

CARBONITE[®]

an **opentext**[™] company

Carbonite Server Backup

vSphere Recovery Agent 9.1

User Guide



© 2022 Carbonite, Inc. All rights reserved.

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

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

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

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

Carbonite, Inc.
Two Avenue de Lafayette
Boston, MA 02111
www.carbonite.com

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

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

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

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

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

Version History

Version	Date	Description
1	August 2022	Initial guide provided for vSphere Recovery Agent 9.1.

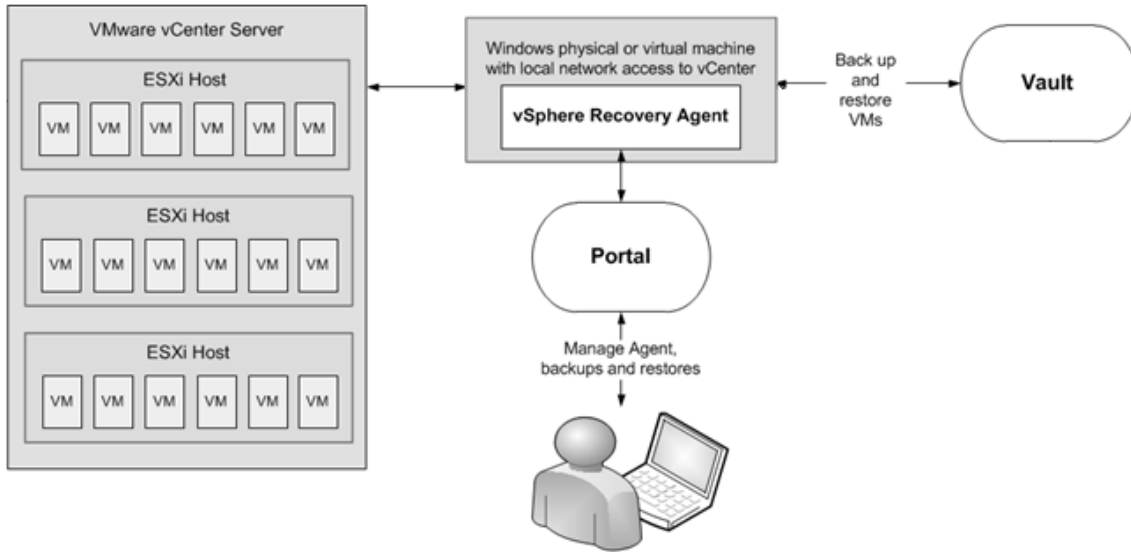
Contents

1 Introduction to the vSphere Recovery Agent	5
2 Prepare for a vSphere Recovery Agent installation	7
2.1 Portal for managing a vSphere Recovery Agent	7
2.2 Vaults for vSphere Recovery Agent backups	7
2.3 Recommended vSphere Recovery Agent deployment	7
2.4 Requirements for specific vSphere Recovery Agent features	8
2.5 vSphere Recovery Agent ports	10
2.6 vSphere Recovery Agent limitations and best practices	11
3 Install, upgrade or uninstall the vSphere Recovery Agent	12
3.1 Install the vSphere Recovery Agent	12
3.2 Install the vSphere Recovery Agent in silent mode	13
3.3 Upgrade the vSphere Recovery Agent	14
3.4 Upgrade the vSphere Recovery Agent in silent mode	14
3.5 Uninstall the vSphere Recovery Agent	14
3.6 Uninstall the vSphere Recovery Agent in silent mode	15
4 Configure a vSphere Recovery Agent	16
4.1 Change vCenter or ESXi host information for a vSphere Recovery Agent	18
4.2 Change the CBT Setting for a vSphere Recovery Agent	19
4.3 Enter backup verification settings for a vSphere Recovery Agent	19
4.4 Change the Portal registration for a vSphere Recovery Agent	20
4.5 Add vault settings	21
4.6 Add a description	23
4.7 Add retention types	23
4.8 Configure bandwidth throttling	25
5 Add a vSphere backup job	27
5.1 Application-consistent backups on vSphere VMs	31
5.2 Backup verification for vSphere VMs	32
5.3 Log file options	33
5.4 Encryption settings	33
6 Run and schedule backups and synchronizations	35
6.1 Schedule a backup	35
6.2 Maximum number of restore points for a job	39
6.3 Specify whether scheduled backups retry after a failure	40
6.4 Run an ad-hoc backup	41
6.5 Synchronize a job	43

7	Resolve certificate failures and potential threats	44
7.1	Resolve certificate failures	44
7.2	Manage potential ransomware threats	45
8	Restore vSphere data	48
8.1	Restore vSphere VMs	48
8.2	Restore a vSphere VM within minutes using Rapid VM Restore	51
8.3	Restore files, folders and database items using a vSphere Recovery Agent	58
8.4	Restore data to a replacement computer	60
8.5	Restore data from another computer	62
8.6	Advanced restore options	63
9	Delete jobs and computers, and delete data from vaults	64
9.1	Delete a backup job without deleting data from vaults	64
9.2	Delete a backup job and delete job data from vaults	65
9.3	Cancel a scheduled job data deletion	67
9.4	Delete a computer without deleting data from vaults	68
9.5	Delete a computer and delete computer data from vaults	69
9.6	Cancel a scheduled computer data deletion	71
9.7	Delete specific backups from vaults	71
10	Monitor computers, jobs and processes	74
10.1	Monitor backups and computers using the Current Snapshot	74
10.2	View computer and job status information	75
10.3	View an unconfigured computer's logs	77
10.4	View current process information for a job	78
10.5	Monitor backups using email notifications	80
10.6	View the Backup Verification Report	84
10.7	Schedule the Daily Status Report	86
10.8	View a job's process logs and safeset information	89
10.9	View and export recent backup statuses	92
11	Carbonite Server Backup Support	93
11.1	Contacting Carbonite	93


1 Introduction to the vSphere Recovery Agent

The vSphere Recovery Agent (VRA) provides data protection for VMware vSphere environments. 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:

- Application-consistent backups. Beginning in version 8.82, the VRA 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 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 backup verification status of Windows VMs in the Backup Verification report in Portal 9.00 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 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](#).

2.1 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.

2.2 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.

2.3 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.

2.4 Requirements for specific vSphere Recovery Agent features

To use specific VRA features, check the following requirements:

- 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](#).

2.4.1 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 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.

2.4.2 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 .	
vSphere environment		
ESXi hosts	<p>Each ESXi host must have the Software iSCSI Adapter installed and bound to a network port group that the VRA can reach.</p> <p>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.</p> <p><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.</p>	<p>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.</p> <p>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.</p> <p><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.</p>
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	
	<p>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.</p> <p>When you migrate a VM to permanent storage, the destination datastore can be local, iSCSI, vSAN or NFS storage.</p>	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.

Component	Rapid VM Restore requirement	Backup verification requirement
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.

2.4.2.1 Enable the Rapid VM Restore feature on a vault

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

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

To enable the Rapid VM Restore feature on a vault:

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

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

2.4.3 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 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

2.6 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.

3 Install, upgrade or uninstall 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. See [Configure a vSphere Recovery Agent](#).

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

3.1 Install the vSphere Recovery Agent

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**.
 - 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**.

3.2 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 (C:\Program Files\Carbonite 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

3.3 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**.

3.4 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
```

3.5 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.

3.6 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
```

4 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](#).

4.1 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**.

4.2 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**.

4.3 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](#).

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**.

4.4 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.

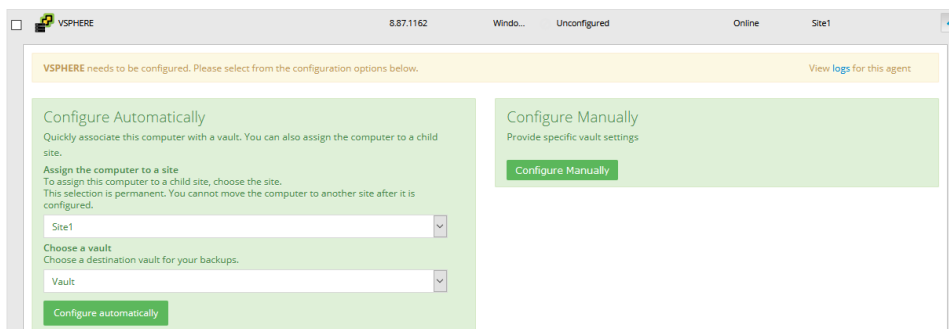
To change the Portal registration for a vSphere Recovery Agent:

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. See [Uninstall the vSphere Recovery Agent](#).
3. Reinstall the VRA. When prompted to register the agent with Portal, enter the new Portal registration information. See [Install, upgrade or uninstall 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**.

4.5 Add vault settings

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

When adding vault settings for a VRA, 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 VRA, Admin users can select any vault profile from the policy. Regular users can only select policy vault profiles that are also assigned to them.

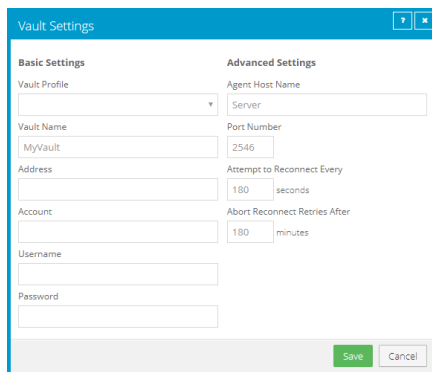
If a policy is not assigned to a VRA, Admin users can select any vault profile in the site. Regular users can only select vault profiles that are assigned to them.

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

To add vault settings:

1. On the navigation bar in Portal, click **Computers**.
2. Find the VRA for which you want to add vault settings, and click the VRA 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 VRA 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**.

4.6 Add a description

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

To add a description:


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

The Computers page shows registered computers.

2. Find the VRA 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 VRA.



5. Click **Save**.

4.7 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 VRA where a policy is not assigned.

If a policy is assigned to a VRA, 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 VRA 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 VRA, 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.

Retention Type

Retention Details

Name:
Daily

Backup Retention
30 days

Number of Backup Copies to Keep
30

Create archived copies

Keep Archives For
365 days

Save Cancel

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**.

4.8 Configure bandwidth throttling

Bandwidth settings include:

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

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

If you edit a VRA'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 VRA, 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:

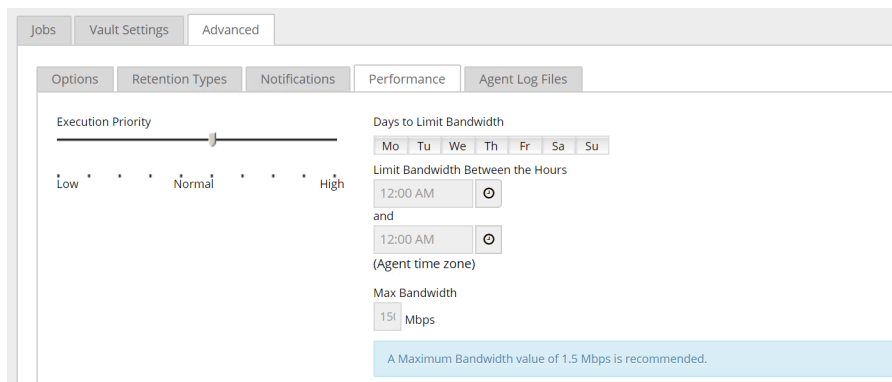
1. On the navigation bar, click **Computers**.
2. Find the VRA for which you want to configure bandwidth throttling, and click the row to expand its view.

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 VRA, you cannot add or change values on the Performance tab. Instead, bandwidth settings can only be modified in the policy.

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



4. Click **Save**.

5 Add a vSphere backup job

You must add vault settings and vSphere environment information before you can add a backup job. See [Configure a vSphere Recovery Agent](#).

You can enable the following options in a vSphere backup job:

- **Application-consistent backups.** Beginning in version 8.82, the VRA 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 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](#).

Note: A backup snapshot only contains VM data that is written to disk. If file writes are still in memory when a backup runs, some data might not be captured, even if the backup is application-consistent. In an application consistent backup, pending transactions are written to disk for some applications but not the entire file system. To force the operating system to flush data in memory to disk, you can run a third-party utility before a backup.

- **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](#).

To back up the data, you can run the backup job manually or schedule the backup job to run. See [Run and schedule backups and synchronizations](#).

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 New 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.

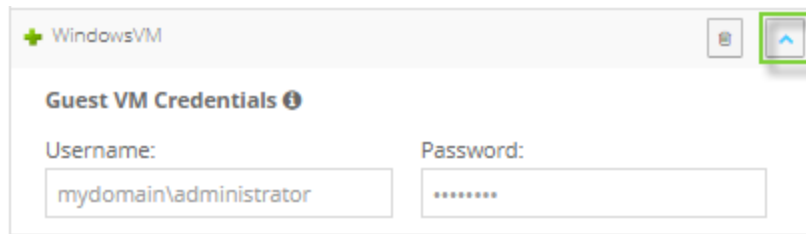
5. In the Create New Job dialog box, specify the following information:
 - In the **Name** box, type a name for the backup job.
 - In the **Description** box, optionally type a description for the backup job.
 - In the **Destination** list, select the vault where you want to save the backup data.
A vault 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 perform application-consistent backups by doing one of the following:
- To perform crash-consistent backups of VMs in the backup job, clear the **Enable Application Consistent Backups** check box.
 - To perform application-consistent backups of SQL Server, Exchange, SharePoint, and Active Directory installed on Windows VMs in the backup job, do the following:
 - a. Select the **Enable Application Consistent Backups** check box.
 - 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 Agent 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.

8. 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.

9. 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](#).

10. 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 and synchronizations](#). If you do not want to create a schedule at this time, click **Cancel**.

5.1 Application-consistent backups on vSphere VMs

Beginning in version 8.82, the vSphere Recovery Agent can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows virtual machines (VMs) in vSphere environments.

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 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.

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 vSphere Recovery Agent 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.

5.2 Backup verification for vSphere VMs

Beginning in version 9.00, the vSphere Recovery Agent 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 version 9.00 or later. See [View the Backup Verification Report](#).

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.

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](#).

5.3 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.

5.4 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.

6 Run and schedule backups and synchronizations

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](#).

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 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.

For 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](#).

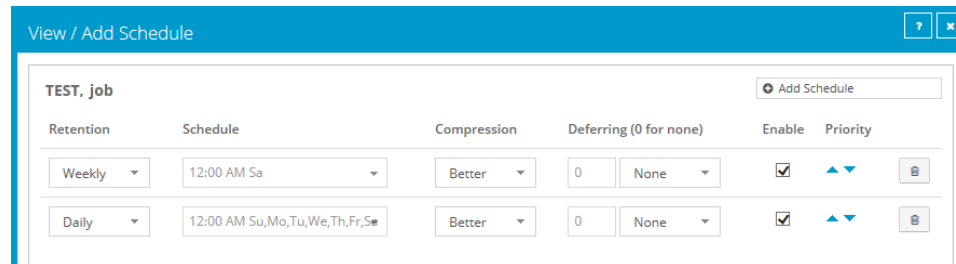
6.1 Schedule a backup

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.

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.



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 backup job to run daily or monthly:

- Do one of the following:
 - On the navigation bar, click **Computers**. Find the VRA 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.

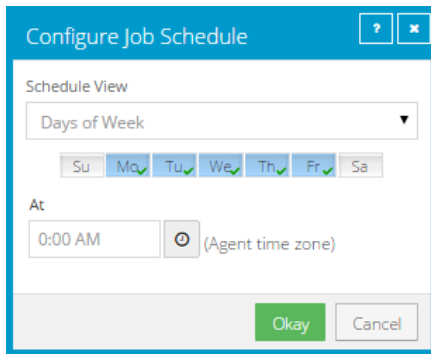
The retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline. See [Add retention types](#).

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

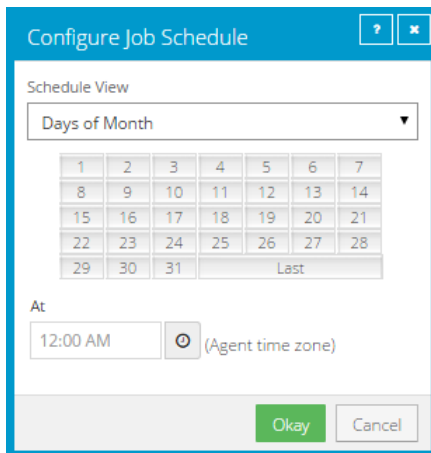
The Configure Job Schedule dialog box opens.

7. 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.




- 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.



- 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.



8. Click **Okay**.

The new schedule appears in the Schedule box.

9. In the **Compression** list, click a compression level for the backup data. Compression levels optimize Compression levels optimize the amount of data stored vs. the backup speed.

10. 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.


11. To run the job on the specified schedule, select the **Enable** check box near the end of the row.

12. 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.

13. 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](#).



14. If an Automatic Retry for Scheduled Backups section appears at the bottom of 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](#).
15. Click **Save**.

6.2 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.

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.

6.3 Specify whether scheduled backups retry after a failure

You can specify whether scheduled backups 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.

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 VRA 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:
 - 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: [Up] [Down]

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

Save Cancel

3. Click **Save**.

6.4 Run an ad-hoc backup

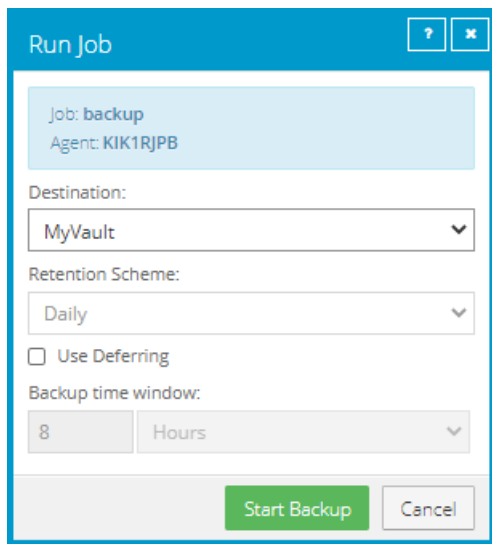
After a backup job is created, you can run the backup at any time, even if the job is scheduled to run at specific times.

To run an ad-hoc backup:

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

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

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



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

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. Do one of the following:

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

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

8. Click **Start Backup**.

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

9. If you want to stop the backup, click **Stop**.
10. To close the Process Details dialog box, click **Close**.

6.5 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.
- 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 VRA 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**.

7 Resolve certificate failures and potential threats

Beginning in version 8.87, the VRA can check vault TLS certificates when they try to connect to vaults. If a VRA reports a certificate failure, you must investigate and resolve the certificate failure before backups and restores can continue. See [Resolve certificate failures](#).

Beginning in version 9.10, the VRA can check for potential ransomware threats on VMs when running backup jobs. If a VRA detects a potential threat, you must investigate and resolve the potential threat. See [Manage potential ransomware threats](#).

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:

- A VRA tries to connect to a version 8.60 vault where certificate pinning is enabled. Beginning with vSphere Recovery Agent 8.87, when a VRA 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 VRA previously connected to the vault. If the public key of the vault certificate is different, the VRA reports a certificate failure and will not connect to the vault.
- A 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

Beginning with 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 VRA 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 and export recent backup statuses](#).
- 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 vSphere backup job. See [Restore vSphere data](#) and [Delete specific backups from vaults](#).

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 VM, you can sign in to the VM in your environment and investigate whether it is infected with ransomware. An Admin user in Portal can then manage the threat:

- If the 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 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, 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.

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**. 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](#).
 - If you investigated or addressed the potential threat and are sure that the 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 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](#)
- [Restore files, folders and database items using a vSphere Recovery Agent](#)

8.1 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.

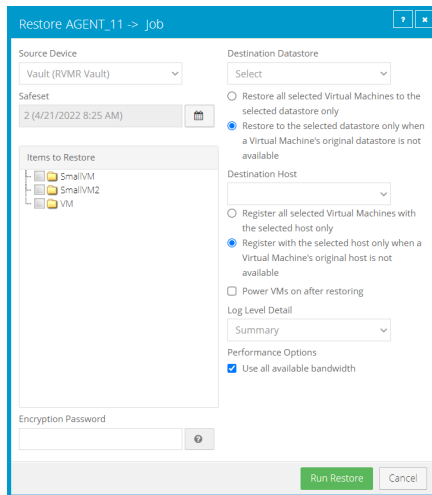
To restore vSphere VMs:

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**.



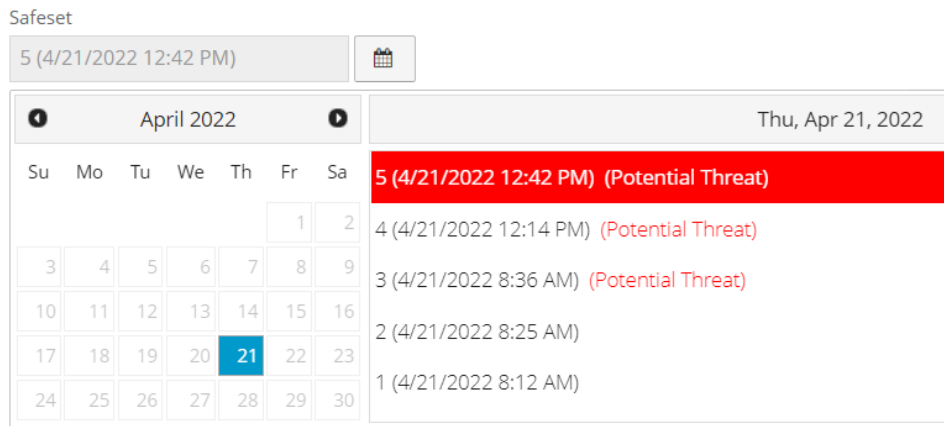
6. Click **Continue**.

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.

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**.

8.2 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](#).

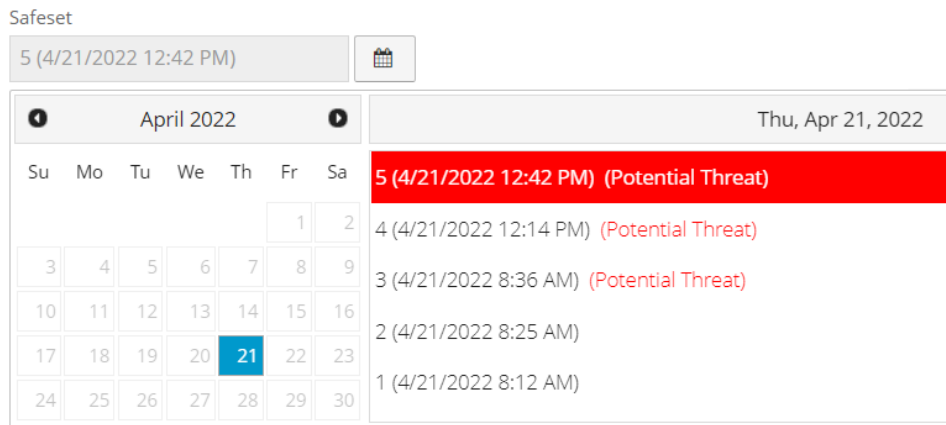



6. Click **Continue**.

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**.

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.

8.2.1 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

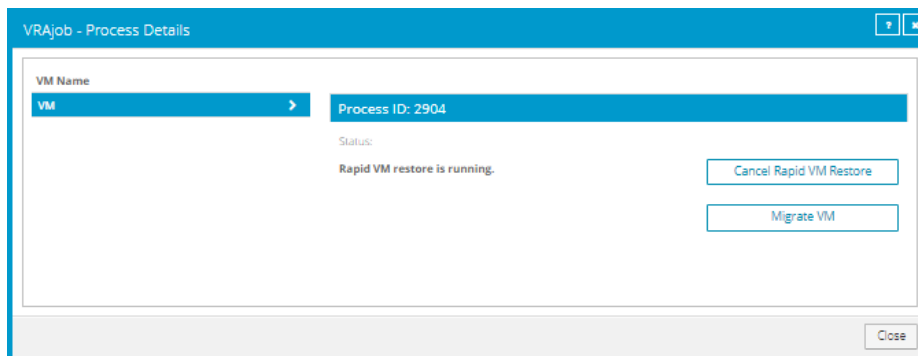
migrate the VM using Portal. The VM might be deleted when you cancel the Rapid VM Restore process.

- 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.

8.2.2 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.

8.3 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.

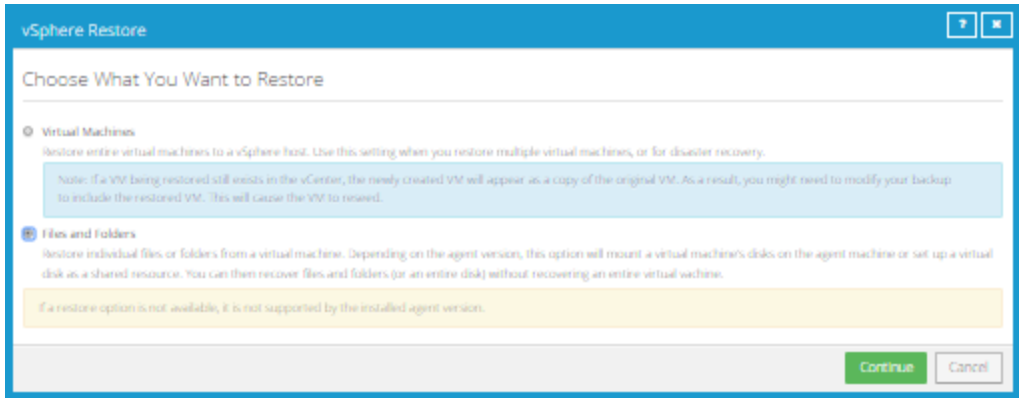
Note: Files and folders on the mounted drives will be accessible to anyone on the VRA system, including non-Admin users. If you are concerned about security, secure the Agent machine and prevent users from logging in to the machine locally.

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**.

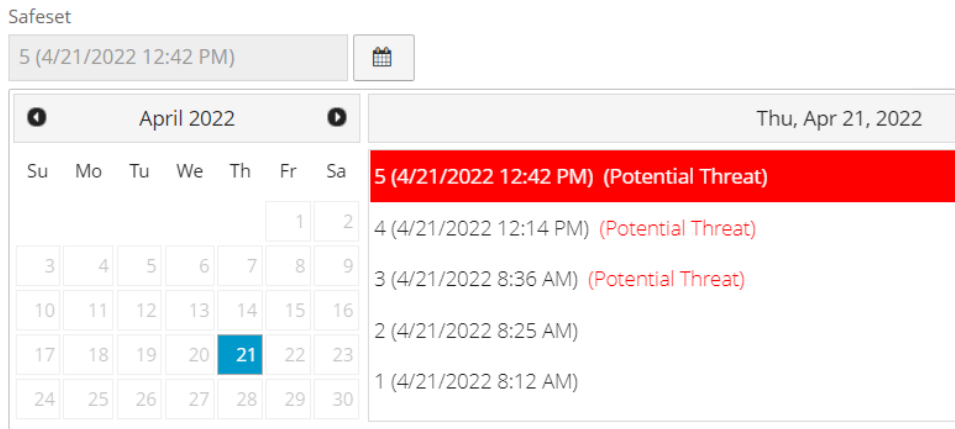



6. Click **Continue**.

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**.

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*.

8.4 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.
- 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](#).

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 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](#).

8.5 Restore data from another computer

You can restore some or all of a computer's backed up data to another (similar) computer.

To restore data from another computer, you can redirect data from a backup job on the vault to a different computer.

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 are often compatible. Operating systems that share similar origins 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 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.

8.6 Advanced restore options

When restoring vSphere VMs, you can specify the following options:

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.

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 Delete jobs and computers, and delete data from vaults

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

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

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

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

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

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

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

- 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.

9.1 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.

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](#).

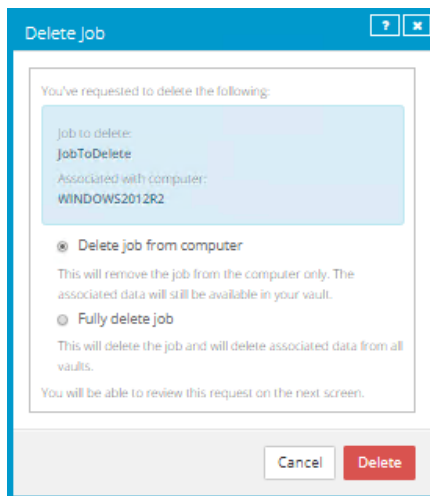
To delete a backup job without deleting data from vaults:

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 **Delete job from computer** 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**.

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

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.

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



9.3 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.

- Find the computer with the scheduled job data deletion that you want to cancel, and expand its view by clicking its row.
- Click the **Jobs** tab.
- 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.

- 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.



9.4 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.

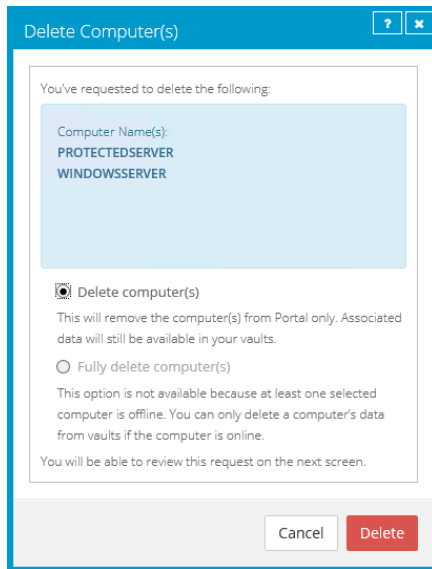
If a computer is deleted from Portal in this way, the data can still be restored using the *Restore from Another Computer* procedure.

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:

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

9.5 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. 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*.

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.

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

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

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

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

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

To delete a computer and delete computer data from vaults:

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

The Computers page shows registered computers.

2. Select the check box for each computer that you want to delete.

3. In the **Actions** list, click **Delete Selected Computer(s)**.

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

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

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

IMPORTANT: To permanently delete unnecessary data from vaults and reduce billing, you must select **Fully delete computer(s)**. If you select **Delete computer(s)**, 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.

9.6 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.

9.7 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. Backups that cannot be deleted (e.g., because a deletion request is scheduled for the job or computer) cannot be selected.

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**.

10 Monitor computers, jobs and processes

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

- **Current Snapshot.** The Current Snapshot provides total numbers of backups and computers in various categories in your site, and allows you to navigate to more detailed information. See [Monitor backups and computers using the Current Snapshot](#).
- **Computers page.** The Computers page shows status information for computers and their jobs. See [View computer and job status information](#). You can also access logs for unconfigured computers from this page. See [View an unconfigured computer's logs](#).
- **Process Details dialog box.** This dialog box shows information about all running, queued and recently-completed processes for a job. See [View current process information for a job](#).
- **Email notifications.** To make it easier to monitor backups, users can receive emails when backups finish or fail. See [Monitor backups using email notifications](#).
- **Backup Verification Report.** Beginning with the vSphere Recovery Agent (VRA) 9.00 and Portal 9.00, the Backup Verification report indicates whether Windows VMs can be restored from vSphere backups. See [View the Backup Verification Report](#).
- **Daily Status Report.** Beginning with VRA 9.10 and Portal 9.10, the Daily Status report indicates whether a potential ransomware threat was detected during a vSphere backup. See [Schedule the Daily Status Report](#).
- **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 and export recent backup statuses](#).

10.1 Monitor backups and computers using the Current Snapshot

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

To monitor backups and computers using the Current Snapshot:

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

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

- **Backups Requiring Attention** — Number of backup jobs where the last backup attempt failed, completed with errors, did not back up any files, reached a license limit, was cancelled or had a

potential ransomware threat.

- **Missed Backups** — Number of backup jobs that have not run for seven days.
 - **Backups With Warnings** — Number of backup jobs where the last backup attempt completed with warnings, was deferred, was deferred with warnings or was skipped. This category also includes backup jobs that have never run.
 - **Computers Requiring Reboot** — Number of computers with a pending reboot.
 - **Offline Computers** — Number of computers that are not currently in contact with Portal. A computer can be offline if it is turned off, if the Agent has been uninstalled from the system, or if the system no longer exists.
 - **Computers Scheduled for Deletion** — Number of computers that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.
 - **Computers With Certificate Failures** — Number of computers reporting a 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.

10.2 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.

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](#).

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.

Beginning in vSphere Recovery Agent 8.87, to reduce schedule overloads, the VRA skips 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.

-  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
-  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](#).

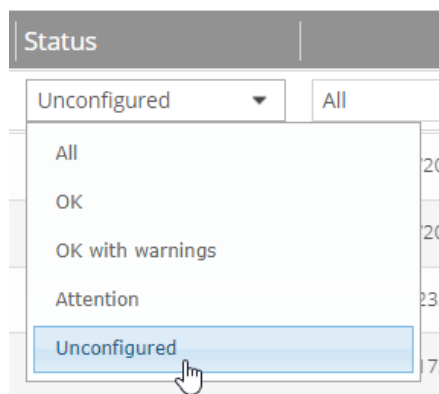
10.3 View an unconfigured computer's logs

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:

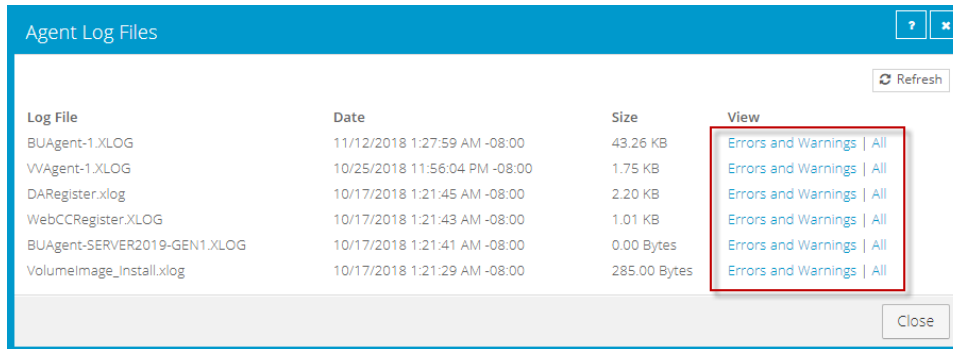
1. On the navigation bar, click **Computers**.

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.

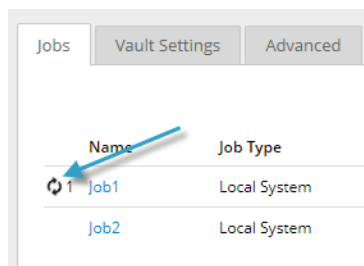
10.4 View current process information for a job



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:

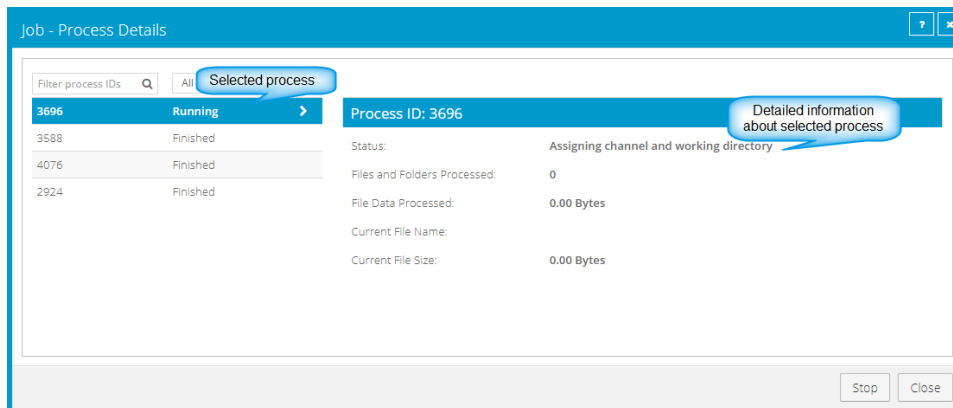
1. While a backup, restore, Rapid VM Restore, or synchronization is running, do one of the following:
 - On the Computers page, on the Jobs tab, click the Process Details symbol  or Rapid VM Restore symbol  beside the job name.



- On the Monitor page, click the Process Details symbol  or Rapid VM Restore symbol  beside the job name.



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**.



10.5 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 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. See [Set up email notifications for encryption password changes](#).

10.5.1 Set up email notifications for backups on a computer

To set up email notifications for a computer:

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

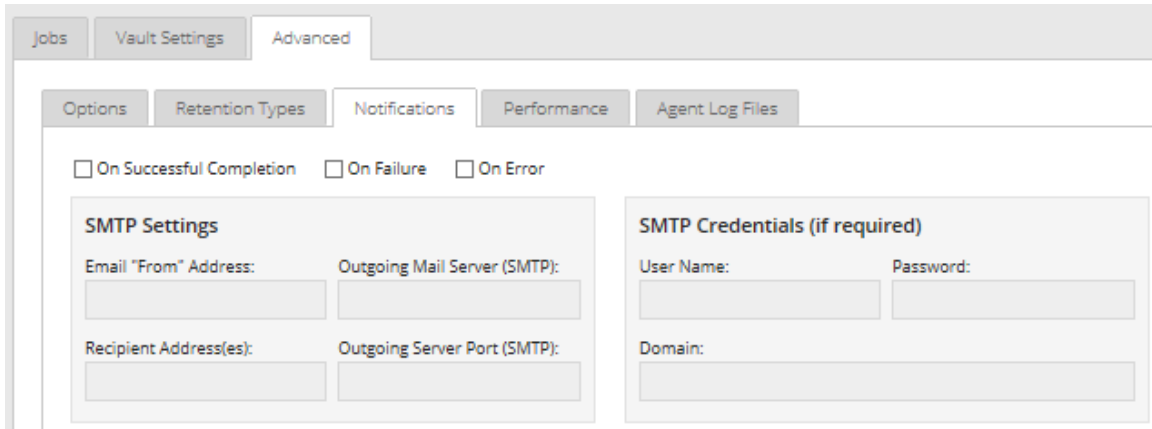
The Computers page shows registered computers.

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.



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**.

10.5.2 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.

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.

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.

4. Click **Update notifications**.

10.5.3 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**.

10.5.4 Set up email notifications for 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**.

10.6 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.

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 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.
 - Email the report data to one or more recipients. Data can be emailed in Adobe Acrobat (.pdf).

- Schedule the report to be emailed to one or more recipients. Data can be emailed in Adobe Acrobat (.pdf).

Note: You cannot export or email the Backup Verification Report in comma-separated values (.csv) or Microsoft Excel (.xls) format.

10.7 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 has a close button in the top right corner. The dialog is divided into two columns of input fields. The left column contains: "To" (test@test.com), "At" (1:00 AM), and "Time Zone" (UTC-05:00 Eastern Time (US & Can)). The right column contains: "Subject" (Daily status), "Report Name" (Daily status), "Attachments" (CSV), and a checked checkbox for "Exclude Successful events".

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.



5. Click **Okay**.

10.7.1 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 or run ad-hoc

Daily Status Report Data Column	Description
Event Status	<p>Status of the event. Possible values include:</p> <ul style="list-style-type: none"> • 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.
Potential Threat	<p>Indicates whether a potential ransomware threat was detected during the backup. Possible values include:</p> <ul style="list-style-type: none"> • 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.

Daily Status Report Data Column	Description
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
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.

10.8 View a job's process logs and safeset information

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

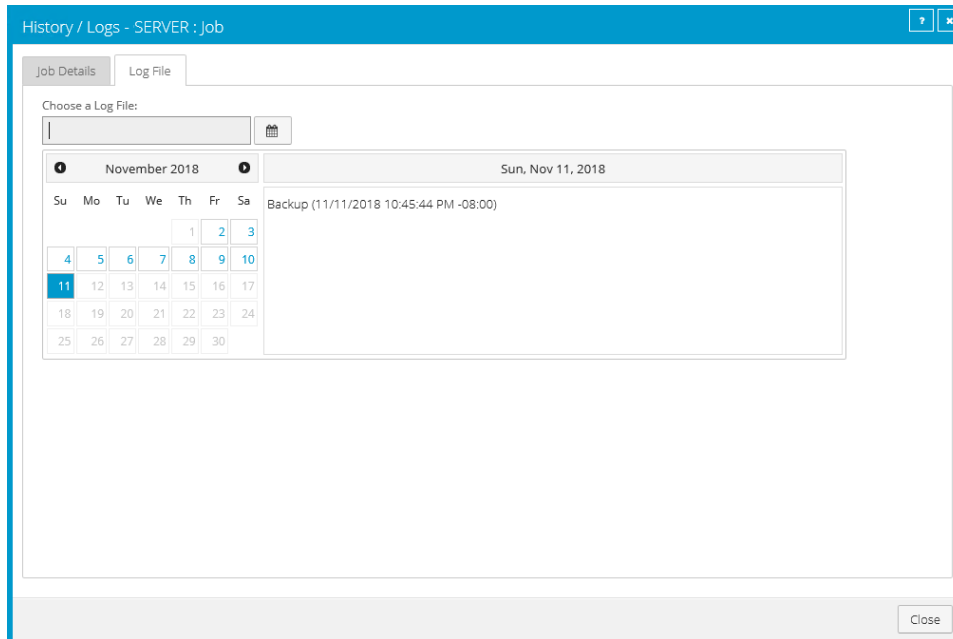
You can also view information about safesets created for the job. A safeset is an instance of backup data on the vault.

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.

3. To view log files for a job, do one of the following:
 - In the job's **Select Action** menu, click **History / Logs**.
 - In the **Last Backup Status** column, click the job status.


The History / Logs or Logs window lists the most recent backups, restores and synchronizations 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.

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.



10.9 View and export recent backup statuses

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

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


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

To view and export recent backup statuses:

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

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

2. To change which backup statuses appear on the page, click the views list at the top of the page, and then click the view that you want to apply.
3. To view information for a job or computer on the Computers page, click the name of an online computer or job.
4. To view the job's logs in the History/Logs window, click the job's last backup status.
5. To export backup status information from the page, click the **Export** box. In the list that appears, click one of the following formats for the exported data file:
 - CSV (comma-separated values)
 - XLS (Microsoft Excel)
 - PDF (Adobe Acrobat)



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

11 Carbonite Server Backup Support

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

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

What can we help you with?

Search

Popular Searches

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

11.1 Contacting Carbonite

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

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



CREATE A
SUPPORT CASE



CHAT WITH A
REPRESENTATIVE



CALL A SUPPORT
REPRESENTATIVE

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

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

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