRA is protecting a vCenter and you later want to migrate the VM to permanent storage, select a host that can access the permanent datastore.

- Do one of the following:
 - To restore the VM with its power on, select the **Power on the VM** option.
 - To restore the VM powered off, clear the **Power on the VM** option.

You might want to restore the VM with its power off, for example, so you can verify or change the VM settings before powering it on.

- Do one of the following:
 - To connect the VM to the network, select **Connect to Network**.
 - To restore the VM without network connectivity, clear **Connect to Network**.


You might want to restore the VM without network connectivity, for example, if you are restoring the VM to a vCenter that does not have the original network. You can then verify the VM settings before connecting the VM to the network.

12. Click **Run Restore**.

If you are restoring from a safeset that has a potential threat, a warning message appears. To cancel the restore, click **Cancel**. To continue restoring the volume, click **Continue**.

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

The restored VM appears in the vSphere environment. You can access the VM and begin using it.



13. Do one or more of the following:

- To close the Process Details dialog box, click **Close** in the dialog box. If you close the Process Details dialog box without canceling the Rapid VM Restore, the VM remains in the vSphere environment.
- To reopen the Process Details dialog box, find the VM's VRA backup job on the Computers page or Monitor page. Click the Rapid VM Restore symbol that appears beside the VRA job name: .
- To permanently restore the VM by migrating it to permanent storage, see [Migrate a vSphere VM restored using Rapid VM Restore to permanent storage](#).

IMPORTANT: You cannot migrate the VM to permanent storage if the VRA is protecting a single ESXi host that is not managed by vCenter Server.

- To remove the VM from the vSphere environment, click **Cancel Rapid VM Restore** in the Process Details dialog box.

8.7.2.1 Migrate a vSphere VM restored using Rapid VM Restore to permanent storage

When you first use Rapid VM Restore to restore a vSphere VM, the VM is dependent on the VRA and vault, and is intended for temporary use.

To restore the VM permanently, use Portal to migrate the VM to permanent storage. If the VM is powered on, you can continue to use the VM during the migration. After migration, the VM is independent from the VRA and vault, and its disks are restored with their original formats (e.g., thin- or thick- provisioned).

IMPORTANT: On an ESXi host that is not managed by vCenter Server, Rapid VM Restore can be used to verify that VMs were backed up correctly, but cannot be used to restore VMs permanently. An ESXi server that is not part of a vCenter does not have the capabilities required to migrate VMs to permanent storage.

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


When migrating a VM that was restored using Rapid VM Restore to permanent storage, we recommend the following:

- Before running a migration, back up the VM that was restored using Rapid VM Restore. See [Best practice: Back up vSphere VMs restored using Rapid VM Restore](#). You cannot back up a VM while it is being migrated, or migrate a VM while it is being backed up.
- Use Portal to migrate a VM to permanent storage rather than using the vSphere Client or Web Client. When migrating a VM to permanent storage, Portal ensures that all disks are migrated and converted to their original formats. If you try to migrate a VM to permanent storage without using Portal but do not migrate all disks and convert them to their original formats, you will not be able to migrate the VM using Portal. The VM might be deleted when you cancel the Rapid VM Restore process.
- Do not perform more than six migrations at one time, even if the migrations are distributed across hosts in the vSphere environment.
- During a migration, do not power off the VM from within the guest operating system or you might be locked out of the VM until the migration is complete. While a VM is being migrated, you cannot power on, power off, or suspend the VM using the vSphere client.

To migrate a VM restored using Rapid VM Restore to permanent storage:

1. Check that the VM is in the state that you want during the migration: powered on, powered off, or suspended.
2. If the Process Details dialog box is not open for the VM's Rapid VM Restore process, find the VM's VRA backup job on the Computers page or Monitor page. Click the Rapid VM Restore symbol that appears beside the VRA job name: 


The Process Details dialog box lists Rapid VM Restores that are running from the selected backup job.



3. If more than one VM appears in the VM Name list, select the VM that you want to migrate.
4. Click **Migrate VM**.

IMPORTANT: The Migrate VM button is not available if you are restoring the VM to a single ESXi host that is not managed by vCenter Server. You cannot permanently restore a VM using Rapid VM Restore if the VRA is protecting a single ESXi host that is not managed by vCenter Server.

The Migration Settings dialog box appears.



5. In the **Select Destination Datastore** list, select the permanent datastore for the VM.

The list includes datastores that are accessible from the host selected for the Rapid VM Restore, but does not include the temporary datastore selected for the Rapid VM Restore.

6. Click **Start Migration**.

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

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

When the VM is migrated to the permanent datastore, the following Status message appears in the Process Details dialog box: *VM has been migrated*. At this point, the VM is permanently restored and is no longer dependent on the VRA and vault. The Rapid VM Restore process ends and the Rapid VM Restore symbol  no longer appears beside the job name on the Computers or Monitor page.

8.7.2.2 Best practice: Back up vSphere VMs restored using Rapid VM Restore

To prevent data loss, we highly recommend backing up vSphere virtual machines (VMs) that are restored using Rapid VM Restore. When a VM is first restored using Rapid VM Restore, it is dependent on a running Rapid VM Restore process and connections to the VRA and vault. If the connection is lost to the VRA or vault, the VM could be lost.

We also recommend backing up a VM immediately before migrating it, in case a problem occurs during the migration. You cannot back up a VM while it is being migrated, or migrate a VM while it is being backed up.

If you restore a VM and the original VM still exists in the vSphere environment, the VM will be restored as a copy of the original VM. You must modify your backup job to include the restored VM.

If you restore a VM and the original VM no longer exists in the vSphere environment, the VM will be restored with the same unique identifier (UUID) as the original VM. The restored VM will be backed up by the existing job, although the first backup might take longer than expected.

In a disaster recovery situation, if multiple VMs from the same backup job no longer exist in the vSphere environment, restore all missing VMs using Rapid VM Restore before running the backup job. If you run the job when only some of the VMs have been restored, the backup will skip the missing VMs and they will reseed when the backup job next runs.

8.7.3 Restore files, folders and database items using a vSphere Recovery Agent

You can restore files and folders from protected Windows VMs using the vSphere Recovery Agent (VRA).

During a file and folder restore, volumes from the selected VM are mounted as drives on the machine where the VRA is running. You can then:

- Share some or all of the mounted drives so that users can access files and folders from other machines.
- Sign in to the VRA machine and copy files and folders from the mounted drives.


Note: Files and folders on the mounted drives will be accessible to anyone on the VRA system, including non-Admin users. If you are concerned about security, secure the Agent machine and prevent users from logging in to the machine locally.

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

Note: You cannot restore specific files and folders from disks that are encrypted using Bitlocker or from Linux VMs.

To restore files, folders and database items using a vSphere Recovery Agent:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the vSphere environment with the VM that you want to restore, and expand its view by clicking the row.
3. Click the **Jobs** tab.
4. Find the backup job with the VM that you want to restore, and click **Restore** in the job's **Select Action** menu.
5. In the Choose What You Want to Restore dialog box, select **Files and Folders**.




6. Click **Continue**.

The Restore dialog box appears. If a potential ransomware threat was not detected when running the backup job, the most recent safeset for the job appears in the Safeset box.

If a potential ransomware threat was detected when running the backup job, a calendar with a list of safesets appears. "Potential Threat" appears beside each safeset where a potential ransomware threat was detected.


Note: If you are restoring data as described in [Restore data to a replacement computer](#) or [Restore data from another computer](#), "Potential Threat" does not appear for any safesets even if a potential threat was detected during a backup in the original vSphere environment.



7. To restore from an older safeset, if a calendar with a list of backups does not already appear, click the **Browse Safesets** button.  In the calendar, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset from which you want to restore.

8. In the **Items to Restore** box, select the check box for the VM with files or folders that you want to restore.

If a potential ransomware threat was detected on a VM, "Potential Threat" appears beside the VM name.



9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button.
10. In the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will automatically unshare. The **Idle time** can range from 2 to 180 minutes.

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

11. To use all available bandwidth for the restore, select **Use all available bandwidth**.

To ensure the best possible performance for your restore, we recommend selecting **Use all available bandwidth**.

12. Click **Run Restore**.

If you are restoring from a safeset that has a potential threat, a warning message appears. To cancel the restore, click **Cancel**. To continue restoring the volume, click **Continue**.

Volumes from the selected VM are mapped as drives on the machine where the VRA is running, and are available in a RestoreMount folder on the VRA machine.

13. (Optional) To allow access to the backup data from other servers, do one of the following on the machine where the VRA is running:

- Share one or more of the mapped drives.
- Share one or more directories from the RestoreMount folder.

14. Do one or both of the following:

- Copy files and folders that you want to restore from the mapped drives or shares.
- Use the Granular Restore for Microsoft Exchange and SQL application to find and restore items from Exchange and SQL Server database backups on the mapped drives or shares. You can restore Exchange mailboxes and messages to PST files or live databases, export SQL Server database items to live databases, and export SQL Server database items as SQL scripts. See the *Granular Restore for Microsoft Exchange and SQL User Guide*.

8.8 Restore Hyper-V data

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

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

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

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

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

8.8.1 Restore Hyper-V VMs

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

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

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

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

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

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

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

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

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

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

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

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

To restore Hyper-V VMs:

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

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

Virtual Machine	VM ID on Vault	Job	Backup Status	Date	Size	Actions
☐ ProtectedTestVM	HY_B4BHT3AIEFRLXWEPANQDZE		Not Protected			
☐ ProtectedVM	HY_HDKET6T7VBJ8J8B322A8464		Not Protected			
☐ supportedVM	HY_AVTU7JGVQVFNVRMT8TBJ8AQ	hyperjob	Completed	today at 3:33 PM	8.33 MB	Select Action
☐ VM	HY_HPVZBL7N8BETV2NVH2Q37XY	hyperjob	Completed	today at 4:30 PM	8.21 MB	Select Action
☐ VM1	HY_GWUJBYNNFAOXQNZDZNEHNDU		Not Protected			
☐ VM2	HY_VTLBN6ZKFB8ZTYER6MVCZBI		Not Protected			

- 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**. 

You can restore up to 50 VMs at a time.

- If the Choose What You Want to Restore dialog box appears, select **Virtual Machines**, and then click **Continue**.

Note: This dialog box only appears for Hyper-V Agent version 8.84 or later.

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



Virtual Machine	New VM Name	Backup Set	Destination	VM Identity
V1		Select Backup Set	Select Volume	Assign new if original exists
Win2012R2		Select Backup Set	Select Volume	Assign new if original exists

7. Do one of the following:

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

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

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

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

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

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

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

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

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

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

- In the **VM Identity** list, do one of the following:

- To restore the VM with a new GUID, click **Assign new identity**.

Note: If a node is down but has not been evicted from the cluster, you can only restore the VM using the **Assign new identity** option. This prevents the VM from being restored with the same GUID as a VM on the cluster node that is down.

- To restore the VM with its original GUID, click **Restore original identity**.

Note: If a VM with the original GUID exists in the Hyper-V environment, the restored VM will not overwrite the existing VM. Two VMs in a Hyper-V environment can have the same GUID if they are on separate hosts and are not configured for high availability.

- To restore the VM with its original GUID unless a VM with the original GUID exists in the Hyper-V environment, click **Assign new if original exists**. If a VM with the original GUID exists in the Hyper-V environment, the VM is restored with a new GUID.
- If the **Host** list is not shown, click the VM row to expand its view. Do one or more of the following:
 - To specify a host for the restored VM, click a host in the **Host** list.
 - To power on the VM after it is restored, select **Power on VM**.
 - To leave the restored VM powered off, clear **Power on VM**.
 - To connect the restored VM to the network, select **Enable network connectivity**.

If **Enable network connectivity** is selected, and the VM has a network adapter with the same name as a network adapter on the host, the VM will be automatically connected to the network.
 - To restore the VM without network connectivity, clear **Enable network connectivity**.

9. Click **Run Restore**.

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

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

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

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

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

Notes:

- A Hyper-V VM restored using Rapid VM Restore cannot be backed up until it is migrated to permanent storage. If you try to back up the restored VM before it is migrated to permanent storage, the following error occurs: *Unable to backup virtual machine "VMname" [VMID] because it is in RVMR state*. The VM's backup status in Portal is *Failed*.
- When you first restore a VM using Rapid VM Restore, it runs on the host that you select during the restore process. The VM could be lost if it is migrated to another host in a Hyper-V cluster. For this

reason:

- High availability is not enabled for a VM that is running using Rapid VM Restore. If high availability was enabled for the VM when it was backed up, high availability will be enabled for the VM after it is migrated to permanent storage using Portal. See [Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage](#).
- Do not enable high availability on a VM that is running using Rapid VM Restore.
- Only use Portal to migrate a VM that is running using Rapid VM Restore. Do not migrate the VM to another host in a cluster using Hyper-V Manager.
- Rapid VM Restore is available with Hyper-V Agent version 9.0 or later.

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

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

A grid lists available computers.

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

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


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

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

5. Find the VM that you want to restore, and click **Restore** in the job's **Select Action** menu.
6. In the Choose What You Want to Restore dialog box, select **Virtual Machine using Rapid VM Restore**.

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

For complete requirements, see [Hyper-V Rapid VM Restore requirements](#).



7. Click **Continue**.

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

8. To restore from an older safeset, click the **Browse Safesets** button. In the calendar that appears, click the date of the safeset from which you want to restore. In the safeset list to the right of the calendar, click the safeset from which you want to restore. The list only includes safesets on vaults that support Rapid VM Restores.
9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button.
10. In the Restore Settings box, do the following:
 - In the **Restored VM Name** box, type a name for the restored VM.
If you specify the name of a VM that already exists in the Hyper-V environment (e.g., the VM that was backed up), the restored VM will have the following name: *VMname-rvmr-yyyy-Mon-dd--hh-mm-ss*, where *yyyy-Mon-dd--hh-mm-ss* is the date and time when the VM was restored (e.g., VM-rvmr-2019-Nov-27--06-14-09).
 - In the **Select Volume** list, select a volume for writing changes while the VM is running using Rapid VM Restore but is not migrated to permanent storage. The amount of free space is shown for each volume in the list.
 - (Optional) In the Sub-Path box, type the folder path (e.g., *RestoredVMs\VM 1*) on the selected volume for writing changes while the VM is running using Rapid VM Restore.

If you do not specify a path, changes will be written to a folder with the name of the restored VM: *VMname-rvmr-yyyy-Mon-dd--hh-mm-ss*, where *yyyy-Mon-dd--hh-mm-ss* is the date and time when the VM was restored.

- In the **Destination Host** list, select a host for running the restored VM.
- Do one of the following:
 - To restore the VM with its power on, select the **Power on the VM** option.
 - To restore the VM powered off, clear the **Power on the VM** option. Restoring a VM with the power off can be useful if you want to verify or change the VM settings before powering on the VM.

Note: If you are restoring a VM that still exists in the Hyper-V environment, power off the original VM before the restore to avoid conflicts between the original VM and the restored VM.


- Do one of the following:
 - To connect the VM to the network, select **Connect to Network**.
 - To restore the VM without network connectivity, clear **Connect to Network**. Restoring a VM with the power off can be useful when restoring to a Hyper-V environment that does not have the original network. You can then verify the VM settings before connecting the VM to the network.

11. Click **Run Restore**.

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

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

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



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

12. Do one or more of the following:

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

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

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

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

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

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

Notes:

- We recommend using Portal to migrate VMs to permanent storage rather than using Hyper-V Manager. If you restore a VM using Rapid VM Restore and migrate it to a different host and storage using Hyper-V Manager, you will not be able to migrate the VM to permanent storage using Portal.
- A Hyper-V VM restored using Rapid VM Restore cannot be backed up until it is migrated to permanent storage.
- While a VM is being migrated, you cannot power on, power off, suspend or create a checkpoint for the VM using the Hyper-V Manager.


- If high availability was enabled for a VM when it was backed up, high availability will be enabled for the restored VM after it is migrated to permanent storage. However, specific high availability settings (e.g., preferred owner, failover and failback settings) are not applied.

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

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


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

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



2. If the protected VM name appears more than once in the VM Name list, check that the process ID for the VM that you want to migrate (recorded in [Restore a Hyper-V VM within minutes using Rapid VM Restore](#)) is shown in the middle of the dialog box. If the correct process ID is not shown, click another VM name in the VM Name list.
3. Click **Migrate VM**.

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



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

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

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

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

8.8.3 Restore Hyper-V files, folders and database items

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

During a file and folder restore, volumes from a protected VM are mounted in a RestoreMount folder on the server where the Management service is running. The folder is shared, and a UNC path to the share is provided in Portal. You can then access the share from a VM or server with network access to the server, and copy files and folders that you want to restore from the protected VM.

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

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

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

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

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

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

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

A grid lists available computers.

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

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

Virtual Machine	VM ID on Vault	Job	Backup Status	Date	Size	Actions
☐ ProtectedTestVM	HI_EI6H73ABF5RLXWVEPAHQDZE		Not Protected			
☐ ProtectedVM	HI_J4DH76T7YB8JB8322A6464		Not Protected			
☐ supportedVM	HI_A1U7U7GVJVCVFNRYN7BTBBAQ	hyper-vjob	Completed	today at 3:33 PM	8.23 MB	Select Action
☐ VM	HI_HPVZ6L7N6BET72NV6HQ2B7XY	hyper-vjob	Completed	today at 4:30 PM	8.21 MB	Select Action
☐ VM1	HI_GWU2YXNYFAXQZDZ7Z5NEHMDU		Not Protected			
☐ VM2	HI_YTLBW62KPB5ZTYER6MVCZBI		Not Protected			

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

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

5. Do one of the following:

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

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

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

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

Virtual Machine	Backup Set	Idle Time
VM	Select Backup Set	5
VM2	Select Backup Set	5

7. If you are restoring files and folders from multiple VMs, do one of the following:
 - If the VMs are protected with the same encryption password, select the **Use the same password for all VMs** check box. In the **Encryption Password** box, enter the data encryption password.
 - If the VMs are protected by jobs with different encryption passwords, clear the **Use the same password for all VMs** check box.
8. If you are restoring files and folders from multiple VMs, do one of the following:
 - To set the same idle time for each VM, select the **Use the same Idle Time for all VMs** check box. In the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will automatically unshare. The idle time value can be from 2 to 180 minutes.


Note: The drive will not unshare as long as new data is being copied. If you copy the same data from a shared drive more than once, the system could time out because no new data is being read.
 - To set a different idle time for each VM, clear the **Use the same Idle Time for all VMs** check box. You can then set the idle time for each VM in [Step 9](#).
9. For each VM from which you are restoring files and folders, do the following in the VM row:
 - In the **Backup Set** list, click the backup from which you want to restore. If you did not enter an encryption password for all VMs in [Step 7](#), enter the password in the **Encryption Password** box. Click **Apply**.
 - If you did not specify an idle time for all VMs in [Step 8](#), in the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will automatically unshare. The idle time value can be from 2 to 180 minutes.

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

10. Click **Run Restore**.

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

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



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

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

12. Use the UNC path on a VM or server with network access to the server where the Management service is running to do one or both of the following:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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


The Computers page shows registered computers.

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

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

The Configuration mode selection section appears.


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



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

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

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



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

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

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

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

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

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

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

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

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

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

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

Click **Save**. In the confirmation message box, click **Continue**.
12. If required, on the **Advanced** tab, on the **Notifications** tab, enter the SMTP password. See [Set up email notifications for backups on a computer](#).
13. Click **Save**. In the confirmation message box, click **Okay**.
14. Reinstall the Host service on each Hyper-V cluster host or standalone node. See [Install the Hyper-V Agent Host service](#).

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

8.9.1 Hyper-V disaster recovery and migration

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

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

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

- Hyper-V Agent Management service

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

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

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

Component Lost			Recovery Process
Hyper-V Agent Management service	Hyper-V environment (cluster or standalone)	Portal	
✓			1. Install the Hyper-V Agent Management service, and recover configuration information and jobs from the offline Hyper-V Agent. See Recover jobs and settings from an offline Hyper-V Agent. 2. If the IP address of the Hyper-V Agent Management server has changed, and Hyper-V Agent Host services were registered to the Management service using the IP address, do the following: a. Back up the host service log on each host where the Hyper-V Agent Host service is installed. The log is named AgentWorker.XLOG and is saved in a \Data\Logs subfolder in the Host service installation folder. <i>Note:</i> This step is suggested as a precaution, in case the host service log was not uploaded to the Management server. b. Uninstall and then reinstall each Hyper-V Agent Host service, and register each Host service to the Management service.
✓	✓		1. Create a new Hyper-V cluster or standalone host (i.e., Windows server with Hyper-V role). 2. Install the Hyper-V Agent Management service in the new Hyper-V environment, and recover configuration information and jobs from the offline Hyper-V Agent. Enter information for the new Hyper-V environment (rebuilt in Step 1) on the Cluster Credentials tab. See Recover jobs and settings from an offline Hyper-V Agent. 3. Install the Hyper-V Agent Host service on each host, and register it to the Management service. See Install the Hyper-V Agent Host service. 4. Restore VMs. See Restore Hyper-V VMs.

Component Lost			Recovery Process
Hyper-V Agent Management service	Hyper-V environment (cluster or standalone)	Portal	
✓	✓	✓	1. Restore the Portal and its protected database. See the <i>Portal Installation and Administration Guide</i>. 2. Rebuild the lost Hyper-V cluster or standalone host (i.e., Windows server with Hyper-V role). 3. Install the Hyper-V Agent Management service, and recover configuration information and jobs from the offline Hyper-V Agent. Enter information for the new Hyper-V environment (rebuilt in Step 1) on the Cluster Credentials tab. See Recover jobs and settings from an offline Hyper-V Agent. 4. Install the Hyper-V Agent Host service on each host, and register it to the Management service. See Install the Hyper-V Agent Host service. 5. Restore VMs. See Restore Hyper-V VMs.


8.10 Search for files to restore

When you restore data from a Windows, UNC, Linux or UNIX backup job, you can search for files to restore or exclude from the restore.

To search for files to restore:

1. In the Restore dialog box, click the **Search** button. 

The Search for files dialog box appears.



2. In the **Search Text** box, enter the file name to search for. You can include asterisks (*) as wildcard characters.
3. To search for files in a specific folder in the backup, enter the path in the **Search Path** box.
4. To search for files only in the specified folder, clear the **Search all subdirectories** check box.
5. Click **Search**.
The Results box lists files that match the search criteria.
6. In the **Results** box, select files to include or exclude. To select multiple consecutive items, press SHIFT while clicking the first and last items in the list. To select multiple items, press CTRL while clicking the items.
7. Do one of the following:
 - To restore the selected files, click **Include Selected**.
 - To exclude the selected files from the restore, click **Exclude Selected**.

8.11 Filter subdirectories and files when restoring data

When you restore data from a Windows, UNC, Linux or UNIX backup job, you can specify folders and files to restore or not restore from the backup.

By default, when you include a folder in a restore, the folder's subdirectories and files are also included. If you only want to restore some of a folder's subdirectories or files, you can add filters to the inclusion record. For example, you could add a filter so that files in a folder are only restored if they have the .doc or .docx extension.

By default, when you exclude a folder from a restore, the folder's subdirectories and files are also excluded. If you only want to exclude some of a folder's subdirectories or files, you can add filters to the exclusion record. For example, you could add a filter so that files in a folder are only excluded from the restore if they have the .exe extension.

To filter subdirectories and files when restoring data:

1. When restoring data from a Windows, UNC, Linux, or UNIX backup job, view the **Restore Set** box.

Restore Set				
	Folders Filter	Files Filter	Recursive	
+ Docs	e.g., a*, b*	*.*	☒	x edit
+ Documents an...	e.g., a*, b*	*.*	☒	x edit
+ Data	e.g., a*, b*	*.*	☒	x edit

2. If editable fields do not appear for a folder inclusion or exclusion record where you want to filter subdirectories and fields, click the **Edit** button in the folder row. 
3. In the **Restore Set** box, for each included folder where you want to include specific subdirectories or files, do one or more of the following:
 - To include specific subdirectories in the restore, in the **Folders Filter** field, enter the names of subdirectories to include. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only restore subdirectories if their names end with “-current” or start with “2015”, enter the following filter: *-current, 2015*
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To restore specific files, in the **Files Filter** field, enter the names of files to restore. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only restore files if they have the .doc or .docx extension, enter the following filter: *.doc, *.docx For example, to only restore files if they have the .pl extension, enter the following filter: *.pl
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To restore the specified folder but not its subfolders, clear the **Recursive** check box.
 - To restore the folder’s subdirectories, select the **Recursive** check box.
4. In the **Restore Set** box, for each excluded folder where you want to exclude specific subdirectories or files, do one or more of the following:
 - To exclude specific subdirectories from the restore, in the **Folders Filter** field, enter the names of subdirectories to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude subdirectories from a restore if their names end with “-old” or start with “2001”, enter the following filter: *-old, 2001*
Note: Asterisks (*) are the only supported wildcards in filter fields.

- To exclude specific files from the restore, in the **Files Filter** field, enter the names of files to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude files from a restore if they have the .exe or .dll extension, enter the following filter: *.exe, *.dll

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To exclude the specified folder, but not its subdirectories, clear the **Recursive** check box.
- To exclude the folder's subdirectories, select the **Recursive** check box.

5. Click **Run Restore**.

8.12 Restore data to a replacement computer

If you are replacing a system and want to migrate all data to a new computer (e.g., at the end of a lease) or in a disaster recovery situation, you can re-register the new computer with the vault as the old computer, and restore data from the old computer's backups. If the old computer backed up data to multiple vaults, you can use Portal version 8.50 or later to re-register the new computer to multiple vaults.

After you re-register a computer with a vault, you must:

- Edit each existing backup job and enter the encryption password for the backup job. See [Edit a backup job](#).
- Synchronize the jobs before they run successfully. See [Synchronize a job](#).


If you want to restore data to another computer without replacing the existing computer, you can restore data from another computer. See [Restore data from another computer](#).

Note: A different procedure is used to recover a protected Hyper-V environment. See [Recover jobs and settings from an offline Hyper-V Agent](#).


To restore data to a replacement computer:

1. Download and install an agent on the new or rebuilt computer.
2. On the navigation bar, click **Computers**.
A grid lists available computers.
3. Find the replacement computer to which you want to restore the data, and expand its view by clicking the computer row.
4. Click **Configure Manually**.
5. Click the Vault Settings tab.


6. Click **Re-register**.



6. In the Vault Settings dialog box, in the **Vault Profile** list, select the vault where the backup from the original computer was stored.



7. Click **Load Computers**.



8. In the list of computers, click the name of the computer where the data was backed up. Click **Save**.

9. In the confirmation dialog box, click **Yes**.

10. If the original computer backed up data to another vault, repeat [Step 6](#) to [Step 9](#) to download job information from the other vault.

11. After job information is downloaded, click the **Jobs** tab.

You must enter any passwords required for the job, including the encryption password.

For a vSphere Recovery Agent, you must also enter vCenter or ESXi host information on the vSphere Settings tab.

12. Find a job whose data you want to restore, and click **Restore** in the job's **Select Action** menu.

The remaining steps are the same as the steps for regular restores.

Note: If you are restoring data from a Windows Image or Local System job, "Potential Threat" appears in the Restore dialog box for any safesets where a potential threat was detected on the original system. If you are restoring data from a vSphere job, "Potential Threat" does not appear for any safesets in the Restore dialog box even if a potential threat was detected during a backup in the original vSphere environment.

IMPORTANT: After you re-register a computer with the vault, you must enter the encryption passwords for the computer's backup jobs and synchronize the jobs before they run successfully. See [Synchronize a job](#).

8.13 Restore data from another computer

You can restore some or all of a computer's backed up data to another (similar) computer.

Note: A different procedure is used to recover VMs from a protected Hyper-V environment. See [Recover jobs and settings from an offline Hyper-V Agent](#).

To restore data from another computer, you can redirect data from a backup job on the vault to a different computer. If the data was backed up using a plug-in, the destination computer must have the same plug-in installed. If the data was backed up using the Exchange Plug-in, the destination computer must also have Microsoft Exchange installed. If the data was backed up using the SQL Plug-in, the destination computer must also have Microsoft SQL Server installed.

The new computer then downloads information from the vault so that the data can be restored on the new computer. For example:

- Computer A backs up data using Job A
- Computer B restores data from Job A (computer A's data) to Computer B

Alternatively, if you wish to perform a disaster recovery on the same or replacement computer, you can re-register a newly configured computer after installing an operating system and an agent on it. See [Restore data to a replacement computer](#).

In some cases, where data streams are compatible, you may be able to restore to another computer with a similar (but not exactly the same) operating system. Different versions of the same operating system (e.g.,

Windows) are often compatible. Operating systems that share similar origins (e.g., Linux and Solaris) are also acceptable.

To restore data from another computer:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer to which you want to restore the data, and expand its view by clicking the computer row.
3. In the **Job Tasks** menu, click **Restore from Another Computer**.
The Restore From Another Computer dialog box opens.
4. In the **Vaults** list, select the vault where the backup is stored.
5. In the **Computers** list, select the computer with the backup from which you want to restore.
6. In the **Jobs** list, select the job from which you want to restore data.
7. Click **Okay**.

Portal attempts to download information about the selected job. After the job information is downloaded, the job appears on the computer's Jobs tab. You can then continue restoring data as you would in a regular restore.

Note: If you are restoring data from a Windows Image or Local System job, "Potential Threat" appears in the Restore dialog box for any safesets where a potential threat was detected on the original system. If you are restoring data from a vSphere job, "Potential Threat" does not appear for any safesets in the Restore dialog box even if a potential threat was detected during a backup in the original vSphere environment.

If Portal cannot download information about the selected job, the restore cannot continue. This can occur if the vault cannot be reached, job information cannot be retrieved, or a required plug-in is not installed on the destination computer. Make sure that any required plug-in is installed on the destination computer before you try again.

8.14 Advanced restore options

When restoring data, you can specify the following options:

Locked File Options

When restoring data from a Windows, Linux or UNIX local job, you can specify whether to overwrite locked files with restored files with the same names. You can select one of the following options:

- **Yes, overwrite locked files** – Files on the system that are locked during the restore are overwritten by restored files when the system restarts. You must select this option for a system state or system volume restore.

- **No, do not overwrite locked files** – Files on the system that are locked during the restore are not overwritten by restored files with the same name.

Streams

When running a backup, information is collected from your files in various streams. Original data created by a user is called a data stream. Other information, such as security settings, data for other operating systems, file reference information and attributes, are stored in separate streams.

When restoring data from a Windows, Linux or UNIX local job, you can select one of the following options:

- **Restore all streams** – Restores all information streams. This option is recommended if you are restoring files to a system with an identical platform.
- **Restore data streams only** – For cross-platform restores, restores data streams only. This option ensures that conflicts do not arise as a result of system-specific information streams.

Log Options

Select one of the following job logging levels from the list:

- **Files** — Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up.
- **Directory** — Provides less detail than the Files logging level. Provides information about folders that are backed up.
- **Summary** — Provides high-level information, including the vault and Agent version, and backup sizes.
- **Minimal** — Provides high-level information, including the vault and Agent version.

Changing the logging level only affects log files that are created from that point and after. It does not affect previously-created log files.

Note: For Image Plug-in jobs, the selected logging level does not affect the content of the logs.

Performance Options

To use all available bandwidth for the restore, select **Use all available bandwidth**.

Bandwidth throttling settings specify the amount of bandwidth consumed by an Agent for backups and restores. 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.

9 Monitor computers, jobs and processes

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

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

9.1 Monitor recent events using the Status Feed

You can view notifications about recent events in the Status Feed on the Dashboard. Notifications can appear for events such as:

- **Backups and restores.** A backup or restore notification indicates whether the process completed successfully, failed, or completed with problems.


- Backup job changes. A job change notification can appear when a job, schedule or event backup trigger is added, changed, or deleted.
- Agent configuration changes.
- Policy changes. A policy change notification can appear when a policy is added, changed, assigned, or deleted.
- User logins and logouts.
- Custom command changes. A notification can appear when a custom command is added, changed, or deleted.

You can also view messages from other users and from your service provider in the Status Feed.

To monitor recent events using the Status Feed:


1. On the navigation bar, click **Dashboard**.
2. In the Notification Center, click **Status Feed**.

Status Feed notifications appear in the center of the Dashboard, with the most recent notifications at the top of the list.




3. Do one or more of the following:
 - To change which events appear in the Status Feed for your current Portal session, click **Show Filter**. In the **Show Filter** list, click items until a check mark appears beside each item that you want to view, and then click **Hide Filter**.


Note: These Status Feed changes are lost when you sign out from Portal. To save Status Feed settings, see [Change your default page settings](#).



- To view more information about a backup or restore, click the backup or restore notification. The backup or restore log opens in the **History / Logs** window. See [View a job's process logs and safeset information](#).
- To view more information about a job that was created, click the job creation notification. The job opens in the Edit Job dialog box on the **Computers** page. See [Edit a backup job](#).
- To view more information about a job, Agent, or policy change, click the change notification. Information about the change appears in a **Changes** box. To print the information, click **Print**.



- To email a user who signed in or out, click the user's email address in the notification. An email to the user is created in your email client.
- To add a comment under an event notification, click **Comment** under the notification. In the **Comment** box, enter the comment. The comment appears under the event notification.



- To hide a notification in the Status Feed, click **Hide** under the notification.
- To load more notifications at the bottom of the Status Feed, scroll to the bottom of the page.

9.2 Monitor backups and computers using the Current Snapshot

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

To monitor backups and computers using the Current Snapshot:

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

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

- **Potential Threats** — Number of backup jobs where a potential ransomware threat was detected. See [Manage potential ransomware threats](#). If a ransomware scan did not run successfully, the backup job could appear in the "Backups with Warnings" category. Please see the backup logs for more information.
- **Backups Requiring Attention** — Number of backup jobs where the last backup attempt failed, completed with errors, did not back up any files, reached a license limit, was cancelled or had a potential ransomware threat.
- **Missed Backups** — Number of backup jobs that have not run for seven days.
- **Backups With Warnings** — Number of backup jobs where the last backup attempt completed with warnings, was deferred, was deferred with warnings or was skipped. This category also includes backup jobs that have never run.
- **Computers Requiring Reboot** — Number of computers with a pending reboot.
- **Offline Computers** — Number of computers that are not currently in contact with Portal. A computer can be offline if it is turned off, if the Agent has been uninstalled from the system, or if the system no longer exists.
- **Computers Scheduled for Deletion** — Number of computers that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.

- **Computers With Certificate Failures** — Number of computers reporting a vault or vSphere environment certificate failure. See [Resolve certificate failures](#).
 - **Total Computers** — Total number of computers in the site.
 - **Successful Backups** — Number of backup jobs where the last backup attempt completed without errors, warnings, or deferrals.
 - **Jobs Scheduled for Deletion** — Number of jobs that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.
2. To view computers in a particular site, click the sites box in the top right of the Current Snapshot box. In the menu, click the site that you want to view.
Computers in the selected site appear on the Computers page.
 3. To view information about backup jobs or computers in one of the categories, click the category.
If you click **Potential Threats, Backups Requiring Attention, Missed Backups, Backups With Warnings** or **Successful Backups**, backup jobs in the category appear on the Monitor page.
If you click **Computers Requiring Reboot, Offline Computers, Computers Scheduled For Deletion, Computers With Certificate Failures** or **Total Computers**, computers in the category appear on the Computers page.

9.3 Monitor storage usage using Site Usage charts and emailed alerts

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

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

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

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


Support users can also view Site Usage charts.

To monitor storage usage using Site Usage charts:

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

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

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



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

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

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

9.4 View computer and job status information

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

To view computer and job status information:

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

The Computers page shows registered computers.

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

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


-  OK — Indicates that all jobs on the computer ran without errors or warnings. This status also appears if the computer has a new job that is less than 24 hours old that has not run yet.
 -  OK with warnings — Indicates that one or more of the computer's jobs completed with warnings. This status also appears if the computer has a new job that is more than 24 hours old that has not run yet.
 -  Attention — Indicates that one or more of the computer's jobs failed or completed with errors.
 -  Unconfigured — Indicates that no jobs have been created for the computer.
 -  Scheduled for deletion — Indicates that the computer is scheduled for deletion from Portal and from vaults. This status only appears in Portal instances where the data deletion feature is enabled.
 -  Certificate failure — Indicates that the agent is reporting a certificate change.
2. Find the computer for which you want to view status information, and click the row to expand its view.
 3. View the **Jobs** tab.

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

Name	Job Type
 1 job1	Local System
 1 job2	Local System

If a Rapid VM Restore is running for a vSphere Recovery Agent (VRA) job, a Rapid VM Restore symbol  appears beside the job name, along with the number of Rapid VM Restores that are running.

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



Name	Job Type
1_VRAjob	vSphere
1_VRAjob2	vSphere

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

-  Completed — Indicates that the last backup completed successfully, and a safeset was created.
-  Completed with warnings — Indicates that the last backup completed and a safeset was created, but problems occurred during the backup. For example, a warning could indicate that a file or volume that was selected in the backup job was not available for backup. A warning could also indicate that a ransomware scan did not run successfully.
-  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.

-  Skipped — Indicates that a backup was skipped. Backups are sometimes skipped if they are scheduled to run multiple times per day. See [Skipped backups](#).
-  Never Run — Indicates that the backup job has never run.
-  Missed — Indicates that the job has not run for 7 days.
-  Completed with errors — Indicates that the backup completed and a safeset is available for restore, but problems occurred. Typically, this status indicates that not all of the data was backed up. This status can also indicate that a potential ransomware threat was detected.
-  No Files backed up — Indicates that no files were backed up during the last backup attempt
-  Failed — Indicates that the backup failed and no safeset was created.
-  Cancelled — Indicates that the backup was cancelled.
-  Scheduled for Deletion — Indicates that the job is scheduled to be deleted from Portal and job data is scheduled to be deleted from all vaults on the date shown in the Date column. This backup status is only possible in Portal instances where the data deletion feature is enabled. See [Delete a backup job and delete job data from vaults](#).

If **Potential Threat** appears after the status in the Last Backup Status column, a potential ransomware threat was detected while running the backup job. See [Manage potential ransomware threats](#).

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

9.5 View skipped rates and backup status histories

When an agent is backing up data to a Director version 8.60 or later vault, backups that are scheduled to run multiple times per day are skipped in some cases. To determine whether backups were skipped, users can view email notifications, the Computers page and Monitor page, and the Daily Status report. See [Skipped backups](#).

In some Portal instances, users can also view the following skipped rate and backup status history information:

- Skipped rate for a job. If a backup was skipped for a job in the 48 hours before the most recent backup attempt, a skipped rate appears for the job on the Computers page and Monitor page. The skipped rate is the percentage of backups that were skipped in the 48 hours before the last backup attempt, and is calculated using the following formula:

$$\text{jobSkippedRate} = \text{numberOfSkippedBackups} / \text{numberOfBackupAttempts}$$

Where:

- *numberOfSkippedBackups* is the number of backups that were skipped for the job during the 48 hours before the last backup attempt.
- *numberOfBackupAttempts* is the total number of backup attempts for the job during the 48 hour period, including skipped, in-progress, deferred, canceled, failed and completed backups.

If no backups were skipped for a job in the 48 hours before the last backup attempt, or if the last backup attempt occurred more than seven days ago, a skipped rate is not shown for the job.

- Skipped rate for a computer. If a skipped rate is reported for one or more jobs on a computer, the highest skipped rate on the computer appears on the Computers page.
- 48-hour backup status history for a job. If a skipped rate appears for a job on the Computers or Monitor page, you can view the job's backup history for the 48 hours before the last backup attempt. The status history shows the dates and times of backup attempts, and indicates the status of each backup attempt (e.g., skipped, in-progress, completed or failed). You can export the status history in comma-separated values (.csv), Microsoft Excel (.xls) or Adobe Acrobat (.pdf) format.

To view skipped rates and backup status histories, see [View skipped rates and backup status histories on the Computers page](#) and [View skipped rates and backup status histories on the Monitor page](#).

9.5.1 View skipped rates and backup status histories on the Computers page

To prevent schedule overloads, backups that are scheduled to run multiple times per day are skipped in some cases. Users can obtain skipped backup information through email notifications, on the Computers page, and in the Daily Status report. See [Skipped backups](#).

In some Portal instances, users can view skipped backup rates for jobs and computers on the Computers page, and view and export a job's backup status history for the 48 hours before the last backup attempt. For more information, see [View skipped rates and backup status histories](#).

To view skipped rates and backup status histories on the Computers page:

1. Click **Computers** on the navigation bar.

A value appears in the Skipped column for any computer where at least one job has a skipped rate. If more than one job on a computer has a skipped rate, the highest skipped rate appears in the Skipped column.

Note: If the Skipped column does not appear, skipped rates and 48-hour backup status histories are not available in your Portal instance.

Computer	Groups	Description	Policy	Version	Operating System	Status	Skipped	Availability	Site Name
SERVER1	All			9.00.1012	Windows Server 2016 x64	Attention	31%	Online	Site1


2. Find a computer with a value in the Skipped column, and click the computer row to expand its view.

On the Jobs tab, a value appears in the Skipped Rate column for any job where a backup was skipped in the 48 hours before the last backup attempt, and the last backup attempt occurred in the last seven days.

Name	Job Type	Description	Last Backup Status	Date	Skipped Rate	Action
Image	Image		Completed	today at 12:53 PM		Select Action
Image2	Image		Completed	today at 1:22 PM	31%	Select Action

3. To see which backups were skipped in the 48 hours before the last backup attempt for a job, click the job's Skipped Rate value.


The 48-Hour Status History for the job shows the date, time and status (e.g., skipped, in-progress, completed or failed) of each backup attempt.



Date and Time	Status
10/28/2021 14:17:25 GMT-07:00	Completed
10/28/2021 14:08:24 GMT-07:00	Skipped
10/28/2021 13:22:06 GMT-07:00	Completed
10/28/2021 13:08:12 GMT-07:00	Skipped
10/28/2021 12:16:06 GMT-07:00	Completed
10/28/2021 12:08:03 GMT-07:00	Skipped
10/28/2021 11:15:10 GMT-07:00	Completed
10/28/2021 11:08:30 GMT-07:00	Skipped
10/28/2021 10:21:46 GMT-07:00	Completed
10/28/2021 10:08:21 GMT-07:00	Skipped
10/28/2021 09:19:04 GMT-07:00	Completed
10/28/2021 09:08:09 GMT-07:00	Skipped
10/28/2021 08:20:37 GMT-07:00	Completed
10/28/2021 07:18:15 GMT-07:00	Completed
10/28/2021 06:15:20 GMT-07:00	Completed
10/28/2021 05:14:49 GMT-07:00	Completed
10/28/2021 04:14:51 GMT-07:00	Completed
10/28/2021 03:14:47 GMT-07:00	Completed
10/28/2021 02:14:47 GMT-07:00	Completed

If you want to export the status history, click the **Export** box. In the list that appears, click one of the following formats for the exported data:

- CSV (comma-separated values)
- XLS (Microsoft Excel)
- PDF (Adobe Acrobat)



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

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

9.5.2 View skipped rates and backup status histories on the Monitor page

To prevent schedule overloads, backups that are scheduled to run multiple times per day are skipped in some cases. Users can obtain skipped backup information through email notifications, on the Computers page, and in the Daily Status report. See [Skipped backups](#).

In some Portal instances, users can view skipped backup rates for jobs on the Monitor page, and view and export a job's backup status history for the 48 hours before the last backup attempt. For more information, see [View skipped rates and backup status histories](#).

To view skipped rates and backup status histories on the Monitor page:

1. Click **Monitor** on the navigation bar.

A value appears in the Skipped column for any job where a backup was skipped in the 48 hours before the last backup attempt, and the last backup attempt occurred in the last seven days.

Computer	Job Name	Job Type	Last Backup Date	Last Backup Status	Skipped	Last Completed Backup	Backup Size	Site
SERVER1	Image	Volume Image	today at 2:23 PM	Completed	31 %	today at 2:10 PM	20.68 GB	Site1

2. To see which backups were skipped in the 48 hours before the last backup attempt for a job, click the job's Skipped value.


The 48-Hour Status History for the job shows the date, time and status (e.g., skipped, in-progress, completed or failed) of each backup attempt.

Date and Time	Status
10/28/2021 14:17:25 GMT-07:00	Completed
10/28/2021 14:08:24 GMT-07:00	Skipped
10/28/2021 13:22:06 GMT-07:00	Completed
10/28/2021 13:08:12 GMT-07:00	Skipped
10/28/2021 12:16:06 GMT-07:00	Completed
10/28/2021 12:08:03 GMT-07:00	Skipped
10/28/2021 11:15:10 GMT-07:00	Completed
10/28/2021 11:08:30 GMT-07:00	Skipped
10/28/2021 10:21:46 GMT-07:00	Completed
10/28/2021 10:08:21 GMT-07:00	Skipped
10/28/2021 09:19:04 GMT-07:00	Completed
10/28/2021 09:08:09 GMT-07:00	Skipped
10/28/2021 08:20:37 GMT-07:00	Completed
10/28/2021 07:18:15 GMT-07:00	Completed
10/28/2021 06:15:20 GMT-07:00	Completed
10/28/2021 05:14:49 GMT-07:00	Completed
10/28/2021 04:14:51 GMT-07:00	Completed
10/28/2021 03:14:47 GMT-07:00	Completed
10/28/2021 02:14:47 GMT-07:00	Completed

If you want to export the status history, click the **Export** box. In the list that appears, click one of the following formats for the exported data:

- CSV (comma-separated values)
- XLS (Microsoft Excel)

- PDF (Adobe Acrobat)



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

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


9.6 View an unconfigured computer's logs

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


To view an unconfigured computer's logs:

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


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



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



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



4. Do one of the following:

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

The log appears in a new browser tab.


9.7 View current process information for a job

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


You can also view information about running and recent Rapid VM Restore and migration processes for a vSphere Recovery Agent (VRA) job. For more information, see [Restore a vSphere VM within minutes using Rapid VM Restore](#).

To view current process information for a job:

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




- On the Monitor page, click the Process Details symbol  or Rapid VM Restore symbol  beside the job name.



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

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




If you clicked a Rapid VM Restore symbol, the Process Details dialog box lists running and recent Rapid VM Restore and migration processes for the VRA job.

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

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

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



9.8 Monitor backups using email notifications

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. See [Set up email notifications for backups on a computer](#).

In some Portal instances, email notifications are configured centrally for Windows systems with Agent version 8.0 or later, Linux systems with Agent version 8.10a or later, and vSphere Recovery Agent 8.40 or later, instead of separately for each computer. See [Set up email notifications for backups on multiple computers](#).

When email notifications are configured centrally in a Portal instance, admin users can also receive email notifications when the encryption password changes for a backup job or when a potential ransomware threat is detected during a backup. See [Set up email notifications for encryption password changes](#) and [Set up email notifications for potential ransomware threats](#).

9.8.1 Set up email notifications for backups on a computer

To set up email notifications for a computer:

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

The Computers page shows registered computers.

2. Find the computer for which you want to configure email notifications, and click the computer row to expand its view.
3. On the **Advanced** tab, click the **Notifications** tab.

If the Notifications tab does not appear, email notifications for the computer's backups are configured centrally instead of for each computer. See [Set up email notifications for backups on multiple computers](#).

Note: If email notifications were set up for the computer before centrally-configured email notifications were enabled in the Portal instance, the Notifications tab can appear for the computer.

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

The screenshot shows the 'Advanced' settings page in the Server Backup Portal. The 'Notifications' tab is selected. At the top, there are three checkboxes: 'On Successful Completion', 'On Failure', and 'On Error'. Below these are two main configuration areas. The 'SMTP Settings' area contains four text input fields: 'Email "From" Address:', 'Outgoing Mail Server (SMTP):', 'Recipient Address(es):', and 'Outgoing Server Port (SMTP):'. The 'SMTP Credentials (if required)' area contains three text input fields: 'User Name:', 'Password:', and 'Domain:'.

4. Select one or more of the following checkboxes:

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

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

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

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

5. Click **Save**.

9.8.2 Set up email notifications for backups on multiple computers

By default in some Portal instances, Admin users receive emails when backups fail, or are canceled, deferred, missed, skipped or completed. Admin users can select backup statuses for which they want to receive email notifications.

These email notifications are sent for Windows systems with Agent version 8.0 or later, Linux systems with Agent version 8.10a or later, AIX systems with Agent version 9.00 or later and vSphere Recovery Agent 8.40 or later, instead of separately for each computer.

When email notifications are configured centrally in a Portal instance, additional notification email addresses can be specified for each child site.


Note: Email notifications selected in Admin users' profile settings are only sent in English. Email notifications for child site email addresses are supported in multiple languages.

In Portal instances where Admin users do not automatically receive email notifications, notifications must be configured separately for each computer. See [Set up email notifications for backups on a computer](#).

To set up email notifications for backups on multiple computers:

1. When signed in as an Admin user, click your email address at the top right of the Portal page.

The user menu appears.



2. Click **Profile Settings**.

Your user profile appears. If your profile includes an Email Notification Settings section with a list of backup events (e.g., Backup Canceled, Backup Completed, Backup Skipped), you can select events for which you want to receive emails.

If Email Notification Settings do not appear, you must set up notifications separately for each computer. See [Set up email notifications for backups on a computer](#).

If an Encryption Password Changed option appears, you can choose to receive email notifications when encryption passwords change in your site. See [Set up email notifications for encryption password changes](#).

3. In the Email Notification Settings list, select any of the following events for which you want to receive emails:

- Backup Cancelled
- Backup Completed
- Backup Completed with Errors
- Backup Completed with Warnings
- Backup Deferred
- Backup Failed
- Backup Missed
- Backup Skipped

Note: Backups are sometimes skipped if they are scheduled to run hourly or multiple times per day. See [Skipped backups](#).

4. Click **Update notifications**.

9.8.3 Set up email notifications for encryption password changes

In some sites, Admin users can choose to receive emails when job encryption passwords change.

Admin users in a parent site can receive emails when job encryption passwords change in the parent site and in its child sites. Admin users in a child site can receive emails when job encryption passwords change in the child site only.

Super users specify whether Admin users in a site can receive encryption password change emails.


When email notifications are configured centrally in a Portal instance, additional notification email addresses can be specified for each child site.

Note: Email notifications selected in Admin users' profile settings are only sent in English. Email notifications for child site email addresses are supported in multiple languages.

To set up email notifications for encryption password changes:

1. When signed in as an Admin user, click your email address at the top right of the Portal page.

The user menu appears.



2. Click **Profile Settings**.

Your user profile appears. If your profile includes an Email Notification Settings section with an Encryption Password Changed option, you can choose to receive emails when encryption passwords change.

3. In the Email Notification Settings list, select the **Encryption Password Changed** option.
4. Click **Update notifications**.

9.8.4 Set up email notifications for potential ransomware threats

When email notifications are configured centrally in a Portal instance instead of separately for each computer, Admin users can receive emails when potential ransomware threats are detected on Windows servers in Image or Local System backup jobs, and on Windows VMs in vSphere backup jobs. See [Add a Windows Local System backup job](#), [Add a Windows Image backup job](#), [Add a vSphere backup job](#) and [Manage potential ransomware threats](#).

Admin users in a parent site can receive emails when potential threats are detected in the parent site and in its child sites. Admin users in a child site can receive emails when potential threats are detected in the child site.


When email notifications are configured centrally in a Portal instance, additional notification email addresses can be specified for each child site.

Note: Email notifications selected in Admin users' profile settings are only sent in English. Email notifications for child site email addresses are supported in multiple languages.

To set up email notifications for potential ransomware threats:

1. When signed in as an Admin user, click your email address at the top right of the Portal page.

The user menu appears.



2. Click **Profile Settings**.
3. In the Email Notification Settings list, select the **Potential Threats** option.
4. Click **Update notifications**.

9.9 View a job's process logs and safeset information

To determine whether a backup, restore or other process completed successfully, determine whether a safeset was flagged with a potential ransomware threat, or to determine why a process failed, you can view a job's process logs.

Note: When you run an Exchange database restore with the **Start Hard Recovery** option selected, the process of restoring database files is recorded in the process logs. The process of replaying transaction logs into the database is recorded in the Windows Event Viewer.

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 computers.


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

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


SERVER 8.70.9504 Windows Ser... OK with warnings Online Site1					
Jobs Vault Settings Advanced					
Select Job Task					
Name	Job Type	Description	Last Backup Status	Date	Action
BMRjob	Local System		Completed	today at 9:31 AM	Select Action
CloudServerBackup	Local System	This backup protects your entire C drive. It will be backed up to the cloud, per your retention schedule.	Completed	yesterday at 7:32 PM	Select Action
Job	Local System		Completed with warnings	yesterday at 10:45 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** or **Logs**.
 - In the **Last Backup Status** column, click the job status.

The History / Logs or Logs window lists the most recent backups, restores and other processes on the computer.



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




6. To only show errors and warnings in the log, click the **Errors and Warnings Only** option at the top right of the window.

7. 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.

Note: The Job Details tab does not appear for a Hyper-V Agent.

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.



Catalog Number:	00000067		
Location:	192.168.6.59		
Status:	Online		
Backup Time:	11/7/2018 10:45 PM		
Storage Size:	Original: 3.83 GB	Deltized: 91.97 MB	Compressed: 47.02 MB
Backup Type:	Full		
Retention:	Days: 30	Copies: 30	Archive Days: 0
Encrypted:	Yes		
Compressed:	Yes		
Media Type:	Vault		
Tape Label:			
Expiry:	12/7/2018 10:45 PM		

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

You can view recent job backup statuses on the Monitor page in Portal. If a potential ransomware threat was detected when running a job, "Potential Threat" appears beside the job's backup status.

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

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

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

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

Last Backup Date column. You can only schedule a report to be emailed from the All Jobs view on the Monitor page.

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


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

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

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

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

2. To change which job backup statuses appear, click a view or enter filter criteria. For more information, see [Filter records on a page](#).
3. To view information for a job or computer on the Computers page, click the name of a job or online computer.
4. To view a job's logs in the History/Logs window, click the job's last backup status.
5. To export job backup status data from the Monitor page, click the **Export** box. In the list that appears, click one of the following formats for the exported data file:
 - CSV (comma-separated values)
 - XLS (Microsoft Excel)
 - PDF (Adobe Acrobat)



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

6. To email a Job Monitor Export report, do the following when signed in as an Admin or Support user:
 - a. To specify which job backup statuses appear in the report, click a view or enter filter criteria.
 - b. Click the **Email/Schedule** box. In the **Send Report** list that appears, click **Email Once**.
 - c. In the Email Once dialog box, do the following:

- i. In the **To** box, type one or more email addresses for sending the report. Use commas to separate multiple email addresses.
 - ii. In the **Subject** box, type a subject for the report email.
 - iii. In the Attachment list, click one of the following formats for the emailed report:
 - CSV (comma-separated values)
 - Excel (Microsoft Excel)
 - PDF (Adobe Acrobat)
 - d. Click **Okay**.
7. To schedule a Job Monitor Export report to be emailed, do the following when signed in as an Admin or Support user:
 - a. To specify which job backup statuses appear in the scheduled report, enter filter criteria in any column except the Last Backup Date column.
Note: You can only schedule a report to be emailed when the All Jobs view is selected on the Monitor page.
 - b. Click the **Email/Schedule** box. In the **Send Report** list that appears, click **Schedule New Report**.
 - c. In the Email/Schedule dialog box, do the following:
 - In the **To** box, type one or more email addresses for sending the report. Use commas to separate multiple email addresses.
 - In the **Report Name** box, type a name for the scheduled report. This name appears in the **Email/Schedule** list.
 - In the **Subject** box, type a subject for the email.
 - In the **Attachment** list, click one of the following formats for the emailed report:
 - CSV (comma-separated values)
 - Excel (Microsoft Excel)
 - PDF (Adobe Acrobat)
 - d. Do one of the following:
 - To email the report on specific days each week, in the **Frequency** list, click **Daily**. In the day row, select the days when you want to email the report each week.

Frequency

Daily ▼

S	M ✓	T	W ✓	T	F ✓	S



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

Frequency

Weekly ☐

S ☒	M ☐	T ☐	W ☐	T ☐	F ☐	S ☐

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

Frequency

Monthly ☐

1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
Last Day						☒

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

9.10.1 Delete a Job Monitor Export report schedule

Admin users and Support users can delete scheduled Job Monitor Export reports. Scheduled reports appear in the **Email/Schedule** list on the Monitor page.

To delete a report schedule:

1. On the navigation bar, click **Monitor**.
2. Click the **Email/Schedule** box.
Scheduled reports appear in the **Currently Scheduled** list.
3. Click the Delete button beside the report schedule that you want to delete.

9.11 View a Hyper-V VM's backup history and logs

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

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


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

The Computers page shows registered computers.

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


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

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




Virtual Machine	VM ID on Vault	Job	Backup Status	Date	Size	Actions
ProtectedTest01/VM	HY_E4GH73AIBF8RL4WIEPAHVKDCE		 Not Protected			
Protected1/VM	HY_4DK8767TVU8J3B83224846L4		 Not Protected			
supported1/VM	HY_A1V7J5V1VCV1K1YK1R1M78T818AQ	hyper-v-job	 Completed	today at 3:33 PM	8.23 MB	Select Action
VM	HY_HPY26L7N8BETJ2V1EH2Q2B70Y	hyper-v-job	 Completed	today at 4:30 PM	8.21 MB	Select Action
VM1	HY_GWU3YX1YFA0X2NT23NEHMHOU		 Not Protected			
VM2	HY_WTLBW62HF85Q7Y5R4UWY22BI		 Not Protected			

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



- To view a log file for a process on the selected date, click the process. The log file appears.
- 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.
- To view safeset information for a particular VM backup, click the **Backup History** tab. The tab shows information for the VM's most recent backup.

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



9.12 Determine whether an agent has been configured automatically

Beginning with Portal 8.89 and Windows Agent 8.90a, backups can be configured automatically on Windows servers based on job templates. The Windows agent must be installed with a default encryption password and registered to a Portal site where agent auto-configuration is enabled.

When agent auto-configuration is enabled, Portal finds and auto-configures eligible Windows agents every three minutes. When an agent is successfully configured, a backup job for the Windows server is created that has:

- The name, description, settings and schedules specified by the job template selected for the child site.
- A randomly-assigned time for running backups with “Days of Week” and “Days of Month” schedules. By default, the time will be between 8 PM and 8 AM, but a different time window might be specified in your Portal instance.
- The vault profile selected for the child site.
- The default data encryption password specified when the agent was installed.

To determine whether an agent has been configured automatically, you can view the agent on the Computers page in Portal. If the agent has not been configured, an auto-configuration status message appears.


To determine whether an agent has been configured automatically:

1. Sign in to Portal as an Admin user in a child site where agent auto-configuration is enabled, or as a parent site Admin user who can manage the child site.
2. Do one of the following:
 - On the navigation bar, click **Computers**. On the Computers page, find the computer for which you want to view the auto-configuration status. Click the computer row to expand it.
 - On the navigation bar, click **Dashboard**. In the Notification Center, click **What’s New**. In the center of the dashboard, find the notification of the computer and click its **Configure Now** link. The computer row is shown on the Computers page.


Note: If the computer is unconfigured and offline, you cannot expand the computer row. An agent cannot be configured automatically if it is offline after it registers to Portal. If an agent stays offline for seven days after registering to Portal, backups cannot be automatically configured on the server. If the agent comes online after seven days, an "automatic configuration has failed" message appears for the agent.

If a job appears in the computer row, the agent has been configured. If the agent was configured automatically, the job has the name, settings and schedule specified by the job template for the child site.

If a "queued for automatic configuration" message appears in the computer row, the agent is waiting to be configured. Portal will attempt to configure the agent in the next three minutes.




If a "needs to be configured" message appears in the computer row, the agent is not eligible for auto-configuration and you must create a backup job manually. For more information, see [Reasons that agents are not eligible for auto-configuration](#).



If a "queued for automatic configuration but x of 3 attempts have failed" message appears in the computer row, an auto-configuration attempt failed. Portal will attempt to configure the agent again in the next three minutes.


Note: Portal attempts to configure backups on a server three times after an agent is registered to Portal.

If error information is available, the second sentence of the message (shown in a green rectangle, below) describes the issue.



If an "automatic configuration has failed" message appears, three auto-configuration attempts failed. Please create a backup job manually.

If error information is available, the second sentence of the message (shown in a green rectangle, below) describes the issue.



9.12.1 Reasons that agents are not eligible for auto-configuration

If a "needs to be configured" message appears for an agent on the Computers page, the agent is not eligible for auto-configuration. This can occur if:

- The agent is not a Windows agent, or is a Windows agent version that is not supported for auto-configuration. Agent auto-configuration is supported with Windows Agent version 8.90a or later.
- An encryption password was not specified when the agent was installed. See [Install the Windows Agent and plug-ins](#).
- An Image job template is selected for the child site, but the Image Plug-in is not installed with the agent.
- The agent has more than one vault registration (vault setting), or has a vault registration that does not match the vault profile selected for agent auto-configuration in the site.
- The agent has one or more backup jobs. This could occur if a backup job was manually configured before the Agent Configuration task ran in Portal.
- The agent is registered to a site where agent auto-configuration is not fully set up (e.g., a job template is not selected).

IMPORTANT: Auto-configuration must be enabled in the child site when the agent first registers to Portal. An agent will not be automatically configured if you enable auto-configuration after the agent is registered to Portal.

10 View and schedule reports

Admin users and Support users in Portal can view and schedule reports that provide detailed information about backups, vault space usage, and activities in a site. For descriptions of available reports, see [Portal report descriptions](#).

The Daily Status report, which includes backup status information for the previous 24 hours, can be scheduled and emailed to users. See [Daily Status Report](#) and [Schedule the Daily Status Report](#).

Other reports can be viewed in Server Backup Portal. Each report has a default view, which shows data for all protected computers and environments in a site. For some reports, Admin users and Support users can change the report date range, specify which data columns and records to show, and save the resulting report as a customized view. See [View a report](#) and [Save a report view](#). Admin users and Support users can also export or email report data in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format. See [Export a report](#), [Email a report](#) or [Schedule an emailed report](#).

We recommend turning off macros in Microsoft Excel when using Portal, particularly if you export reports in XLS or CSV format and open these reports in Excel.

Note: Beginning in Portal 9.00, the Backup Verification, Daily Status and Custom Command reports appear on the Reports page by default. Additional configuration is required for other reports to appear. For more information or assistance, please contact your service provider or Portal administrator.

10.1 Portal report descriptions

Admin users and Support users can view or schedule the following reports through Portal:

- Backup Verification Report. This report is available beginning with vSphere Recovery Agent (VRA 9.00) and Portal 9.00, and indicates whether each Windows VM in a vSphere environment was backed up correctly and can be restored. See [Backup Verification Report](#).
- Daily Status Report. This report includes backup status information for the previous 24 hours, including missed and skipped backups and running jobs for computers where Agent version 8.0 or later is installed. See [Daily Status Report](#).

The Daily Status report can be scheduled and emailed to users. It cannot be viewed in Portal. See [Schedule the Daily Status Report](#).

- Backup Status Report. This report provides information about the last time each backup job ran during a specified time period, including the amount of data that was backed up. See [Backup Status Report](#).
- Backup Details Report. This report provides detailed information about each time a backup job ran during a specified time period, including the amount of data that was backed up, the amount of data that changed since the last backup, and whether the backup was successful or not. See [Backup Details Report](#).

- Activity Details Report. This report provides information about activities in a site during a specified time period, including backups and restores. See [Activity Details Report](#).
- Usage Summary Report. This report indicates how much data is saved on the vault for each backup job in a site. See [Usage Summary Report](#).

Usage Summary Report information can also be viewed in a configurable chart. In addition to vault space usage, the chart can show activities that could affect space usage on the vault (e.g., backup job changes and computer changes). See [View a Usage Summary Report chart](#).

Note: The Usage Summary Report might not appear for every site.

- Aggregated Usage Summary Report. The Aggregated Usage Summary Report, available in some Portal instances, shows the total amount of data backed up for a site. You can then navigate in a tree structure to see usage data for each child site, if any, and computer. See [View the Aggregated Usage Summary Report](#).

Note: The Aggregated Usage Summary Report is only available in Portal instances that obtain data from billing systems.

- Custom Command Report. This report lists all custom commands for a site. See [Custom Command Report](#).
- Job Monitor Export. This report, which includes job backup status data, is available beginning in Portal 9.20 and can be emailed and scheduled from the Monitor page. See [View, export and email backup statuses on the Monitor page](#).

Note: If only the Backup Verification, Daily Status and Custom Command reports appear on the Reports page in your Portal instance, additional configuration might be required. For more information or assistance, please contact your service provider or Portal administrator.

10.1.1 Backup Verification Report

The Backup Verification Report indicates whether Windows VMs can be restored from vSphere backups. This report is available in Portal version 9.00 or later.

The report shows the results of backup verification processes, available with vSphere Recovery Agent (VRA) version 9.00 or later. When backup verification settings are entered for a VRA and backup verification is enabled for a vSphere backup job, the VRA backs up VMs in the job and then checks whether each Windows VM can be restored from the backup. See [Backup verification for vSphere VMs](#) and [vSphere Rapid VM Restore and backup verification requirements](#).

The following table lists and describes data columns that appear in the Backup Verification Report.

Backup Verification Report Data Column	Description
VM Name	Name of the VM that was backed up <i>Note:</i> If a VM is included in multiple backup jobs where verification is enabled, the VM can appear multiple times in the report. <i>Note:</i> vSphere allows two or more VMs in a vSphere environment to have the same name if each VM is located in its own folder. If multiple VMs have the same name, you cannot differentiate between the VMs in Portal or in the Backup Verification Report. Consider renaming the VMs in this case.
Site	Site of the VRA that backed up the VM
Computer	VRA that backed up the VM
Job	Name of the backup job
Verification Status	Indicates whether the VM was verified and can be restored from the backup. Possible values include: - Completed — The VM was verified and can be restored from the backup. To view a screenshot of the restored VM's login screen, click View in the Screenshot column. - Unsuccessful - Time Out — The VM backup could not be verified within 10 minutes. This can occur, for example, if the host specified in the VRA backup verification settings does not have sufficient memory or storage, if there is a heavy load on the vault, if the VM takes a long time to start, or if VMware Tools are not installed on the VM. - Unsuccessful (See logs) — The VM backup could not be verified. For more information, see the backup log. <i>Note:</i> Rarely, a <i>Not Verified</i> or <i>Unknown</i> status appears in the report. These statuses also indicate that the VM backup could not be verified. If the VM backup could not be verified, you can run a Rapid VM Restore to determine whether the VM can be restored. See Restore a vSphere VM within minutes using Rapid VM Restore. <i>Note:</i> The report only shows the most recent verification status for each VM in a backup job.
Backup Time	Date and time when the VM's backup was committed to the vault. This is the safeset with the VM that the VRA tried to verify. <i>Note:</i> A vSphere backup job can contain multiple VMs.
Screenshot	For each VM with the <i>Completed</i> verification status, you can view a screenshot of the restored VM's login screen by clicking the View link in this column. If the VM backup could not be verified, the value in the column is "Unavailable".

10.1.2 Daily Status Report

The Daily Status Report includes backup status information for the previous 24 hours, including missed and skipped backups and running jobs for computers where Agent version 8.0 or later is installed. Beginning in Portal 9.0, the Daily Status report also indicates whether potential ransomware threats were detected during backups.

Note: The Daily Status report cannot be viewed in Portal. It can only be scheduled and emailed to users. See [Schedule the Daily Status Report](#).

The following table lists and describes data that is available in the Daily Status Report.

Daily Status Report Data Column	Description
Parent Site	Parent company or service provider that owns or manages the Agent
Site	Child company that owns the Agent (if applicable)
Agent	Computer being protected (or where the backup ran)
Job	Backup job
Event Start	Date and time when the backup started
Event Stop	Date and time when the backup ended (if applicable)
Event Safeset	Numeric value of the safeset which was attempted in the backup. If the backup failed, this safeset was not committed to the vault and the next event will have the same safeset number. The safeset number is not incremented until a backup is committed.
Event Type	Indicates whether the backup was scheduled, run ad-hoc or triggered.

Daily Status Report Data Column	Description
Event Status	Status of the event. Possible values include: • Cancelled — The backup was cancelled by a user before it was completed. • In Progress — The backup was in progress when the report was run. • Completed — The backup started and successfully completed. • Completed with Warnings — The backup started and completed but with some warnings. • Completed with Errors — The backup started and completed but with some errors. • Deferred — The backup successfully committed but some data was deferred. • Failed — The backup started but failed to complete. • Vault license limit reached — The backup failed because there were no available licenses on the vault. • No files were backed up — The backup failed because there were no files available to be backed up. • Schedule Disabled — All schedules for the job have been disabled. • Offline — The job's Agent was offline at the time when the report was run. • Missed — The job was scheduled to run but did not run according to its schedule. This could occur if the Agent system was shut down or backup services on the Agent were stopped. • Skipped — The job was scheduled to run multiple times per day but was skipped. See Skipped backups.

Daily Status Report Data Column	Description
Potential Threat	Indicates whether a potential ransomware threat was detected during the backup. Possible values include: • Undefined — The agent does not support threat detection. "Undefined" is the only possible Potential Threat value for an agent that does not support ransomware threat detection. • Disabled — Ransomware threat detection was not enabled in the backup job. • Not Detected — Ransomware threat detection was enabled in the backup job but a potential ransomware threat was not detected during the backup. • Detected — A potential ransomware threat was detected during the backup. See Manage potential ransomware threats. • Not Run — Ransomware threat detection did not run during the backup, even though it was enabled in the backup job. This can occur, for example, if the operating system is not supported for threat detection. No further action is required. • Error — An error occurred during ransomware threat detection during the backup. Please check the logs for more information and try to resolve the issue. In a vSphere backup, an error can occur if the specified VM credentials are not correct, a VM is offline or unresponsive, or VMware Tools is not installed.
Retention	Name of the retention type used for the backup.
Options	Options used for the backup. This is applicable only to some SQL Server and Exchange backup types. For SQL Server, this will indicate if the backup was a Full, Full with transaction logs, or an Incremental only backup. For Exchange, this will indicate if it was a Full or Incremental backup and if the backup also verified the database.
Original Size	Total amount of source data which was included in the backup
Modified Size (Delta)	Amount of changed data protected in this backup
OTW Data Size (Compressed)	Amount of data sent over the wire
Deferred Data Size	Approximate amount of data which was selected to be backed up but could not be completed within the defined backup window and was therefore deferred to a following backup. This generally only occurs on an initial backup or when re-seeding a job when the amount of data to be protected is more than what can be processed and transmitted within a single backup window.
Most Recent Job Status	Most recent status of the job. For jobs which run multiple times in a day, this is the latest result for the job.
Last Event Date	Date and time of the most recent event for the job

Daily Status Report Data Column	Description
Last Committed Safeset Date	Date and time of the last completed backup which committed a safeset to the vault.
Current Safeset	Numeric value of the last completed backup which committed a safeset to the vault.

10.1.3 Backup Status Report

The Backup Status Report provides information about the last time each backup job ran during a specified time period, and indicates how much data was backed up.

The following table lists and describes data columns that are available in the Backup Status Report, and indicates whether each column is required or optional.

Backup Status Report Data Column	Description	Required	Optional
Vault	Name of the vault where the backup data is saved		✓
Site	Site of the backup		✓
Computer	Computer or environment that was backed up	✓	
Job	Name of the backup job	✓	
OS Version	Operating system of the computer or environment that was backed up		✓
Agent Type	Type of Agent that ran the backup		✓
Backup Date	Date and time when the backup started		✓
Deleted At	If the backup data has been deleted, date and time when the safeset was deleted from the vault		✓
Backup Duration	Length of time that the backup ran		✓
Last Backup Original Size (GB)	Amount of data that was backed up from the computer or environment, in gigabytes		✓

10.1.4 Backup Details Report

The Backup Details Report provides detailed information about each time that a backup job ran during a specified time period, including the amount of data that was backed up, the amount of data that changed since the last backup, and whether the backup was successful or not.

The following table lists and describes data columns that are available in the Backup Details Report, and indicates whether each column is required or optional.

Backup Details Report Data Column	Description	Required	Available
Vault	Name of the vault where the backup data is saved		✓
Site	Site of the backup		✓
Computer	Computer or environment that was backed up	✓	
Job	Name of the backup job	✓	
Backup Date	Date and time when the backup started		✓
Safeset	Number of the safeset that was created by the backup		✓
Backup Duration	Length of time that the backup ran		✓
Original Size (GB)	Amount of data that was backed up from the computer or environment, in gigabytes		✓
Delta Size (MB)	Amount of changed data extracted during the backup, before compression, in megabytes		✓
Result	Outcome of the backup session: Successful or Failed	✓	

10.1.5 Activity Details Report

The Activity Details Report provides information about activities in a site during a specified time period. Activities in the report include successful and failed backups.

The following table lists and describes data columns that are available in the Activity Details Report, and indicates whether each column is required or optional.

Activity Details Report Data Column	Description	Required	Available
Vault	Name of the vault involved in the activity		✓
Site	Site of the activity		✓
Computer	Name of the computer or environment that was involved in the activity	✓	
Job	Name of the backup job involved in the activity	✓	
Date	Date and time when the activity started		✓
Activity Type	Activity that occurred (e.g., Backup, Restore)		✓
Duration	Length of time that the activity ran		✓
Result	Outcome of the activity: Successful or Failed	✓	

10.1.6 Usage Summary Report

The Usage Summary Report table view shows how much data is backed up for each job and computer in a site.

Note: The Usage Summary Report might not appear for every site.

You can also view the amount of data backed up in a site, and activities that affect vault usage, in a configurable chart. See [View the Aggregated Usage Summary Report](#).

The following table lists and describes data columns that are available in the Usage Summary Report table view, and indicates whether each column is required or optional.

Usage Summary Report Data Column	Description	Required	Available
Site	Site of the backup		✓
Vault	Name of the vault where the backup data is saved		✓
Computer	Computer or environment that was backed up	✓	
Job	Name of the backup job	✓	
Date	Date when the backup started		✓
Active Safesets	Number of online safesets created by the backup job		✓
Original Size (GB)	Amount of data that is backed up from the computer or environment in gigabytes. The original size is the size of the data before it was compressed during a backup.		✓
Compressed Size (GB)	Amount of data that is backed up for the computer or environment in gigabytes. The Compressed size is the size of the data after it was compressed during a backup. <i>Note:</i> This value is only available in some Portal instances.		✓
Billing ID	Billing ID associated with the vault usage.		✓

10.1.7 Aggregated Usage Summary Report

The Aggregated Usage Summary Report, available in some Portal instances, shows the total amount of data (original size and compressed size) backed up for a site. You can then navigate in a tree structure to see usage data for each child site, if any, and computer. The report also shows the Billing ID for each computer.

This report is only available in Portal instances that obtain data from billing systems. To associate billing data with a Portal site, the customer short name from the vault must be added for the Portal site.

Aggregated Usage Summary Report Data Column	Description
Name	Name of the site or computer associated with the vault usage. A customer short name from the vault appears in this column if the short name is not associated with a Portal site. An "Unidentified" category appears at the bottom of this column if a computer in the billing data is not associated with a Portal site or customer short name.

Aggregated Usage Summary Report Data Column	Description
Original Size (GB)	The size of the data backed up (in gigabytes) before it was compressed during a backup. At the site level, this value is the total amount of data backed up from all computers in the site. At the computer level, this value is the amount of data that is backed up from the computer.
Compressed Size (GB)	The size of the data backed up (in gigabytes) after it was compressed during a backup. At the site level, this value is the total amount of data backed up from all computers in the site. At the computer level, this value is the amount of data that is backed up from the computer.
Billing ID	Billing ID associated with the vault usage.

10.1.8 Custom Command Report

The Custom Command Report lists all custom commands for a site and indicates when they are scheduled. Custom commands are scripts that are saved on a Windows or Linux computer where an Agent is installed and are scheduled to run through Portal.

The following table lists and describes data columns that are available in the Custom Command Report, and indicates whether each column is required or optional.

Custom Command Report Data Column	Description	Required	Available
Agent	Computer with the custom command script		✓
Company	Site of the computer with the custom command		✓
Script	Custom command script name		✓
Schedule	Time and days when the custom command is scheduled to run, and priority of the schedule		✓
Parent company	Parent site, if any, of the computer with the custom command script		✓

10.2 Schedule the Daily Status Report

The Daily Status report includes backup status information for the previous 24 hours, including missed and skipped backups and running jobs for computers where Agent version 8.0 or later is installed. The Daily Status report also indicates whether potential ransomware threats were detected during backups. See [Daily Status Report](#).

This report can be scheduled and emailed to users, but cannot be viewed in Portal. Each scheduled Daily Status Report is listed in the Daily Status Report section on the Reports page.


To schedule the Daily Status report:

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


The Reports page lists available reports.

If you are signed in as a Support user, and the Support Dashboard appears on the Reports page, select a site.

2. In the Daily Status Report section, click **Add New Report**.
3. In the Email/Schedule dialog box, do the following:
 - In the **To** box, type one or more email addresses that will receive the emailed report. Use commas to separate multiple email addresses.
 - In the **Subject** box, type text for the subject line of the report email.
 - In the **Report Name** box, type a name for the scheduled report. This name will appear in the Daily Status Report section on the Reports page.
 - Using the **At** field, specify the time for running and emailing the report each day. In the **Time Zone** list, click the time zone of the specified time.
 - To exclude completed backups from the report, mark the **Exclude Successful events** check box.

The screenshot shows a dialog box titled "Email/Schedule: Daily Status Report". It contains several input fields: "To" (with the value "test@test.com"), "Subject" (with the value "Daily status"), "Report Name" (with the value "Daily status"), "At" (with the value "1:00 AM" and a clock icon), "Time Zone" (with a dropdown menu showing "(UTC-05:00) Eastern Time (US & Can...)", "Attachment" (with a dropdown menu showing "CSV"), and an "Exclude Successful events" checkbox which is checked.

4. If the Email/Schedule dialog box includes a **Sites to include in report** section, do one of the following:
 - To include computers from child sites in the report, along with computers from the parent site, mark the check box for each child site.
 - To only include computers from the parent site in the report, do not mark any child site check boxes.

The screenshot shows a dialog box titled "Sites to include in report". It contains a "Type to filter" input field and a list of sites with checkboxes: "Child2", "child 1", "1000", and "child3". All checkboxes are unchecked. At the bottom, there are "Okay" and "Cancel" buttons.

5. Click **Okay**.

10.2.1 Delete a scheduled Daily Status Report

Admin users and Support users can delete scheduled Daily Status Reports. These reports appear in the Daily Status Report section on the Reports page.

To delete a scheduled Daily Status Report:

1. When signed in as an Admin user or Support user, click **Reports** on the navigation bar.
The Reports page lists available reports.
If you are signed in as a Support user, and the Support Dashboard appears on the Reports page, select a site. For more information,
2. In the Daily Status Report section, click **[delete]** beside the scheduled report that you want to delete.
3. In the confirmation dialog box, click **Yes**.

10.3 View a report

Admin users and Support users can view reports in Portal that provide information about backups, vault space usage, and activities in a site. For information about available reports, see [Portal report descriptions](#).

Each report in Portal has a default view, which shows data for all protected computers and environments in a site. In most reports, an Admin user or Support user can change the report date range, specify which data columns and records to show, and save the resulting report as a customized view. Admin users and Support users can view default report views and any customized report views that have been saved in their sites.

Note: The following reports do not include options described in this procedure:

- Backup Verification Report. You cannot specify a date range or columns in this report or save a report view. You can only export or email the report in .pdf format. See [View the Backup Verification Report](#).
- Daily Status Report. You cannot view the Daily Status Report in Portal; you can only schedule it to run. See [Schedule the Daily Status Report](#).
- Aggregated Usage Summary Report, which is available in some Portal instances, includes a tree structure for viewing data for child sites and individual computers. See [View the Aggregated Usage Summary Report](#).

Note: Most reports in Portal are available on the Reports page. The Job Monitor Export report, available beginning in Portal 9.20, can be emailed and scheduled from the Monitor page. See [View, export and email backup statuses on the Monitor page](#).

To view a report:

1. When signed in as an Admin user or Support user, click **Reports** on the navigation bar.
The Reports page lists available reports. Default and customized views of each report appear under the report name.

If you are signed in as a Support user, and the Support Dashboard appears on the Reports page, select a site.

Note: If an error message appears on the Reports page, short names might not be entered for the site. Short names are used to extract report data, and must be added by a Super user.

2. To reduce the number of views that are listed, do one of the following:
 - To list default report views only, click **System**.
 - To list customized report views only, click **Custom**.
 - To list default reports and reports with specific characters in their names, enter the characters in the **Find a custom view** box.
3. Click a report view.

Note: You cannot click a report name (which appears in bold text). You can only click a report view.

Report data appears on the Reports page. If a scroll bar appears at the right side of the page, you can scroll down to see more records in the report.

The screenshot shows the Reports page interface. At the top, there's a 'Reports' tab and a 'Backup Status Report' dropdown. Below this is a table with columns: Vault, Site, Computer, Job, OS Version, Agent T, Date range box, Backup Duration..., and Last Backup Ori... The table contains 15 rows of data. A 'Filter row' callout points to the first row. A 'Columns box' callout points to the 'Columns' dropdown. A 'Date range box' callout points to the date range filter. A 'Scroll bar' callout points to the vertical scroll bar on the right side of the table.

	Vault	Site	Computer	Job	OS Version	Agent T	Date range box	Backup Duration...	Last Backup Ori...
1	VAULT	SITE	1AGENT2008	27528	Windows Server 2008	Server Agent	2020/05/04 11:51	00:00:10	0.01
2	VAULT	SITE	1AGENT2008	actmkejob	Windows Server 2008	Server Agent		0	
3	VAULT	SITE	1AGENT2008	itgpf	Windows Server 2008	Server Agent		0	
4	VAULT	SITE	1AGENT2008	ghd	Windows Server 2008	Server Agent			
5	VAULT	SITE	1AGENT2008	jh-Tec	Windows Server 2008	Server Agent			
6	VAULT	SITE	1AGENT2008	Public	Windows Server 2008	Server Agent		0	
7	VAULT	SITE	1AGENT2008	Steven_TEST	Windows Server 2008	Server Agent		0	
8	VAULT	SITE	1AGENT2008	TwoChangeEncryption	Windows Server 2008	Server Agent		0	
9	VAULT	SITE	1AGENT2008	swmb01	Windows Server 2008	Server Agent		0	
10	VAULT	SITE	1AGENT2008	swd0ag	Windows Server 2008	Server Agent		0	
11	VAULT	SITE	1AGENT2008	xctkxctkxctkxctkxctk	Windows Server 2008	Server Agent		0	
12	VAULT	SITE	1AGENT2008	job123	Windows Server 2008	Server Agent	2020/03/08 12:38	00:00:11	0
13	VAULT	SITE	1AGENT2008	TwoExclusions	Windows Server 2008	Server Agent	2020/03/02 12:17	00:00:11	0
14	VAULT	SITE	2008AGENT64	SharePointJob	Windows Server 2008	Server Agent	2020/04/02 10:46	00:00:02	0
15	VAULT	SITE	1AGENT2008	TwoProdDemo	Windows Server 2008	Server Agent	2020/06/07 15:06	00:00:07	0

1164 records shown

4. To change the columns that appear, click the Columns box. In the list that appears, do one or more of the following:
 - To hide a column that currently appears, click **[remove]** beside the column name.
 - To display a column that does not appear, click **[add]** beside the column name.

Columns

Vault	[remove]
Site	[remove]
Computer	
Job	
OS Version	[remove]
Agent Type	[remove]
Backup Date	[remove]
Deleted At	[add]
Backup Duration	[remove]
Last Backup Original Size (GB)	[remove]

- To change the report date range, click the date range box. In the list that appears, do one of the following:

- Click a pre-defined date range: Last 24 Hours, Last 7 Days, Last 30 Days, Last 3 Months, or Year to date.

Apr 15, 2020 - May 14, 2020

Last 24 Hours

Last 7 Days

Last 30 Days

Last 3 Months

Year to date

Date Range

- Click **Date Range**, and specify a custom date range. In the **Start date** calendar, click the first day of records to include in the report. In the **End date** calendar, click the last day of records to include in the report. Click **Done**.

Apr 15, 2020 - May 14, 2020

Last 24 Hours

Last 7 Days

Last 30 Days

Last 3 Months

Year to date

Date Range

Start date

End date

Apr		2020					
Su	Mo	Tu	We	Th	Fr	Sa	
				1	2	3	4
5	6	7	8	9	10	11	
12	13	14	15	16	17	18	
19	20	21	22	23	24	25	
26	27	28	29	30			

May		2020				
Su	Mo	Tu	We	Th	Fr	Sa
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31						

Done

6. If you are viewing the Usage Summary Report and a **Display billing usage only** box appears, the report is only showing billable vault storage usage. To include both billable and non-billable vault storage usage in the report, click the box and then click **Display all usage**.
7. If you are viewing the Usage Summary Report and a **Display all usage** box appears, the report is showing all vault storage usage. To only include billable vault storage usage in the report, click the box and then click **Display billing usage only**.
8. To change which data records appear, enter criteria that records must match. In the filter row under the column headings, in each column where you want to apply a filter, do one of the following:
 - In the empty box, type text that records must match.
 - In the list, click the value that records must match.

Records only appear in the report if they match all specified criteria.

9. When the report includes the data that you want to view, you can do any of the following:
 - Save the customized report view. See [Save a report view](#).
 - Export the report data in comma-separated (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format. See [Export a report](#).
 - Email the report data to one or more recipients. Data can be emailed in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf). See [Email a report](#).
 - Schedule the customized report to be emailed to one or more recipients. Data can be emailed in comma-separated (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf). See [Schedule an emailed report](#).

10.4 View the Backup Verification Report

To determine whether Windows VMs can be restored from vSphere backups, Admin users and Support users can view the Backup Verification Report in Portal version 9.00 or later. You can also view the results in Verification logs in Portal 9.30 or later. See [View a job's process logs and safeset information](#).

The report shows the results of backup verification processes, available with vSphere Recovery Agent (VRA) version 9.00 or later. When backup verification settings are entered for a VRA and backup verification is enabled for a vSphere backup job, the VRA backs up VMs in the job and then checks whether each Windows VM can be restored from the backup. See [Backup verification for vSphere VMs](#) and [vSphere Rapid VM Restore and backup verification requirements](#).

Note: Backup verification does not run for vSphere backups that are started by intra-daily schedules.

The report only shows the most recent verification status for each VM in a backup job. If a VM is included in multiple backup jobs where verification is enabled, the VM can appear multiple times in the report. If two VMs with the same name are backed up, you cannot differentiate between the two VMs in the report.

Note: vSphere allows two or more VMs in a vSphere environment to have the same name if each VM is located in its own folder. If multiple VMs have the same name, you cannot differentiate between the VMs in Portal or in the Backup Verification Report. Consider renaming your VMs in this case.

To view the Backup Verification Report:

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

The **Reports** page lists default and customized report views.

If you are signed in as a Support user and the Support Dashboard appears on the Reports page, you need to select a site.

2. In the Backup Verification Report section, click **Table View**.

The report shows Windows VMs in backup jobs where backup verification was enabled. The site name appears in the Name column. If a site is a parent site, a parent site icon () appears beside the site name.

The Verification Status for each VM indicates whether the VM was verified and can be restored from the backup. Possible values include:

- **Completed** — The VM was verified and can be restored from the backup.

To view a screenshot of the restored VM's login screen, click **View** in the Screenshot column.

- **Unsuccessful - Time Out** — The VM backup could not be verified within 10 minutes. This can occur, for example, if the host specified in the VRA backup verification settings does not have sufficient memory or storage, if there is a heavy load on the vault, if the VM takes a long time to start, or if VMware Tools are not installed on the VM.
- **Unsuccessful (See logs)** — The VM backup could not be verified. For more information, see the backup log.

Note: Rarely, a *Not Verified* or *Unknown* status appears in the report. These statuses also indicate that the VM backup could not be verified.

If the VM backup could not be verified, you can run a Rapid VM Restore to determine whether the VM can be restored. See [Restore a vSphere VM within minutes using Rapid VM Restore](#).

3. For a VM with the *Completed* verification status, to view a screenshot of the restored VM's login screen, click **View** in the Screenshot column.
4. To change which data records appear, enter criteria that records must match. In the filter row under the column headings, in each column where you want to apply a filter, do one of the following:
 - In the empty box, type text that records must match.
 - In the list, click the value that records must match.

Records only appear in the report if they match all specified criteria.

5. When viewing the report, you can do any of the following:
 - Export the report data in Adobe Acrobat (.pdf) format. See [Export a report](#).
 - Email the report data to one or more recipients. Data can be emailed in Adobe Acrobat (.pdf). See [Email a report](#).
 - Schedule the report to be emailed to one or more recipients. Data can be emailed in Adobe Acrobat (.pdf). See [Schedule an emailed report](#).

Note: You cannot export or email the Backup Verification Report in comma-separated values (.csv) or Microsoft Excel (.xls) format.

10.5 View the Aggregated Usage Summary Report

In some Portal instances, Admin users and Support users can view the Aggregated Usage Summary Report. This report shows the total amount of data backed up for a site in a specified time period: both the original size (size of the data backed up before it was compressed during a backup) and the compressed size (size of the data backed up after it was compressed during a backup) in gigabytes. The report includes a tree structure where you can navigate to see usage data for each child site, if any, and each computer. The report also shows the Billing ID for each computer.

This report is only available in Portal instances that obtain data from billing systems. To associate billing data with a Portal site, the customer short name from the vault must be added for the Portal site. If a customer short name is not associated with a Portal site, the short name from the vault appears in the report instead of the Portal site name. If a computer in the billing data is not associated with a Portal site or customer short name from the vault, the computer appears in an "Unidentified" category.

To view the Aggregated Usage Summary Report:

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

The **Reports** page lists default and customized report views.

If you are signed in as a Support user, and the Support Dashboard appears on the Reports page, you need to select a site.


2. In the Usage Summary Report section, click **Aggregated Usage Summary Report**.

Note: If this report name does not appear, the Aggregated Usage Summary Report is not available in your Portal instance.


The report shows the total amount of data backed up from all computers in the site (original and compressed size) in the most recent billing period. The site name appears in the Name column. If a site is a parent site, a parent site icon () appears beside the site name.

3. To change the report date range, click the date range box. In the list that appears, do one of the following:

- Click a pre-defined date range: Last 24 Hours, Last 7 Days, Last 30 Days, Last 3 Months, or Year to date.



- Click **Date Range**, and specify a custom date range. In the **Start date** calendar, click the first day of records to include in the report. In the **End date** calendar, click the last day of records to include in the report. Click **Done**.



- Click the site name in the Name column to expand the tree view.
- If the site has child sites, the report shows each site and the total amount of data backed up from computers in each site.

To view the amount of data backed up from each computer in a site, click the site name in the Name column to expand the tree view.

If an information icon  appears beside a name in the Name column, the name is either:

- A customer short name from the vault. A customer short name appears when the short name is not associated with a Portal site.
- "Unidentified". This category appears when one or more computers in the billing data are not associated with a Portal site or customer short name. For assistance with computers in this category, please contact Support.

For more information, point to the information icon beside the name to view a tooltip.

If the tree view for a site is fully expanded, the report shows the amount of data backed up from each computer in the site, and shows each computer's billing ID.

- When viewing the report, you can do any of the following:
 - Export the report data in comma-separated (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format. See [Export a report](#).


- Email the report data to one or more recipients. Data can be emailed in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf). See [Email a report](#).
- Schedule the customized report to be emailed to one or more recipients. Data can be emailed in comma-separated (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf). See [Schedule an emailed report](#).

10.6 Switch to another report view


When viewing a report in Portal, an Admin user or Support user can quickly switch to another report view.

To switch to another report view:

1. When viewing a report on the Reports page, click the report view list.



Default and customized views of each report are listed under each report name. You cannot click a report name (which appears in bold text). You can only click a report view.



2. In the report list, click the report view that you want to show.

The report view appears on the Reports page.

10.7 Save a report view


After viewing a report in Portal, changing the report date range, and specifying which data columns and records to show, an Admin user or Support user can save the customized report view.

After a report view is saved, the view name appears on the Reports page for all Admin users in the site, and all Support users in the Portal instance.

Note: You cannot save a view for the Backup Verification Report or Aggregated Usage Summary report.


To save a report view:

1. View a report. Specify a report date range, and data columns and records to show in the report. See [View a report](#).



2. Click the **Save View** box.

The box becomes editable.



3. Type a name for the report view, and then press the Enter key.

The customized report view now appears on the Reports page.

10.8 Delete a customized report view

Admin users and Support users can delete customized report views. The default view of a report cannot be deleted.

To delete a customized report view:

1. When signed in as an Admin user or Support user, click **Reports** on the navigation bar.
The Reports page lists default and customized report views.
2. Click **[delete]** beside the customized report view that you want to delete.

10.9 Export a report

Admin users and Support users can export data from a report in comma-separated values (.csv), Microsoft Excel (.xls) or Adobe Acrobat (.pdf) format. The exported report file (named Report.csv, Report.xls or Report.pdf) is downloaded to the user's computer.

Note: The Backup Verification Report can only be exported in Adobe Acrobat (.pdf) format.

We recommend turning off macros in Microsoft Excel when using Portal, particularly if you export reports in XLS or CSV format and open these reports in Excel.

Reports in Microsoft Excel (.xls) and Adobe Acrobat (.pdf) format are formatted using the site's logo, color, and custom text.

Admin users and Support users can also email report data in these formats to one or more recipients. See [Email a report](#) and [Schedule an emailed report](#).

To export a report:

1. View a report.
2. Specify a date range, and data columns and records to show in the report. See [View a report](#).
Note: You cannot specify a date range or columns to show in the Backup Verification Report.
3. Click the **Export** box. In the list that appears, click one of the following formats for the exported report data file:
 - CSV (comma-separated values)
 - XLS (Microsoft Excel)
 - PDF (Adobe Acrobat)

Note: The Backup Verification Report can only be exported in Adobe Acrobat (.pdf) format.

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

10.10 Email a report

Admin users and Support users can email a report in comma-separated values (.csv), Microsoft Excel (.xls) or Adobe Acrobat (.pdf) format. The report file (named Report.csv, Report.xls or Report.pdf) is sent as an email attachment to one or more specified recipients.

Note: The Backup Verification Report can only be emailed in Adobe Acrobat (.pdf) format.

Note: Most Portal reports can be emailed from the Reports page. The Job Monitor Export report, available beginning in Portal 9.20, can be emailed from the Monitor page. See [View, export and email backup statuses on the Monitor page](#).

Reports in Microsoft Excel (.xls) and Adobe Acrobat (.pdf) format are formatted using the site's logo, color, and custom text.


Admin users and Support users can also schedule reports to be emailed to one or more recipients. See [Schedule an emailed report](#).

To email a report:

1. View a report.
2. Specify a date range, and data columns and records to show in the report. See [View a report](#).

Note: You cannot specify a date range or columns to show in the Backup Verification Report.

3. Click the **Email/Schedule** box. In the **Send Report** list that appears, click **Email Once**.



4. In the Email Once dialog box, do the following:
 - In the **To** box, type one or more email addresses for sending the report. Use commas to separate multiple email addresses.
 - In the **Subject** box, type a subject for the report email.
 - In the **Attachment** list, click one of the following formats for the emailed report:
 - CSV (comma-separated values)
 - Excel (Microsoft Excel)
 - PDF (Adobe Acrobat)

Note: The Backup Verification Report can only be emailed in Adobe Acrobat (.pdf) format.

 - In the **Date Range** list, click the date range for the emailed report:
 - Selected Dates. If this option is selected, the emailed report uses the date range currently specified for the report in Portal.
 - Yesterday
 - Last 24 Hours
 - Last 7 Days
 - Last 30 Days
5. Click **Okay**.

10.11 Schedule an emailed report

Admin users and Support users can schedule reports to be emailed in comma-separated values (.csv), Microsoft Excel (.xls) or Adobe Acrobat (.pdf) format. A report file (named Report.csv, Report.xls or Report.pdf) is sent as an email attachment to one or more specified recipients on the specified days at the specified time.

Note: The Daily Status Report can only be emailed in .csv format. See [Schedule the Daily Status Report](#).

Note: The Backup Verification Report can only be emailed in Adobe Acrobat (.pdf) format.

Note: Most Portal reports can be scheduled from the Reports page. The Job Monitor Export report, available beginning in Portal 9.20, can be scheduled from the Monitor page. See [View, export and email backup statuses on the Monitor page](#).

Reports in Microsoft Excel (.xls) and Adobe Acrobat (.pdf) format are formatted using the site's logo, color, and custom text.


We recommend turning off macros in Microsoft Excel when using Portal, particularly if you email reports in XLS or CSV format and open these reports in Excel.

Scheduled reports appear in the Email/Schedule list on the Reports page.

To schedule an emailed report:

1. View a report.
2. Specify a date range, and data columns and records to show. See [View a report](#).
3. Click the **Email/Schedule** box. In the **Send Report** list that appears, click **Schedule New Report**.

Note: You cannot specify a date range or columns to show in the Backup Verification Report.



4. In the Email/Schedule dialog box, do the following:
 - In the **To** box, type one or more email addresses for sending the report. Use commas to separate multiple email addresses.
 - In the **Report Name** box, type a name for the scheduled report. This name appears in the **Email/Schedule** list.
 - In the **Subject** box, type a subject for the email.
 - In the **Attachment** list, click one of the following formats for the emailed report data file:
 - CSV (comma-separated values)
 - Excel (Microsoft Excel)
 - PDF (Adobe Acrobat)

Note: The Backup Verification Report can only be emailed in Adobe Acrobat (.pdf) format.

- In the **Date Range** list, click the date range for the emailed report:

- Yesterday
- Last 24 Hours
- Last 7 Days
- Last 30 Days

5. Do one of the following:

- To email the report on specific days each week, in the **Frequency** list, click **Daily**. In the day row, select the days when you want to email the report each week.

Frequency

Daily 

S	M 	T	W 	T	F 	S



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

Frequency

Weekly 

S 	M	T	W	T	F	S



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

Frequency

Monthly 

1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
Last Day 						



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

10.11.1 Delete a backup status report schedule

Admin users and Support users can delete backup status report schedules from the Monitor page.

To delete a Monitor page report schedule:

1. On the navigation bar, click **Monitor**.
The Monitor page shows recent backup statuses for jobs in your site.
2. Click the **Email/Schedule** box.
Scheduled backup status reports appear in the **Currently Scheduled** list.
3. Click the Delete button beside the report schedule that you want to delete.

10.12 View a Usage Summary Report chart

To see how much data is backed up for computers in a site, Admin users and Support users can view and configure a Usage Summary Report chart. The chart also shows activities that affect vault space usage (e.g., backup job changes).

The Original Size chart view of the Usage Summary Report shows the amount of data backed up for computers in a site before the data is compressed.

In some Portal instances, a Compressed Size chart view also appears. This chart shows the amount of data backed up for computers in a site after the data was compressed.

The amount of data backed up in a site can also be viewed in the Usage Summary Report table view. See [Usage Summary Report](#).

Note: You cannot export or email data from a Usage Summary Report chart. You can only export or email data from the Usage Summary table view.

To view a Usage Summary Report chart:

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

The **Reports** page lists default and customized report views.


If you are signed in as a Support user, and the Support Dashboard appears on the Reports page, you need to select a site.

2. In the Usage Summary Report section, click a chart view.

Chart View (Original Size) shows the amount of data backed up for computers and environments in a site before the data was compressed.

In some Portal instances, **Chart View (Compressed Size)** is also available. This chart view shows the amount of data backed up for computers and environments in a site after the data was compressed.

The vertical axis scale and units change according to the amount of data. The vertical axis does not start at zero (0).



In the Select box, a check mark appears beside each item that is included in the chart. In some Portal instances, you can select vaults, sites and computers in the Select box. In other Portal instances, you can only select sites and computers, and data in the chart is combined across vaults for each site and computer.

3. To change which usage data is shown in the chart, do one or more of the following in the **Select** box, and then click **Apply Filter**:
 - To exclude an item from the chart, click the item name to clear the check mark.
 - To remove all items from the chart, click **None**. The check mark is cleared from every item name.
 - To add an item to the chart, click the item name so that a check mark appears.
 - To add all items to the chart, click **All**. A check mark appears beside every item name.
4. To change the chart date range, click the date range box. In the list that appears, do one of the following:
 - Click a predefined date range:
 - Last 24 Hours
 - Last 7 Days
 - Last 30 Days
 - Last 3 Months
 - Year to date

Jul 15, 2018 - Aug 14, 2018 

Last 24 Hours

Last 7 Days

Last 30 Days

Last 3 Months

Year to date

Date Range 

- In the date slider below the chart, drag the date range start and end markers to change the selected date range.

AugustSeptember 20October 2017November 2017

- Click **Date Range**, and specify a custom date range. In the **Start date** calendar, click the first day of records to include in the report. In the **End date** calendar, click the last day of records to include in the report. Click **Done**.

Sep 15, 2017 - Oct 14, 2017 

Last 24 Hours

Last 7 Days

Last 30 Days

Last 3 Months

Year to date

Date Range 

Start date

End date

Su	Mo	Tu	We	Th	Fr	Sa
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30

Su	Mo	Tu	We	Th	Fr	Sa
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31				

Done

- If a Display usage box appears beside the date range box above the chart, do one of the following:
 - To only show billable storage usage in the chart, click **Display billing usage only**.
 - To include both billable and non-billable storage usage in the chart, click **Display all usage**.

Display billing usage only 


Display billing usage only

Display all usage

Note: The Display usage box only appears in Portal instances where the Compressed Usage chart is available.


6. To view additional information on the chart, click the **Display event types** box. In the list, do one or more of the following:
 - To view job changes, mark the **Job Changes** check box.
 - To view computer changes, mark the **Computer Changes** check box.
 - To view other configuration changes, mark the **Other Config Changes** check box.
 - To only view vault space usage on the chart, mark the **Show none** check box.
 - To view job changes, computer changes, and other configuration changes, mark the **Show all** check box.

A circle appears on the chart for each event type that is marked in the list.



7. To view precise information about vault disk space usage and any displayed event types on a particular date, click the dot for the date on the chart.

A box shows the amount of disk space used on the vault, and detailed information about any displayed event types.



11 Server Backup Support

If you have a question about 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 Server Backup solutions you may need. We highly recommend searching here first for the quickest answers to your questions.

What can we help you with?

Search

Popular Searches

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

11.1 Contacting Support

If you need live assistance from a qualified support agent, 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 Support in the [Knowledge Base](#).



CREATE A
SUPPORT CASE



CHAT WITH A
REPRESENTATIVE



CALL A SUPPORT
REPRESENTATIVE