

CARBONITE[®]

an **opentext**[™] company

Carbonite Server Backup

Linux Agent and Oracle Plug-in 9.2

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 with Linux Agent 9.20.

Contents


1 Introduction to the Linux Agent	5
2 Install the Linux Agent	7
2.1 Install and verify Relax-and-Recover for Linux BMR backups	10
2.2 Install or upgrade the Linux Agent in silent mode	11
2.3 Register the Linux Agent with Relax-and-Recover and enable BMR backups	13
2.4 Upgrade the Linux Agent	13
2.5 Change the Portal registration for a Linux Agent	15
2.6 Uninstall the Linux Agent	16
3 Configure the Linux Agent	17
3.1 Add vault settings	17
3.2 Add a description	19
3.3 Add retention types	19
3.4 Configure bandwidth throttling	21
3.5 Resolve certificate failures	22
4 Add a Linux backup job	24
4.1 Add the first backup job for a Linux server	27
4.2 Add an NFS backup job	28
4.3 Log file options	30
4.4 Encryption settings	30
4.5 Advanced backup options	31
4.6 Filter subdirectories and files in backup jobs	32
5 Run and schedule backups and synchronizations	34
5.1 Schedule a backup	35
5.2 Schedule a backup to run multiple times per day	38
5.3 Maximum number of restore points for a job	42
5.4 Specify whether scheduled backups retry after a failure	42
5.5 Run an ad-hoc backup	43
5.6 Synchronize a job	44
6 Restore Linux files and folders	46
6.1 Restore ACLs	48
6.2 Restore data to a replacement computer	49
6.3 Restore data from another computer	50
6.4 Advanced restore options	51
6.5 Filter subdirectories and files when restoring data	52
6.6 Search for files to restore	53

7 Restore a Linux system from a BMR backup	55
8 Restore a Linux system without a BMR backup	59
8.1 Hardware requirements	59
8.2 Software requirements	59
8.3 Recovery steps	60
8.4 Recovery problems	61
9 Back up and restore Oracle databases using the Oracle Plug-in	62
9.1 Install the Oracle Plug-in for Linux	62
9.2 Add an Oracle database backup job	63
9.3 Restore Oracle databases	67
9.4 Uninstall the Oracle Plug-in for Linux	68
10 Delete jobs and computers, and delete data from vaults	70
10.1 Delete a backup job without deleting data from vaults	70
10.2 Delete a backup job and delete job data from vaults	71
10.3 Cancel a scheduled job data deletion	73
10.4 Delete a computer without deleting data from vaults	74
10.5 Delete a computer and delete computer data from vaults	75
10.6 Cancel a scheduled computer data deletion	77
10.7 Delete specific backups from vaults	78
11 Monitor computers, jobs and processes	80
11.1 Monitor backups and computers using the Current Snapshot	80
11.2 View computer and job status information	81
11.3 View skipped rates and backup status histories	83
11.4 View an unconfigured computer's logs	86
11.5 View current process information for a job	87
11.6 Monitor backups using email notifications	89
11.7 View a job's process logs and safeset information	93
11.8 View and export recent backup statuses	95
12 Carbonite Server Backup Support	96
12.1 Contacting Carbonite	96

1 Introduction to the Linux Agent

The Linux Agent backs up data on Linux systems, and restores data from the backups.

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



The Linux Agent can back up:

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

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

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

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

2 Install the Linux Agent

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

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

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

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

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

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

To install the Linux Agent:

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

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

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

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

```
cd packageName
```

5. Run the following command to start the installation:

```
./install.sh
```

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

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

```

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

                          Installing Backup Agent

Installation directory? [/opt/BUAgent] _

```

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

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

```

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

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

```

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

```

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

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

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

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

Select language: [en-US] _

```

You are then prompted to choose the data encryption method.

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

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


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

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

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

10. Do one of the following:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

For BMR backup support with Linux Agent 9.20, Relax-and-Recover (rear) version 2.6 must be installed on the Linux system. This section describes how to install this open-source tool and verify the installation. For more information, see the Relax-and-Recover website: <http://relax-and-recover.org/>

Note: Relax-and-Recover version 2.5 was required for BMR backups with Linux Agent 8.83. Before upgrading the Linux Agent from version 8.83 to version 9.20, we recommend uninstalling Relax-and-Recover version 2.5 and performing a fresh install of Relax-and-Recover version 2.6, as described in this procedure.

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

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

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

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

1. Ensure that the following Relax-and-Recover requirements are installed on the Linux system:
 - bash
 - mkisofs or genisoimage
 - mingetty

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

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

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

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

```
rear -V
```

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

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

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

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

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

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

2.2 Install or upgrade the Linux Agent in silent mode

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

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

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

Linux Agent installation parameters

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

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

Linux Agent registration options

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

```
wccAddress=ADDRESS_OF_AMP_SERVER
```

```
wccPort=PORT_OF_AMP_SERVER # Defaults to 8086
```

```
wccLogin=PortalUserName
```

```
wccPassword=PortalPassword
```

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

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

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

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

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

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

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

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

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

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

2.3.1 Disable BMR backups

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

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

To disable Linux BMR backups:

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

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

2.4 Upgrade the Linux Agent

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

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

Relax-and-Recover version 2.5 was required for BMR backups with Linux Agent 8.83. Before upgrading the Linux Agent from version 8.83 to version 9.20, we recommend uninstalling Relax-and-Recover version 2.5

and performing a fresh install of Relax-and-Recover version 2.6. See [Install and verify Relax-and-Recover for Linux BMR backups](#).

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

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

To upgrade the Linux Agent:

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

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

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

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

```
cd packageName
```

4. Run the following command to start the upgrade:

```
./install.sh
```

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

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

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

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

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

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

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

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

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

2.5 Change the Portal registration for a Linux Agent

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

The Agent is restarted when you change the Portal registration.

To change the Portal registration for a Linux Agent:

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

```
./register
```

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

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

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

The Agent is restarted and the Portal registration is changed.

2.6 Uninstall the Linux Agent

To uninstall the Linux Agent:

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

```
./uninstall.sh
```

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

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

3 Configure the Linux Agent

After a Linux agent is installed and registered with Portal, you can configure settings for the agent. Settings include:

- Vault connections. Vault connections provide vault information and credentials so that the agent can back up data to and restore data from the vault. See [Add vault settings](#).
- Description for the protected computer. The description appears for the agent on the Computers page in Portal. See [Add a description](#).
- Retention types. Retention types specify how long backups are kept on the vault. See [Add retention types](#).
- Amount of bandwidth consumed by backups and restores. See [Configure bandwidth throttling](#).
- Email notifications, so that users receive emails when backups complete, fail, or have errors. See [Monitor backups using email notifications](#).

If an agent reports a certificate failure, you must resolve the certificate failure before backups and restores can continue. See [Resolve certificate failures](#).

3.1 Add vault settings

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

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

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

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

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

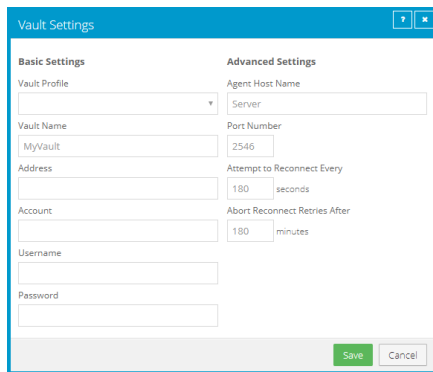
To add vault settings:

1. On the navigation bar in Portal, click **Computers**.
2. Find the Agent for which you want to add vault settings, and click the Agent 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 Agent on the vault.
- **Port Number.** Port used to connect to the vault. The default port is 2546.
- **Attempt to Reconnect Every.** Specifies the number of seconds after which the agent should try to connect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 30 to 1800 seconds.
- **Abort Reconnect Retries After.** Enter the number of minutes after which the agent should stop trying to reconnect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 60 to 720 minutes. If the Agent cannot connect to the vault successfully in the specified time, the backup or restore fails.

6. Click **Save**.

3.2 Add a description

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

To add a description:

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


The Computers page shows registered computers.

2. Find the Agent 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 Agent.



5. Click **Save**.

3.3 Add retention types

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

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

You cannot add, change or delete retention types for intra-daily schedules. For intra-daily schedules, you must choose one of two intra-daily retention types that are available beginning in Portal 8.88. See [Schedule a backup to run multiple times per day](#).

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

To add a retention type:

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

The Computers page shows registered computers.

2. Find the Agent for which you want to add a retention type, and click the row to expand its view.

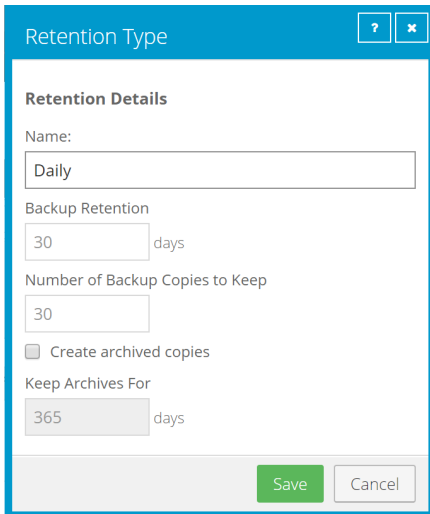
If the Configure Manually box appears, click **Configure Manually**. The Configure Manually box appears for some computers where a backup job has not been created.

3. On the Advanced tab, click the **Retention Types** tab.

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

4. Click **Create Retention Type**.

The Retention Type dialog box appears.



5. Complete the following fields:

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

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

6. Click **Save**.

3.4 Configure bandwidth throttling

Bandwidth settings include:

- Maximum bandwidth (upper limit), in megabits per second, to be consumed by the Agent for backups and restores. For example, if three jobs are running at the same time on the same computer, each job gets 1/3 of the specified maximum bandwidth.
- Period of time during the day that throttling is in effect. Only one time window can be specified. Outside the window, no throttling takes place.
- Days of the week that throttling is in effect.

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

If you edit an Agent’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 an Agent, 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:

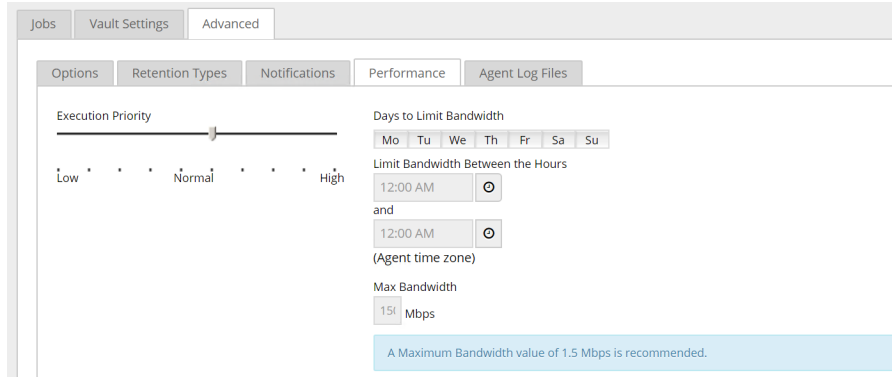
1. On the navigation bar, click **Computers**.
2. Find the Agent for which you want to configure bandwidth throttling, and click the row to expand its view.

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 Agent, 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**.

3.5 Resolve certificate failures

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

A certificate failure can occur when a Linux agent tries to connect to a Director version 8.60 vault where certificate pinning is enabled. Beginning with Linux Agent 8.90, when a Linux agent tries to connect to the vault (e.g., to run a backup or restore), it checks whether the public key of the vault's TLS certificate is the same as when the agent previously connected to the vault. If the public key of the vault certificate is different, the agent reports a certificate failure and will not connect to the vault.

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

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

To resolve certificate failures:

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

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

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

4 Add a Linux backup job

After a Linux system is added in Portal, you can create backup jobs for the system.

You can create backup jobs for files and folders that are saved locally on the computer. The backup job specifies which folders and files to back up, and where to save the data. You can also create a backup job for files and folders that are saved on mounted NFS shares. See [Add an NFS backup job](#).

You can also create Bare Metal Restore (BMR) backup jobs that can be used to restore entire Linux systems. A Linux BMR backup includes an .iso file for starting the destination system and running a restore, and a backup in the vault that includes all required system volumes and files.

IMPORTANT: We recommend creating only one BMR backup job for each Linux system. If you create and run multiple BMR backup jobs, the resulting .iso file might not be usable.

Note: By default, a Linux BMR job backs up all data on the system. You can exclude folders from the BMR backup. However, if you exclude any of the following required folders, the exclusion will be ignored when the backup job runs: /bin; /boot; /etc; /lib; /lib64; /root; /usr/bin; /usr/lib; /usr/lib64; /usr/share; /usr/sbin

Note: On a server where Oracle or another database is running, we recommend that you shut down database services when running a BMR job. Alternatively, on a server where Oracle is running, you can exclude database directories from the BMR job and set up a separate Oracle Plug-in job for the database. Otherwise, database data might be inconsistent after it is restored.

A symbolic link (also called a symlink or soft link) consists of a special type of file that serves as a reference to another file or directory. During a backup, a symbolic link gets backed up with the timestamp of the link. Restoring a symbolic link sets its modification date and time to the date and time of the restore (rather than the date and time of the symbolic link when it was backed up).

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

To add a Linux backup job:

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

The Computers page shows registered computers.

2. Find a Linux system, and expand its view by clicking the computer row.

If a backup job has not been created for a Linux computer, the system can attempt to create a backup job automatically. See [Add the first backup job for a Linux server](#).

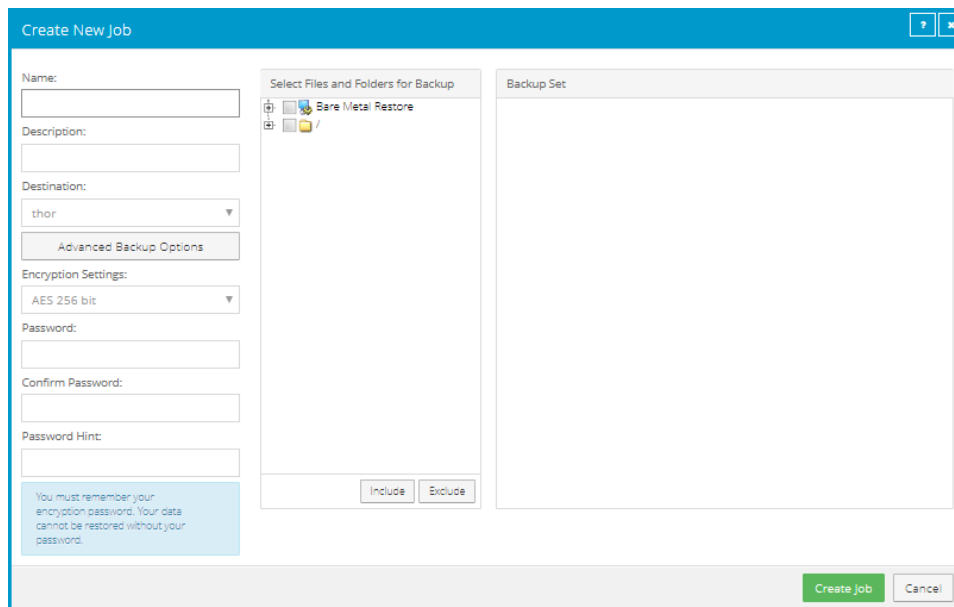
3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the Jobs tab. See [Add vault settings](#).

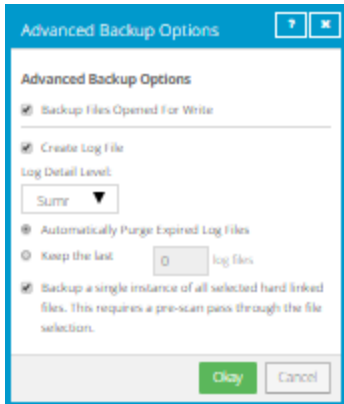
4. In the **Select Job Task** menu, click **Create New Local System Job**.

5. In the Create New Job dialog box, specify the following information:

- In the **Name** box, type a name for the backup job.
- In the **Description** box, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.
A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.



6. To change log file or other backup options, click **Advanced Backup Options**. In the Advanced Backup Options dialog box, select options and then click **Okay**. For more information, see [Log file options](#) and [Advanced backup options](#).




7. In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the folders and files that you want to include in and exclude from the backup:

- To create an ISO file that contains the volumes and files that are needed to boot up the system, and to back up all system data, select **Bare Metal Restore**, and then click **Include**.

By default, a Linux BMR job backs up all data on the system. You can exclude folders from the backup but, if you exclude a required folder, the exclusion will be ignored when the backup job runs and a message will appear in the log file.

When a Linux BMR job runs, it creates a boot ISO file named Bare_Metal_Restore_Image.iso in the root directory (/). The file is overwritten every time a BMR job runs. We recommend reserving a minimum of 1 GB of space in the root file system for the .iso file when you first run a BMR backup job.

- To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#)
 - To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
 - To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 
8. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups and synchronizations](#).

4.1 Add the first backup job for a Linux server

Portal can automatically create a backup job for a Linux computer that does not have a backup job. An automatically-created job backs up everything from the root, and is scheduled to run every night.

After a job is automatically created, you can change the job settings, if desired. For example, you can specify different directories to back up or change the schedule for running the job.

A valid vault profile must be available before Portal can automatically create a backup job.

After a job is created, you can change the job settings, if desired. For example, you can specify different folders to back up or change the schedule for running the job.


To add the first backup job for a Linux server:

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

The Computers page shows registered computers.

2. Find a Linux computer, and expand its view by clicking the computer row.

If a backup job has not been created for the computer, the Configure Manually box appears. If a backup job has not been created for the computer and at least one vault profile is available, the Configure Automatically box also appears.



3. Do one of the following:


- To create a backup job manually, click **Configure Manually**. See [Add a Linux backup job](#).
- To automatically create a backup job for the computer, do the following:

- a. In the **Create a password** and **Confirm your password** boxes, enter an encryption password.

IMPORTANT: Your encryption password is required for restoring your data, so be sure to store it somewhere safe. If you forget the password, you will not be able to restore your data. The password is not maintained anywhere else and cannot be recovered.

- b. In the **Password hint** box, enter a hint to help you remember the encryption password.
- c. If the **Assign the computer to a site** list appears, choose a site for the computer.

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



- d. If more than one vault is available, choose a vault from the **Choose a vault** list.
- e. Click **Configure automatically**.

If the configuration succeeds, a backup job appears for the computer.

If the automatic job configuration fails, do the following:

- i. Click **Configure Manually**.
- ii. On the Vault Settings tab, click **Add Vault**.
- iii. In the Vault Settings dialog box, enter vault information and credentials.
- iv. Create a backup job manually. See [Add a Linux backup job](#).

4.2 Add an NFS backup job

After a Linux system is added in Portal, you can create a backup job for files and folders that are saved on mounted NFS shares. The backup job specifies which folders and files to back up, and where to save the data.

NFS servers must share their exports in order to make them available to client systems. If you want to perform a mount-point backup or restore, the NFS server must be available, and it must provide sufficient privileges to your client system. Also, the NFS must be mounted on your client system at the time of the backup or restore.

Note: If you restore an NFS backup, and the NFS mount does not exist, the restore will proceed as if it were a local restore. It will put the data on the local disk (with a similar path that is local) without using a mount-point (NFS) path. It will not indicate a “failure”.

NFS does not export extended attributes from remote file systems. On Linux NFSv3 clients, remote file system ACLs will be presented as standard Linux ACLs if possible. NFSv4 clients will present remote file system ACLs as native NFSv4 ACLs, but the Agent will protect them as extended attributes.

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

To add an NFS backup job:

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

The Computers page shows registered computers.

2. Find a Linux system, and expand its view by clicking the computer row.

In some Portal instances, if a backup job has not been created for a Linux computer, the system can attempt to create a backup job automatically.

3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab.

4. In the **Select Job Task** menu, click **Create New NFS Files Job**.


5. In the Create New Job dialog box, specify the following information:

- In the **Name** box, type a name for the backup job.
- In the **Description** box, optionally type a description for the backup job.
- In the **Destination** list, select the vault where you want to save the backup data.

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.

- In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

6. In the **Select Files and Folders for Backup** box, do one or more of the following until the **Backup Set** box shows the folders and files that you want to include in and exclude from the backup:

- To add one or more folders or files to the backup job, select the check box for each folder or file, and then click **Include**. The included folders or files appear in the **Backup Set** box. If you include a folder, the backup job includes all of the folder's subdirectories and files by default. If you do not want to back up all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
- To exclude one or more folders or files from the backup job, select the check box for each folder or file, and then click **Exclude**. The excluded folders or files appear in the **Backup Set** box. If you exclude a folder, all of the folder's subdirectories and files are excluded from the backup job by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files in backup jobs](#).
- To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the folder or file record. 

7. Click **Create Job**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups and synchronizations](#).

4.3 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.

The following log file options are also available:

- **Create log file.** If this check box is selected, the system generates log files for each job. Log files can contain the start-connect-completion and disconnect times, file names (i.e., the names of the files that were copied during backup), and any processing errors.
- **Automatically purge expired log files.** If this check box is selected, the log file associated with a backup is automatically deleted when the backup has been deleted from the vault. Backups are typically deleted from the vault according to retention types. See [Add retention types](#).
- **Keep the last <number of> log files.** Specifies the number of log files to keep for a backup job. When the specified number is reached, the oldest log file for a backup job will be deleted to make space for the newest one.

Note: You must choose either the **Automatically purge expired log files** option or the **Keep the last <number of> log files** option. When a backup job runs, log files are removed according to the specified option. Log files are not removed when a backup job is synchronized.

4.4 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.

4.5 Advanced backup options

When you create or edit a Linux backup job, the following options are available in the Advanced Backup Options dialog box.

Back up a single instance of all selected hard linked files

A hard link is a reference, or pointer, to data on a storage volume. More than one hard link can be associated with the same data. Hard-linked files cannot cross disk boundaries and only exist on the same disk.

If the **Back up a single instance of all selected hard linked files** option is selected, only one copy of the data is backed up, along with all hard links. When the data is restored, both the data (with a new inode) and the hard links are restored. When this option is selected, a pre-scan process is required. The pre-scan reads through the file system, gets each inode and stores it in a map. The larger the file system, the more memory this map requires and the more time it takes to process. However, the resulting backup size is smaller.

If the **Back up a single instance of all selected hard linked files** option is not selected, the data is backed up separately for each hard link. When the data is restored, the hard-link relationship will not be re-established. Each file will be restored individually and applications depending on hard links may not be automatically restored. When this option is not selected, the backup is faster but the total backup size is larger.

4.6 Filter subdirectories and files in backup jobs

When you include and exclude folders in a backup job, the folder's subdirectories and files are also included or excluded by default.

If you only want to back up some subdirectories or files in a folder, you can add filters to the inclusion record. For example, you could add a filter so that files in a folder are only backed up if they have the .pl extension.

If you only want to exclude some subdirectories or files in a folder from a backup job, you can add filters to the exclusion record. For example, you could add a filter so that files in a folder are only backed up if they have the .mpg extension.


If a policy is assigned to a computer, you can add filters from the policy to a folder inclusion or exclusion record.

Filters in a backup job are applied when the job runs. New subdirectories and files that match the filters are automatically backed up or excluded when the job runs.

To filter subdirectories and files in a backup job:

1. When creating or editing a backup job, view the **Backup Set** box.

Backup Set			
	Folders Filter	Files Filter	Recursive
+ usr		*	✓
+ etc		*	✓

2. If editable fields do not appear for a folder inclusion or exclusion record where you want to filter subdirectories and files, click the **Edit** button in the folder row. 



3. In the **Backup Set** box, for each included folder where you want to include specific subdirectories or files, do one or more of the following:

- To include specific subdirectories in the backup job, in the **Folders Filter** field, enter the names of subdirectories to include. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only include subdirectories in a backup if their names end with "-current" or start with "2015", enter the following filter: *-current, 2015*

Note: Asterisks (*) are the only supported wildcards in filter fields.

- To include specific files in the backup job, in the **Files Filter** field, enter the names of files to include in the backup. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only include files in a backup if they have the .pl or .sh extension, enter the following filter: *.pl, *.sh

Note: Asterisks (*) are the only supported wildcards in filter fields.

- If a policy is assigned to the computer, to apply filters from the policy to the folder inclusion record, click the **Apply Policy Filters** button. 
 - To back up the specified folder, but not its subdirectories, clear the **Recursive** check box.
 - To back up the folder's subdirectories, select the **Recursive** check box.
4. In the **Backup Set** box, for each excluded folder where you want to exclude specific subdirectories or files, do one or more of the following:
- To exclude specific subdirectories from the backup job, in the **Folders Filter** field, enter the names of subdirectories to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude subdirectories from a backup if their names end with "-old" or start with "2001", enter the following filter: *-old, 2001*
- Note:* Asterisks (*) are the only supported wildcards in filter fields.
- To exclude specific files from the backup job, in the **Files Filter** field, enter the names of files to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude files from a backup if they have the .mpg or .gif extension, enter the following filter: *.mpg, *.gif
- Note:* Asterisks (*) are the only supported wildcards in filter fields.
- If a policy is assigned to the computer, to apply filters from the policy to the folder exclusion record, click the **Apply Policy Filters** button. 
 - To exclude the specified folder, but not its subdirectories, clear the **Recursive** check box.
 - To exclude the folder's subdirectories, select the **Recursive** check box.
5. Click **Create Job** or **Save**.

5 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](#).

To help you meet your recovery point objectives (RPOs), when Linux Agent 8.90 or later is backing up data to a Director version 8.60 or later vault, you can also schedule a backup job to run multiple times per day, as often as hourly. See [Schedule a backup to run multiple times per day](#).

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 computers with Linux Agent version 8.60 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](#).


5.1 Schedule a backup

After creating a backup job, you can add one or more schedules for running the job on specific days of the week or month. You can create complex schedules for a job by creating multiple schedules. For example, you can schedule a backup job to run at midnight every Friday, and schedule the job to run at 8 PM on the first day of every month.

Note: Beginning in Portal 8.88, when Linux Agent 8.90 or later is backing up data to a Director version 8.60 vault, you can also schedule a backup job to run multiple times per day, as often as hourly. See [Schedule a backup to run multiple times per day](#).

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 Agent with the backup job that you want to schedule, and click the row to expand its view. On the Jobs tab, find the job that you want to schedule. In its **Select Action** menu, click **View/Add Schedule**.
 - Create a new backup job. The View/Add Schedule dialog box appears when you save the job.
- In the View/Add Schedule dialog box, click **Add Schedule**.
A new row appears in the dialog box.
- In the new schedule row, in the **Retention** list, click a retention type.

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

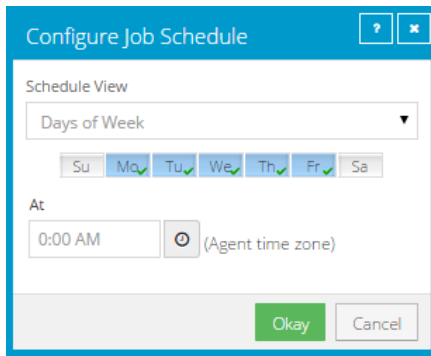
The 24-Hours and 48-Hours retention types are only available for intra-daily schedules. See [Schedule a backup to run multiple times per day](#).

6. 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:

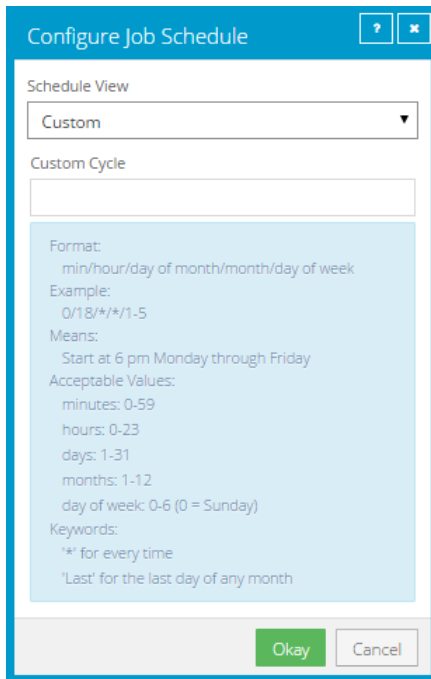
- To run the backup on specific days each week, select **Days of Week** in the **Schedule View** list. Select the days when you want to run the job. Then use the **At** field to specify the time when you want to run the job.



- 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.



Note: If **Intra-daily** appears in the **Schedule View** list, you can also schedule the backup to run multiple times each day. See [Schedule a backup to run multiple times per day](#).

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 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.

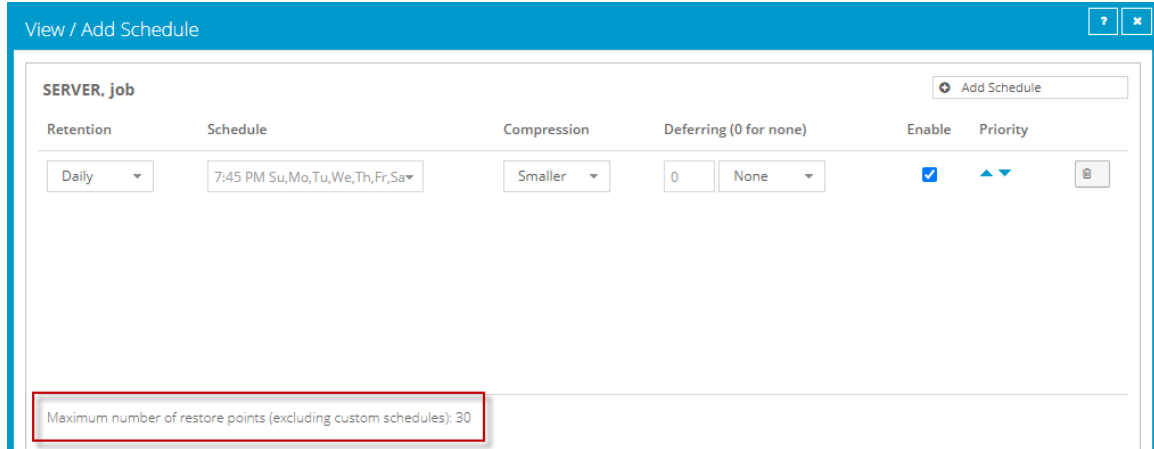
Note: When deferring is used, the backup job does not back up any new data after the specified amount of time, even if some data is not backed up. Changes to data that was previously backed up will be backed up, regardless of the amount of time specified.

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.

- 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](#).



- 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](#).
- Click **Save**.

5.2 Schedule a backup to run multiple times per day

Beginning in version 8.90, when the Linux Agent is backing up data to a Director version 8.6x vault, you can schedule the backup job to run multiple times per day by creating an intra-daily schedule using Portal 8.88 or later.

Note: To schedule a backup job to run on specific days of the week or month, see [Schedule a backup](#).

Each backup job can have one intra-daily schedule. If the job has other schedules, the intra-daily schedule has the lowest priority and is at the bottom of the schedule list. If a job is scheduled to start at exactly the same time by an intra-daily schedule and another schedule, the job only runs once and the retention type of the other schedule (e.g., daily or monthly) is applied to the resulting safeset.

When you create an intra-daily schedule for a backup job, you can choose one of two retention types:

- 24-Hours.** With this retention type, each backup is kept for at least 24 hours and at least one backup with this retention type is stored online.
- 48-Hours.** With this retention type, each backup is kept for at least 48 hours and at least one backup with this retention type is stored online.

Other retention types are not available for intra-daily schedules. You cannot add, change or delete retention types for intra-daily schedules.

When you schedule a backup, the View/Add Schedule dialog box shows the maximum number of restore points that could result from the job's current schedules and retention types. You can then change your schedules, if required. See [Maximum number of restore points for a job](#).

To reduce schedule overloads, backups that are scheduled by intra-daily schedules are skipped in some cases. See [Skipped backups](#).

To schedule a backup job to run multiple times per day:

1. Do one of the following:
 - On the navigation bar, click **Computers**. Find the computer with the backup job that you want to schedule, and click the row to expand its view. On the Jobs tab, find the job that you want to schedule. In its **Select Action** menu, click **View/Add Schedule**.
 - Create a new backup job. The View/Add Schedule dialog box appears when you save the job.

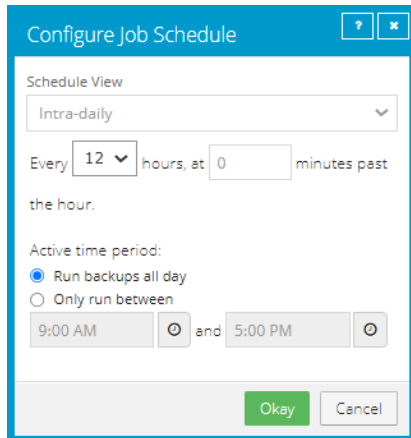
2. In the View/Add Schedule dialog box, click **Add Schedule**.

A new row appears in the dialog box.

3. In the new schedule row, click the arrow in the **Schedule** box.

IMPORTANT: To create an intra-daily schedule, you must select **Intra-daily** in the Schedule box before selecting a retention type.

4. In the Configure Job Schedule dialog box, do the following:
 - a. In the **Schedule View** list, select **Intra-daily**.



- b. In the **Every x hours** list, click the frequency for running the job. You can schedule the job to run every 1, 2, 3, 4, 6, 8 or 12 hours.
- c. In the **at y minutes past the hour** box, type the number of minutes after the hour when you want to run the job. For example, enter 15 to run the job at 15 minutes past each hour when the job runs.
- d. In the Active time period area, do one of the following:

- To run the job at the specified frequency for the full 24 hour period, click **Run backups all day**.
- To run the job according to the intra-daily schedule for only part of each 24-hour day period, click **Only run between**. Click the first clock icon and specify the start of the time period for running backups at the specified frequency. Click the second clock icon and specify the end of the time period for running backups at the specified frequency.

e. Click **Okay**.

If the job has other schedules, the intra-daily schedule appears at the bottom of the schedule list and has the lowest priority. The priority of the intra-daily backup schedule cannot be changed.

5. In the **Retention** list, click one of the following retention types:

- **24-Hours**. With this retention type, each backup is kept for at least 24 hours and at least one backup with this retention type is stored online.
- **48-Hours**. With this retention type, each backup is kept for at least 48 hours and at least one backup with this retention type is stored online.

Other retention types are not available for intra-daily schedules.

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

The Configure Job Schedule dialog box opens.

7. Click **Okay**.

The new schedule appears in the Schedule box.

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

9. Do one of the following:

- To allow the backup job to run without a time limit, click **None** in the Deferring list.
- To specify a maximum amount of time that the backup job can run, click **Minutes** or **Hours** in the **Deferring** list. In the adjacent box, type the maximum number of minutes or hours that the job can run.

Note: When deferring is used, the backup job does not back up any new data after the specified amount of time, even if some data is not backed up. Changes to data that was previously backed up will be backed up, regardless of the amount of time specified. For deferral behavior for specific backup types, see [Run and schedule backups and synchronizations](#).

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

11. 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](#).

12. In the Automatic Retry for Scheduled Backups section at the bottom of the View / Add Schedule dialog box, specify whether scheduled backups should retry after a failed backup. See [Specify whether scheduled backups retry after a failure](#).
13. Click **Save**.

5.2.1 Skipped backups

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

To reduce schedule overloads when a backup job runs multiple times per day, backups are skipped when:

- An agent starts a backup that is scheduled by an intra-daily schedule, and a backup is already running for the job.
- An agent contacts a Director version 8.60 or later vault to start a backup that is scheduled by an intra-daily schedule, and the vault is busy with high-priority maintenance for the job data.

If email notifications are configured centrally in a Portal instance, Admin users can receive an email when a backup is skipped. See [Set up email notifications for backups on multiple computers](#). When the last backup status reported for a job was "Skipped", this Last Backup Status appears for the job on the Computers page and Monitor page. See [View computer and job status information](#) and [View and export recent backup statuses](#). The Daily Status report also shows skipped backups.

In some Portal instances, users can also see skipped rates and 48-hour backup status histories for jobs. See [View skipped rates and backup status histories](#).

Best practices: Reducing the number of skipped backups

If you notice that some backups are skipped frequently, you can make changes to the backup job, backup schedule, or servers to ensure reliable backups. For example, you could:

- Reduce the frequency of the scheduled backups.
- Reduce the size of the job.
- Add system resources (e.g., RAM, CPU, Storage IO) on the server where the agent is running. While the resources on a server might be sufficient for backing up and restoring data periodically, the resources might not be sufficient to run backups multiple times per day.
- Add system resources to the vault server.

5.3 Maximum number of restore points for a job

Beginning in Portal version 8.88, when you schedule a backup job, the View/Add Schedule dialog box shows the maximum number of restore points that could result from the job's current schedules and retention types. The maximum number of restore points, or backups in the vault, is updated when you add or change a schedule row so you can understand the impact of your schedule changes and make additional changes, if required.

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.

5.4 Specify whether scheduled backups retry after a failure

You can specify whether scheduled backups automatically retry if they do not run successfully.

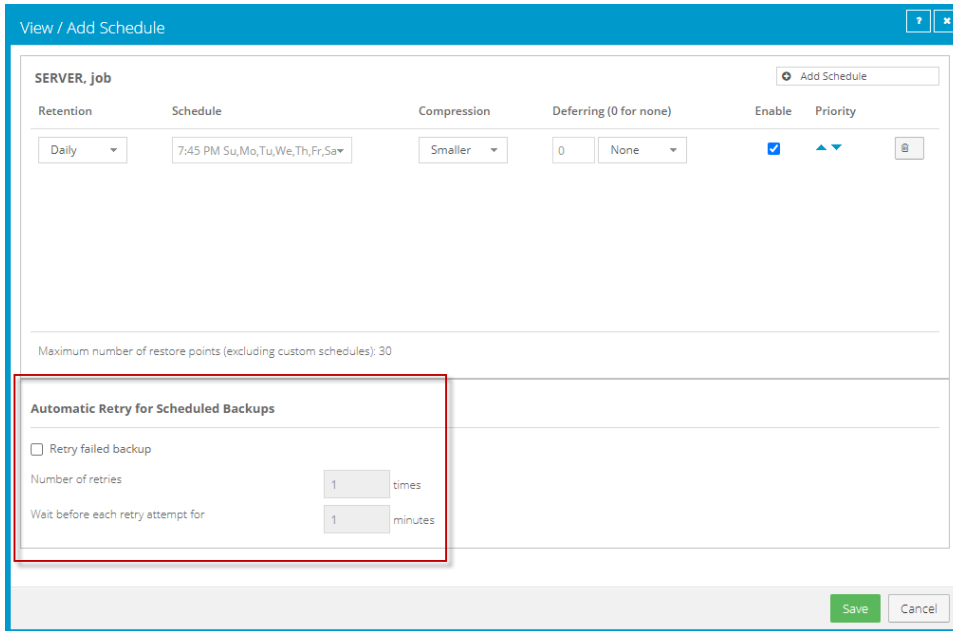
You can also specify how many times a scheduled backup should retry after a failed attempt, and specify the amount of time between retries.

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 Agent 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.



3. Click **Save**.

5.5 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 Agent 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**.
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**.

5.6 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 Agent with the job that you want to synchronize. Expand its view by clicking its row.

3. Click the **Jobs** tab.
4. Find the job that you want to synchronize, and click **Synchronize** in its **Select Action** menu.

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

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

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

6 Restore Linux files and folders

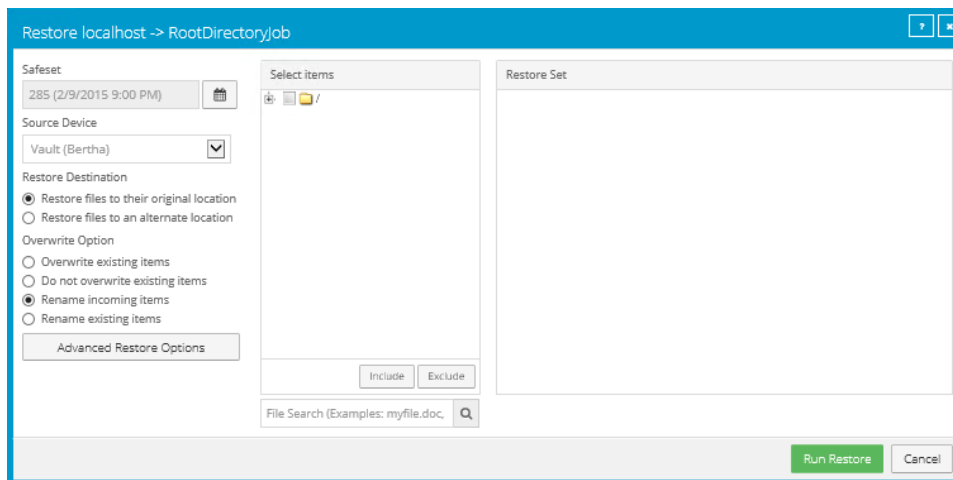
After backing up data from a Linux computer, you can restore files and folders from the backup.


You can also restore Linux systems from BMR backups. See [Restore a Linux system from a BMR backup](#).

To restore Linux files and folders:

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



The Restore dialog box shows the most recent safeset for the job.



5. To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:
 - To restore data from an older safeset, click the calendar button. In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
 - To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the Select Folder dialog box, select the directory where the files are located, and click **Okay**.

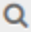

SSI files are full backups exported from the vault or backed up to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

6. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
7. Select a Restore Destination option.
 - To restore files and folders to the location where they were backed up, select **Restore files to their original location**.
 - To restore files and folders to a different location, select **Restore files to an alternate location**. Click the folder button.  In the Select Folder dialog box, select the location where you want to restore, and click **Okay**.
8. Select an Overwrite Option. This option specifies how to restore a file, folder or symbolic link to a location where there is a file, folder or symbolic link with the same name.
 - To overwrite the existing item with the restored item, select **Overwrite existing items**.

Note: If you try to restore multiple files with the same name to an alternate location and select **Overwrite existing items**, only the last file restored will remain. Other files with the same name will be overwritten.

IMPORTANT: Using Agent version 8.70 or later, if you select **Overwrite existing items** and restore a file that has the same name as a folder in the restore location, the file will overwrite the folder. The folder and all of its contents will be removed.
 - To skip restoring the item that has the same name as an item in the destination location, select **Do not overwrite existing items**.
 - To add a numeric extension (e.g., .0001) to the *restored* item name, select **Rename incoming items**. For example, if you restore a file named “filename.txt” to a location where there is a file with the same name, an extension is added to the **restored** file name (e.g., “filename.txt.0001”).
 - To add a numeric extension (e.g., .0001) to the *existing* item name, select **Rename existing items**. For example, if you restore a file named “filename.txt” to a location where there is a file with the same name, an extension is added to the *existing* file name (e.g., “filename.txt.0001”). The name of the restored file is “filename.txt”.
9. To change locked file, data streams, log detail level or bandwidth options, click **Advanced Restore Options**. Specify settings in the Advanced Restore Options dialog box, and click **Okay**. See [Advanced restore options](#).
10. In the **Select Items** box, do one or more of the following until the **Restore Set** box shows the folders and files that you want to restore:

- Select the check box for each folder and file that you want to restore, and then click **Include**. The **Restore Set** box shows the included folders and files. If you include a folder, all of the folder's subdirectories and files are restored by default. If you do not want to restore all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).
 - To exclude one or more folders or files from the restore, select the check box for each folder or file, and then click **Exclude**. The **RestoreSet** box shows the excluded folders and files. If you exclude a folder, all of the folder's subdirectories and files are excluded from the restore by default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).
 - To search for files to restore or exclude from the restore, click the **Search** button.  In the **Search for files** box, enter search criteria and select files. See [Search for files to restore](#). Click **Include Selected** or **Exclude Selected**. The **RestoreSet** box shows the included or excluded files.
 - To remove an inclusion or exclusion record from the **Restore Set** box, click the Delete button beside the folder or file record. 
11. Click **Run Restore**.
- The Process Details dialog box shows the restore progress and indicates when the restore is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).
12. To close the Process Details dialog box, click **Close**. If the restore is running, it will continue to run.

6.1 Restore ACLs

You can back up and restore Access Control Lists (ACLs). The following behaviors can occur when you restore ACLs on a Linux server.

ACLs control the access of users or groups to particular files. Similar to regular file permissions (e.g., owner, group, world), ACLs are tracked by the ID of the user/group. ACLs provide access-control granularity beyond regular file permissions, and unlike regular permissions, they are not always enabled.

ACL implementations might differ by variety of Linux, and by the type of file system. Not all ACL implementations are "portable" (i.e., ACLs on one OS/file system may be incompatible with ACLs on another OS/file system). In addition, you might need to enable ACL support on a partition before you can configure it.

If you attempt to restore ACLs to an incompatible system (e.g., a file system that does not support ACLs), the ACLs will not be restored. An error message will appear in the backup log.

If you restore to a compatible system (e.g., the original system, or a different system with the same variety of Linux), ACLs will also be restored.

Since ACLs are associated with user and group IDs, you will observe the following on a compatible system:

- If the group, user names, and IDs on the restored system match those of the original system, the ACLs will be associated with the same user name as on original system.
- If the group, user names, and IDs on the restored system do not match those on the original system, the ACLs will be associated with a different user or group name compared to the original system.
- If the group or user name ID does not exist on the restored system, the ACLs will be associated with the user ID or group ID respectively. Therefore, browsing ACLs on these files will show user/group IDs as opposed to user/group names.

6.2 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.

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.

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](#).

6.3 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 (e.g., Linux and Solaris) are also acceptable.

To restore data from another computer:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer to which you want to restore the data, and expand its view by clicking the computer row.
3. In the **Job Tasks** menu, click **Restore from Another Computer**.
The Restore From Another Computer dialog box opens.
4. In the **Vaults** list, select the vault where the backup is stored.
5. In the **Computers** list, select the computer with the backup from which you want to restore.
6. In the **Jobs** list, select the job from which you want to restore data.

7. Click **Okay**.

Portal attempts to download information about the selected job. After the job information is downloaded, the job appears on the computer's Jobs tab. You can then continue restoring data as you would in a regular restore.

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.

6.4 Advanced restore options

When restoring data, you can specify the following options:

Locked File Options

When restoring data from a local job, you can specify whether to overwrite locked files with restored files with the same names. You can select one of the following options:

- **Yes, overwrite locked files** – Files on the system that are locked during the restore are overwritten by restored files when the system restarts. You must select this option for a system state or system volume restore.
- **No, do not overwrite locked files** – Files on the system that are locked during the restore are not overwritten by restored files with the same name.

Streams

When running a backup, information is collected from your files in various streams. Original data created by a user is called a data stream. Other information, such as security settings, data for other operating systems, file reference information and attributes, are stored in separate streams.

When restoring data, you can select one of the following options:

- **Restore all streams** – Restores all information streams. This option is recommended if you are restoring files to a system with an identical platform.
- **Restore data streams only** – For cross-platform restores, restores data streams only. This option ensures that conflicts do not arise as a result of system-specific information streams.

Log Options

Select one of the following job logging levels from the list:

- **Files** — Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up.
- **Directory** — Provides less detail than the Files logging level. Provides information about folders that are backed up.

- **Summary** — Provides high-level information, including the vault and Agent version, and backup sizes.
- **Minimal** — Provides high-level information, including the vault and Agent version.

Changing the logging level only affects log files that are created from that point and after. It does not affect previously-created log files.

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.

6.5 Filter subdirectories and files when restoring data


When you restore data, you can specify folders and files to restore or not restore from the backup.


By default, when you include a folder in a restore, the folder’s subdirectories and files are also included. If you only want to restore some of a folder’s subdirectories or files, you can add filters to the inclusion record. For example, you could add a filter so that files in a folder are only restored if they have the .pl extension.

By default, when you exclude a folder from a restore, the folder’s subdirectories and files are also excluded. If you only want to exclude some of a folder’s subdirectories or files, you can add filters to the exclusion record. For example, you could add a filter so that files in a folder are only excluded from the restore if they have the .sh extension.

To filter subdirectories and files when restoring data:

1. When restoring data, view the **Restore Set** box.




2. If editable fields do not appear for a folder inclusion or exclusion record where you want to filter subdirectories and fields, click the **Edit** button in the folder row. 
3. In the **Restore Set** box, for each included folder where you want to include specific subdirectories or files, do one or more of the following:

- To include specific subdirectories in the restore, in the **Folders Filter** field, enter the names of subdirectories to include. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only restore subdirectories if their names end with “-current” or start with “2015”, enter the following filter: *-current, 2015*
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To restore specific files, in the **Files Filter** field, enter the names of files to restore. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only restore files if they have the .pl extension, enter the following filter: *.pl
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To restore the specified folder, but not its subdirectories, clear the **Recursive** check box.
 - To restore the folder’s subdirectories, select the **Recursive** check box.
4. In the **Restore Set** box, for each excluded folder where you want to exclude specific subdirectories or files, do one or more of the following:
- To exclude specific subdirectories from the restore, in the **Folders Filter** field, enter the names of subdirectories to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude subdirectories from a restore if their names end with “-old” or start with “2001”, enter the following filter: *-old, 2001*
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To exclude specific files from the restore, in the **Files Filter** field, enter the names of files to exclude. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only exclude files from a restore if they have the .pl extension, enter the following filter: *.pl
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To exclude the specified folder, but not its subdirectories, clear the **Recursive** check box.
 - To exclude the folder’s subdirectories, select the **Recursive** check box.
5. Click **Run Restore**.


6.6 Search for files to restore

When you restore data, you can search for files to restore or exclude from the restore.

To search for files to restore:

1. In the Restore dialog box, click the **Search** button. 

The Search for files dialog box appears.



2. In the **Search Text** box, enter the file name to search for. You can include asterisks (*) as wildcard characters.
3. To search for files in a specific folder in the backup, enter the path in the **Search Path** box.
4. To search for files only in the specified folder, clear the **Search all subdirectories** check box.
5. Click **Search**.
The Results box lists files that match the search criteria.
6. In the **Results** box, select files to include or exclude. To select multiple consecutive items, press SHIFT while clicking the first and last items in the list. To select multiple items, press CTRL while clicking the items.
7. Do one of the following:
 - To restore the selected files, click **Include Selected**.
 - To exclude the selected files from the restore, click **Exclude Selected**.

7 Restore a Linux system from a BMR backup

You can restore entire Linux servers from Bare Metal Restore (BMR) backups. A BMR backup includes:

- An .iso file for starting the destination system and running the restore. The .iso file is created on the source system during a BMR backup and is backed up to the vault
- A backup in the vault of all folders and files that are required for the system. By default, a Linux BMR backup includes all folders and files from the root (/), although some files can be excluded.

When restoring a Linux server from a BMR backup, the destination machine must have:

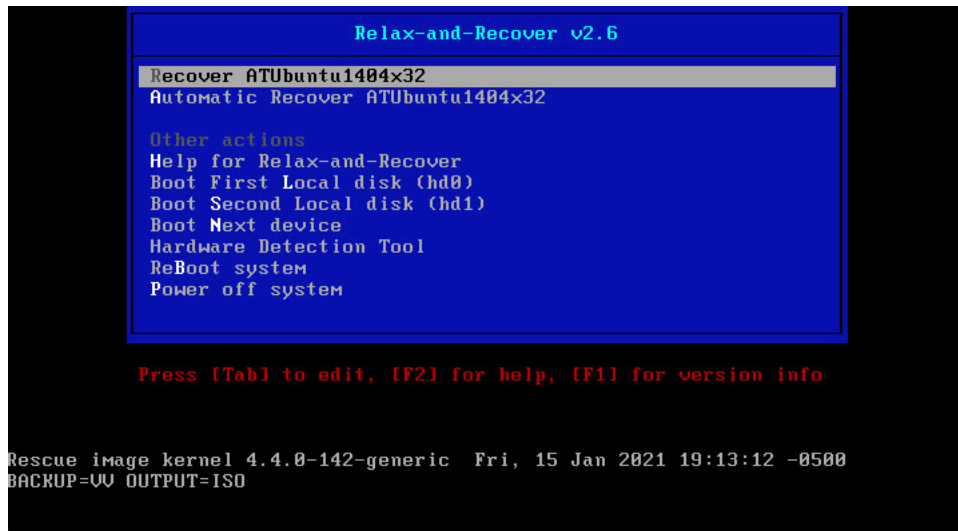
- At least 4 GB of RAM.
- The same boot type (BIOS or UEFI) as the source system, and compatible hardware.
- Hard drives that are the same size or larger than drives on the source system.
- A connection to the network, so that it can communicate with the vault.

Note: Restores are not supported to systems with different types of firmware.


To restore a Linux system from a BMR backup:

1. Do one of the following:
 - If the source system is still available, copy the `/Bare_Metal_Restore_Image.iso` file from the root directory (/) of the source system to another machine.
 - Use the *Restore from another computer* procedure to restore the `/Bare_Metal_Restore_Image.iso` file from the Linux BMR backup to another machine. See [Restore data from another computer](#).
2. Create a bootable USB device, CD or DVD from the `Bare_Metal_Restore_Image.iso` file and mount it on the destination system.
3. Boot the destination system from the bootable file.
4. Do one of the following:
 - If the Relax-and-Recover screen appears, select **Recover *sourceSystemName*** and then press Enter. This screen appears if the protected system is BIOS-based.

Do not select the Automatic Recover *sourceSystemName* option or the system will not start successfully.



- If the following screen appears, select **Relax-and-Recover (no Secure Boot)** and then press Enter. This screen appears if the protected system is UEFI-based.



5. At the login prompt, log in as root.

Note: If a login prompt does not appear initially, press Enter.


```
Relax-and-Recover 2.5 / 2019-05-10
Relax-and-Recover comes with ABSOLUTELY NO WARRANTY; for details see
the GNU General Public License at: http://www.gnu.org/licenses/gpl.html

Host localhost.localdomain using Backup UU and Output ISO
Build date: Thu, 26 Mar 2020 14:42:11 -0500

CentOS Linux 7 (Core)
Kernel 3.10.0-1062.18.1.el7.x86_64 on an x86_64

SSH fingerprint: 2048 SHA256:1at0v5avfn2Qim+kodeTBIi5cgXSC6oeWaabSe0oQAE root@localhost (RSA)

localhost login:
```

- (Optional) To ensure that the connection to the vault is active, ping the vault IP address. If there are network connection issues, change the network settings.
- Enter the following command:

```
./bmragent
```
- On the Vault Login screen, enter information for connecting to the vault where the BMR backup is saved.

In the Address field, enter the vault IP address or fully qualified domain name (FQDN). In the Port field, enter the port number for connecting to the vault (2546, by default). In the Account, Username and Password fields, enter an account and credentials used for the backup.

```
Vault Login:
=====
Address      :
Port Number  : 2546
Account      :
Username     :
Password     :

'ESC' : quit

-----
Enter information about the vault where the BMR backup is saved.
Address: IP address or FQDN
Default port: 2546

=====
```

9. On the Protected Servers screen, press the up and down arrow keys to select the protected server to restore, and then press Enter.
10. On the Job List screen, press the up and down arrow keys to select the BMR job to restore, and then press Enter.
11. On the Safeset List screen, press the up and down arrow keys to select the backup to restore, and then press Enter.
12. On the Start Restore screen, enter the encryption password for the BMR backup job.
On the CONFIRM line, press the right arrow key to choose yes, and then press Enter.
13. If the destination system is larger than the protected system, a *Confirm the recreated disk layout or go back one step* prompt appears. Press Enter to select the default option.
If an error occurs at this stage, please go to <http://relax-and-recover.org/support/> for Relax-and-Recover support.
14. If the destination system has a larger disk than the protected system, a *Confirm restored config files are OK or adapt them as needed* prompt appears. Press Enter, and then press Enter again to select the default options.

The system restore begins. The restore time depends on the size of the backup.

If the restore takes a long time, the screen might go blank. To refresh it, press Enter.

When the restore is finished, a *Completed the bare metal restore* message appears, followed by the RESCUE prompt.

```
Finished recovering your system. You can explore it under '/mnt/local'.
Exiting rear recover (PID 772) and its descendant processes ...
Running exit tasks
Completed the bare metal restore.
RESCUE localhost:/opt/BUAgent # _
```

15. Run the following command to view the restore log:
`./xlogcat jobName/RSTyyyyymmdd-hhmmss.XLOG | tail -n 25`
Where *jobName* is the name of the BMR job from which you restored the system, and *yyyyymmdd-hhmmss* is the date and time of the restore.
Review the restore log. Check that the restore completed with no errors and that the restored system size is correct
16. Restart the system.
Depending on the source system platform and configuration, the system could restart once or twice automatically.
17. Log in to the restored system with credentials from the protected system, and verify that the system is working.

8 Restore a Linux system without a BMR backup

You can restore entire Linux systems from BMR backups. See [Restore a Linux system from a BMR backup](#).

If a Linux system is not protected by a BMR backup, you must recover the system using techniques described in this section. This section also describes the minimum resources required to rebuild a file system to its state at the last system backup.

The basic recovery procedure is:

1. Install the minimal operating system, including networking.
2. Install and configure the Agent.
3. Restore the backed up system state, programs, and data using the Agent.
4. Perform M-restore maintenance.
5. Verify the restore.

Before performing a recovery, ensure that your hardware configuration is sufficient for the programs, data, and system state of the protected system.

8.1 Hardware requirements

It is crucial for local storage on the system to be sufficient for a full restore of programs, system state, and data. Otherwise, the restore will fail, and your system may be left in an indeterminate state.

If any configuration files for your operating system depend on specific identifiers of installed hardware (such as the MAC address of a network card), ensure that this information is noted, as the values may be different than when the system was backed up using the Agent.

Note: When performing a complete system restore (DR), you need to ensure there is ample disk space for the creation of large recovery logs from our Agent and other possible logging or auditing from the operating system. Using file level logging on a system containing a large file system can generate a large log, which can potentially fill up the available or allocated disk space. If the logs are on the same partition as the root file system, this may prevent the OS from booting.

8.2 Software requirements

Ensure that the appropriate installation media is available. The minimum system software includes:

- Installation media identical to that installed on the original system.
- Any necessary OS patches to install the Agent, as described in the installation instructions for the Agent on the OS.
- Agent Installation media identical to that installed on the original system.

8.3 Recovery steps

This section describes the steps to perform a system recovery.

Install the minimal operating system

Follow the instructions in your operating system manual and installation media to install a minimal operating system.

- When prompted to partition your drive(s), ensure that the partitions are large enough to restore to; they should be at least as large as the original partitions.
- If restoring over the network, TCP/IP network services must be installed and configured appropriately, and there must be a connection between the system and the backup vault.
- If restoring from a directory on disk, there must be sufficient disk space to handle all the restored data.

Install and configure the Agent

1. Install the Agent for your operating system.
2. Configure the Agent. Re-register the Agent to the vault where the data was backed up.
3. Synchronize the job to ensure that local copies of job catalogs are created.

Restore the backed up system

1. Start a restore.
2. Select the files you wish to restore. The Agent will restore most files to their original locations and protect against many known restore problems (for file systems mounted in their default locations), but some files may cause unpredictable results if restored. These files vary and may generally be restored to alternative locations without problems.
3. Ensure that the files are not being restored to a file system that is mounted read-only.

Note: The Agent will prevent recovery of files to critical locations, but not all critical locations are necessarily detected.

When the recovery procedure is complete, the process of verifying the integrity of the restore can commence.

Perform post-recovery maintenance

If any modifications to the configuration of the restored system are required after restore, these should be performed now. Known post-restore maintenance steps are noted below.

Verify the recovery

Once the restore procedure is complete, determine if the recovery is complete and correct. The listing and testing of the jobs should be performed as part of the systems recovery planning. The specific jobs to be performed for verification depend on the application environment and the system's importance.

Once the system is restored, the integrity of the recovery must be verified. The test can be as simple as placing a duplicate file in a different directory structure and testing for any differences within the file. Then, confirm that the file can be opened using a known application and that you are able to send e-mail to a known address. It can also be as complex as completing an SQL query on a known database set.

Whatever the test, both the list and the test itself must be planned and executed during normal system operation.

8.4 Recovery problems

Should any of the recovery jobs fail, consider these questions:

- Was the system restored using the same version of OS?
- What possible differences were there in the hardware or software settings that could have affected the recovery?
- Were any errors reported in the error log file?
- Were all the necessary drivers installed?
- Were the applicable OS patches added?
- Was there sufficient disk space to handle all of the restored data?

9 Back up and restore Oracle databases using the Oracle Plug-in

The Oracle Plug-in is an add-on to the Linux Agent that allows you to perform database backups on Oracle databases.

The Plug-in is installed with the Agent on the database host.

A user, typically a DBA, configures the backup using Portal or the legacy Windows CentralControl. A user can schedule a backup of the database, at which time the Agent (with the help of the Oracle Plug-in) will send database information to the Director vault.

The Oracle Plug-in provides ARCHIVELOG-based, non-RMAN backups of whole online database instances. All non-temporary tablespaces and instance parameter files are automatically backed up.

Full and partial databases are restored through normal user-managed Oracle recovery mechanisms.

Agents specify databases using Oracle Service Names. They do not require script-level or backup-level ORACLE_HOME customization.

Database passwords are encrypted for enhanced security over script-based methods.

Limitations

- Only local, single-instance, disk-based databases are backed up.
- Database clusters are not backed up.
- Raw devices are not backed up.
- Remote databases are not backed up.
- The database must run in ARCHIVELOG mode, and the user under which the backup is configured must have SYSDBA privileges.

9.1 Install the Oracle Plug-in for Linux

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

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

```
SELECT banner
   FROM v$version

SELECT version
   FROM v$instance
```

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

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

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

To install the Oracle Plug-in for Linux:

1. Download the Oracle Plug-in for Linux tar.gz installation package on the machine where you are installing the Plug-in.

2. Run the following command to extract files from the installation package:

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

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

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

```
cd packageName
```

4. Run the following command to start the installation:

```
./install.sh
```

5. Follow the installation instructions on the screen.

9.2 Add an Oracle database backup job

When the Linux Agent and Oracle Plug-in are installed on a computer, you can create a backup job for one or more Oracle databases. The backup job specifies which databases to back up, and where to save the backup data. You must also specify credentials for the Agent to use to connect to the Oracle server.

The Oracle Plug-in performs what Oracle Corporation deems an “inconsistent” whole database backup, requiring that the database be run in ARCHIVELOG mode. During a live backup, any changes to the database will be written to archived logs. The database administrator should ensure that the database is in ARCHIVELOG mode.

The Oracle Plug-in backs up redo and archive logs that are created while the database backup job is running. For example, if an Oracle database backup job runs from 22:00 to 01:00 each day, the plug-in backs up redo and archive logs that are created between 22:00 and 01:00. To back up logs that are created after the Oracle database backup job runs, we recommend running a Local System job at another time each day. Using the Local System, you will be able to recover the database to a point in time that is later than the time when the Oracle database backup job ran.

To ensure that archived log files do not take up too much disk space on your system, the Oracle Plug-in can delete archived redo logs after a successful backup. This functionality is available with the Oracle Plug-in for Linux Agent version 8.60 or later. If you specify that archived logs should be deleted after a backup, ensure that the logs are backed up using a Local System job.

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

To add an Oracle database backup job:

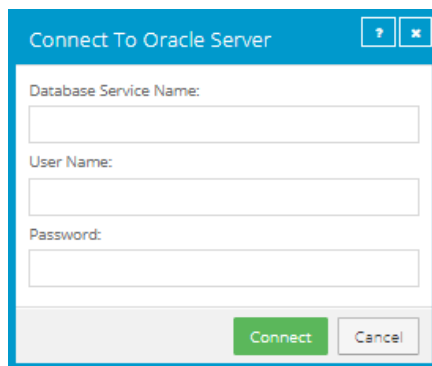
1. On the navigation bar, click **Computers**.

The Computers page shows registered computers.

2. Find a computer with the Oracle Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the Jobs tab.

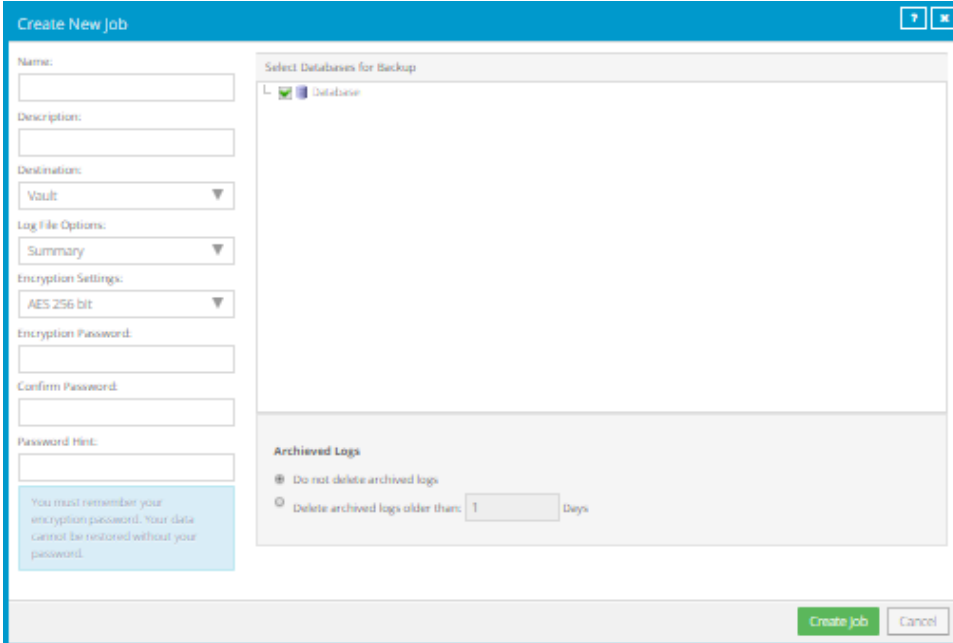
4. In the **Select Job Task** menu, click **Create New Oracle Job**.
5. In the Connect to Oracle Server dialog box, specify the following information:
 - In the **Database Service Name** box, type the service name of the database that you want to back up.
 - In the **User Name** box, type the name of a user who has sysdba privileges.
 - In the **Password** box, type the password for the specified user.



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

A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
 - In the **Log File Options** list, select the level of detail for job logging. For more information, see [Log file options](#).
 - For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
 - In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also

enter a password hint in the **Password Hint** box.



8. In the **Select Databases for Backup** box, select the database to back up.
9. Do one of the following:
 - To leave Oracle archived redo logs on the system, click **Do not delete archived logs**.
 - To delete Oracle archived redo logs after a successful backup, click **Delete archived logs older than [...] days**. Enter the number of days after which archived logs can be deleted.
10. Click **Save**.

The job is created, and the View/Add Schedule dialog box appears. Now you can create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups and synchronizations](#).

9.2.1 About Oracle backups

The Oracle Plug-in for the Linux Agent performs what Oracle Corporation deems an “inconsistent” whole database backup, requiring the database to run in ARCHIVELOG mode. During a live backup, any changes to the database will be written to archive logs. The DBA should ensure that the database is in ARCHIVELOG mode:

```
SELECT log_mode
FROM v$database
```

The value ARCHIVELOG should return. Otherwise, follow the normal Oracle procedure for putting the database in ARCHIVELOG mode. This is typically:

```
> shutdown normal
> startup mount
> alter database archivelog;
> archive log start
> alter database open
```

In Oracle, this is done directly from SQL*Plus. You can also put the database in ARCHIVELOG mode when you initially set it up. Alternatively, you can use the Enterprise Manager GUI or other DBA tools.

No tablespaces can be in backup mode before a backup job starts. You can verify this with:

```
SELECT d.file_name, b.status
FROM dba_data_files d, v$backup b
WHERE b.file# = d.file_id;
```

If any files display with ACTIVE status, the backup job will not start.

Note: The Agent leaves the database in an appropriate state when a backup completes successfully.

Before you can use the Oracle Plug-in to create backup jobs, a license must be available on the vault. See the vault operations manual for more information.

9.2.2 How the backup works

When a backup starts, the Oracle Plug-in for the Linux Agent iterates through all non-TEMPORARY tablespaces (including ONLINE, OFFLINE, and READONLY tablespaces). Each ONLINE tablespace will enter ARCHIVELOG mode (which creates a snapshot of the tablespace's files). The tablespace's component files will be backed up. When the backup of an ONLINE tablespace's files finishes, the tablespace will return to normal mode.

After all of the tablespaces have been backed up, the Plug-in flushes any pending redo logs, and also backs up the generated archive logs. These logs will always be new files.

The instance control files are backed up as binary files, as well as TRACE log entries. The instance parameter files (**init**<ORACLE_SID>.ora and/or **spfile**<ORACLE_SID>.ora, depending on the version and configuration of Oracle) and the Oracle password file are also backed up.

Note: OS and Oracle Configuration files that are not instance-specific (such as **kernel parameters**, **tnsnames.ora**, **sqlnet.ora** and **listener.ora**) are not backed up by the Plug-in. You can back these up using an ordinary file-based Agent.

9.2.3 Table of backup information

Before you perform Oracle database backup or restore processes on a Linux server, be sure that you have all information such as names, locations, passwords, etc., that the wizard will request. You can use the following table for reference.

System Requirement	Customer/User Supplied Value	Comments
New Job Name	Job Name =	Name of job to communicate with an Agent that has the Oracle Plug-in
Backup Source Type	Oracle	Choose Oracle from the dropdown menu
Oracle Options (database to back up, and database account information)	Database Service Name * = User Name = Password =	Validates the fields, and allows connection to the database. In Portal, set the Database Service Name to the Database Instance from Oracle (rather than the Instance Name from Oracle). In Windows CentralControl, set the Oracle Service Name to the Database Instance from Oracle (rather than the Instance Name from Oracle).
Encryption type	Encryption type = Password = Password Hint =	If you select a type, you must supply a password
Logging options	Create log file = Y/N Log detail level = Keep or purge log files = Number of logs to keep =	
Schedule		You can run backup jobs immediately, or through a schedule. You can optionally use the scheduling wizard.
Destination vault	Vault Name = Network Address =	Choose from the dropdown list of Director vaults

* If you connect to a database that listens on a port other than the default, the format for the Database Service Name is **service name:port number** (for example, **orcl:1523**).

9.3 Restore Oracle databases

You might need to restore a full database, or restore a system from the ground up (“bare metal”): installing the OS, applications, and then the full database (plus any transaction logs) on a new system.

If there is an Oracle backup and a full-system backup, restore the system (putting back the contents of ORACLE_HOME – specifically the database installation). You may safely exclude the data files and archive logs that are backed up by the Plug-in.

Finally, restore the Oracle backup, and copy the required components to the appropriate directories. Follow the standard user-managed Oracle recovery procedure outlined in the appropriate OS Oracle Backup and Recovery Guide (available on the Oracle website).

An Oracle restore process is performed by a Database Administrator. Briefly, the steps are:

1. Shut down the database.
2. Restore the files using **Restore to an Alternate Location**.
3. If the files have been renamed, you must change them back to their original file names (i.e., control files).
4. If necessary, reset the control information for the database.
5. Start and recover the database.
6. Re-open the database for use.

The Plug-in does not do table-level restores.

9.3.1 Guidelines for restoring

Note: This section provides general database recovery guidelines. For detailed recovery procedures, please refer to documentation from Oracle.

Note: For a full disaster recovery (in which the full database instance is restored), be careful when you recover the database because the Oracle Plug-in for Linux does not back up TEMPORARY tablespaces.

Start the database recovery with an explicit PFILE or SPFILE reference:

```
SQL> STARTUP mount PFILE='path-to-pfile\initSIDNAME.ora'
```

It may be necessary to take the temporary tablespace files offline:

```
SQL> ALTER DATABASE DATAFILE 'path-to-datafile' OFFLINE
```

Restore the database as usual, but when you open it after recovery, use this command:

```
SQL> ALTER DATABASE OPEN NORESETLOGS
```

TEMPORARY tablespaces should be dropped, the data files for the temporary tablespaces should be removed, and the TEMPORARY tablespaces should be recreated (this may include the default TEMP tablespace).

At this point, the database can be closed normally and restarted (with RESETLOGS, for example).

Note: Oracle parameter files are backed up to a different directory by default.

9.4 Uninstall the Oracle Plug-in for Linux

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

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

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

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

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

10 Delete jobs and computers, and delete data from vaults

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

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

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

When deleting job data from vaults, there is a 72-hour waiting period before the data deletion request is sent to vaults. 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.

10.1 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**.

10.2 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.



10.3 Cancel a scheduled job data deletion

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

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


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

To cancel a scheduled job data deletion:

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

The Computers page shows registered computers.

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



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

5. Click **Yes**.

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

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



10.4 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.

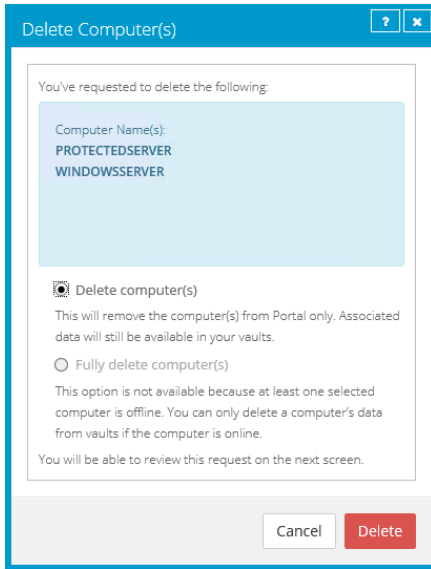
If a computer is deleted from Portal in this way, the data can still be restored using the *Restore from Another Computer* procedure.

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:

1. When signed in as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Select the check box for each computer that you want to delete.
3. In the **Actions** list, click **Delete Selected Computer(s)**.
4. If the data deletion feature is enabled in your Portal instance, a Delete Computer(s) dialog box appears.

To delete the computer without deleting data from vaults, click **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**.

10.5 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.



10.6 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 Selected 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.

10.7 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**.

11 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](#).
- **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](#).

11.1 Monitor backups and computers using the Current Snapshot

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

To monitor backups and computers using the Current Snapshot:

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

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


- **Backups Requiring Attention** — Number of backup jobs where the last backup attempt failed, completed with errors, did not back up any files, reached a license limit, was cancelled or had a potential ransomware threat.
- **Missed Backups** — Number of backup jobs that have not run for seven days.
- **Backups With Warnings** — Number of backup jobs where the last backup attempt completed with warnings, was deferred, was deferred with warnings or was skipped. This category also includes backup jobs that have never run.
- **Computers Requiring Reboot** — Number of computers with a pending reboot.

- **Offline Computers** — Number of computers that are not currently in contact with Portal. A computer can be offline if it is turned off, if the Agent has been uninstalled from the system, or if the system no longer exists.
 - **Computers Scheduled for Deletion** — Number of computers that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.
 - **Computers With Certificate Failures** — Number of computers reporting a certificate failure. See [Resolve certificate failures](#).
 - **Total Computers** — Total number of computers in the site.
 - **Successful Backups** — Number of backup jobs where the last backup attempt completed without errors, warnings, or deferrals.
 - **Jobs Scheduled for Deletion** — Number of jobs that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.
2. To view computers in a particular site, click the sites box in the top right of the Current Snapshot box. In the menu, click the site that you want to view.
Computers in the selected site appear on the Computers page.
 3. To view information about backup jobs or computers in one of the categories, click the category.
If you click **Potential Threats, Backups Requiring Attention, Missed Backups, Backups With Warnings** or **Successful Backups**, backup jobs in the category appear on the Monitor page.
If you click **Computers Requiring Reboot, Offline Computers, Computers Scheduled For Deletion, Computers With Certificate Failures** or **Total Computers**, computers in the category appear on the Computers page.

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

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

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

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

-  Skipped — Indicates that a backup was skipped. Backups are sometimes skipped if they are scheduled to run multiple times per day. See [Skipped backups](#).
-  Never Run — Indicates that the backup job has never run.
-  Missed — Indicates that the job has not run for 7 days.
-  Completed with errors — Indicates that the backup completed and a safeset is available for restore, but problems occurred. Typically, this status indicates that not all of the data was backed up.
-  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](#).

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

11.3 View skipped rates and backup status histories

Beginning with Linux Agent 8.90, when a Linux agent is backing up data to a Director version 8.60 or later vault, backups that are scheduled to run multiple times per day are skipped in some cases. To determine whether backups were skipped, users can view email notifications, the Computers page and Monitor page, and the Daily Status report. See [Skipped backups](#).

- Skipped rate for a job. If a backup was skipped for a job in the 48 hours before the most recent backup attempt, a skipped rate appears for the job on the Computers page and Monitor page. The skipped rate is the percentage of backups that were skipped in the 48 hours before the last backup attempt, and is calculated using the following formula:

$$\text{jobSkippedRate} = \text{numberOfSkippedBackups} / \text{numberOfBackupAttempts}$$

Where:

- *numberOfSkippedBackups* is the number of backups that were skipped for the job during the 48 hours before the last backup attempt.
- *numberOfBackupAttempts* is the total number of backup attempts for the job during the 48 hour period, including skipped, in-progress, deferred, canceled, failed and completed backups.

If no backups were skipped for a job in the 48 hours before the last backup attempt, or if the last backup attempt occurred more than seven days ago, a skipped rate is not shown for the job.

- Skipped rate for a computer. If a skipped rate is reported for one or more jobs on a computer, the highest skipped rate on the computer appears on the Computers page.
- 48-hour backup status history for a job. If a skipped rate appears for a job on the Computers or Monitor page, you can view the job's backup history for the 48 hours before the last backup attempt. The status history shows the dates and times of backup attempts, and indicates the status of each backup attempt (e.g., skipped, in-progress, completed or failed). You can export the status history in comma-separated values (.csv), Microsoft Excel (.xls) or Adobe Acrobat (.pdf) format.

To view skipped rates and backup status histories, see [View skipped rates and backup status histories on the Computers page](#) and [View skipped rates and backup status histories on the Monitor page](#).

11.3.1 View skipped rates and backup status histories on the Computers page

To prevent schedule overloads, backups that are scheduled to run multiple times per day are skipped in some cases. Users can obtain skipped backup information through email notifications, on the Computers page, and in the Daily Status report. See [Skipped backups](#).

In some Portal instances, users can view skipped backup rates for jobs and computers on the Computers page, and view and export a job's backup status history for the 48 hours before the last backup attempt. For more information, see [View skipped rates and backup status histories](#).

To view skipped rates and backup status histories on the Computers page:

1. Click **Computers** on the navigation bar.

A value appears in the Skipped column for any computer where at least one job has a skipped rate. If more than one job on a computer has a skipped rate, the highest skipped rate appears in the Skipped column.

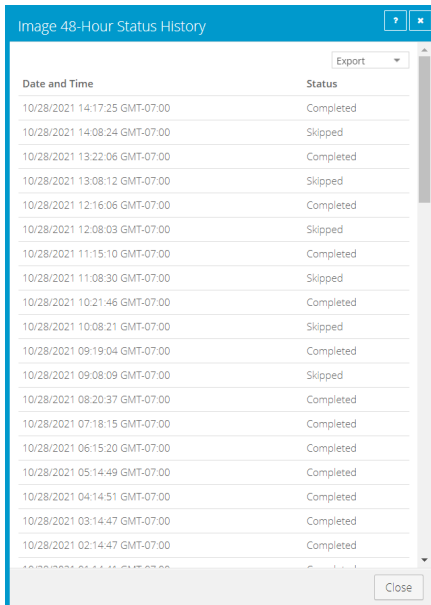
Note: If the Skipped column does not appear, skipped rates and 48-hour backup status histories are not available in your Portal instance.

2. Find a computer with a value in the Skipped column, and click the computer row to expand its view.

On the Jobs tab, a value appears in the Skipped Rate column for any job where a backup was skipped in the 48 hours before the last backup attempt, and the last backup attempt occurred in the last seven days.


3. To see which backups were skipped in the 48 hours before the last backup attempt for a job, click the job's Skipped Rate value.

The 48-Hour Status History for the job shows the date, time and status (e.g., skipped, in-progress, completed or failed) of each backup attempt.



If you want to export the status history, click the **Export** box. In the list that appears, click one of the following formats for the exported data:

- CSV (comma-separated values)
- XLS (Microsoft Excel)
- PDF (Adobe Acrobat)



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

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

11.3.2 View skipped rates and backup status histories on the Monitor page


To prevent schedule overloads, backups that are scheduled to run multiple times per day are skipped in some cases. Users can obtain skipped backup information through email notifications, on the Computers page, and in the Daily Status report. See [Skipped backups](#).

In some Portal instances, users can view skipped backup rates for jobs on the Monitor page, and view and export a job's backup status history for the 48 hours before the last backup attempt. For more information, see [View skipped rates and backup status histories](#).

To view skipped rates and backup status histories on the Monitor page:

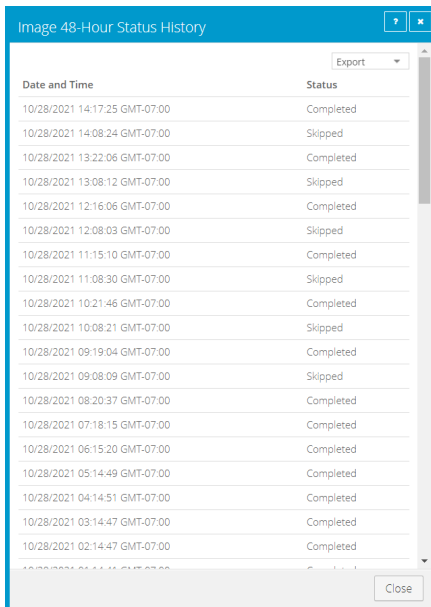
1. Click **Monitor** on the navigation bar.

A value appears in the Skipped column for any job where a backup was skipped in the 48 hours before the last backup attempt, and the last backup attempt occurred in the last seven days.




2. To see which backups were skipped in the 48 hours before the last backup attempt for a job, click the job's Skipped value.

The 48-Hour Status History for the job shows the date, time and status (e.g., skipped, in-progress, completed or failed) of each backup attempt.



If you want to export the status history, click the **Export** box. In the list that appears, click one of the following formats for the exported data:

- CSV (comma-separated values)
- XLS (Microsoft Excel)
- PDF (Adobe Acrobat)



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

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

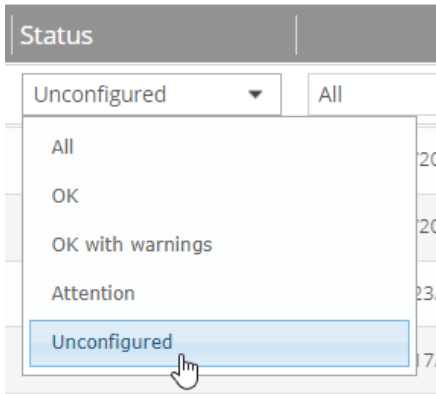
11.4 View an unconfigured computer's logs

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

To view an unconfigured computer's logs:

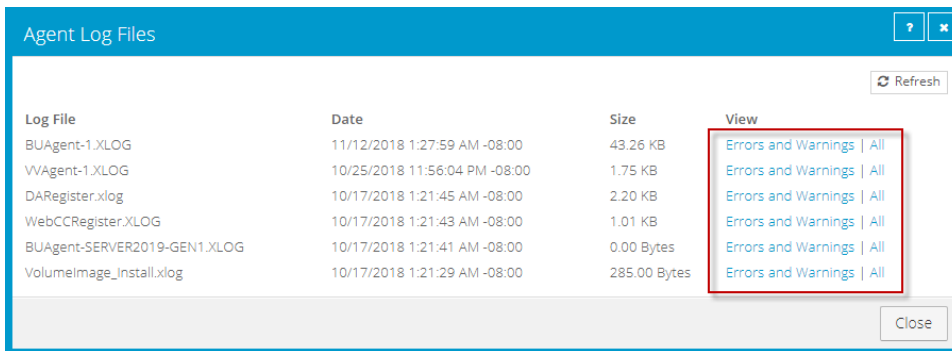
1. On the navigation bar, click **Computers**.

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



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

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




4. Do one of the following:
 - To only view errors and warnings in a log, click **Errors and Warnings** for the log.
 - To view an entire log, click **All** for the log.


The log appears in a new browser tab.

11.5 View current process information for a job


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

To view current process information for a job:


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



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




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.



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

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

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**.



11.6 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 Linux systems with Agent version 8.10a 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](#).

11.6.1 Set up email notifications for backups on a computer

To set up email notifications for a computer:

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

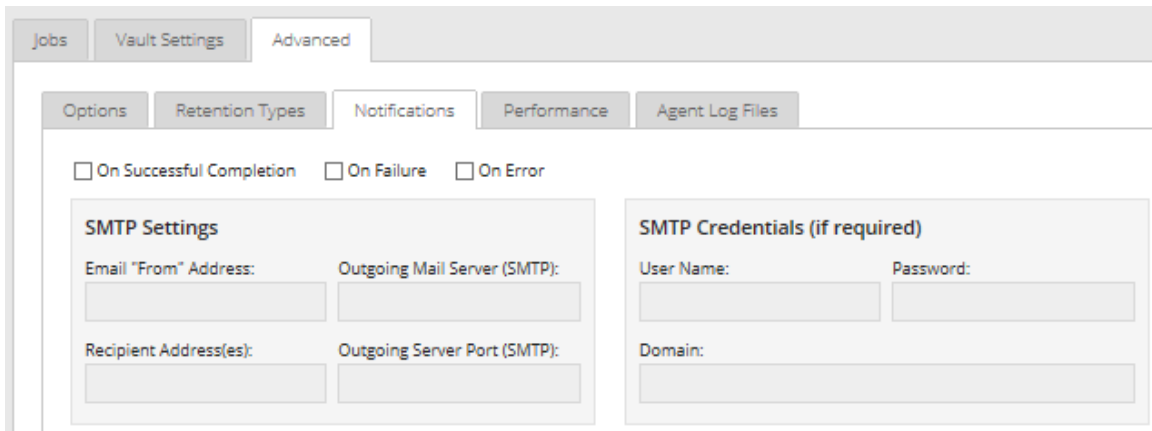
The Computers page shows registered computers.

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**.

11.6.2 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. See [Skipped backups](#).

4. Click **Update notifications**.

11.6.3 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**.

11.7 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.

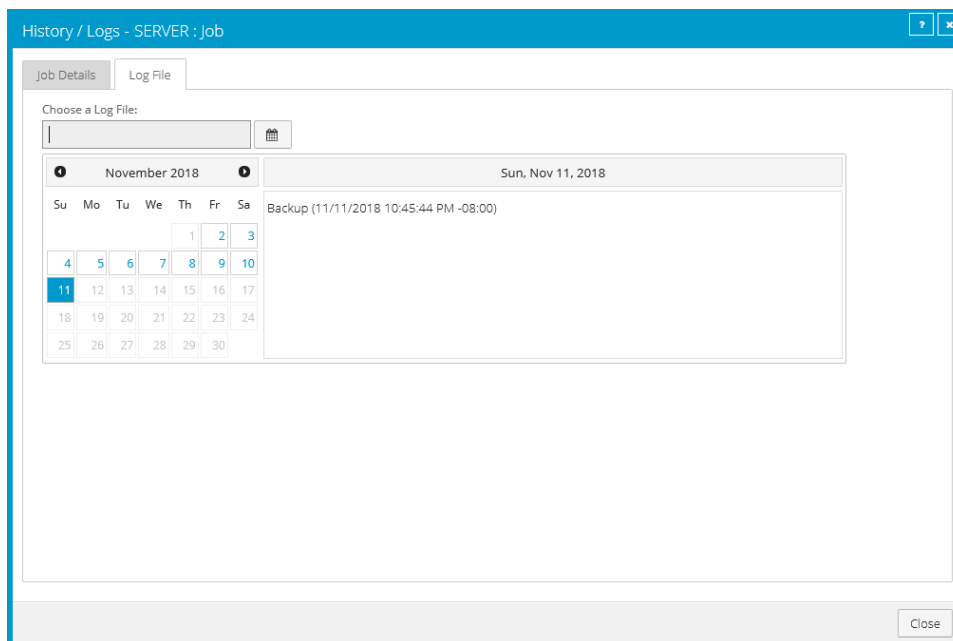
2. Find the computer for which you want to view logs, and click the row to expand its view.

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:

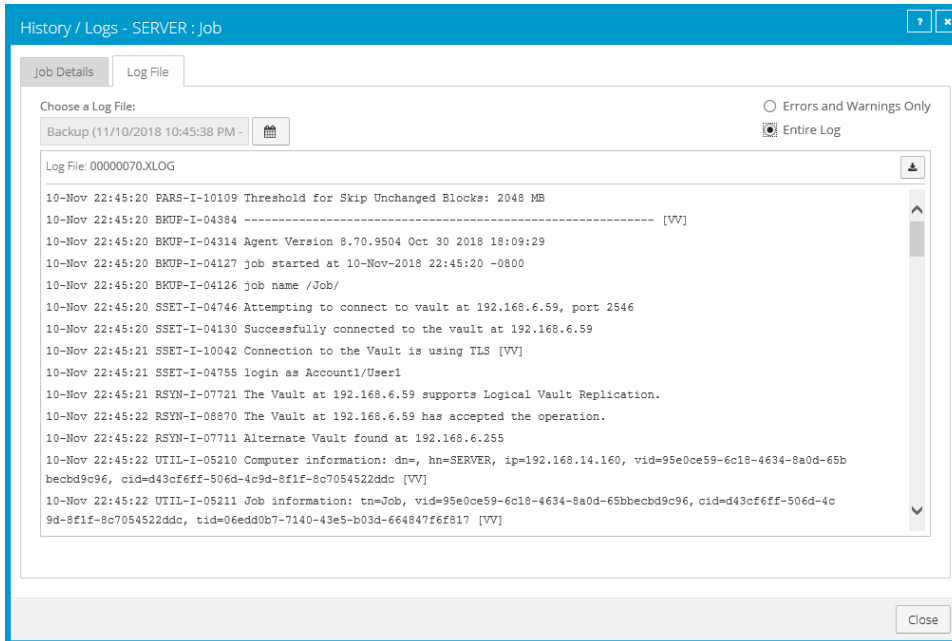
- In the job's **Select Action** menu, click **History / Logs**.
- In the **Last Backup Status** column, click the job status.

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




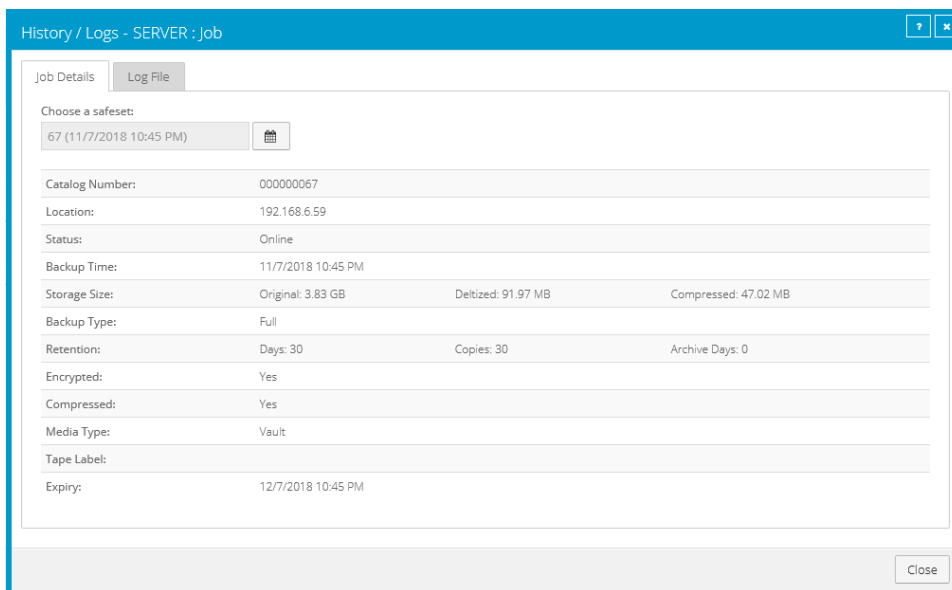
4. To view processes for a different day, click the calendar button. 📅 In the calendar that appears, click the date of the log that you want to view.
5. In the list of processes on the selected date, click the process for which you want to view the log.

The window shows the selected log.



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

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



11.8 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.

12 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](#)

12.1 Contacting Carbonite

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

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



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

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.