



Carbonite Server Backup

Linux Agent and Oracle Plug-in 8.6

User Guide



© 2018 Carbonite, Inc. All rights reserved.

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

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

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

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

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

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

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

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

Document History

Version	Date	Description
1	March 2018	Initial guide provided for Linux Agent 8.6x.

Contents


1	Introduction to the Linux Agent	5
2	Install the Linux Agent	6
2.1	Install the Linux Agent in silent mode	8
2.2	Upgrade the Linux Agent.....	10
2.3	Change the Portal registration for a Linux Agent.....	11
2.4	Change the language for Linux Agent messages.....	11
2.5	Uninstall the Linux Agent	11
3	Configure the Linux Agent.....	13
3.1	Add vault settings.....	13
3.2	Add a description	15
3.3	Add retention types.....	15
3.4	Configure bandwidth throttling	17
4	Add a Linux backup job.....	19
4.1	Add the first backup job for a Linux computer	21
4.2	Add an NFS backup job.....	22
5	Run and schedule backups and synchronizations.....	28
5.1	Schedule a backup.....	28
5.2	Specify whether scheduled backups retry after a failure	31
5.3	Run an ad-hoc backup.....	32
5.4	Synchronize a job	33
6	Restore Linux files and folders	34
6.2	Restore ACLs.....	38
6.3	Restore data to a replacement computer.....	39
6.4	Restore data from another computer.....	41
6.5	Advanced restore options	42
7	System recovery.....	45
7.1	Hardware requirements.....	45
7.2	Software requirements	45

- 6.3 Recovery steps.....45
- 8 Back up and restore Oracle databases using the Oracle Plug-in..... 48**
 - 8.1 Install the Oracle Plug-in for Linux48
 - 8.2 Add an Oracle database backup job.....49
 - 8.3 Restore Oracle databases.....53
 - 8.4 Uninstall the Oracle Plug-in.....54
- 9 Monitor computers, jobs and processes..... 55**
 - 9.1 View computer and job status information55
 - 9.2 View an unconfigured computer’s logs.....56
 - 9.3 View current process information for a job58
 - 9.4 Monitor backups using email notifications60
 - 9.5 View a job’s process logs and safeset information62
 - 9.6 View and export recent backup statuses65
- 10 Carbonite Server Backup Support 66**
 - 10.1 Contacting Carbonite.....66

1 Introduction to the Linux Agent

Carbonite Server Backup Agent for Linux 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 Carbonite Server Backup 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 the Linux system.
- System files required for recovering the operating system, including registry and boot files.

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

The Linux Agent is available as a 64-bit application and a 32-bit application. The 64-bit version of the Linux Agent can only be installed on 64-bit systems. The 32-bit version of the Linux Agent can only be installed on 32-bit systems.

To run the installation kit, you must have root privileges.

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.

- A target system running a supported version of Linux. For supported platforms, see the Linux Agent release notes.
- Root privileges on the target system are required to install the Agent.
- Sufficient disk space for the new installation, and later job activities. The installation program will determine whether there is enough disk space for the installation to continue. If the available disk space is insufficient, the installation directory will roll back to its original state.

To install the Agent:

1. Download the 64-bit or 32-bit Linux Agent tar.gz installation kit on the machine where you are installing the Agent.

Note: Download the correct installation kit for your system. You can only install the 64-bit Linux Agent on a 64-bit system, and you can only install the 32-bit Linux Agent on a 32-bit system.

2. Extract files from the installation package. To do so, run the following command:

```
> tar -xzf packageName.tar.gz
```

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

The following screenshot shows files that are extracted to the Linux Agent folder.

3. Run the following command to start the installation:

```
> ./install.sh
```

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

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

```

I HAVE READ THE TERMS AND CONDITIONS OF THIS AGREEMENT, I UNDERSTAND THE CONTENT
, AND I AGREE THAT I WILL ABIDE BY THE TERMS SET FORTH HEREIN. I UNDERSTAND THAT
IF I DO NOT AGREE WITH THE FOREGOING, THEN I WILL NOT BE PERMITTED TO USE THE L
ICENSED SOFTWARE.

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] _

```

5. Press **Enter** to accept the installation directory.

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) _

```

6. Enter **Y** to create the BUAgent directory.
7. When prompted to select a language, 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] _

```

8. When prompted to register to Portal, enter **Y**.
9. At the Portal address prompt, enter the Portal address.
10. At the Portal connection port prompt, enter the Portal connection port. The default value is 8086.
11. At the Portal username prompt, enter the Portal username for registering the Agent.
12. At the Portal password prompt, enter the password for the Portal user specified in Step [11](#).

```
Do you wish to register to the Portal? (Y|n) y
What is the Portal address? ("-" to cancel) portalprod.corp.ssv.com
What is the Portal connection port? [8086] ("-" to cancel)
What is your Portal username? ("-" to cancel) test@test.com
What is your Portal password? ("-" to cancel)
Registering with portalprod.corp.ssv.com:8086 using login of test@test.com
```

The next step in the installation is to choose the encryption method. 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 with the Agent. Using the external encryption library can degrade Agent performance.

```
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
```

13. 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**.

The installation proceeds. When complete, a message appears and the Agent will be running.

After installation, the installation log (Install.log) is located in the installation directory.

```
Starting Agent: /etc/rc.d/init.d/ovagent start quiet

Installation complete. Agent started successfully.

Install finished at 10:41:04 2017.11.23

The installation has been recorded in /opt/BUAgent/Install.log.
copy EULA to installation directory.
```

2.1 Install 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
-----------	-------------

Parameter	Description
<code>-shutdown -s</code>	Force the Agent to shutdown, if running.
<code>-force -F</code>	Force the installation; skip the initial free space check.
<code>-defaults -D</code>	Use the default values for installation.
<code>-force-defaults</code>	Force the installation using the defaults (assumes <code>-s</code> and <code>-F</code>).
<code>-web-registration=off</code> <code>-W-</code>	Turns off Portal registration.
<code>-web-registration=file</code> <code>-W=file</code>	Attempts to register to Portal with the values found in the FILE. See Linux Agent registration options .
<code>-quiet -Q</code>	Quiet install; does not echo output to the screen. If user interaction is required in quiet mode, the install will fail unless <code>-force-defaults</code> is specified.
<code>-log=NAME -L=NAME</code>	Writes the installation log to the specified file NAME.
<code>-lang=NAME -l=NAME</code>	Selects NAME 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., <code>fr</code> , <code>fr-FR</code> , and <code>fr_FR</code> 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 <code>en-US</code> [English (US)]. If not specified, the language will be guessed from your system value of <code>"en_US.UTF-8"</code> .
<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>-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.

2.2 Upgrade the Linux Agent

Before you upgrade the Agent, ensure that your system meets the minimum 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.

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 64-bit or 32-bit Linux Agent tar.gz installation kit on the machine where you are installing the Agent).

Note: Download the correct installation kit for your system. You can only install the 64-bit Linux Agent on a 64-bit system, and you can only install the 32-bit Linux Agent on a 32-bit system.

2. Extract files from the installation package by running the following command:

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

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

3. Run the following command to start the upgrade:

```
> ./install.sh
```

4. If a message states that VVAgent is running, enter Y to stop the Agent.
5. At the Installation directory prompt, enter the installation directory.

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.

6. At the Select language prompt, enter the language for Agent messages. The default language is English [en-US].
7. If a message states that you are already registered to a Web-based Agent Console server, 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.
8. 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**.

The upgrade proceeds. When complete, a message appears, and the Agent daemon will be running.

2.3 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 Web-based Agent Console address prompt, enter the Portal address.
5. At the Web-based Agent Console connection port prompt, enter the Portal connection port. The default value is 8086.
6. At the Web-based Agent Console username prompt, enter the Portal username for registering the Agent.

The Agent is restarted and the Portal registration is changed.

2.4 Change the language for Linux Agent messages

When you install a Linux Agent, you can specify a language for Agent messages. You can also change the Agent language at any time.

To change the language for Linux Agent messages:

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

```
> ./set_language
```
2. At the Select language prompt, enter one of the following values:
 - de-DE (German)
 - en-US (English)
 - es-ES (Spanish)
 - fr-FR (French)

2.5 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 Carbonite Server Backup 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. See [Configure bandwidth throttling](#).
- Email notifications, so that users receive emails when backups complete, fail, or have errors. See [Monitor backups using email notifications](#).

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.

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

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

To add vault settings:

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

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

The Vault Settings dialog box appears.

4. Do one of the following:

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

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

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

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

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

- **Agent Host Name.** Name to use for the computer on the vault.
- **Port Number.** Port used to connect to the vault. The default port is 2546.
- **Attempt to Reconnect Every.** Specifies the number of seconds after which the Agent should try to connect to the vault, if the vault becomes unavailable during a backup or restore.

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


6. Click **Save**.

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



5. Click **Save**.

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.

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.

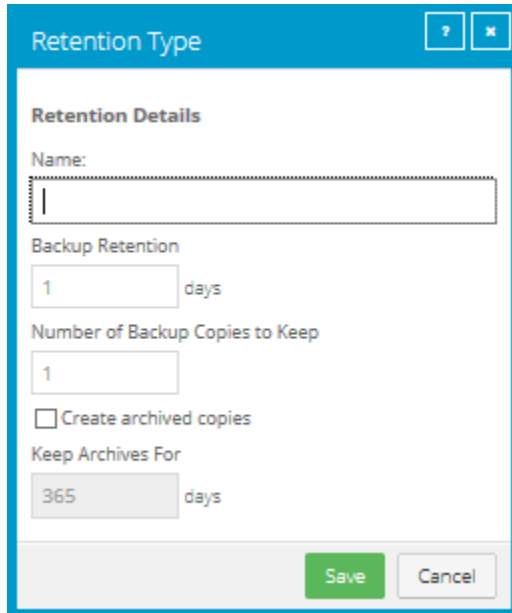
To add a retention type:

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

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

4. Click **Create Retention Type**.

The Retention Type dialog box appears.



5. Complete the following fields:

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

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

6. Click **Save**.

3.4 Configure bandwidth throttling

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

Bandwidth settings include:

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

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

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

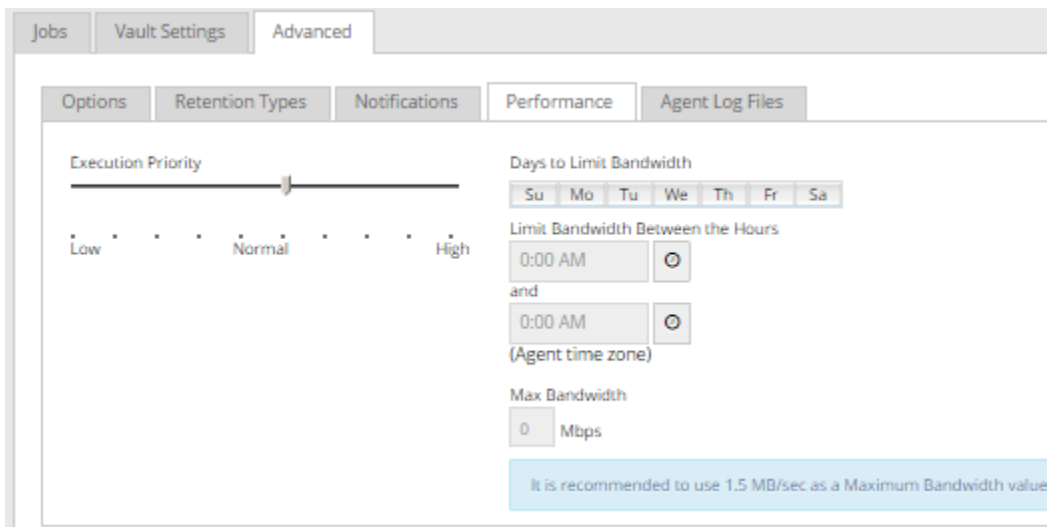
If a policy is assigned to a computer, bandwidth throttling settings cannot be modified on the Computers page. Instead, settings can only be added or modified in the policy.

To configure bandwidth throttling:

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

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

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



4 Add a Linux backup job

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

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 computer](#).

3. Click the **Jobs** tab.

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

4. In the **Select Job Task** menu, click **Create 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 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 computer

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.

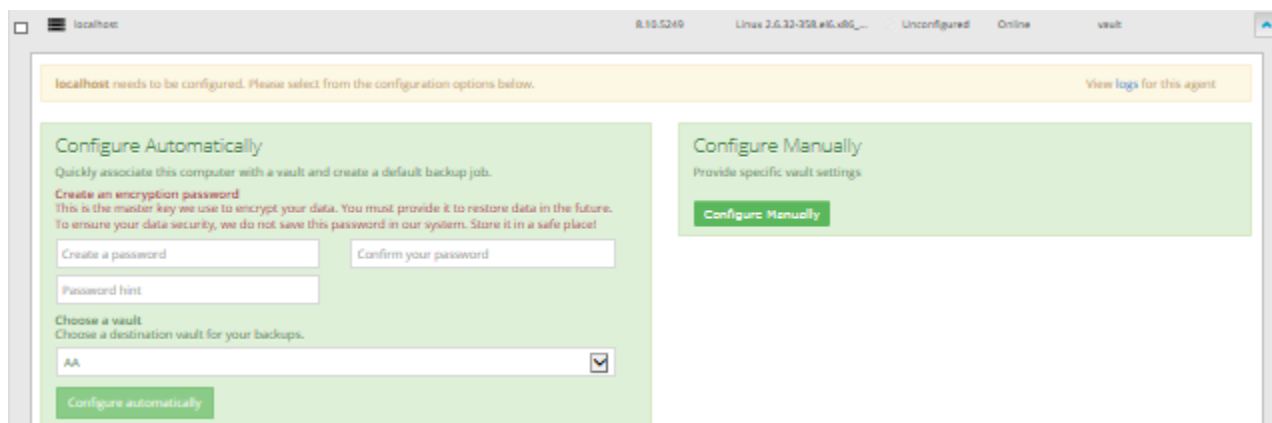
To add the first backup job for a Linux computer:

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:

- i. 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.
- ii. In the **Password hint** box, enter a hint to help you remember the encryption password.
- iii. 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.



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

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

If the automatic job configuration fails, do the following:

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

4.2 Add an NFS backup job

After a 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.

1. Find a 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.

2. Click the **Jobs** tab.

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

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

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


- In the **Name** box, type a name for the backup job.
- In the **Description** box, 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.

5. 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. 

6. 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](#).

4.2.1 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.

4.2.2 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.

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.

4.2.3 Advanced backup options

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

Back up files opened for write

If the **Backup files opened for write** option is selected, files are backed up if they are open for writing or shared reading during the backup. Files that are open for exclusive writes cannot be backed up.

When this option is selected, inconsistencies in the backup can occur if an open file is modified during the backup process.

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 of the 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.

4.2.4 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 .doc or .docx extension.

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

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

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

To filter subdirectories and files in a backup job:

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

	Folders Filter	Files Filter	Recursive	
+	C:	**	<input checked="" type="checkbox"/>	
+	Documents an...	<input type="text" value="e.g., a*, b*"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
+	ProgramData	<input type="text" value="e.g., a*, b*"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

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


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

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

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

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

- 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 .exe or .dll extension, enter the following filter: *.exe, *.dll
- 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 **Apply Now** to consolidate and simplify records in the **Backup Set** box, if changes need to be applied.
6. 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.

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 the remaining data.

- 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 volume of data sent to the vault against the speed of processing. 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 running a backup, you can view logs to check whether the backup completed successfully.

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 automatically.

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

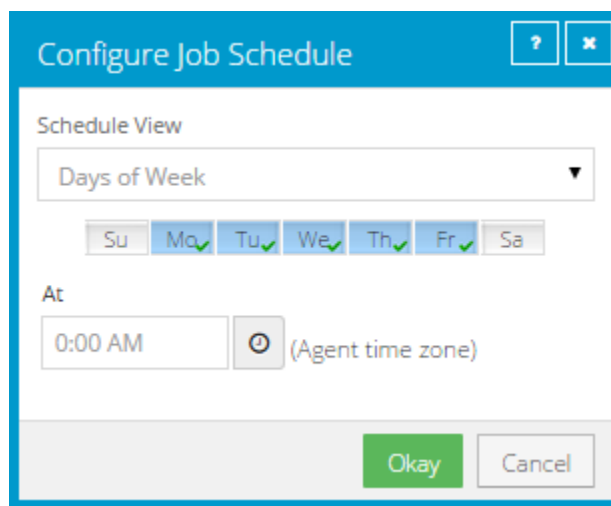
If a job is scheduled to start at exactly the same time by multiple schedules, the job only runs once at the scheduled time. If the jobs have different retention types, the retention type of the schedule that is highest in the list is applied to the resulting safeset. For example, in the following screenshot, the job is scheduled to

run at 12 AM each Saturday with the Weekly retention type, and every day at 12 AM with the Daily retention type. 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 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.

To schedule a backup:

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 computer 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, in the **Retention** list, click a retention type.
4. In the **Schedule** box, click the arrow.
The **Configure Job Schedule** dialog box opens.
5. In the **Configure Job Schedule** dialog box, do one of the following:
 - To run the backup on specific days each week, 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.

Configure Job Schedule

Schedule View

Days of Month

1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	Last			

At

12:00 AM (Agent time zone)

Okay Cancel

- 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.

Configure Job Schedule

Schedule View

Custom

Custom Cycle

Format:
min/hour/day of month/month/day of week
Example:
0/18/*/*/1-5
Means:
Start at 6 pm Monday through Friday
Acceptable Values:
minutes: 0-59
hours: 0-23
days: 1-31
months: 1-12
day of week: 0-6 (0 = Sunday)
Keywords:
** for every time
"Last" for the last day of any month

Okay Cancel

- Click **Okay**.

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

- In the **Compression** list, click a compression level for the backup data. Compression levels optimize the volume of data sent against the speed of processing.

8. 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.
9. To run the job on the specified schedule, select the **Enable** check box near the end of the row.
10. 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](#).
11. Click **Save**.

5.2 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 computer for specifying automatic retry settings, and click the computer 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.

The screenshot shows a 'View / Add Schedule' window for a backup job named 'ACCEL, Test1'. The window has a blue header with a question mark and a close button. Below the header, there's a table with columns: Retention, Schedule, Compression, Deferring (0 for none), Enable, and Priority. The 'Retention' is set to 'Daily', 'Schedule' is '12:00 AM Mo,Tu,We,Th,Fr,Sa,Su', 'Compression' is 'Smaller', 'Deferring' is '0', 'None', 'Enable' is checked, and 'Priority' is set to a default value. Below the table, there's a section titled 'Automatic Retry for Scheduled Backups' which is highlighted with a red box. This section contains a checked checkbox for 'Retry failed backup', a 'Number of retries' field set to '1' with the unit 'times', and a 'Wait before each retry attempt for' field set to '1' with the unit 'minutes'. At the bottom right of the window, there are 'Save' and 'Cancel' buttons.

3. Click **Save**.

5.3 Run an ad-hoc backup

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

To run an ad-hoc backup:

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

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

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

5. 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.

6. 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](#).

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

5.4 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 reregistered computers. See [Restore data to a replacement computer](#).
- Before restoring data from jobs that are backed up to a Satellite vault and replicated to the cloud or another vault.
- To rebuild a delta (.dta) file for a job. If an error message in a log file says that the delta mapping file is corrupt, delete the delta (.dta) file from the job folder on the protected computer and then synchronize the job to rebuild the delta file.

To synchronize a job:

1. On the navigation bar, click **Computers**.
The **Computers** page shows registered computers.
2. Find the computer with the job that you want to synchronize. Expand its view by clicking its row.
3. Click the **Jobs** tab.
4. Find the job that you want to synchronize, and click **Synchronize** in its **Select Action** menu.

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

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

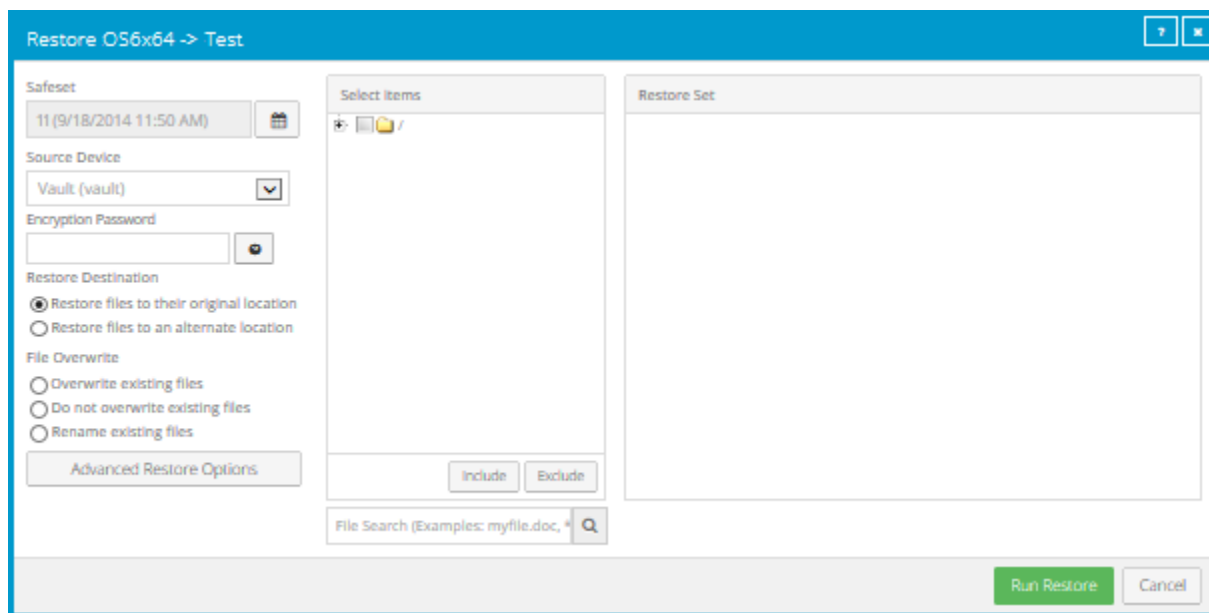
6 Restore Linux files and folders



After backing up data from a Linux computer, you can restore files and folders from the 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 a **File Overwrite** option. This option specifies how to restore a file to a location where there is a file with the same name.
 - To overwrite existing files with restored files, select **Overwrite existing files**.

Note: If you try to restore multiple files with the same name to an alternate location and select **Overwrite existing files**, only the last file restored will remain. Other files with the same name will be overwritten.
 - To add a numeric extension (e.g., .0001) to a *restored* file name, select **Do not overwrite existing files**. 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 an *existing* file name, select **Rename existing files**. 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 continues to be “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 **Restore Set** box shows the excluded folders and files. If you exclude a folder, all of the folder’s subdirectories and files are excluded from the restore by

default. If you do not want to exclude all of the subdirectories and files, you can add filters. See [Filter subdirectories and files when restoring data](#).

- To search for files to restore or exclude from the restore, click the **Search** button.  In the **Search for files** box, enter search criteria and select files. See [Search for files to restore](#). Click **Include Selected** or **Exclude Selected**. The **Restore Set** box shows the included or excluded files.
- To remove an inclusion or exclusion record from the **Restore Set** box, click the Delete button beside the folder or file record. 

Click **Apply Now** to consolidate and simplify records in the **Restore Set** box, if changes need to be applied.

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.

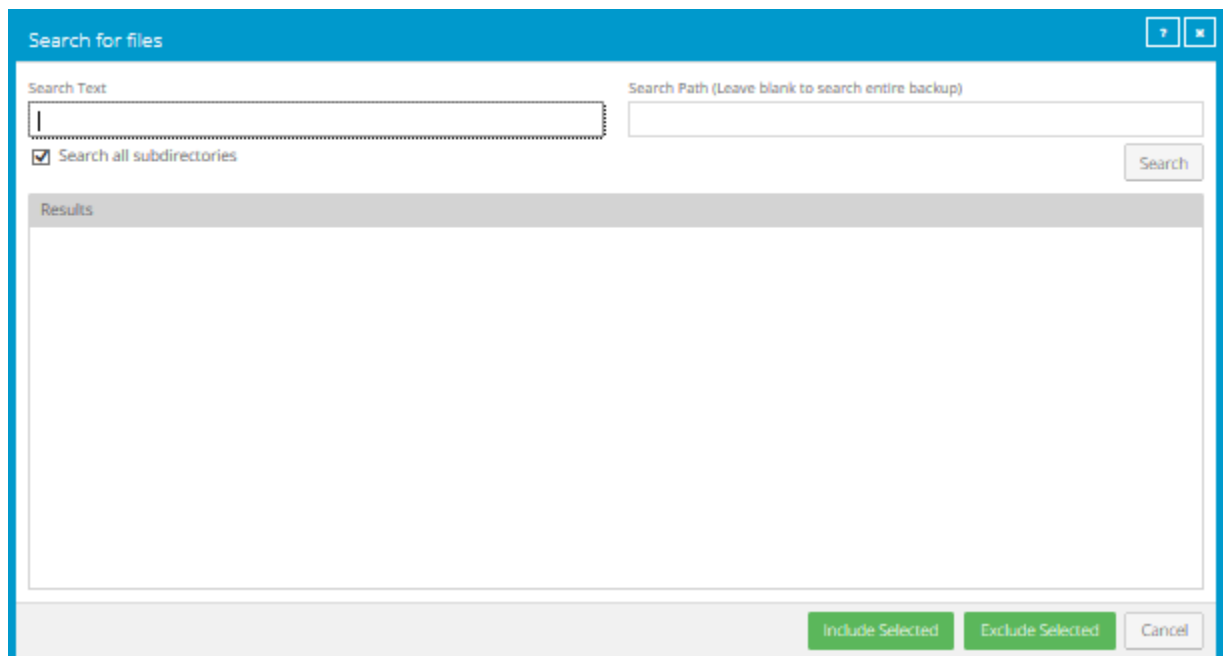
6.1.1 Search for files to restore

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

To search for files to restore:

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

The **Search for files** dialog box appears.



2. In the **Search Text** box, enter the file name to search for. You can include asterisks (*) as wildcard characters.

3. To search for files in a specific folder in the backup, enter the path in the **Search Path** box.
4. To search for files only in the specified folder, clear the **Search all subdirectories** check box.
5. Click **Search**.

The **Results** box lists files that match the search criteria.

6. In the **Results** box, select files to include or exclude. To select multiple consecutive items, press SHIFT while clicking the first and last items in the list. To select multiple items, press CTRL while clicking the items.
7. Do one of the following:
 - To restore the selected files, click **Include Selected**.
 - To exclude the selected files from the restore, click **Exclude Selected**.

6.1.2 Filter subdirectories and files when restoring data

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


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

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

To filter subdirectories and files when restoring data:

1. When restoring data from a backup job, view the **Restore Set** box.

Restore Set					
	Folders Filter	Files Filter	Recursive		
+	Docs	e.g., a*, b*	*.*	<input checked="" type="checkbox"/>	<input type="button" value="x"/> <input type="button" value="edit"/>
+	Documents an...	e.g., a*, b*	*.*	<input checked="" type="checkbox"/>	<input type="button" value="x"/> <input type="button" value="edit"/>
+	Data	e.g., a*, b*	*.*	<input checked="" type="checkbox"/>	<input type="button" value="x"/> <input type="button" value="edit"/>

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

- To include specific subdirectories in the restore, in the **Folders Filter** field, enter the names of subdirectories to include. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only restore subdirectories if their names end with “-current” or start with “2015”, enter the following filter: *-current, 2015*
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To restore specific files, in the **Files Filter** field, enter the names of files to restore. Separate multiple names with commas, and use asterisks (*) as wildcard characters. For example, to only restore files if they have the .doc or .docx extension, enter the following filter: *.doc, *.docx
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 .exe or .dll extension, enter the following filter: *.exe, *.dll
Note: Asterisks (*) are the only supported wildcards in filter fields.
 - To exclude the specified folder, but not its subdirectories, clear the **Recursive** check box.
 - To exclude the folder’s subdirectories, select the **Recursive** check box.
5. Click **Apply Now** to consolidate and simplify records in the **Restore Set** box, if changes need to be applied.
6. Click **Run Restore**.

6.2 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.

6.3 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.

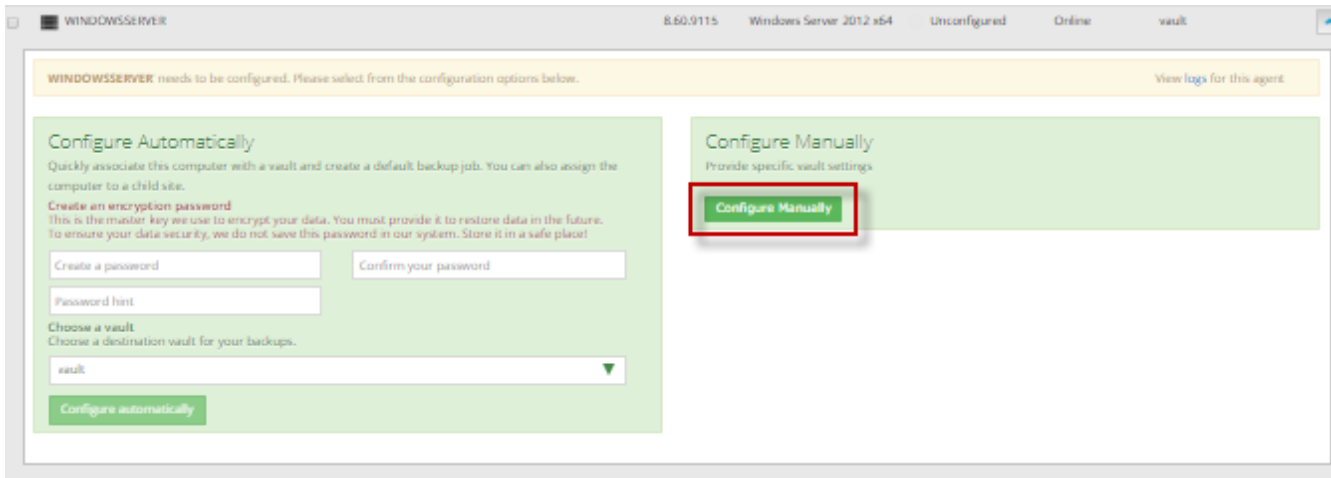
After you re-register a computer with a vault, you must synchronize existing backup jobs before they run successfully. See [Synchronize a job](#).

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.

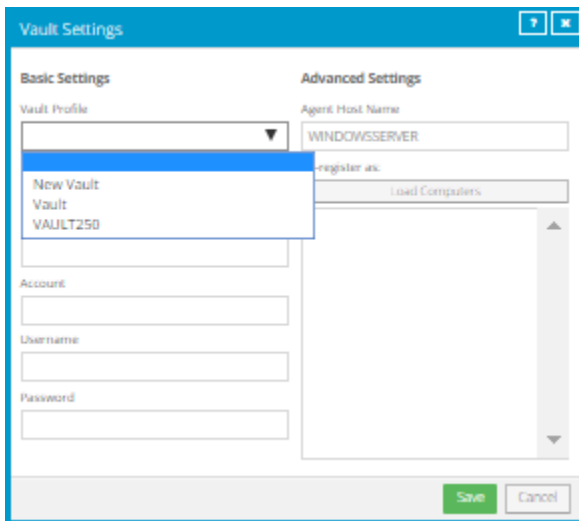
If the following messages appear, a backup job has not been created for the computer. Click **Configure Manually**.



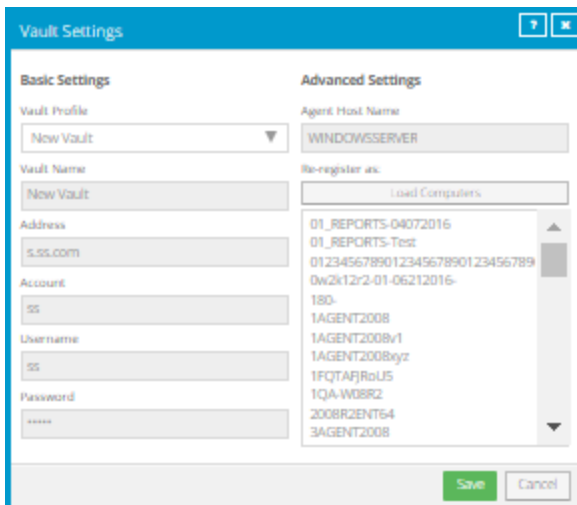
4. Click the **Vault Settings** tab.
5. Click **Re-register**.



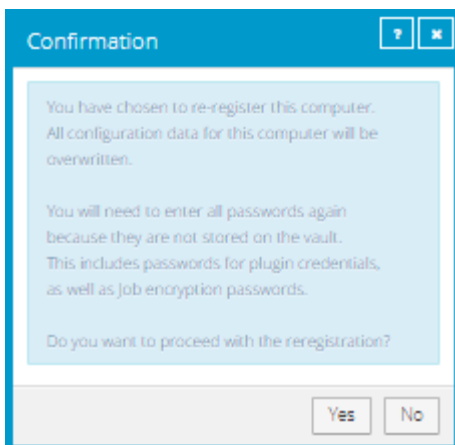
6. In the **Vault Settings** dialog box, in the **Vault Profile** list, select the vault where the backup from the original computer was stored.



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. After job information is downloaded, click the **Jobs** tab.
11. Find a job whose data you want to restore, and click **Restore** in the job's **Select Action** menu.

During a restore, you must enter any passwords required for the job, including the encryption password. The remaining steps are the same as the steps for regular restores.

Note: After you re-register a computer with the vault, you must synchronize existing backup jobs before they run successfully. See [Synchronize a job](#).

6.4 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. If the data was backed up using a plug-in, the destination computer must have the same plug-in installed. If the data was backed up using the Exchange Plug-in, the destination computer must also have

Microsoft Exchange installed. If the data was backed up using the SQL Plug-in, the destination computer must also have Microsoft SQL Server installed.

The new computer then downloads information from the vault so that the data can be restored on the new computer. For example:

- Computer A backs up data using Job A
- Computer B restores data from Job A (computer A's data) to Computer B

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.

6.5 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 from a local job, you can select one of the following options:

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

Log Options

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

- **Files** — Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up.
- **Directory** — Provides less detail than the Files logging level. Provides information about folders that are backed up.
- **Summary** — Provides high-level information, including the vault and Agent version, and backup sizes.
- **Minimal** — Provides high-level information, including the vault and Agent version.

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

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

Bandwidth settings include:

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

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

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

7 System recovery

The purpose of this chapter is to illustrate techniques for recovering a file system. The procedures provided describe the minimum resources and information required to rebuild the file system to its state at the last system backup. The recovery procedure can be performed from a backup disk or directly from a vault.

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 post-restore maintenance.
5. Verify the restore.

Prior to performing a recovery, ensure that your hardware configuration is at least sufficient to hold the programs, data, and system state previously installed on the system.

7.1 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.

7.2 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.

6.3 Recovery steps

This section describes the steps to perform a system recovery.

7.2.1 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.

7.2.2 Install and configure the Agent

1. Install the Agent for your operating system.
2. Configure the Agent. Reregister 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.

7.2.3 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.

7.2.4 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.

7.2.5 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.

7.2.6 Recovery problems

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

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

8 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.

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.¹

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.

8.1 Install the Oracle Plug-in for Linux

To protect Oracle databases, you can install the Oracle Plug-in with the Linux Agent.

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

For supported platforms and database versions, see the Oracle-Plug-in for Linux release notes.

You can determine which version of Oracle you have 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.

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

To install the Oracle Plug-in for Linux:

1. Download the Oracle Plug-in for Linux tar.gz installation package.
2. Extract the files from the package. To do so, type the following, where *PackageName* is the name of the Oracle Plug-in installation package:

```
> # cd /tmp
> # tar xvf PackageName.tar
```

3. Type the following, where *PackageName* is the name of the Oracle Plug-in installation package:

```
> # cd PackageName.xxxx
```

4. Run the installation script:

```
> # ./install.sh
```

5. Follow the installation instructions on the screens.

8.2 Add an Oracle database backup job

After a Windows computer with the Oracle Plug-in is added in Portal, you can create a backup job for one or more Oracle databases. The backup job specifies which databases to back up, and where to save the backup data. You must also specify credentials for the Agent to use to connect to the Oracle server.

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

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 the Windows Agent or Linux Agent version 8.60 or later. If you specify that archived logs should be deleted after a backup, ensure that the logs are backed up using a Local System or Image job.

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

To add an Oracle database backup job:

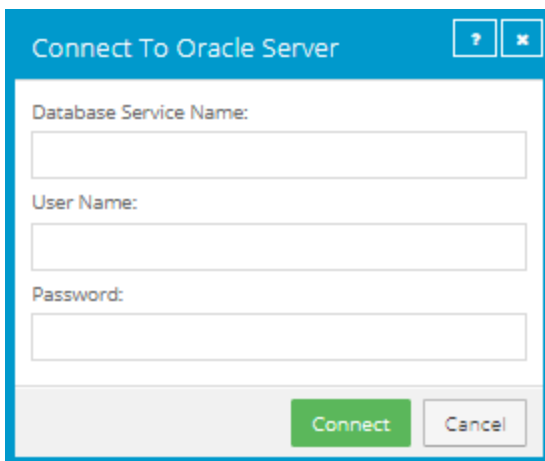
1. On the navigation bar, click **Computers**.

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](#).

8.2.1 About Oracle backups

The Oracle Plug-in 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.

8.2.2 How the backup works

When a backup starts, the Oracle Plug-in 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.

8.2.3 Table of backup information

Before you perform Oracle database backup or restore processes, 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

System Requirement	Customer/User Supplied Value	Comments
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 <i>Database Instance</i> from Oracle (rather than the <i>Instance Name</i> from Oracle). In Windows CentralControl, set the Oracle Service Name to the <i>Database Instance</i> from Oracle (rather than the <i>Instance Name</i> from Oracle).
Encryption type	Encryption type = Password = Password Hint =	You must supply an encryption 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 Directors (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**).

8.3 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.

8.3.1 Guidelines for restoring

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

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

```
SQL> STARTUP 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.

8.4 Uninstall the Oracle Plug-in

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.

9 Monitor computers, jobs and processes

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

- Computer page. The Computer page shows status information for protected 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 computers logs](#).
- Process Details dialog box. This dialog box shows information about all running, queued and recently-completed processes for a job. See [View current process information for a job](#).
- 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 jobs 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](#).

9.1 View computer and job status information

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





To view computer and job status information:


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



The Computers page shows registered Agents.

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

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




-  OK — Indicates that all jobs on the computer ran without errors or warnings.
 -  OK with warnings — Indicates that one or more of the computer's jobs completed with warnings.
 -  Attention — Indicates that one or more of the computer's jobs failed or completed with errors.
 -  Unconfigured — Indicates that no jobs have been created for the computer.
2. Find the Agent for which you want to view logs, and click the row to expand its view.
 3. View the **Jobs** tab.

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






Name	Job Type	Description
 AppAware	Image	
 FilesAndFolders	Local System	

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

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

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

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

-  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.
-  Failed — Indicates that the backup failed and no safeset was created.
-  Cancelled

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

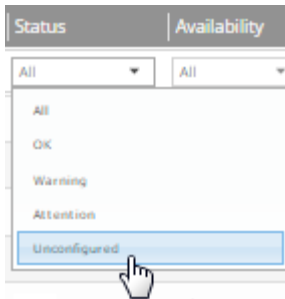
9.2 View an unconfigured computer's logs

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

To view an unconfigured computer's logs:

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

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

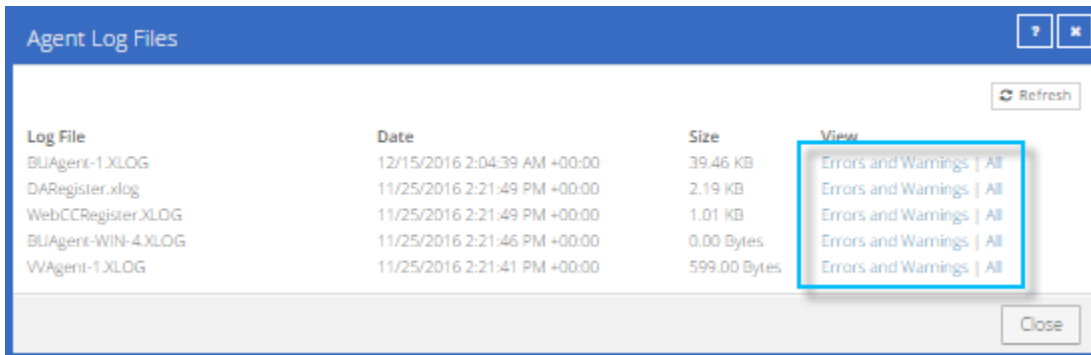


- Find the unconfigured computer, and expand its view by clicking the computer row.



- 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.



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

The log appears in a new browser tab.

```


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


```

9.3 View current process information for a job


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

To view current process information for a job:


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




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



The **Process Details** dialog box lists processes that are running, queued and recently completed for the job. Detailed information is shown for the process that is selected on the left side of the dialog box.



- To view information about a different process, click the process on the left side of the dialog box. Detailed information for the process is shown at the right side of the dialog box.
- To show only some processes in the dialog box, do one of the following in the status list:
 - To only show queued processes, click **Launched**.
 - To only show processes that are waiting for user action, click **Operator Request**.
 - To only show processes that are in progress, click **Running**.
 - To only show completed processes, click **Finished**.
 - To only show processes that are finishing, click **Finalizing**.



9.4 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](#).

9.4.1 Set up email notifications for backups on a computer

To set up email notifications for a computer:

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

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

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.

Note: Email notifications can be sent using CRAM-MD5, AUTH LOGIN and AUTH PLAIN authentication.

4. Click **Save**.

9.4.2 Set up email notifications for backups on multiple computers


By default in some Portal instances, Admin users receive emails when backups fail, or are cancelled, deferred, missed or completed. Admin users can select backup statuses for which they want to receive email notifications. These email notifications are sent for Linux systems with Agent version 8.10a or later, instead of separately for each computer.

For other computers, and 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), you can select events for which you want to receive emails.



If Email Notifications Settings do not appear, you must set up notifications separately for each computer. See [Set up email notifications for backups on a computer](#).

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
4. Click **Update notifications**.

9.5 View a job's process logs and safeset information

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

You can also view information about safesets created for the job. A safeset is an instance of backup data on the vault. For most Agents, one safeset is created by each successful backup.

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

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

The Computers page shows registered Agents.
2. Find the Agent for which you want to view logs, and click the row to expand its view.

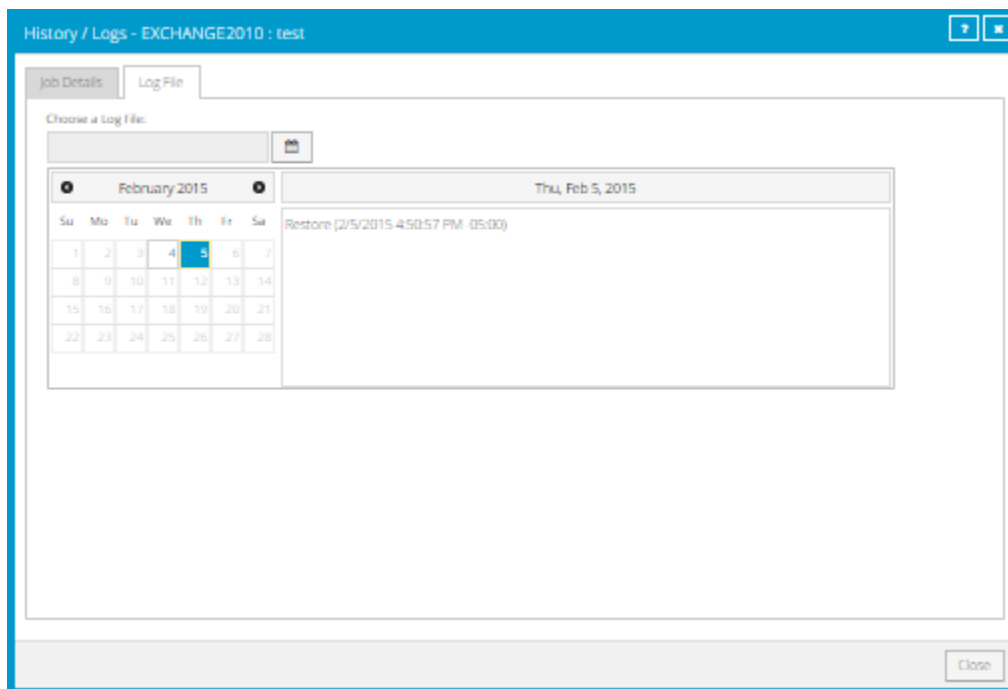
On the **Jobs** tab, the **Last Backup Status** column shows the status of each backup job.


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

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


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

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





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

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



5. To only show errors and warnings in the log, click the **Errors and Warnings Only** option at the top right of the window.
6. 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.



9.6 View and export recent backup statuses

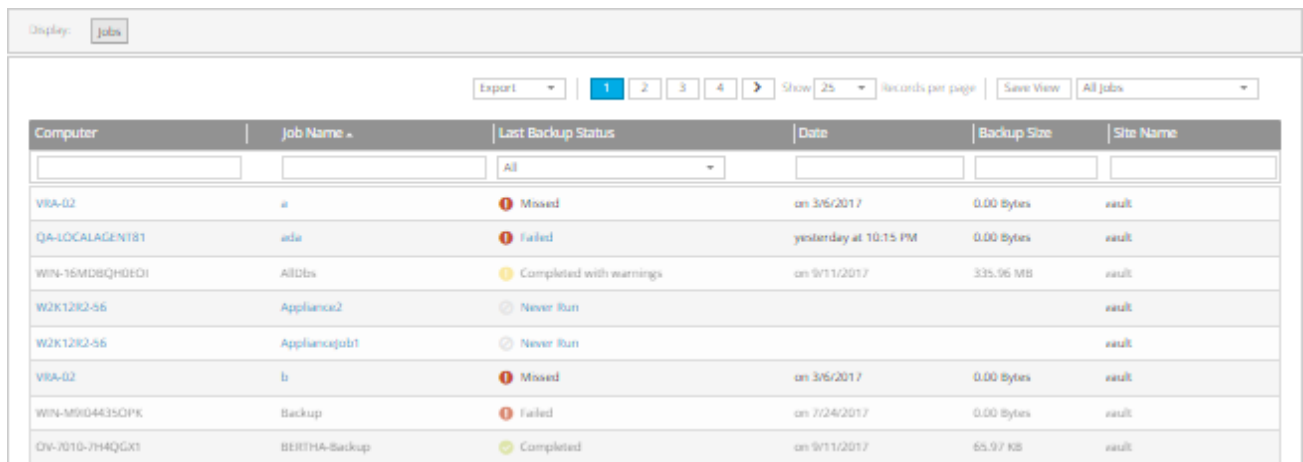
You can view recent backup statuses for computers on the Monitor page in Portal. You can also export the information in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format.

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

To view and export recent backup statuses:


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

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



Computer	Job Name	Last Backup Status	Date	Backup Size	Site Name
VRA-02	a	Missed	on 3/6/2017	0.00 Bytes	naul
QA-LOCALAGENT81	ada	Failed	yesterday at 10:15 PM	0.00 Bytes	naul
WIN-16MD8QH0ED1	AllDBs	Completed with warnings	on 9/11/2017	335.96 MB	naul
W2K12R2-56	Appliance2	Never Run			naul
W2K12R2-56	ApplianceJob1	Never Run			naul
VRA-02	b	Missed	on 3/6/2017	0.00 Bytes	naul
WIN-M9104435DPK	Backup	Failed	on 7/24/2017	0.00 Bytes	naul
DN-7010-7H4QGK1	BERTHA-Backup	Completed	on 9/11/2017	65.97 KB	naul

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



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

10 Carbonite Server Backup Support

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

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

What can we help you with?

Search


Popular Searches

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

10.1 Contacting Carbonite

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

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



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

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

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