



Carbonite Server Backup Windows CentralControl 8.8

Operations Guide



© 2020 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	April 2020	Initial guide provided for Windows CentralControl 8.8x.

Contents

1	Welcome to Windows CentralControl	6
1.1	Product Overview	6
1.2	Workspaces and Jobs	8
1.3	Safesets / Backups	8
1.4	Security	9
1.5	CentralControl Passwords	9
1.6	Encryption	9
1.7	Seeding	10
1.8	Additional Documentation	11
2	Getting Started with Windows CentralControl.....	12
2.1	Permission Requirements	12
2.2	Locking Configuration Files.....	12
2.3	Default Installation Directories	13
2.4	Installing Windows CentralControl	13
2.5	About the Setup Maintenance Wizard.....	13
2.6	Repairing an Installation.....	14
2.7	Uninstalling CentralControl	14
2.8	Managing Workspaces	15
2.9	Creating a Workspace	15
2.10	Renaming a Workspace	15
2.11	Opening a Workspace	15
2.12	Encrypting a Workspace.....	16
2.13	Setting Workspace Options	16
2.14	Deleting a Workspace.....	17
3	Managing Agents.....	18
3.1	Adding an Agent	18
3.2	Registering an Agent	19
3.3	Privileges for Agent Configuration	21
3.4	Existing Agents	21
3.5	Propagating Settings	26

3.6	Working with Groups of Agents	29
4	Managing Vault Profiles.....	33
4.1	Adding a Vault Profile.....	33
4.2	Modifying a Vault Profile.....	35
4.3	Copying a Job to Another Vault.....	36
5	Creating and Managing Jobs	37
5.1	Creating a Job	37
5.2	Backup Source Type	38
5.3	Destination/Vault	39
5.4	New Job Name.....	39
5.5	Data Sources.....	40
5.6	Encryption	45
5.7	Local Catalog Files	45
5.8	Log Files	45
5.9	Finished Screen/Options	50
5.10	Existing Jobs.....	50
5.11	Monitoring Jobs.....	52
6	Creating and Managing Schedule Entries.....	54
6.1	Schedule List.....	54
6.2	Creating a New Schedule Entry	54
6.3	Working with Schedule Entries	57
6.4	Scheduler Log Files	59
6.5	Effects of Time Changes	59
6.6	Time Zones	60
7	Backing up Data	61
7.1	Seeding.....	61
7.2	Seeding with Deferring.....	61
7.3	Backup Wizard.....	62
7.4	System State Objects.....	62
7.5	Locked/Open Files	63
7.6	Back Up After Reregistration.....	63

- 8 Restoring Data 64**
 - 8.1 Restoring from Conventional Backups64
 - 8.2 Restoring from vSphere Agent Backups.....73

- 9 Resolving Common Errors..... 78**
 - 9.1 Email Notification Not Received.....78
 - 9.2 Failed to Authorize – Insufficient Privileges78
 - 9.3 Failed to Connect.....79
 - 9.4 I/O Device Error.....79
 - 9.5 Limit on Number of Shared Memory Segments.....79
 - 9.6 VVAgent Unexpected Shutdown on Unix.....80
 - 9.7 System Trouble after Active Directory Restore.....80
 - 9.8 Configuration File Missing Info on Reregistration.....81

- 10 Additional Information 82**
 - 10.1 Recursive Backups82
 - 10.2 User and System State Data82
 - 10.3 Active Directory Restores.....84
 - 10.4 Backing Up/Restoring Event Log Databases.....87
 - 10.5 Backing Up/Restoring Terminal Service License Databases.....87
 - 10.6 Using the Exchange Plug-In87
 - 10.7 Using the SQL Server Plug-In93
 - 10.8 Using the Cluster Plug-In101
 - 10.9 Using the Oracle Plug-in103
 - 10.10 Encryption, Compression and OTW105

- 11 Carbonite Server Backup Support 107**
 - 11.1 Contacting Carbonite.....107

1 Welcome to Windows CentralControl

You can use Windows CentralControl to configure, monitor, and manage all backup and recovery activities for computers where Agents are installed.


IMPORTANT: To take advantage of new features and important security fixes, we highly encourage using Portal to manage Agents instead of Windows CentralControl.

You can configure jobs to back up selected files, databases, or complete systems. Backups can occur on a regularly scheduled basis, and you can also run them immediately ("ad hoc"). In addition, you can receive regular email regarding the status (success or failure) of your backups.

The material in this information resource is intended for people who are responsible for installing, configuring, and using CentralControl.

1.1 Product Overview

This diagram shows the relationship between data protection products:



The Vault, Agents and CentralControl are the primary components of the data protection application. You use these components to back up from your network computers to a local or remote vault, and restore data from the vault. These applications protect your data on site or off site (in the cloud, for example).

1.1.1 CentralControl

You can use Windows CentralControl to manage computers that run Agent software. You use CentralControl to configure Agents, jobs, scheduling, monitoring, and to restore backup data. During backups, data moves from the Agent computer to the vault. Backup data does not go through CentralControl. You need a vault account and user name in order to back up to a vault.

After you configure the Agent and create schedules, backups occur automatically. You do not need to run CentralControl continuously. Use CentralControl to configure the Agent, check the progress of backups, and view error logs.

You do not require a license to use CentralControl. You must have Administrator privileges on the Agent computers to create or modify backups, and manage recoveries.

1.1.2 Agents

The Agent is an application that runs on the host system as a local component service.

After an Agent is installed, use Windows CentralControl and the Agent to back up data from the computer to a vault. The applications provide an automated, unattended method for protecting your data.

Normally each computer that you back up must meet these requirements:

- The Agent software must be installed and running.
- The computer and the Agent must be connected to a network.
- The Agent must be able to access a vault.

The Agent scans the server for changed data and delivers those changes to a remote electronic vault. Changes are then added to the backup on the vault to complete a full backup of the server. You can manage and control many Agents through one CentralControl application. An Agent can have many jobs, but Agent names must be unique on the vault.

Each backup is considered a full backup because you can restore a full set of data from any backup version (safeset). This is easier than doing an initial restore, and then restoring additional incremental versions. The Agent runs on the computer as a background service. It starts automatically when the system starts.

1.1.3 Plug-ins

Plug-ins for creating application-consistent backups are available with some Agents. For more information, see the Agent documentation.

To see a list of plug-ins installed with an Agent, highlight the Agent, and click **File > Properties > Get Status**.

Plug-ins are optional, and require extra licenses. When you install the Agent, you can also install the plug-ins (and activate them later).

1.1.4 Vaults and Licensing

A vault is a computer system that receives a user's backup data from a LAN or the Internet and securely stores the data on disk. See also [Creating a Vault Connection](#).

Vaults use a quota system to control Agent licensing. When an Agent connects to a vault, the vault automatically supplies the license. Licenses are required for most Agents, plug-ins, and other Agent products that use a vault. You can create a backup job without a license, but you will get a warning message.

Without a valid license, the backup will not run. You can, however, run a recovery. Previously licensed Agents are unaffected by the failure. Contact your service provider to purchase additional licenses.

1.2 Workspaces and Jobs

A [workspace](#) in Windows CentralControl is a uniquely named, user-defined environment. In this environment, specific computers that run the Agent software are represented and defined, and jobs are created for those Agents.

[Jobs](#) define what data to protect and how to do it. This can involve file selection, filters, compression, encryption, and retention settings.

A job always belongs to only one Agent. Job names are unique within an Agent.

1.3 Safesets / Backups

Each time a backup is successfully performed, a safeset is sent to the vault. The safeset contains the data protected in that backup session. The safeset also carries the associated backup type, data options, and retention settings. This information is used to help manage the safeset until it (optionally) expires. The list of safesets is maintained on the vault (if this is the destination), as well as on the Agent.

Over time, a pool of safesets will develop based on the values specified in the retention settings. This safeset pool makes up the backup history for a specific job. The safeset pool, when stored on a vault, is dynamic. For instance, each time a new backup is performed or a safeset expires, the pool changes. Typically this happens automatically after each backup is performed. On occasion, such as before a recovery, this information may require synchronization with the local Agent.

Depending on the destination of your backup, you may or may not see the safeset file locally. Each backup operation creates a safeset image (.ssi file) into which the contents of selected disk files, directories or volumes are placed. The contents reflect the job settings for the backup operation.

1.3.1 Safeset Properties

To view safeset properties, double-click a catalog entry in the **Safesets** folder. The **Safeset Properties** screen opens, displaying this information:

- **Job** - The name of the job that created this safeset
- **Catalog number** - The entry number (e.g., **00000002**) within the catalog
- **Location** - The place (e.g., vault name or IP address) where the safeset is stored
- **Status** - The current status (e.g., **Online**)

- **Backup time** - The date and time the backup was created. This is the local time for the Agent/Server, rather than vault time.
- **Storage size** - The size of the backup in its original, deltized, and compressed forms
- **Backup type** - The type of the backup (e.g., **Full**)
- **Retention** - The retention settings based on days online, copies online and days archived
- **Encrypted** - A **Yes** value indicates that the backup was encrypted
- **Compressed** - A **Yes** or **No** value indicating whether or not the backup was compressed
- **Media type** - The media used (e.g., **Disk**)
- **Expiry** - The user-defined date on which the safeset expires

1.4 Security

Data security can be applied in several places:

- Data going to/from the Agent and vault is encrypted by the system. You might choose to turn off this Over The Wire (OTW) encryption in a secure network.
- You can encrypt the workspace to prevent unauthorized access.
- Data that is stored on the vault is encrypted.

1.5 CentralControl Passwords

Agent password: Used by CentralControl to communicate with an Agent.

Vault password: Used by an Agent to communicate with a vault. It is created by the vault operator and given to a user. If a user loses or forgets the password, the vault operator can issue a new one. There can be many different users with different passwords in the same vault account.

Workspace password: Protects all of the other passwords and access to information through CentralControl. It is set at the workspace level. If the password is lost, you must recreate the workspace and settings.

Encryption password: Secures data stored on a vault. You must remember the encryption password in order to restore data from the vault.

1.6 Encryption

Encrypting settings in a job specify the encryption type used for backup data.

You must specify a password for encrypted backup data. The password is case-sensitive. To recover files, you need to remember the encryption password that was entered when the files were backed up.

Important: If you forget the encryption password for a job, you cannot restore data from the job.

If you change the encryption options of an existing job, it will force a reseed of the backup data. The next backup session will take longer than previous delta backups, and will increase the amount of data stored on the vault.

1.7 Seeding

"Seeding" is the process of getting the initial full copy of the data to the vault. After the vault has received an initial full copy, the built-in DeltaPro technology reduces subsequent traffic so that further backup procedures complete in a shorter period of time, and use less disk space on the vault. Note that each subsequent backup is still considered a full backup.

Seeding the first backup

There are ways to create the initial seed for a backup job: perform a full backup over the internet or a local network within your backup time window, or use Removable Media Seeding.

Each technique has its own advantages. Determine the best approach for your system based on available resources, communications bandwidth and scheduling requirements.

1.7.1 Creating a Seed Using the Backup Time Window

You can seed your backup by performing a full backup after creating a **Backup time window**. This setting is used to specify a time period in which a backup can run.

Any files that are not backed up within the specified window will be deferred until the next backup. This window can be used to help a full backup grow over a period of days.

Establish your window when you create a backup. Using the Job Wizard, set up the time window on the Options screen. After your job has been created, you can change this setting on the Advanced tab (in Job Properties). Enter the number of hours or minutes after which all new backup files will be deferred.

After the backup has run a few times, you can look at the summary log to determine if your backup time window is sufficient to complete a typical full backup with the Delta processing option.

An advantage to this technique is that you do not permanently need removable media connected to the Agent system. In addition, there is no need to transport media to your vault.

A disadvantage is that the seeding process can take several days to complete if your communication bandwidth is limited. During this time, you have only partially completed backups that may or may not suffice in a crisis. In this case, continue regular local disk backups until your system is fully seeded.

1.7.2 Removable Media Seeding

If your communications bandwidth is limited, it may be faster to perform removable media seeding. This is the process of directing your first full backup to a disk or other type of removable media. The media is then transported to the vault where the vault processes it and initializes your backup data. Subsequent backups can be performed over the network connection to the vault directly.

To create the seed:

1. Create your backup job as if you were going to back up over the network directly to the vault.
2. Before defining your backup schedule, execute an ad-hoc (immediate) backup.
3. When the Backup Wizard starts, temporarily redirect your backup by choosing **Alternate safeset location** for the destination. This will not affect your permanent settings for the job.

1.8 Additional Documentation

Release notes for this product are available from your service provider. Release notes contain the most current information about the product, including an overview of new features, any known defect (bug) fixes incorporated since the last release, and descriptions of any known issues.

Windows CentralControl includes online help for many screens. To access the help, press **F1**, or select **Help** on the toolbar.

2 Getting Started with Windows CentralControl

The topics in this chapter describe the steps required to install the CentralControl application. See the release notes for supported operating systems.

The CentralControl installation kit is available in a self-extracting executable file.

2.1 Permission Requirements

This table shows the Microsoft permissions required by CentralControl for installation and use. The ADMINISTRATOR setting is needed for installations that require registry configuration.

Program	Function	Minimum Permission
CentralControl	Installation	BACKUP OPERATOR
	Using the CentralControl GUI	BACKUP OPERATOR is sufficient for accessing non-mapped drives. ADMINISTRATOR rights are needed to access remote mapped drives.
Agent	Installation	ADMINISTRATOR

2.2 Locking Configuration Files

The directories in which the Agent and CentralControl program files and configuration files are stored can be locked down to prevent reading and writing by non-privileged users.

The only users with permission are:

- .\Administrators group
- .\BackupOperators group
- .\LocalSystem user

Locking the directory tree prevents non-privileged users from reading sensitive information in the configuration files.

Note: With the default “locked down” installation, non-privileged users cannot run CentralControl, since CentralControl is normally installed in the same directory as the other Agent programs and configuration files. (Scheduled backup jobs can still run normally.)

2.3 Default Installation Directories

By default, the installer puts the files for fresh installations in C:\Program Files (x86)\<service name>\CentralControl on a 64-bit computer.

All of the CentralControl executables are then located in a single directory, but have their own subdirectories in which to run. Requirements (such as locking down for Admin users only) do not affect the other applications.

The installer will not remove the top-level directory unless it is empty at the end of an uninstall.

2.4 Installing Windows CentralControl

To install Windows CentralControl:

1. Double-click the self-extracting executable file to start the installation process.
The Welcome screen appears. Click **Next**.
2. On the View Notes screen, you can view and/or print support and contact information. Click **Next** to proceed with the installation.
3. For a new installation, the Software License Agreement screen appears. Read the agreement, and then click **ACCEPT** and **Next**.
The Choose Destination Location screen appears. The default installation location is "C:\Program Files (x86)\Carbonite Server Backup\CentralControl".
4. To install the application in the default directory, click **Next**.
To install the application in another location, click **Browse**. Choose a location, click **OK**, and click **Next**.
The Authentication Information screen opens.
5. Enter the credentials for a user with Backup Operator or Administrator privileges. Click **Next**.
The Desktop Shortcut screen appears.
6. To add a shortcut icon to your desktop, choose **Yes**. Click **Next** to continue installing.
7. On the InstallShield Wizard Complete screen, you can choose to directly launch the CentralControl application. Click **Finish**.

2.5 About the Setup Maintenance Wizard

You can modify, repair/upgrade, or uninstall your CentralControl software by double-clicking the installation file (located on your computer, CD or the Web).

If the installer detects that you have the same version of CentralControl on your system, the Setup Maintenance wizard presents you with the choices that follow. Otherwise, an **Upgrade** prompt appears, allowing you to install a newer version of CentralControl.

Modify: Add or remove specific components.

Repair: Fix the current version of your application.

Uninstall: Partially or totally remove the application.

If you try to run the installation, and the installer detects that you already have a newer version of the CentralControl application on your system, the Setup Maintenance wizard will advise you, and then stop.

2.6 Repairing an Installation

When the installation program is launched, it searches your computer for previously-installed versions of the CentralControl application. If the same version of the application is located, you can modify or repair (or uninstall) the software. If the existing CentralControl is an older version, you will be prompted to upgrade.

To repair the CentralControl application:

1. Double-click the self-extracting executable file.
2. The process starts, and a Welcome page appears. Click **Next** to continue.
3. The View Notes screen appears. Click **Next** to continue.
4. Select **Repair**, and click **Next**. A confirmation message appears. Click **OK** to repair, or **Cancel** to return to the selection screen.
5. When the repairs complete, a Maintenance Complete screen appears. Click **Finish**.

Note: You may be prompted here to restart your computer.

2.7 Uninstalling CentralControl

To uninstall CentralControl:

1. Double-click the self-extracting executable file, or follow the process for uninstalling through the Windows Control Panel.
2. The Welcome page appears. Click **Next**.
3. If the View Notes screen appears, click **Next** to proceed. Select **Uninstall** and click **Next**. When you are prompted to confirm, click **OK**.
4. When the Uninstallation Type screen appears, select **Total Uninstall** or **Program Files only**. **Total Uninstall** removes all traces of the application from your system. **Program Files only** leaves job configuration files and log files on your computer for future use.

If you choose **Program Files only**, the process of uninstalling continues when you click **Next**. If you choose **Total Uninstall**, you are prompted. If you choose **Yes**, the process of totally uninstalling continues. If you choose **No**, only the program files are uninstalled.

5. Once the process of uninstalling is complete, click **Finish**.

2.8 Managing Workspaces

You use workspaces to organize your Agent connections into logical groups. For example, you can create a workspace for individual company departments. Within each workspace, you add Agents and jobs.

The file type used to identify workspace files is Vault Workspace (**.vws** file extension). By default, these files are saved in the installation directory.

When you open CentralControl for the first time, the default workspace is called **MyWorkspace**. It is recommended that you change this name.

Also see [Creating a Workspace](#).

2.9 Creating a Workspace

To create a new workspace:

1. Click **File** and select **New Workspace**.
2. Click on **Workspace (Untitled)** in the left pane of the screen.
3. Click **File** and select **Save Workspace As**.
4. The **Save As** dialog box will open. Enter a name for the workspace in the **File name** field.
5. Click **Save**.

2.10 Renaming a Workspace

To rename a workspace:

1. Select the workspace in the left pane of CentralControl.
2. Click **File** and select **Save Workspace As**.
3. Enter a name for the workspace in the **File name** field.
4. Click **Save**.

2.11 Opening a Workspace

To open an existing workspace:

1. Click **File** and select **Open Workspace**.
2. Browse to the location of the workspace file. For the default locations of workspace (VWS) files, see [Managing Workspaces](#).
3. Select the file, and click **Open**.

2.12 Encrypting a Workspace

To encrypt a workspace:

1. Select a workspace in the left pane of CentralControl.
2. Select **File > Workspace Password**.
3. Enter your existing password in the **Old password** field.
4. Select an encryption type from the **Encryption type** menu.
5. Enter a new password in the **New password** field. If you forget this password, you must recreate the workspace.
6. Enter the new password in the **Confirm password** field.
7. Click **OK**.

2.13 Setting Workspace Options

To set workspace options:

1. Select a workspace in the left pane of Windows CentralControl.
2. Select **Tools > Options**.
3. Edit the fields in the **Options** dialog box. For field descriptions, see [Fields for Workspace Options](#).
4. Click **OK**.

2.13.1 Fields for Workspace Options

These fields appear in the workspace **Options** dialog box:

Field	Description
Automatically reload last workspace on startup	Automatically open the current workspace when you close and then reopen CentralControl.
Auto-refresh display for selected agent every <...> minutes	Frequency for automatically refreshing the Agent information on display. You can enter values from 60 to 1440.
Update progress display every <...> seconds	Frequency for refreshing the Process Information screen. You can enter values from 5 to 10000.
Return maximum <...> files and directories	Maximum number of files and directories to display at one time. You can enter values from 10 to 10000000.
Default text viewer	Default viewer for logs and text-based files.

2.14 Deleting a Workspace

To delete a workspace:

1. In Windows Explorer, browse to the location of the workspace file. For the default location of workspace (VWS) files, see [Managing Workspaces](#).
2. Right-click the workspace file, and select **Delete**.

3 Managing Agents

In a CentralControl workspace, an Agent represents an individual computer. The Agent program residing on a computer enables configuration and status communication between the system and CentralControl.

Before communication starts, the Agent program must be installed on each computer to be managed by CentralControl. The Agent uses local permission information to protect it from unauthorized access.


To configure and manage an Agent system from CentralControl, you must connect through an account on the Agent system, and you must have sufficient privileges to use the account.

Also see [Creating a New Agent](#) and [Fields for Agent Properties](#).

3.1 Adding an Agent

To create an Agent:

1. Open CentralControl.
2. Right-click a workspace in the left pane, and select **New Agent**.



3. Complete the fields on the **Agent Properties** screen. For field descriptions, see [Fields for Agent Properties](#).

4. Click **Get Status** to test the Agent settings.

If the DNS or IP information is incorrect, the message Failed to connect to <...> appears. If the authorization information is incorrect, the message Failed to authorize user () or user () possesses insufficient privilege appears. Contact your service provider.

5. If you are satisfied with the settings, click **OK**.
6. Click **OK** again.

3.1.1 Fields for Agent Properties

See also [Creating a New Agent](#).

These fields appear on the **Agent Properties** screen:

Field	Description
Description	Agent profile name
Network address	IP address or Domain Name System (DNS) name for the Agent computer. For the vSphere Agent , use the IP address that appeared after vCenter registration.
Default ports	Default ports used to communicate with the Agent computer
Custom ports	Custom ports used to communicate with the Agent computer
User name	User name for accessing the Agent
Password	Password for accessing the Agent. This field is case-sensitive.
Save password	Allows you to save the password for the Agent
Domain	Name of the Windows domain on which the Agent computer is installed. If you are not using a domain name, you can leave this as an empty field.

3.2 Registering an Agent

Before you can start backing up data, you need to register the Agent with an assigned vault. You supply parameters such as your user ID, password, and vault address (obtained from vault personnel) so that the vault has information indicating that the Agent computer is valid. This connection to the vault registers that computer (with its Agent) with the vault. The registration has succeeded if the vault allows you to connect and proceed.

The vault then sends an ID back to the Agent computer. This ID is a unique Agent computer identifier. This part happens without your intervention or knowledge.

As you create jobs, you register them with the vault. There can be multiple jobs associated with a vault registration, but an Agent computer only needs to register with the vault once.

Now the Agent has information about the vault, and the vault has information about the Agent. This information is kept on the Agent side and on the vault side. If something happens to the Agent computer (e.g., it crashes), when you do a disaster recovery, the vault information about the Agent is sent back to the Agent computer. This is a reregistration to recover all of the Agent's parameters, such as schedules and job files.

During reregistration:

1. Provide the same vault address, user ID, and password as the initial registration.
2. The vault will provide a list of computers that it has associated with that ID.
3. Select the computer that you want to replace in order to recover that Agent's job information from the vault.

3.2.1 Editable Host Names

When you first use an Agent to connect to a vault, you must register the Agent with the vault using a unique Agent name. This allows the vault to recognize that Agent and computer for future backup and recovery.

If, for conflict reasons, you need to use a different host name, click **Advanced** when you reach the Authentication screen. This opens the Agent Host Name screen. Enable **Change Agent host name**, enter the new name, and click **OK**.

Note: You can only do this during initial registration with the vault. After configuration, you cannot change this name.

If you lose the system and need to recreate it, you must [reregister](#). (You can select the name that you gave to the previous system, but you cannot modify it. You can change the authentication information, but not the host name.)

3.2.2 Registering for the first time

To register an Agent (i.e., computer) with a vault, select **Tools > Agent Configuration > Vaults > New**. The Vault Configuration Wizard will open.

You will need to provide authentication information. This information is case-sensitive. Normally the system name is used to register with a vault.

The System Name of the Agent server is the DNS name of that machine. It must be registered with the vault to allow that machine access to the vault, under that account. But what happens if two machines, say in different locations, have the same name? Both of them cannot be registered on the same vault because the names must be unique across the vault.

Changing the Agent host name

If, for conflict reasons, you need to use a different host name, click **Advanced** when you reach the Authentication screen. This opens the Agent Host Name screen. Enable **Change Agent host name**, enter the new name, and click **OK**.

You can only do this during initial registration with the vault. After configuration, you cannot change this name.

This only changes the alias name of the Agent computer on the vault so that the two machines can be distinguished. The real System Name of the computer does not change.

3.2.3 Reregistration/Edit

Through **Agent Configuration > Vaults > Edit > Authentication**, you can change the Account, User name, and Password on a registered computer, but not the Agent Host Name (it is display-only here).

To create a new vault connection and reregister a previously registered computer, use **Agent Configuration > Vaults > New**. The Vault Configuration Wizard will open.

Select **Re-register previously registered computer** to reregister a previously registered computer with the vault. This option is not active when you try to create a new vault destination from within the New Job Wizard.

3.3 Privileges for Agent Configuration

To configure Agents on different operating systems (e.g., Windows, Unix), you require different access privileges.

3.3.1 Windows Agent

This Agent requires users to have BACKUP OPERATOR or ADMINISTRATOR privileges (to run from the scheduler or CentralControl).

To specify the domain property of the Agent, type its domain name in the **Domain** field. You can leave the domain empty if any of these apply:

- You are not specifying a Windows Agent.
- You belong to the same domain as the Windows Agent system.
- Your network does not use a domain controller.

3.3.2 Unix Agent

The Unix Agent requires full root group (GID 0) access privileges to run. Installation and configuration require authentication by users with root group privileges. Installation with "su <command>" is not recommended. If the Agent is not run as the root user, the user who runs it must be part of the root group.

The Agent Service account must also belong to this group.


3.4 Existing Agents

3.4.1 Agent Configuration

After you create an Agent and provide its properties, you can configure the Agent. Note that you must have a valid Agent connection (via IP or host name) and a valid authorization to that system.

Agent configurations are specific to an Agent, and they affect all jobs for that Agent. You can manually set or modify the settings through Agent configuration.

For an existing Agent, click **Tools > Agent configuration**. The Agent Configuration screen opens:



3.4.2 Vaults Tab

Working with Existing Vaults

If you have an existing vault configured, it will show in the list of servers here.

For information about creating a new vault, see [Creating a Vault Connection](#).

Editing a Vault Entry

Note: If an Agent is registered to Portal, you cannot change settings on the Vault tab.

1. On the **Vaults** tab of the Agent Configuration screen, select a vault.
2. Click **Edit**. The **Vault Entry** screen opens.

Deleting a Vault Entry

Before you can delete a vault entry, it must be empty of all jobs.

1. Highlight the target vault.
2. Click **Delete**.
3. Click **Yes** to confirm the deletion.

If there are still jobs associated with this vault, a screen will list the jobs associated with the vault entry, and a message will indicate that you cannot proceed until you have deleted all of the Jobs.

On the screen, **Copy to Clipboard** allows you to copy the list of job names to another place, such as a text file or an email. This gives you a record of the vaults that must be deleted. This can be helpful if the list is large.

3.4.3 Retentions Tab

You can manually define and set retentions through this tab. A retention includes the following

- **Days Online** is the number of days a safeset is kept on the vault before expiring.
- **Copies Online** indicates how many copies of the backup safeset are stored online.
- **Days Archived** indicates if and how long the data will be stored offline.

Note that **Days Online** and **Copies Online** work together. Both conditions need to be met before any backups are deleted.

For example, if both **Days Online** and **Copies Online** are 7, there can never be less than 7 backups. If **Days Online** goes over 7, there will still be 7 copies. If **Copies Online** goes below 7, there will still be 7 days' worth. It must be over 7 days with more than 7 copies before any are expired.

Creating a New Retention

To create a new retention for the Agent:

1. On the Retentions tab, click **New**. The Retention Wizard opens.
2. On the Welcome screen, click **Next**.
3. On the Retention Name screen, enter a name for the retention. Retention names can be up to 32 alphanumeric characters in length. No spaces are allowed. Names are not case-sensitive. Underscores (_) and hyphens (-) may be used.
4. If desired, enter a number from 1 to 9999 in the **days** box. This number defines the number of days a safeset is to be kept on the Vault Server before expiring. Once the expiry date is reached, a safeset is automatically deleted.
5. If desired, enter a number from 1 to 999 in the **copies** box. This number defines the minimum number of copies of a particular safeset (as defined by a job) that will be maintained online. It functions in a first in, first out manner. Once the number of copies is exceeded, the oldest copy will automatically expire or be deleted. This process occurs until the actual number of copies matches the definition. This setting does not apply to archived data.
6. If desired, enter a number from 365 to 9999 in the box for the number of days to archive. This number defines the number of days a safeset will be maintained on removable media at the vault before it expires and is deleted.
7. On the Finish screen, click **Finish**.

To save the retention, you must click **OK** on the Retentions tab before you leave the tab.

Editing a Retention

To manually edit a retention for the Agent:

1. On the Retentions tab, select a retention by clicking on its line.
2. Click **Edit**. This opens the Retention screen.
3. In the **Retention name** field, you can change the name of the retention. Retention names can be up to 32 alphanumeric characters in length. No spaces are allowed. Names are not case-sensitive. Underscores (_) and hyphens (-) may be used.
4. In the **Settings** pane, the number of days and the number of copies can be modified. The number of days can be from 1 to 9999 and the number of copies can be from 1 to 999.
5. In the **Archive** pane, choose whether you want to archive your data. Data can be archived from 365 to 9999 days. If you would like to archive data, discuss this with your service provider.

Note: Changing a retention definition affects only the retention settings and its application to future (yet to be created) safesets.

You may also want to create a new retention. A retention is defined through the Retention Wizard. Activate this wizard by clicking **New** on the Retentions tab.

3.4.4 Notification Tab

You can receive email notification after a successful and/or failed backup procedure.

To configure email notification:

1. Select an Agent, and click **Tools > Agent configuration**.
2. Click the **Notification** tab.
3. Enable or disable the **Send e-mail** options as you prefer.
4. For the **E-mail from address**, enter the address from which the notification is sent. This can be any valid address.
5. For the outgoing mail server, enter the network address of the SMTP server that will send the email.
6. Enter the recipient email addresses, separated by commas.
7. Click **OK**.

Explanation of Notifications

Send email on successful completion: There may be warnings in the log file, but you can still recover any file from this safeset.

Send email on failure: No backup occurred. You cannot recover any files from the safeset.

Send email on error: There are errors in the log file. You cannot recover the files that have the errors, but you can restore other files from the safeset.

To see whether or not the email notification has succeeded, check the log file.

Note: If you enable one, two or all three conditions, you will only receive one email for each backup session. Enable all three if you want to see all conditions.

SMTP Credentials

In some Agent and CentralControl versions, you can see the **SMTP Port** and **SMTP Credentials** fields.

In the outgoing **SMTP Port** field, enter the port number assigned to your SMTP server. The default is port 25.

Depending on your SMTP server security, you might need to enter authentication to allow the Agent to access the server to send email. Enter a valid **User name**, **Password** and (if required) a **Domain** name. This does not get checked or tested until an actual attempt to send an email.

3.4.5 Advanced Tab

This tab provides **Execution priority** and **Bandwidth** settings.

Execution priority

Drag the slider to select your preferred **Execution priority** (for Windows and Unix operating systems). This affects the priority of the Agent program on the computer running the backup or recovery.

High is the highest priority, and **Low** takes the lowest CPU usage priority. The **Normal** (middle) value is usually adequate, depending on the number of concurrent applications that run on the computer during the assigned job.

Bandwidth

This allows you to control the amount of network bandwidth consumed by the Agent during backup and recovery operations. Controlling bandwidth in this way is often called *bandwidth throttling*.

For example, you might want restrictions on daytime ad-hoc backup/recovery so that online users are not affected, and then unlimited usage at night so that scheduled backups will run as fast as possible.

This configuration includes:

- Maximum bandwidth (upper limit), in kilobits per second, to be consumed during each backup and recovery
- Period of time during the day that throttling is in effect. Only one time window can be specified, and outside the window, no throttling takes place.
- The days of the week that throttling is in effect

The time period will appear in 24-hour format, or AM/PM format, depending on how the clock is set on the system running the Agent.


The setting is determined at the start of the job. The setting is adjusted dynamically if the job crosses a throttling window boundary.

The throttling setting applies globally to all backup and recovery jobs, but it may be overridden (for ad-hoc backups, for example) by temporarily resetting the bandwidth to **Use all available bandwidth**. Restore operations can override the setting through the **Use all available bandwidth** setting in **Advanced Restore Options**.

Bandwidth throttling applies across multiple jobs running on a single Agent. So each job running on the same Agent will share the allocated bandwidth equally.

3.4.6 vCenter Tab

You can reach the **vCenter** tab by selecting a vSphere Agent and clicking **Tools > Agent configuration > vCenter**.



The screenshot shows the 'Agent Configuration' dialog box with the 'vCenter' tab selected. The dialog has a title bar with a close button (X) and a tabbed interface with tabs for 'Vaults', 'Retentions', 'Notification', 'Advanced', and 'vCenter'. The 'vCenter' tab is active and contains the following fields:

- vCenter Credentials:**
 - vCenter address: 122.222.22.22
 - User name: Admin
 - Example: domain\username
 - Password: [masked with dots]
 - Test button
- Change Block Tracking:**
 - Enable Change Block Tracking (CBT) for Virtual Machines during backup

vCenter Credentials

Provide the same credentials that were used to register the Agent with the vCenter server. Click **Test** to verify that the credentials are correct.

If the credentials change within vCenter, you must make the same changes here. Click **Test** to verify that any changed credentials are correct.

Change Block Tracking

Backups always use CBT for VMs that have it enabled.

If you clear the **Enable Change Block Tracking (CBT) for Virtual Machines during backup** checkbox, this does not disable CBT for systems that already have it enabled. It will only stop the application from enabling CBT for future VMs.

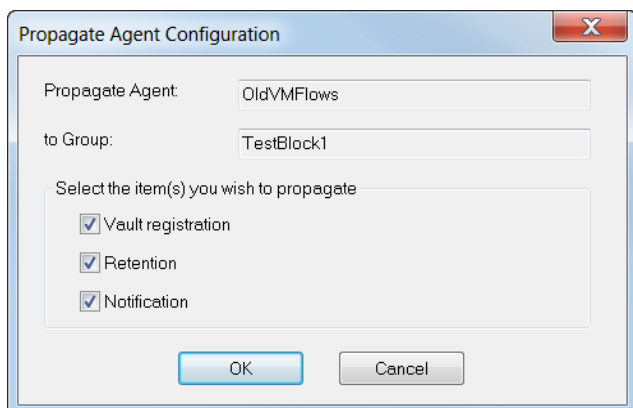
In other words, clearing this checkbox does not stop the Agent from using CBT for backups if the VM has it turned on.

3.5 Propagating Settings

Propagation allows you to configure an Agent or job within a group, and then copy the settings to other Agent/job groups. For information about creating groups, see [New Groups](#).

3.5.1 Propagate Agent Configuration

If you select an Agent within a group, you can choose **Propagate Agent Configuration**. This allows you to configure one Agent, and then copy its information to all of the Agents in the group.



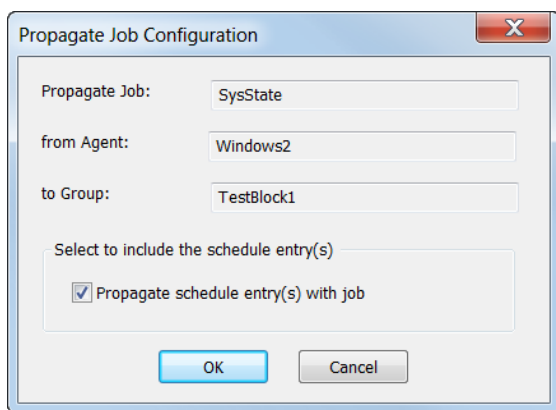
The screen displays the name of the Agent from which to propagate, and the name of the group to which you are propagating.

Select the items that you want to propagate, and click **OK**. A progress screen will open.

3.5.2 Propagate Job Configuration

When you select a job in an Agent group, and right-click on it, you can choose **Propagate Job Configuration**.

Propagate Job Configuration is similar to Propagate Agent Configuration. It shows the name of the job from which you are propagating, the name of the Agent that has the job, and the name of the group to which you are propagating. Optionally you can propagate schedules for that job to the others in the group.




Click **OK** to continue. A Propagation Progress screen will open.

Note: Job propagation is supported only on file-based jobs on similar operating system platforms, such as Windows to Windows, or Linux to Linux. You cannot propagate a Windows-based Agent job to a Unix-based Agent, for example. Also, job propagation is not supported for Agent plug-ins or vSphere Agent jobs.

3.5.3 Propagation Progress

When you click **OK** from Propagate Job Configuration, the progress screen opens.



To view the log file, click **View Log**.

3.5.4 Propagation Logs

A log file is created for each propagation session. It gives you information about the CSV file, and whether each propagation was successful or not.

Example

```
30-Nov 15:49 AGNT-I-04314 Agent Version 6.90.4354 Nov 29 2011 18:14:43
30-Nov 15:49 GCFG-I-07803 propagating started at 30-NOV-2011 15:49:45.85 -0500
30-Nov 15:49 GCFG-I-07804 group name: TestBlock1, base agent name: Windows2, Version 6.90.4280
30-Nov 15:49 GCFG-I-07807 Propagation settings:
30-Nov 15:49 GCFG-I-07829 Propagate job configuration: Yes
30-Nov 15:49 GCFG-I-07832 started propagating job configuration, base job name: Testing
30-Nov 15:49 GCFG-I-07834 Propagating job configuration to agent Windows1, Version 6.72.1072
30-Nov 15:49 GCFG-I-07850 Checking if job Testing exists
30-Nov 15:49 GCFG-I-07854 job does not exist, will register with base vault
30-Nov 15:49 GCFG-I-07855 propagating job vault setting to target agent
30-Nov 15:49 GCFG-I-07898 successfully registered agent Windows1 with the vault 192.168.2.191
30-Nov 15:49 GCFG-I-07857 propagating job retention setting to target agent
30-Nov 15:49 GCFG-I-07862 started registering job with base vault 192.168.2.191
30-Nov 15:49 GCFG-I-07895 successfully registered job Testing with the vault 192.168.2.191
30-Nov 15:49 GCFG-I-07836 Successfully propagated base job configuration to agent Windows1, Version 6.72.1072
```

3.6 Working with Groups of Agents

You can organize your Agents into groups. Normally, when you create (add) a single new Agent, you enter this information:

- Description
- Network address
- Port
- User name
- Password
- Domain

When you first create a group to logically hold Agent names, it is empty. Instead of adding the names individually, as if they were new, the group function allows you to add a “block” of them using a CSV text file that you create.

Later you can propagate the Agent configuration (vault registration, retention and notification), job, and schedule configuration to all of the Agents in that group.


Notes/Rules about Agents and groups:

- The Agent must either exist first, to be moved into the group, or be created by the Agent Import function in the group.
- If you move an Agent (with Add/Delete Agents from group) from the workspace to a group, it will still exist in the workspace.
- If you move an Agent from a group to the workspace, it will move it if it does not exist there. If it does exist there, the move does not occur (that is, nothing happens).
- You may appear to have duplicate Agents in the workspace and/or a group. These are only the Agent Information Description fields. The other information about network address, port, user name, password and domain must be unique. So, if you have two Agent configurations with the same name (Description) but different address, port, user or password, and you move one into or out of a group, it will overwrite the one with the same name. There is no prompt, and the information will be changed.
- You cannot move an Agent directly from one group to another. You need to first move it to the workspace, then to the other group.
- Deleting a group only deletes the Agents in that group, even if they are also duplicated in the workspace, or another group.
- If you delete Agents in the workspace, it does not affect any Agents in a group. They will still be there, even if they are duplicates (copies).

- If you want to totally delete an Agent, you must delete it from the workspace and the group (if it exists in both places).
- Similar to when you add a new Agent manually, adding Agents to a group does not check the validity of the network address, passwords, etc. Only when you do a “Check Status” or click on the Agent icon does the application try to connect to that Agent.

3.6.1 New Groups

A new group can be created from the **File** menu when a workspace is selected. Or it can be created by right-clicking the workspace.



You will be prompted for a new group name. This can be up to 32 alphanumeric characters. The name is not case-sensitive, but must be unique with respect to other groups (with no leading or trailing blanks). It can have the same name as a single Agent in the workspace, or an Agent in another group.

The new group will appear in the left pane of CentralControl. Right-click on the group name to expand the view to see all the Agents listed.

You can also delete the group name only, without affecting any existing Agents. You can refresh the list, and you can add or delete Agents from the group (that is, move Agents in and out of the group).

3.6.2 Deleting Groups

Deleting a group removes the Agent group name entry from your local workspace. All Agent names in the group are also deleted. If the names were also in the workspace or in another group, they do not get deleted from there. Any data on the Agent computer (the server that gets backed up) is not affected.

3.6.3 Importing Agents to a Group

When you right-click on a group name, you can choose **Import Agents** from the menu. Select this option to use the lines of a CSV file to populate the **New Agent** screen.

The CSV file allows you to automate the process of adding many Agents for which much of the information is duplicated (while allowing for any differences to be entered manually).

Manually create and save a text file. The **Import Agents** function will read the file and populate new Agents with the information from there. The first field (**Network Address**) is required. If other fields are missing, the **Agent Properties** screen will prompt you for values.

The fields are positional. If you do not enter one, but you want to enter the next one, you still need the commas to indicate the order/position of that field. See [CSV File Format](#).

Import Agent reads the CSV file sequentially. If there are any errors, they will be reported in a log file in the **Logs** folder under the group name.

If any of the fields in the CSV file are not supplied, the **Import Agent** process will stop and prompt you for a response.

3.6.4 CSV File Format

The CSV file is created with a name of your choosing. It can have from one to six fields per line, representing (in this order):

- Network address
- Description
- Port
- User name
- Password
- Domain

Notice that in the CSV file the Network Address comes first, and the Description comes second. The Network Address is required, where the other five fields are not always required.

Each line represents one Agent to be added to the group. You may “skip over” field entries after the first one in a line, or not mention them at all. In that case you will be prompted for the entry on the Imported Agent Properties screen.

You do not need to enter a final Enter key (CR/LF) after the last entry. If you do, it will show as an error in the log: “Can’t find an address in line <#>.”, which can be ignored in this case.

The system expects to have a password. If you skip over that field, it will prompt you for a password, even if you want the password to be null. You may use the “Use as default for all the following imported Agents” box to automatically assign a null password to all the following Agents.

A typical CSV line in a file might be:

```
192.168.5.211, Agent_1 in Group_B, 808, User211, Password, corp.company.com
```

or

```
192.168.5.211
```

or

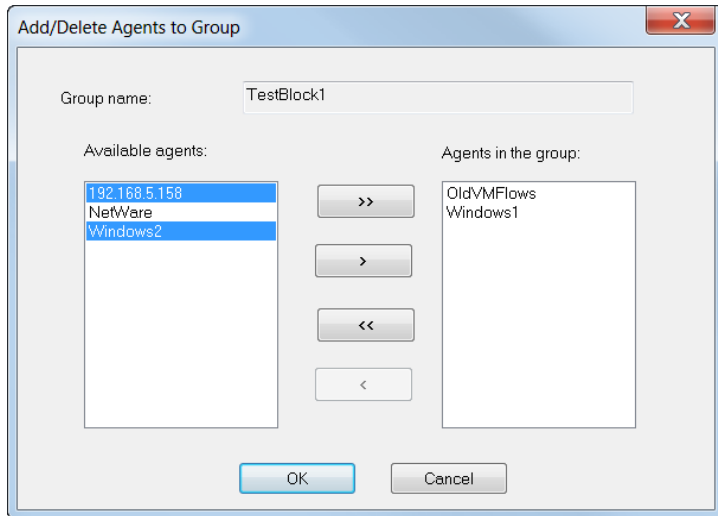
```
192.168.5.211, Agent_1 in Group_B, 808, User211
```

or

```
192.168.5.211, , 808, , , corp.company.com
```

3.6.5 Adding and Removing Agents from Groups

Open the Add/Delete Agents to Group screen by right clicking, or by using the **F2** key when a group is selected in CentralControl. This screen allows you to move one or more Agents to or from the group.



The << and >> buttons will move all of the Agents across from one side to the other, regardless of what is selected. The < and > buttons will move one or more selected Agents across. You can select multiple Agents by clicking on them while holding down the **Ctrl** key.

4 Managing Vault Profiles

If an Agent is not registered to Portal, you can use CentralControl to add and modify vault profiles for the Agent.

If an Agent is registered to Portal, the Agent's vault profiles are read-only in CentralControl. You can view the Agent's vault profiles in CentralControl, but must use Portal to add and edit vault settings.

4.1 Adding a Vault Profile

Highlight an Agent, and select **Tools > Agent configuration > Vaults**.

The Agent's vault profiles, if any, appear on the Vaults tab of the Agent Configuration dialog box.

If the Agent is registered to Portal, the following message appears on the Vaults tab: *Vault profiles for this Agent are managed through Portal*. If this message appears, you can view the Agent's vault profiles in CentralControl, but must use Portal to add vault settings.

If you can manage the Agent's vault profiles using CentralControl, click **New**.

The Vault Configuration Wizard opens. On the Welcome screen, click **Next**. The Vault Registration screen opens.

If this is a new vault, you must register it as a new computer.

You can also create a new vault connection for a previously registered computer. Select **Re-register previously registered computer** to reregister the computer with the vault. This will retrieve all jobs from the vault, so you might need to wait a minute or longer.

To change the Agent host name, see [Editable Host Names](#).

4.1.1 Profile Name

Enter a profile name for the new vault. This name will be used for all interactions with the newly created vault.

4.1.2 Vault Network Addresses

Enter a network host name or an IP address.

For a multi-home server, you can enter multiple network addresses. In other words, if the vault has multiple network cards and IPs, you can list the addresses here in order of priority.

If you want to change priorities within the list, there are buttons to help you.

4.1.3 Ports

Backup software benefits from long, secure sessions of activity.

Because network transports such as TCP do not provide sufficiently stable connectivity, Backup Restore Transport Protocol (BRTTP) supplies network connectivity resumption support. In the event of a network failure, a backup or recovery job can resume its functions from where it left off.

In the **New ports** field, enter the port values at which the new vault will listen. The pane below the field lists any currently configured ports.

The default port for new installations is **2546**. For upgrades, port **807** is used. Should you create multiple ports, port connections will be attempted in descending order as listed in the pane.

The software can be configured to use the HTTP (**80**) and FTP (**21**) ports instead of the traditional BRTTP port (**807** or **2546**). These ports are often open to outbound connectivity with remote systems because many companies permit browsers on desktops.

Alternatively, any other port can be used, but you need your service provider's approval. The highest allowable port entry is **65536**.

If the port you select is used by another service, a message notifies you, and offers the opportunity to continue or use another port.

Click **Next** to proceed to the **Connection Settings** screen.

Note: If you want to add a new port, you might need to discuss the options with your service provider.

4.1.4 Connection Settings

The default connection/reconnection settings have already been optimized for most network situations. If you wish to change them, follow these steps:

In the **Try to reconnect every** field, enter the length of the time intervals after which your system should try to re-establish connection with the server. This value should range from 3 to 3600.

In the **Stop reconnection attempts after** field, enter the number of minutes after which your system should stop trying to reconnect before terminating the job. This value should range from 1 to 1440.

Note: In both fields above, enter whole numbers (with no fractions, decimals or commas).

Over The Wire Encryption Settings: Check this box to enable OTW encryption for transmission to and from the vault. The default mode is "enabled" using AES encryption. This setting applies to all jobs that use this profile.

You can turn the encryption off if security is not an issue. This gives you slightly more speed/efficiency for backup and recovery operations.

Click **Next** to continue.

4.1.5 Authentication

On the Authentication screen, enter the account, user name and password as supplied by your service provider. Note that the password is case-sensitive.

If you need to change the Agent host name (to avoid conflicts, for example), click the **Advanced** button. This opens the Agent Host Name screen. Enable **Change Agent host name**, and modify the name in the **Agent Host Name** field. This information is case-sensitive. Also see [Editable Host Names](#).

Click **Next**. On the following screen, click **Finish**.

4.2 Modifying a Vault Profile

To modify an existing vault:

On the **Vaults** tab of the Agent Configuration screen, select a vault.

If the Agent is registered to Portal, the following message appears on the Vaults tab: *Vault profiles for this Agent are managed through Portal*. If this message appears, you can view the Agent's vault profiles in CentralControl, but must use Portal to modify the vault profiles.

If you can manage the Agent's vault profiles using CentralControl, click **Edit**.

The Vault Entry screen opens.

The Vault Entry screen provides fields that contain the existing vault settings. In most cases, you can change the settings in the fields.

4.2.1 General Tab

Also see [Modifying a Vault](#).

Vault name: Use the existing name, or enter a new one.

Network address(es): Use the existing addresses, change them, or add to them. Use semicolons to separate multiple addresses. All of the addresses must apply to the same vault.

4.2.2 Authentication Tab

Change or use the existing credentials as displayed. Note that the password is case-sensitive.

Agent Host Name is read-only here.

4.2.3 Connectivity Tab

The default connection/reconnection settings have already been optimized for most network situations. If you wish to change them, follow these steps:

In the **Retry up to** field, enter the amount of time after which your system should stop trying to reconnect before terminating the job.

In the **Retry every** field, enter the length of the time intervals after which your system should try to re-establish connection with the server.

Note: In both fields above, you must enter whole numbers (with no fractions, decimals or commas).

Over The Wire Encryption: Check this box to enable OTW encryption for transmission to and from the vault. The default mode is "enabled" using AES encryption. This setting applies to all jobs that use this profile.

You can turn the encryption off if security is not an issue. This gives you slightly more speed/efficiency for backup and recovery operations.

4.2.4 Port Tab

The pane here lists any currently configured ports. To add ports at which the vault should listen, click **New**.

See [Ports](#) for additional information.

4.3 Copying a Job to Another Vault

The vault operator may determine that a job needs to move to another vault because of space problems, failing hardware, or other technical or logistical problems.

If a job is copied (moved) on the same vault, it is not obvious because the unique Job ID is retained. But a job on a new vault will generate a new Job ID.

So if a job is copied to another vault, you must do two things from CentralControl to allow the job to work properly:

- Reregister the Agent (to recover the jobs)
- Manually change the IP address (and user name/password for the vault registration, if necessary) to point to the new vault

5 Creating and Managing Jobs

Each Agent manages a collection of jobs. A job defines default parameters to associate with a backup or recovery. Job parameters can include file selections, compression and encryption settings, log file locations, and advanced settings.

[To create a job](#), select (highlight) an Agent, and choose **File > New Job**.

Alternatively, right-click an Agent, and select **New Job**.

You can see the properties of an existing job through **File > Properties**. You must have a valid job selected.

5.1 Creating a Job

Once you have created an Agent, you can create a job. A job is initially made using the New Job Wizard. Right-click on an Agent. Choose **New Job**. This starts the New Job Wizard. On the **Welcome** screen, click **Next**.

Following are the main steps for job creation. They are described further in this chapter, and also in the appropriate Agent guides.

1. **Welcome** screen: This may not display if you have previously selected **Skip this screen in the future**. Click **Next**.
2. **Backup Source Type** screen: Choose the appropriate source type from the menu. *For simplicity*, these condensed instructions are based on choosing **Local Drive Only**. Click **Next**.
3. Select a destination vault, or click **New** to set up a new vault. Click **Next**.
4. Enter a **job name**. This can be up to 30 characters in length.
5. **Source** screen: This is where you select files and directories to back up. You can also include **System State** and other items that appear. If you select **Data Files**, and click **Add**, you are shown your local drives. Include or exclude as many directories and files as you want. Click **OK** when you complete a selection, choose more sources if you want, and finally click **Next**.

Note: You may see a **Bare Metal Restore** option in this list. You need a license to run this type of backup. See [BMR with System Restore](#) for more information.

6. **Options** screen: Choose options for the job.
 - Quick file scanning. Quick file scanning reduces the amount of data read during the backup process. Any file streams that are deemed unchanged since the last backup are skipped. To read files in their entirety, disable the Quick file scanning option.
 - Disable deferring. Normally, when the backup time window expires, all jobs that have not yet been backed up are put on hold until the next backup time window. This process is known as deferring. When deferring is disabled for a specific job (i.e., the Disable deferring option is turned on), the job is backed up in its entirety, even if this means extending the backup beyond the prescribed backup time window.

- Backup time window. You can set a time frame in which a backup occurs. Any new files are deferred until the next backup if they are not backed up within the specified window. Adjust the Backup time window to allow a full backup grow over a period of days or even weeks. The Backup time window setting must be at least 15 minutes, and no more than 48 hours.

If you prefer, disable deferring, and let the backup run from start to finish. This allows the backup to run (but more slowly) during times when other users need the bandwidth.

7. **Encryption** screen: For new backup jobs, AES (256-bit) is the **Encryption type** for data on the vault. Enter a password and password hint.

Note: If an existing job uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES, None), 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.

Important: To recover data, you must remember the encryption password. If you forget the encryption password for a job, you cannot restore data from the job.

8. **Log Options** screen: Choose the level of logging you want, and the number of logs to keep.
9. **Finished** At the end of the process, you can exit from the wizard, or you can schedule or run the job immediately.

5.2 Backup Source Type

The Backup source type menu might include **Local Drive Only**, **Network UNC Share**, **VMware vSphere**, and specific plug-in source types, depending on the plug-ins installed on the Agent. Additional options might also display depending on the source type that you select.

Select the source of the data that will be backed up by this job. Your selection will affect the dialogs that follow.

5.2.1 Local Drive Only

Local Drive Only allows you to include **Data Files**, and, optionally, **Bare Metal Restore**, **System State**, **RSM database**, and **Event logs**. You will only be able to back up from the drives that your local computer can access (for example, C:).

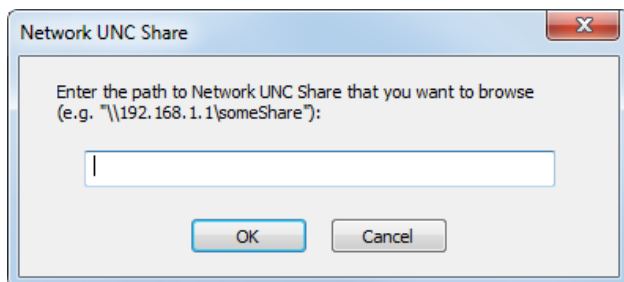
5.2.2 Network UNC Share for Windows Agents

The **Network UNC Share** source type allows you to browse and back up network resources via UNC paths (for Windows Agents only). The UNC share will be the root for the browse, instead of a device letter. Using UNC gives you more flexibility than mapped network drives for remote files and directories.

After you have chosen a destination and a job name, you must provide network account information that will enable the backup process to "Run As" those credentials.

This source type only allows you to select files and directories.

When you reach the Source screen, the share appears as the root of the hierarchy, rather than a drive letter. When you click **Add**, you are prompted for the path to the share that you want (for example: \\197.168.1.2\MyShare). You must add at least one item.



If the path or network credentials are not valid, an error message will result.

Now you can select files and directories to include/exclude, as well as other options.

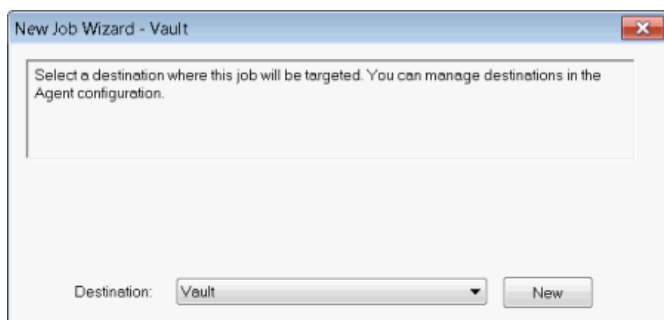
5.2.3 VMware vSphere

The **VMware vSphere** source type allows you to include virtual machines, virtual disks, and template files in [vSphere Agent](#) backups.

5.3 Destination/Vault

Select a destination for the backup. Click **Next**.

You can also set up a new vault from here.



5.4 New Job Name

On the New Job Name screen, enter a name for the job. You can also add a job description if the Agent is version 6.9 or later.

The name must be 1-30 characters in length. It must consist of letters (A-Z and a-z), numbers (0-9) and/or `_`, `-`, `$` (underscore, hyphen, dollar sign).

The following names cannot be used as job names in connection with an Agent: **PRN**, **CON**, **LPT1**, **LPT2**, **LPT3**, **LPT4**, **COM1**, **COM2**, **COM3**, **COM4**, **NUL**, **AUX**, **Register**, and **Global**.

(**Register** and **Global** each result in a message that the job name already exists, even though you did not create it yourself. The other prohibited names result in messages that those job names are reserved by the system.)

5.5 Data Sources

5.5.1 Conventional Agents

5.5.1.1 Source Screen

Select an item to back up. You must select at least one item.

As available, you can select:

- **Data Files** to back up files. When you click **Add**, the **Include/Exclude** screen opens. This allows you to choose which files to back up.
- **Bare Metal Restore** to create a BMR-type backup
- **System State** to add system-specific information to your backup.
Note: Vault profiles, schedules, and jobs that are associated with a physical node in a cluster can have the system volume and System State backed up and restored.
- **RSM database** to add information about the Removable Storage Manager
- **Event logs** to back up the system logs
- **VMware vSphere** to back up virtual machines. When you click **Add**, the **Include/Exclude** screen opens. This allows you to choose which VMs to back up.

Plug-in types might be preselected (e.g., for [SQL Server](#) or [Exchange](#)), depending on workflows and what is installed on the Agent.

5.5.1.2 Include/Exclude

For information about the **Include/Exclude** screen for Virtual Machines, click [here](#).

For information about **Include/Exclude** screens for Exchange plug-ins, click [here](#).

If you select **Data Files** on the **Source** screen, the **Add** button opens the **Include/Exclude** screen. From here, browse for files or directories and select them. To include a file or directory, highlight it and click **Include** (or double-click it). To exclude a file or directory, highlight it and click **Exclude**.

You can use the **Ctrl** and **Shift** keys to help you make your selections. You must include at least one file in your backup or recovery.

If you choose an entire directory, the **Confirm Include** screen appears. Select **Recursive** if you want to include/exclude all subdirectories and files.

Otherwise, filter to specify the files that you want. Use filtering formats similar to the following examples.

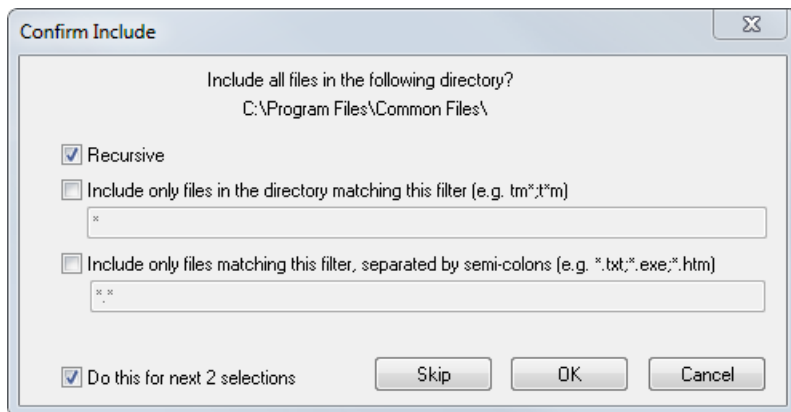
Syntax	Meaning
sample.txt	Include the file called sample.txt
*.txt	Include all .txt files
.	Include all files in the selected directory

If you enter a list of filters, use semicolons to separate them:

```
*.txt;*.exe;*.htm
```

Note: Wildcard searches do not work for file names that contain semicolons.

The **Confirm Include** and similar screens allow you to apply their settings to some or all of the directories that you have chosen:



After you finish making your Include/Exclude selections, click **OK**. The **Source** screen returns. Click **Next** to continue.

Exclusion Tree Scanning

The selection with the most detail takes precedence. This allows for nesting of multiple inclusions and exclusions. In these cases, the filtering avoids traversing excluded directories.

Folder/Container Wildcards

Wildcards are supported on file folders for use during backup only.

* (asterisk) - signifies a wildcard string up to the next separator character.

? (question mark) - signifies a single wildcard character.

. (period) signifies a recursive directory.

Examples

*.txt - selects all files that have the txt file type.

C:\.*.doc - selects all files on drive C that have the doc file type.

t*.* - selects all files whose names begin with the letter t.

Note: vSphere Agents do not support multiple filters in a horizontal sequence (e.g., *.*.txt; *.*.exe; *.*.htm). To apply several filters, use the **Confirm VMs to Include** screen several times (adding only one filter each time).

In Windows, files or directories that contain reparse points display in red.

5.5.1.3 File Backup Options

If you select **Data Files > Options**, you can choose options for file backups.

Back up files opened for write

This option allows you to back up files that are open while the backup is in progress. Files that are open for write, and also opened for shared read, can be backed up. However, files open only for exclusive write cannot be backed up.

The danger in backing up files that are open is that you are not guaranteed to get a copy of the file at that time. Therefore, an open file might be modified during the backup process, producing inconsistencies in the backup copy.

Suppress archive bit processing

When a file is created or modified, an archive attribute (depending on the OS) can sometimes be placed in the file. For some programs, the archive bit indicates that the file requires backup. By default, the application clears the archive bit when it backs up a file.

If you use other backup programs that rely on archive bits, enable **Suppress archive bit processing**. This leaves the archive bit intact.

UNIX options

If you choose to back up a single instance of the selected hard-linked files, the backup is slower, but smaller, as only the actual data (one copy) and the hard links are backed up. During recovery, the data and all the hard links are restored.

If you do not enable the **Back up a single instance** option, the backup is faster, but the backup size may be larger. This is because the data and each link are copied separately, and restored separately. The hard links are not reestablished in this case.

5.5.1.4 System State – Windows

If **System State** is available for selection, you can choose it to add system-specific information to your backup.

Note: Vault profiles, schedules, and jobs that are associated with a physical node in a cluster can have the System Volume (SYSVOL) folder and System State backed up and restored.

If available, you can select **RSM database** to add information about the Removable Storage Manager.

Objects that are backed up in System State backups can include:

- Registry
- COM+ Class Registration database
- Boot files
- Windows System Files
- Performance Counter
- .NET Framework
- IIS Metabase
- Certificate Services database
- Active Directory
- SYSVOL folder
- Cluster database and Quorum disk
- Terminal Services database
- Volume quotas
- WMI-specific state and data
- DHCP database
- WINS database
- Windows System Files

Including the system files allows you to:

- Recover from corruption of system files
- Recover from accidentally uninstalling service packs
- Perform a bare metal restore

Including the system files also allows you to return to the state at backup time without needing to reinstall the OS from the installation kit, and then needing to install each service pack separately.

You should include the system files in a backup whenever you modify the operating system.

5.5.1.5 Event Logs

Event log databases store events that are viewed by using the Windows Event Viewer program. The Event log service, located in the Event Viewer, records events in the system, security and application logs. These events include services starting and stopping, users logging on and off and accessing resources.

5.5.1.6 IIS Metabase

The backing up of IIS involves the backup of the IIS Metabase. The IIS Metabase is a database similar in structure to the Windows Registry. The IIS Metabase is optimized for IIS. It provides hierarchical storage and fast retrieval of IIS configuration properties for websites, virtual directories, FTP Sites, SMTP and NNTP sites.

If you have IIS on your server, you should see this source option.

5.5.2 vSphere Agent

5.5.2.1 ESX Operations and Credentials

Credentials

Enter your credentials, and then click **Next**.

The credentials are verified when you use the Source screen, and also when the backup job actually runs.

5.5.2.2 Source Screen

Click **Add**. The **Include/Exclude** screen will open, allowing you to choose which items to back up.

You must select at least one item.

5.5.2.3 Include/Exclude VMs

On the **Source** screen, click **Add** to open the **Include/Exclude** screen. Expand **Virtual Machines** to browse for items to select. (If you wish, you can select the entire **Virtual Machines** server.)

To include an item, highlight it and click **Include** (or double-click it). To exclude an item, highlight it and click **Exclude**.

You can use the **Ctrl** and **Shift** keys to help you make your selections. You must include at least one item in your backup or recovery.

If you try to include the entire **Virtual Machines** server, the **Confirm VMs to Include** screen opens. You can include all of the VMs, or you can filter them according to their names. To see examples of the filtering scheme, click [here](#).

Notes

- Wildcard searches do not work for VM names that contain semicolons.
- vSphere Agents do not support multiple filters in a horizontal sequence (e.g., `*.txt;*.exe;*.htm`). To apply several filters, use the **Include** button and **Confirm VMs to Include** screen several times (adding only one filter each time).

After you finish making your **Include/Exclude** selections, click **OK**. The names of the VMs that you have chosen will appear in the **Virtual Machine** pane of the **Source** screen.

5.6 Encryption

For new backup jobs, AES (256-bit) is the Encryption type for data on the vault. Enter a password and password hint. The password is case-sensitive.

If an existing job uses another encryption type (e.g., AES 128-bit, Blowfish, DES, Triple DES, None), 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.

To recover data, you must remember the encryption password. If you forget the encryption password for a job, you cannot restore data from the job.

Note: If you change your password or any other encryption options, the next backup will reseed.

5.7 Local Catalog Files

Local catalog files are created automatically. A catalog file is used to view and analyze your backup files. Also, it is used to select your backup files for recovery without needing to load a tape or connect to a remote server.

A local catalog is created in a safeset subdirectory in the product's working directory. Its name is `xxxxxxx.cat`, where the `x` characters represent the number of the backup. The first backup is created with the number **0000001.cat**.

Note: A catalog cannot exceed the system limit of 4 GB. This represents a catalog with 2 to 3 million files, or more.

5.8 Log Files

Each time a backup or recovery runs, a job log is created (unless logging is turned off). These logs contain different levels of detail, depending on the job settings. Log files can help you to investigate backup and recovery failures.

You can configure jobs to generate log files containing start-connect-completion and disconnect times, file names (i.e., the names of all files that were copied during a backup process), and any processing errors.

The activities of a successful backup job are recorded in a file called **zzzzzzz.XLOG**, where the `z` characters form the number of the backup (i.e., the safeset number). For example, the first backup will generate a log file named **0000001.XLOG**. Activities of a failed backup will be noted in a file called **BACKUP.XLOG**. This log is useful for troubleshooting.

A successful recovery log is named according to **RSTYYYYMMDD-HHMMSS.XLOG**. That is, the date and time are used as the file name to make the name unique.

Note: Older Agents use **.log** for the file extension.

5.8.1 Creating Job Log Files

You can set your log file options when you create a job. You can also modify these options through the **Log** tab (**Job Properties**).

More detailed logging creates larger log files, and is normally used only for troubleshooting problems.

Changing the detail level only affects log files that are created from that point on. It does not affect any previously created log files.

To create log files from the Job Wizard:

1. On the **Log** screen in the Job Wizard, select **Create log file**.
1. To conserve disk space, the application allows you to specify how much information to display in the log file. From the **Log detail level** menu, choose **None**, **Files**, **Directories** or **Summary** to specify the amount of detail required in your log files. For example, the **Files** option would be used for a recovery, but the **Directories** and/or **Summary** options would be used for backups.
2. To further conserve disk space, the application allows you to specify the number of log copies to save. Select **Automatically purge expired log files only** to have all files saved. To specify a specific number of files to keep, enter a number in the **Keep the last <X> log files** field.
3. Click **Next** and complete the Job Wizard.

Note: For recovery jobs performed through Agents, full recovery logs with file-level logging are generated, regardless of which log options are selected.

5.8.2 Viewing Log Files

You can view log files by opening the **Logs** folder. Log folders are automatically created for each job. Log files can be identified by relating the file name to the action performed.

Backup log files are renamed to the catalog number created if the backup completes successfully (without a serious error). If the backup job fails, its log is called **BACKUP.XLOG**. Double-click the log file to open it for viewing.

Process	File	Description
Backup	BACKUP.XLOG	Log of backup job. Renamed to <catnumber>.XLOG upon backup completion without serious errors.
Restore	RESTORE.XLOG	Recovery log. Renamed to RSTYYYYMMDD-HHMMSS.XLOG upon a successful recovery.
Synch	SYNCH.XLOG	Log of Synchronize operation

The default log viewer is called LogViewer. It can display the logs in a number of languages. This is independent of which language was used for the Agent installation.

Older Agents do not use the **.XLOG** format, so you cannot view their logs in multiple languages.

5.8.3 Service Logs

When a Windows Agent runs, a **VVAgent-#** file is created. The number sign (#) is replaced by the numerals 1 through 4, in rotation. Logging switches to the next file when the current file reaches about 32 KB in size.

The older logs are kept for reference only. The logs, which are kept on the Agent computer in the Agent installation directory, contain service start and stop information, and any errors (e.g., an error about an inability to start a scheduled backup).

This means that there are a maximum of four log files maintained called **VVAgent-1**, **VVAgent-2**, **VVAgent-3**, and **VVAgent-4**. Each program appends to a log until the log file reaches its maximum size. At that point, the program will write to the next file in the series. The old data there is removed first. When the fourth file has been filled, **VVAgent-1** is used again.

The log files can be viewed from the rightmost pane (double-click on one of the file names) when you have an Agent selected.

5.8.4 Automatically purge expired log files only

You have a choice of either automatically purging expired log files, or keeping a selected number of them before they get deleted (oldest one first).

Automatic purging will delete the log file associated with a backup when that backup (safeset) is deleted.

5.8.5 Keep the last <number of> log files

You can select a number here, whereby the system will keep a certain number of log files for the backup. When that number is reached, the oldest log file will be deleted to make room for the newest one.

Note: The default choice (to automatically purge expired log files) means that when a safeset expires, the log files will also be deleted. If you want to delete log files before the safeset expires, you can choose to keep the last “X” number of days of logs. In both cases you cannot keep the logs for longer than the safeset unless you copy them manually.

5.8.6 Example email log

This is an example of a “good” log file emailed from a normal backup to show the meaning of the lines in the log. These email notices are optional. The actual (full) log file is found under the job name in CentralControl.

Note: Line numbers have been added for this example only. Normally there are no line numbers.

01. Agent: EXAMPLE1
02. Job: BACKUP01
03. Retention: daily
04. Job start time: 05-Oct-2010 20:30:14 -0400
05. Job end time: 05-Oct-2010 21:46:01 -0400
06. Elapsed Time: 01:15:47
07. SafeSet: 00000389

08. Errors encountered: 0
09. Warnings encountered: 0
10. Files/directories examined: 491,960
11. Files/directories filtered: 40,973
12. Files/directories deferred: 0
13. Files/directories backed-up: 450,967
14. Files backed-up: 425,820
15. Directories backed-up: 25,147
16. Data stream bytes processed: 85,664,629,239 (79.78 GB)
17. All stream bytes processed: 85,752,048,987 (79.86 GB)
18. Pre-delta bytes processed: 6,649,699,821 (6.19 GB)
19. Deltized bytes processed: 830,521,523 (792.04 MB)
20. Compressed bytes processed: 527,557,613 (503.11 MB)
21. Approximate bytes deferred: 0 (0 bytes)
22. Reconnections on recv fail: 0
23. Reconnections on send fail: 0

01. This is the name that you (i.e., the Administrator) have assigned to this particular computer that will be backing up.

02. This is the name that you have assigned to a job on this computer (Agent). There may be more than one job associated with an Agent. Each job will create separate email notices and log files.

03. This is the retention schedule that was chosen for this job.

04-05. The job's start/end time comes from the Agent. CentralControl and/or the vault may be in different time zones. (In this example, -0400 shows the GMT offset.)

06. Elapsed time (start to end) is the total time that the backup took.

07. Each successive backup creates a new safeset number, starting from 00000001. This does not necessarily mean that there are X number of safesets on the vault (X is 389 in this example). Depending on retention settings, there may be fewer.

08-09. Number of warnings and errors encountered. As well, the actual log files will list the warnings and errors. Different Log detail levels will give you different log files. With None, you will still get a short log with warnings and errors. With Files and/or Directories, you will see all the files and directories used for your backup, plus the warnings and errors. With Summary, you will see the totals for your backup job, plus the warnings and errors. Note that a job may have warnings and errors, but still complete. Also note that if a job is not successful (that is, it fails for some reason), the backup is not completed. There is never a "partial" backup. It must be completed and committed before it can be considered successful.

10-11. Files/Directories examined/filtered: You select the files and directories that you want to back up. Your selection may include nesting of multiple inclusions and exclusions. In these cases, the filtering will avoid traversing excluded directories.

12. Files/Directories deferred: If you have enabled deferring, and chosen a backup time window, the backup will be deferred after that time if it has not completed. It will continue at the next backup time (usually according to a schedule). This line shows how many files and directories are waiting to be backed up.

13-15. This shows the total number of files/directories that were backed up.

16-17. Data stream/All stream bytes processed: Depending on the operating system, some files include other streams along with the data stream (e.g., Security streams, Extended Attributes, and Access Control Lists). You cannot include or exclude these at backup time, but you might be able to exclude them from a recovery.

Note: If this backup is an initial seed, the numbers in lines 16 and 17 will be approximately the same. That is, all files will be examined, and all blocks considered “new” (no deltas). (See the partial log example that follows.)

18. Pre-delta bytes processed: The Agent has examined (with Quick File Scanning) all the files and this number is the total size of all files that have been flagged as changed by QFS.

19. Deltized bytes processed: If QFS finds any files that have changed (size, date, attributes), the Agent examines them further, looking for changed blocks. Only those changed (delta) blocks (32 KB each) will be included in the current backup.

20. These deltized bytes are then compressed and sent to the vault. This number could be the same or smaller than “All stream bytes processed”.

21. Bytes deferred: If the backup could not complete in the time allocated, the remaining portion (bytes deferred) will be included in the next backup time window. Note that if backups keep getting deferred, with more and more data not being backed up, your system may not be protected properly. You should allow sufficient time and resources for the backups to complete.

22-23. Reconnections on Receive/Send Fail: If you temporarily lose communication with the vault while a backup is running, the system will recover and continue, provided that the backup is not too lengthy, and there are not too many failures. You can change parameters to decide how long and how many times to wait before a backup should fail.

This partial log (see above) comes from a seed (initial) backup, so there are no deltas:

16. Data stream bytes processed:	108,715 (106.16 KB)
17. All stream bytes processed:	108,903 (106.35 KB)
18. Pre-delta bytes processed:	108,903 (106.35 KB)
19. Deltized bytes processed:	108,903 (106.35 KB)
20. Compressed bytes processed:	66,320 (64.76 KB)

5.9 Finished Screen/Options

The last screen in the New Job Wizard gives you three choices regarding what to do with your new job.

Run the job immediately

The backup job will be saved, and the Backup Wizard will be launched. This is also called an “ad hoc” backup, a single backup for a specific reason. The job is saved, though, and may be reused or scheduled.

You can also select scheduled jobs and run them ad hoc (immediately).

Schedule the job

This opens the Schedule List screen so that you can schedule the job. See [Schedule List](#) for more information.

Just exit from this wizard

This option allows you to leave the New Job Wizard, but retain the job. You can return afterwards to run it directly, or schedule it for later.

If you choose this option, remember to click **Finish** to save your job.

5.10 Existing Jobs

5.10.1 Existing Jobs - General tab

The **Job Properties** function allows you to change parameters that were previously created for a job.

To use it, select a job, and then choose **File > Properties**, or press **F2**. The **General** tab of the **Job Properties** screen will open.

Destination pane

This displays the name of the vault with which this job is registered. It is the vault that will receive the backup.

Job description pane

You can add a job description for Agent versions 6.9 and later.

Options pane

Retention

You can choose a retention from the menu.

Compression type

Compression allows you to optimize the volume of data sent versus the speed of transmission. In some cases it might be better to take the time and CPU cycles to compress the data before sending it at a slower rate, as opposed to not compressing it and sending it at a faster rate.

Also, compression reduces the space required to store the data on the vault.

For Agent versions 7.5 and later, the following compression types are available:

- **Faster:** Minimizes the amount of time that is required for backing up the data.
- **Smaller:** Minimizes the size of the backup data, but can take longer to process the data.

For previous Agent versions, jobs can have the following compression types:

- **None:** Do not compress any data.
- **Minimum:** Minimize CPU consumption, possibly at the expense of a larger safeset size.
- **Normal:** Balance CPU consumption against safeset size.
- **Better:** Minimize safeset size, possibly at the expense of extra CPU consumption.
- **Maximum:** Always minimize safeset size, regardless of CPU consumption.

5.10.2 Advanced tab

This tab allows you to work with options such as **Quick file scanning**, **Disable deferring**, and **Backup time window**.

Disable deferring is not available for SQL legacy jobs.

5.10.3 Log tab

This tab allows you to select options for creating and retaining [job logs](#).

5.10.4 Credentials tab

Credentials

Enter your credentials, and then click **Next**.

The credentials are verified at this time, and also when the backup job actually runs.

5.10.5 Source tab

To specify what to back up, select an item.

As available, you can select:

- **Data Files** to back up files. When you click **Add**, the **Include/Exclude** screen opens. This allows you to specify which files to back up.
- **Bare Metal Restore** to create a BMR-type backup
- **System State** to add system-specific information to your backup.
Note: Vault profiles, schedules, and jobs that are associated with a physical node in a cluster can have the system volume and System State backed up and restored.
- **RSM database** to add information about the Removable Storage Manager
- **Event logs** to back up the system logs

- **VMware vSphere** to back up virtual machines. When you click **Add**, the **Include/Exclude** screen opens. This allows you to choose which VMs to back up.
- Plug-in items (e.g., for Exchange data)

5.10.6 Network UNC Share tab

You must provide network account information that will enable the backup process to "Run As" those credentials.

See also [Network UNC Share for Windows Agents](#).

5.10.7 Encryption tab

For new backup jobs, AES (256-bit) is the Encryption type for data on the vault.

If an existing job uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES, None), 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.

Note: If you change any encryption options, the next backup will reseed. This will occupy more space on the vault because all of your job data will again be sent to the vault using the new encryption password. (In other words, you will not have a deltized set of data.) Your retention settings will determine when older versions of the job (with the previous passwords) will expire. When older jobs expire, their space on the vault becomes available.

You must remember all of your passwords (and the backups to which they belong) because there is no way to retrieve them from the system (Agent or vault).

The **Password Hint** field is available for Agent versions 6.9 and later. Any password hint that you may have entered when you backed up only applies to the last password that you used. If you changed passwords before that, their hints will not display, so you must correctly remember those previous passwords.

5.11 Monitoring Jobs

After a backup or recovery has started, CentralControl begins to monitor the progress of the job. The **Process Information** screen indicates the activity being performed.

When you monitor a backup or recovery process, file names display with a percent-completed indicator. (This applies to large files only.)

Completion status is provided so that the application can detect whether the job is successful or not.

5.11.1 Viewing Process Information

The **Process Information** screen (Job status) can be viewed while a scheduled job is running. To view the process information while a job is running:

1. Expand the target Agent.
2. Open the **Processes** folder.

3. Click on the target job. This opens the **Process Information** screen. The progress of the job can be viewed from here.

5.11.2 Stopping a Job

You can halt a job while it is running. To stop a job while it is running:

1. Expand the target Agent.
2. Open the **Processes** folder.
3. Select the target job. The **Process Information** screen opens.
4. Click **Stop Process**.

If this does not work (i.e., the process does not stop gracefully), you can click **Terminate Process** if necessary.

5.11.3 Removing an Entry from the Processes Folder

When a command completes, its status displays in the **Processes** folder, and the **Stop Process** button changes to **Delete Entry**. Jobs are visible through the **Processes** folder for up to an hour after their completion. This provides an opportunity to check the command's final status without opening log files.

All entries are automatically removed after an hour, but can be manually deleted at any time. To remove an entry:

1. Expand the target Agent.
2. Open the **Processes** folder.
3. Click on the target job. The **Process Information** screen opens.
4. Click **Delete Entry**.

6 Creating and Managing Schedule Entries

In order to perform regular automated commands, CentralControl provides an integrated scheduling feature. Backups and synchronizations can be scheduled for any or all of your Agent jobs from the CentralControl application. New schedules are created using the Schedule Wizard. This allows many useful predefined options that support the most common scheduling scenarios.

Additionally, customized commands and time schedules can be entered to help support more sophisticated scheduling requirements. Existing schedules can be modified from the Schedule Entry screen.

Each Agent system monitors the schedule file for changes, and prepares to initiate scheduled commands at the specified times. This program is started or loaded at system start time on all Agent systems.

To start, click **Tools > Schedule Entries**, or right-click on the Agent.

6.1 Schedule List

Also see [Working with Schedule Entries](#).

The Schedule List shows a summary of currently schedule entries. To view or modify an entry, double-click the entry, or select the entry and click **Edit**. Each entry is summarized using these headings:

- **Command:** Name of command scheduled to run
- **Job:** Name of job for this command
- **Retention:** Name of retention for this command
- **Time:** Scheduled time to run this command
- **Days:** Scheduled days to run this command

6.2 Creating a New Schedule Entry

You can add a schedule entry when you create a job. You can also add a schedule entry directly to the Schedule file.

1. With the target Agent highlighted, click the Schedule file. The Schedule List screen opens, displaying all scheduled jobs.
2. Click **New** to launch the Schedule Wizard. The wizard guides you through the process of creating a schedule for a specific job.

6.2.1 Command Screen

The **Command** screen of the Schedule Wizard offers a selection of actions to schedule:

- **Backup** - Create a schedule to back up data for a specified job.
- **Synchronize** - Create a schedule to update local catalogs with server information for a specified job.

- **Custom command** - Advanced users can create a schedule to run custom commands. Custom commands can be any third party executable.

For Windows, such a command could be: `vv Backup Jobname
</param=parameter.files>`

6.2.2 Job List

The **Job List** screen of the Schedule Wizard shows the list of jobs in your current workspace.

Select the job on which you wish to perform a command, and click **Next**.

6.2.3 Options Screen

Quick file scanning reduces the amount of data read during the backup process. Any file streams unchanged since the last backup are skipped. If this option is not selected, files are read in their entirety.

You may also enter **Backup time options**. This is the duration of time you want your backup to run. The default backup window is eight hours. The backup will terminate at the end of this time period, regardless of whether the process is complete or not. Any files not backed up are deferred until the next scheduled backup.

When **Disable deferring** is selected for a specific job, the job is backed up in its entirety, even if this means extending the backup beyond the prescribed backup time window.

Retention

The number of days online defines the number of days a safeset is kept on the vault. Once the expiry date is reached, a safeset is deleted. Online storage is used to store data so that it is readily accessible. When data is stored online, recovery operations are initiated immediately without any operator intervention on the vault side.

The number of copies online indicates how many copies of the backup safeset are stored online. It functions in a first in/first out manner. Once the number of copies is exceeded, the oldest copies are automatically deleted until the actual number of copies matches the definition.

The number of days archived indicates how long the data is stored offline. Archive storage is used to store data offline for long periods of time. Typically, only long-term data archives are stored offline. The parameters for archived data are from 365 to 9999 days.

Assuming that at least one successful backup has been completed for the job, there will never be less than one copy of its safeset online. This is true even if all retention settings for 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.

Note that **Days Online** and **Copies Online** work together. Both conditions need to be met before any backups are deleted.

For example, if both Days and Copies are 7, then there can never be less than 7 backups. If Days goes over 7, there will still be 7 copies. If Copies goes to less than 7, there will still be 7 days' worth. It must be over 7 Days with more than 7 Copies before any are expired.

6.2.4 Command Cycle

The **Command Cycle** screen of the Schedule Wizard allows you to define the frequency of your scheduled command. You can choose **Weekly**, **Monthly**, or create a **Custom** cycle.

[Weekly](#) means that the entry will run on the same selection of days every week. [Monthly](#) means that the entry will run on the same selection of days every month.

If you choose **Custom** and click **Next**, the [Custom Date/Time](#) screen will open.

6.2.5 Weekly and Monthly Screens

For a **Weekly** cycle, select the days of the week on which the entry should run.

For a **Monthly** cycle, enter the days of the month on which the entry should run.

The **Start Time** field defines the time of day for the scheduled command to commence. To schedule more entries on the same day, create a separate schedule entry at a different time of the day.

The software displays time in 12-hour mode (AM/PM) or 24-hour mode. The mode displayed depends on the time mode set on your computer. To change the time mode on your computer, select **Control Panel > Regional Settings > Time**.

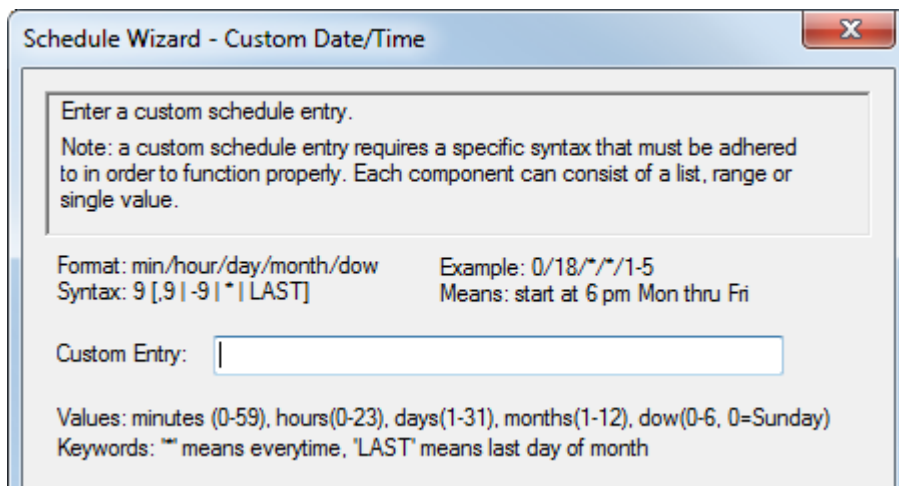
- For 24-hour mode, use **HH:mm**, where **HH** is 0-23 and **mm** is 0-59.
- For 12-hour mode, use **hh:mm**, where **hh** is 0-12 and **mm** is 0-59. Then select **AM** or **PM**.

6.2.6 Custom Date/Time Screen

You can specify a custom schedule by entering time and date information. A custom schedule entry requires a specific syntax. Use the following format to enter your custom schedule data: min/hour/day/month/dow (where dow means days of the week).

Values: minutes (0-59), hours (0-23), days (1-31), months (1-12), dow (0-6, 0=Sunday)

Keywords: * means every time, and **LAST** means the last day of the month



This **Custom Entry** runs the command at 6 pm, Monday through Friday:

```
0/18/*/*/1-5
```

6.2.7 Schedule Wizard - Finish

Click **Finish** to complete the schedule configuration. To further change the settings, click **Back**.

You can also change the schedule after creating it. See [Working with Schedule Entries](#).

6.3 Working with Schedule Entries

To edit a schedule:

1. Select an Agent and open its **Schedule** file.
2. Select the target schedule entry.
3. Click **Edit**. The Schedule Entry screen opens, providing tabs for quick adjustments to the entries.

To enable a schedule entry:

1. Select a disabled entry from the Schedule List.
2. Click **Enable**. The entry will display a 'scheduled' symbol.
3. Click **OK**.

To disable a schedule entry:

1. Select an enabled entry from the Schedule List.
2. Click **Disable**. The entry will display a 'halted' symbol.
3. Click **OK**.

To remove a schedule entry:

1. Select an existing schedule from the list.
2. Click **Remove**. You will be prompted to make sure that you want to permanently remove this schedule.
3. Click **OK**.

6.3.1 Command Tab

This tab allows you to make quick adjustments to a selection of the [Command Screen](#) options.

6.3.2 Day/Time Tab

This tab allows you to make quick adjustments to a selection of options from the [Weekly and Monthly Screens](#).

6.3.3 Parameters Tab

This tab allows you to make quick adjustments to a selection of fields from the [Options Screen](#) in the Schedule Wizard.

6.3.4 Prioritizing Scheduled Entries

The Scheduler uses a prioritization scheme that permits you to override entries under certain circumstances. If the following conditions are met by the schedule entries, the Scheduler will only execute the highest prioritized entry:

- Both commands must start with "VV" or "VV.EXE".
- Both command verbs after the "VV" or "VV.EXE" must be the same (for example, "Backup").
- Both job names must be the same.
- Both execution times must be the same.

Note: The "VV.exe" component on the Agent performs backup and recovery functions to the vault.

The priority of the entry is determined by the order in which the entries reside in the Schedule.cfg file. The first entry is given a higher priority over the entries that follow it and meets the conditions specified above. If two entries conflict, the entry that is higher on the list takes priority. The entry with lesser priority does not run and is rescheduled. See the examples that follow. To move an entry up, use the Up button. To move an entry down the list, use the Down button.

Note: For simplicity, the following examples do not include all schedule parameters that are available, including:

- /quickscan=[yes|no]
- /deferafter=<value>
- /type=[full|incremental|...]

- /compression=[smaller|faster]
- /delta=[yes|no]

Example 1

The following entries meet all prioritization conditions even though the retention parameters are different (day versus week): both end with some form of "vv", the verb is the same "Backup" and the job name is the same "full". When the two backups conflict (this occurs if the third day of the month is ever a Monday, Wednesday or Friday), the second one is rescheduled to its next available time.

```
0/1/3/*/* vv Backup full /retention=day
```

```
0/1/**/1,3,5 c:\agent\vv.exe Backup full /retention=week
```

Example 2

The following entries meet all prioritization conditions as both end with some form of "vv": the verb is the same "Backup", and the job name is the same "system". When the two backups conflict (which occurs on the LAST day of every month), the second entry is rescheduled to its next available time.

```
0/23/LAST/*/* vv Backup system /retention=year
```

```
0/23/*/*/* vv Backup system /retention=week
```

6.4 Scheduler Log Files

The Scheduler creates a log file during its execution. As commands are scheduled and run, each activity is logged into one of four **VVAgent-#** log files located in the CentralControl directory.

The Scheduler will write to the first file until it reaches about 32 KB in size. At this point, it will write to the second file. When all four files reach capacity, the scheduler deletes the first file and begins writing the log again. The current log file version is the one with the latest modification date.

6.5 Effects of Time Changes

The Scheduler checks the server clock and determines whether or not it has been changed by more than 75 minutes. A change forward or backward within 75 minutes would typically be as a result of Daylight Savings time or a simple clock adjustment.

Example 1

If there is a change greater than seventy-five minutes, all scheduled entries are rescheduled to their next available times.

Example 2

A job is scheduled to run at 1:15 a.m. At 1:00 a.m., the server clock is changed to 2:00 a.m. The job scheduled for 1:15 a.m. then runs immediately. This happens because the time change is within 75 minutes and constitutes a daylight savings time change or a clock adjustment.


Example 3

A scheduled job runs at 1:30 a.m. At 2:00 a.m. the server clock is changed to 1:00 a.m. The job will not run twice. When the schedule file is reloaded, it only updates the next scheduled time of each existing entry.

6.6 Time Zones

Scheduled jobs use the time of the local machine (Agent). If CentralControl is in one time zone, and the Agent is in another, the starting time of the job depends on the local machine's clock. The log files and **Last Modified** dates also show the local machine's time.

These times also display the offset (difference) from Greenwich Mean Time. In the example, **-0500** means that the time zone is 5 hours behind GMT, which is defined as zero.



The screenshot shows the MyWorkspace.vws application window. The left pane displays a tree view of the workspace structure, including folders for vSphere4, vSphereJob1, vSphereRest, Windows1, Processes, Backup, Safesets, Logs, ED_Regular, and NetWare. The right pane displays a table of backup entries.

Backup Age	Safeset	Date and Time	Type
New	00000001	07-NOV-2011 13:08:00.00 -0500	Full
New	00000002	07-NOV-2011 13:23:00.00 -0500	Full
New	00000003	09-NOV-2011 09:32:00.00 -0500	Full
New	00000004	09-NOV-2011 09:33:00.00 -0500	Full (delta)

7 Backing up Data

When you select a job, you can run a backup through the **Actions** menu, or by clicking the **Backup** icon on the toolbar.

A backup is a process that copies data to a vault or disk. A backup may be scheduled or run “ad hoc”, which is a backup not on a scheduled time. Usually, it is a one-time backup, for a specific reason. You can either run a regularly scheduled backup job as a special ad hoc backup, or create a new job to be run immediately.

An initial seed copies, in whole, all files and directories to create the first full backup. This includes user data, system data, and applications/programs. Subsequent “deltized” backups are also regarded as full backups. They only send the changes since the previous backup, but you can do a complete recovery from a deltized backup.

Note: It is not recommended that you back up the Agent or CentralControl installation directories. Exclude these directories from backup jobs.

7.1 Seeding

Seeding is the process of getting an initial backup to a vault. From then on, any changes to the data are applied to that initial backup as deltas (changes). This means that every backup is as useful as a full backup. You do not need to apply differential or incremental changes to allow recovery.

Seeding may be accomplished across the network (LAN/WAN), or you may want to save your data on media that you physically transport to the vault.

7.2 Seeding with Deferring

The Agent allows you to specify time windows for network reconnection and backup jobs.

Depending on how these values are set, you may see errors about failed deferrals because a vault process is no longer available.

Stop reconnection attempts after... specifies the maximum time over which your system should try to re-establish connection with the vault before terminating/failing the job. This value also specifies how much time the vault should wait for the Agent before it terminates the job from the vault side.

Try to reconnect every... specifies the number of seconds to wait between consecutive attempts to re-establish connection with the vault.

Backup time window defines a window of time in which to perform the backup procedure. Choosing this option will limit the backup time window to the period of time specified. Any new files that are not completely backed up within the available window will be deferred until the next backup session.

7.3 Backup Wizard

The Backup Wizard helps you through the process of choosing destinations and additional options.

Before starting a backup, you must correctly configure an Agent and a job.


When you start the Backup Wizard, you may not see the Welcome screen. You bypass the screen if the **Skip this screen in the future** option is enabled. To stop bypassing, click **Back** on the second screen, and clear the option.

7.3.1 Destination

You can back up to a vault that was configured for this job, or an **Alternate safeset location**.

If you choose a destination and click **Back Up Now**, the backup will start immediately. Otherwise you can click **Next** to go to the **Options** screen.

If you choose **Alternate safeset location**, you can enter the location information into the field that appears.

Alternatively, click the  button that appears so that you can browse. Find a suitable directory, select it, and click **OK**.

Click [here](#) for information about recovery destinations.

7.3.2 Job Options

The [Retention](#) setting specifies the length of time to store the backup.

You can also specify **Quick file scanning**, **Compression type**, and **Backup time options** (as applicable).

Click [here](#) for information about specific backup options for the SQL Server plug-in.

Click [here](#) for information about specific backup options for Exchange DR.

7.3.3 Finish Screen

Click **Finish** to start the backup immediately. Otherwise, click **Cancel** or **Back**.

7.4 System State Objects

If you include “System State” in your Windows job, the Windows Registry and Active Directory will be backed up.

7.4.1 Windows Registry

The Windows Registry (or Local Registry) stores virtually all of the custom data that Windows uses. Every time a program starts and every time Windows performs an operation, the Registry fills in all of the variables with your computer’s custom settings.

7.4.2 Active Directory

Active Directory is a Windows directory service designed for distributed computing environments. It allows organizations to centrally manage and share information on network resources while acting as the central authority for network security.

In addition to providing comprehensive directory services in a Windows environment, Active Directory is designed to be a consolidation point for isolating, migrating, centrally managing, and reducing the number of directories that companies require.

7.5 Locked/Open Files

Locked files are busy (in use) by the system or another application.

To ensure that open or locked files are backed up properly, VSS is used.

7.6 Back Up After Reregistration

If you have reregistered a computer, the job credentials are intentionally not moved to the new computer. This means that you cannot use that job for a backup until you have successfully re-entered the proper credentials. These include the encryption password, if applicable.

You must edit the job, enter the proper credentials, and save.

8 Restoring Data

8.1 Restoring from Conventional Backups

You can restore from any conventional backup that you have created. Restores must be run manually from the GUI. If a backup has a password for encryption, you must know the password. It is not kept anywhere on the system.

You can restore a complete file system, or one or more files or directories.

You can search, select, and restore from multiple backup catalogs in a single Restore workflow. Earlier versions of CentralControl only permitted the search and selection of files from a single backup catalog.

Restores are not scheduled. You can run many restore operations at the same time. Each one starts a new process, which you can monitor.

8.1.1 Operating System Independence

You can recover data independent of the operating system. However, there are limitations for this sort of action.

Specifically you can only recover the data stream and basic file information attributes, such as creation date, modification date, etc. Recovery of security, extended attributes, and alternate data streams is not supported.

8.1.2 Restoring from CD, USB, and Other Media

The Agent allows you to restore directly from CD, DVD, and other media (such as USB drives), without needing to copy safesets to the hard drive first.

There are several ways to restore safesets on the media:

- One backup/safeset image (SSI) on a single CD/DVD/USB
- Backups that are divided into many SSI files with a fixed length (but the whole set fits on a single medium)
- CDs/DVDs/USBs that contain a single backup spanning more than one medium

Choose **Alternate safeset location**, and then browse to the folder containing the SSI file. The SSI file on the medium must correspond correctly with the safeset number under **Restore from**. If not, CentralControl will show an error when the restore operation begins, saying that the medium is not correct.

When the restore operation begins, it will request that a certain CD/DVD/USB be mounted in the drive if there are no media mounted. If additional media are required, you will be prompted. You will not be prompted if the requested SSI is on a CD/DVD/USB that is already in use (i.e., mounted). If there are multiple SSI files on the same medium, you will need to make a selection.

8.1.3 Restore Wizard

You use the Restore Wizard to complete these tasks:

- Select the source for restore operations
- Enter the password if the backup is encrypted
- Select items to restore
- Enter the destination
- Select other options

Select a job that has at least one safeset. Click the **Restore** icon. This starts the Restore Wizard.

8.1.3.1 Selecting a Source

(For information related to Exchange DR *share options*, click [here](#).)

Choose a source from which to restore: **Vault** or **Alternate safeset location**.

If you choose **Vault**, you must select a single safeset or a range of safesets.

Restore from a safeset or range of safesets

All safesets that have been created through backups, and that are still valid, are shown here.


For a single safeset, you must know (from the job, date and safeset properties) which safeset you want to recover.

For a range of safesets, use the **from** and **to** lists. Note that the **from** date must be earlier than the **to** date. Also, older safesets restore before newer ones.

You cannot select **System State** when you restore from multiple safesets.

Alternate safeset location

If you choose **Alternate safeset location**, you can enter the location information into the field that appears.

Alternatively, click the  button that appears so that you can browse. Find a suitable directory, select it, and click **OK**.

8.1.3.2 Encryption Password

To restore data from a job, you must enter the encryption password for the job.

The password is not recorded anywhere on the backup, or the system. It is your responsibility to retain this password. Note that the password is case-sensitive.

8.1.3.3 Select Items to Restore

You can select some or all of the items that were backed up by the job.

Add and Include/Exclude

CentralControl distinguishes between full and partial directory backups. If the entire directory has not been backed up, hidden file data such as security and reparse information is not backed up.

- To include items, click **Add**, select them, and click **Include**.
- To exclude items within your selection, highlight them, and click **Exclude**.
- If you select an entire directory, a confirmation screen appears. Click [here](#) for more information.

In some cases, a **Search** button is available. If you are not sure which directory your files are in, click **Search** to [scan the backup catalog](#) for the files that you want to recover.

System State

This selection allows you to recover all system files that were optionally included in the backup. The names of all System State files that are eligible for recovery appear in the lower pane.


System State recovery guidelines:

- System State data can only be restored on a local computer.
- You must restore all the System State data (Registry, COM+ Class Registration Database, Boot Files, etc.) that was backed up.
- Vault profiles, schedules and jobs that are associated with a physical node in a cluster can have the system volume and System State backed up and restored.
- System State restores should not be done on a computer that has a system root location different from the root location where the backup was done.
- Ideally, System State restores should be performed on the same computer on which the backup occurred.
- To restore System State on a different computer, the new system must have the same or higher number of disk drives, and similar hardware configuration as the original computer.

8.1.4 Backup Catalog

8.1.4.1 Searching the Backup Catalog

The **Search** feature within the Restore wizard allows wildcard searches using the * (asterisk), ? (question mark), and . (period) characters.



In your search, you can use the asterisk to substitute for a string of data. For example, *.DAT or BOB.* could be used to search for BOB.DAT.

You can use the question mark to substitute for one character or multiple characters. For example, BO?.DAT could be used to search for BOB.DAT. Additionally, *.S?? could be used to search for BOB.S01.

You can use the period to signify a recursive directory. For example, use C:\.*.doc to search for all files on drive C that have the doc file type.

To further restrict your search, enter a path in the **Search path** field.

Click **OK** to start your search.

8.1.4.2 Selecting from Search Results

After you enter your search criteria and click **OK**, the **Include/Exclude Search Result** screen will open. You might need to wait more than a few seconds for this to happen.

To include a file or directory, highlight it and click **Include**. To exclude a file or directory, highlight it and click **Exclude**. You must include at least one item in your recovery set.

The search can return huge numbers of items, so you will sometimes receive message prompts allowing you to "break" the search and display only the items that have been found so far.

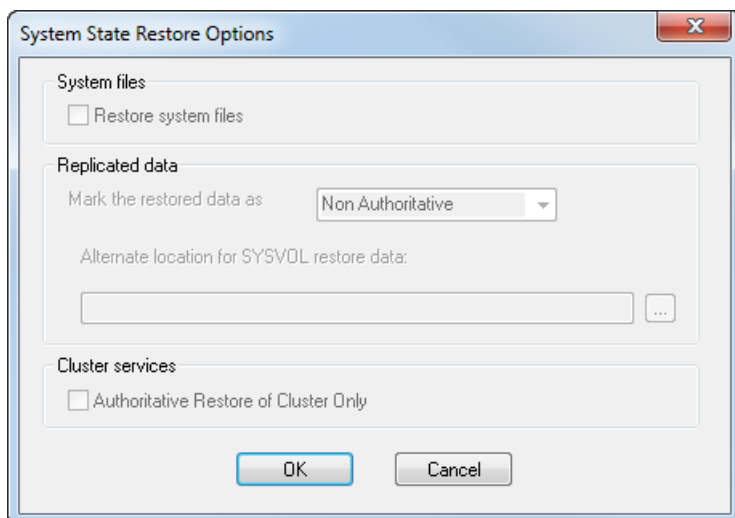
8.1.4.3 System State Restore Options

This screen allows you to select how to restore your System State. For example, you may want to recover:

- To a previous version of the System State
- From a system failure (identical hardware)
- Active Directory (authoritative)
- Active Directory (non-authoritative)
- Exchange server
- DNS/DHCP server
- Cluster server

System files:

You can turn the **Restore system files** option on or off. It is on by default if your backup included system files.

**Replicated data:**

You can mark the selected data as **Non Authoritative**, **Authoritative**, or **Primary**. If you are restoring Windows 2008 or later, the list will only include **Authoritative** and **Non Authoritative**.

If **Authoritative** is selected, the next field, **Alternate location for SYSVOL restore data**, is active.

Cluster services:

You can turn **Authoritative Restore of Cluster Only** on or off if your backup included cluster services.

8.1.5 Locations, Overwriting and Renaming

You can choose locations for restoring your items, as well as options for overwriting and renaming.

Destination options

Original or alternate locations: If you choose original, you can overwrite existing items.

To restore to an alternate location, enable the corresponding option. Next, enter location information into the field provided. Otherwise, click **Browse**. Find a suitable directory, select it, and click **OK**.

Similarly, to restore to an alternate database, enable the **Restore to an alternate Exchange database** option. Enter location information into the field provided. Otherwise click **Browse**. Find a suitable database, select it, and click **OK**.

Note: You can restore a backup from a UNC path to the original location. But if you attempt to restore to an alternate path, it will fail and log an error.

Suggestions related to overwriting and renaming:


- Rename files coming in from your recovery so that they don't conflict with existing ones.

- Rename files that already exist, so that they don't conflict with restored ones.
- Do not restore existing files by overwriting the existing ones.
- Overwrite existing files with the restored ones.

Note: The Agent will add serial extensions to renamed files. These extensions start as .0001, then .0002, .0003, etc.

8.1.5.1 Advanced Restore Options

Here are the default settings for file-based restores:



Overwriting locked files

If you choose to overwrite locked files, these files will be replaced after restarting the machine where the data is being restored.

Restore all streams or data streams only

CentralControl stores the information from your files in various streams. The original data created by you is called the data stream. Other information, such as security settings, alternate data for other operating systems and file reference information, is stored in separate streams.

All streams recovers all information streams. This setting is recommended, provided that you are restoring files onto a system with an identical operating system. For cross-platform recovery, select **Data streams only** - This setting ensures that conflicts do not arise as a result of system-specific information streams.

Log detail level

The default setting is to create a log whose detail level is **Summary**.

When you perform a complete system recovery ("Disaster Recovery") on a Unix platform, make sure that you have enough disk space for large recovery logs.

In addition to Agent logging, entries for logging or auditing could come from the OS. File-level logging on a large file system can generate a very large log, which could consume the available or allocated disk space. If the logs occupy the same partition as the root file system, this could prevent the OS from starting.

Use all available bandwidth

By default, restore operations use all available bandwidth in order to run as fast as possible. If you prefer otherwise, clear this option. This allows the *bandwidth settings that were configured for the Agent* to apply to the restore.

Disable prescanning - Unix

This option disables prescanning, so that any hard links are not searched for files, and only the links themselves are backed up. On a restore, the links are not followed for their files.

8.1.6 Restoring from Another Computer

(For corresponding information about vSphere Agents, click [here](#).)

Normally a recovery would go to the same computer where the backup was performed. But if that computer is damaged, or unavailable, and only some of the data needs to be recovered before a replacement system can be built, you can restore using **Restore from Another Computer**.

If you are replacing a computer after the loss of a system, use the reregister workflow to assume the identity of the previous system. See the section about reregistering for more information.

There are limitations regarding the operating systems that can successfully transfer data using this workflow. For example, different versions of the same operating system can be compatible. Operating systems that are part of the same family, or share similar origins (such as Linux and Solaris), can also be compatible.

Restore from Another Computer allows you to redirect a recovery from the (original) job to a different computer. It downloads the job to an alternate system so that a recovery can be performed.

If you already have a job with the same name as the job you want to use, this is acceptable. A new job name is automatically generated when the job is downloaded. The new name is a combination of the original job name, name of the original computer, and a key. The new job contains the original job name since this information is required to perform the restore. The restore also records the original job name in the restore log.

1. Select an Agent or job.
2. Click **Actions > Restore from another computer**. The **Restore From Another Computer** screen will open.
3. Select the **Vault**, **Computer**, and **Job** from which to restore.
4. Click **Next**.

The software will attempt to download information about the job that you have selected. If the job has not produced a usable backup, the download will fail.

Other things can cause the download to fail:

- Job is not from a compatible system
- Vault cannot be reached
- Catalog file cannot be retrieved

If the download fails, select another job, or try again later.

After a successful download, the recovery continues in a way that is similar to restoring from the original computer – select a safeset, select items to restore, etc.

If the job being restored requires credentials, cancel the **Restore from Another Computer** procedure after the job downloads. Select the newly downloaded job, and edit the job configuration to add the new credentials. Then start a regular recovery.

You will now have a “new” job in your list of jobs, which came from the other Agent. This job can only be used for restores. It does not do backups. When you finish restoring, remove this job.

Note: This procedure should not be used to restore from jobs originally registered on this system. If you have lost a job configuration on this system and need to recover it, use the vault reregistration workflow to obtain your job configuration again.

8.1.7 Synchronizing

This is the process of connecting the current job to the remotely defined vault, updating the local Agent with vault information, and then disconnecting from the vault. The list of the safesets and their current status information is used to display status for backup and recovery.

Immediately after you have done a backup, the safeset status shows as **Full (delta)**. This indicates that the Agent does not have an updated status list from the vault. That is, it does not have an indication of how many safesets are currently online. To get the current status, you need to synchronize with the vault.

Normally a Synchronize operation is scheduled to run automatically. It may also be run manually.

A Synchronize operation produces a log file called SYNC in the logs for the job.

8.1.7.1 Synchronizing manually after reregistering

With newer Agents, a backup also copies the catalog information to the vault. This is the structure of directories and files from the Agent machine. This information is encrypted, along with the backup data. The encryption password is not stored anywhere on the Agent or vault.

Previous Agent versions did not keep catalog information in the backup. They rebuilt it (dynamically) from the files in the backup.

If you restore using the Restore wizard, you are prompted for the encryption password, which is used to unencrypt the catalog data and user data.

But if you choose to manually run Synchronize before the recovery, there is no associated password to recover the catalog file. The Synchronize operation needs to recover the catalog information from the vault to the new Agent.

Under **Job > Properties > Encryption**, you must know and enter the password that was used to back up this job so that the Synchronize operation can correctly read the catalog information.

If you do not supply the password, the Synchronize operation might succeed, but it will run slowly. This is because it will need to build new catalog information as it retrieves the data files, instead of simply reading the catalog file.

8.1.8 Bare Metal Restores

A Bare Metal Restore (BMR) refers to a computer system that starts from an “original, new” state, with nothing (i.e., operating system, applications, or data) stored on its disks. This may be the case after a system hardware crash, where new disks are installed.

The Restore process then recreates everything on the system. (This is not the same as reinstalling everything. With that, you lose customization and parameters for applications, systems, and service packs.)

Bare Metal Restore is a licensed option that captures the information necessary for the System Restore (SR) application. SR is a separate application used for bare metal recovery of systems in case of a disaster.

There are two scenarios for restoring your backups from a vault:

- Restore everything using network communication lines. Reinstall your operating system and then the Client Agent/CentralControl. With the Client Agent, recover your data using the Synchronize option to retrieve your catalogs from the vault.
- Request to have your backup sent to disk, and delivered, along with a copy of the CentralControl and Agent software. After restoring (reinstalling) your operating system and CentralControl software, you can recover directly from local media.

BMR with System Restore

System Restore (SR) is a separate, licensed application used for bare-metal restores. BMR with System Restore can completely restore a Windows system (Windows Server 2008 or later).

This is different from using the Agent and System State to do a bare-metal restore. See the System Restore documentation for more information.

Note that encrypted volumes (BitLocker, TrueCrypt, etc.) are not supported for BMR jobs.

You would want to perform a BMR with System Restore, for example, if a system has crashed, and the disk has been replaced. Now you want to recover all the system and user data back to that disk.

Reinstalling the OS, applications, and then the data is possible, but you may not be able to recreate the exact state of the system that you would get with a restore of a full-drive backup that included data files, System State, and System Files.

Note: Bare Metal Restore-type backups/safesets (which can be used by SR to perform a bare-metal restore for disaster recovery) can also be used via CentralControl to restore individual files and folders. This is similar to how you use regular safesets.

8.2 Restoring from vSphere Agent Backups

8.2.1 Restore Type

To restore from a vSphere Agent backup:

1. Select an Agent.
2. Select a job.
3. Click the **Restore** icon. The Restore From a Backup workflow will open.
4. Choose what you want to restore:

Virtual Machines

This restore type allows you to recover an entire virtual machine to your vCenter.

You can also choose a datastore and host for restoring and registering the VM.

Virtual Disks

This restore type allows you to recover individual Virtual Disk (VMDK) files to a datastore in your vCenter.

You can choose the datastore and a folder within (or create a new folder within the datastore) for restoring the Virtual Disk files.

Restoring Virtual Disk files does not restore whole VMs. The VMs will not be registered.

Files and Folders

This restore type allows you to select and restore specific files and folders from a VM, rather than restoring an entire VM or individual virtual disks.

After selecting a VM and virtual disk whose data you want to restore, you set up the disk as a shared resource (a "share"). Then you mount the disk on the machine where you want to restore the files.

5. After you choose a restore type, click **Next**.

8.2.2 vSphere Source Selection

Select a source (usually a vault) from which to restore, and choose an item from the **Safeset** menu.

(To adjust the options for virtual disk shares, click [Advanced Share Options](#).)

Restoring from an alternate location

If your safeset has been exported to a disk and shipped to you for offline recovery, use the **Source** menu to select **Alternate safeset location**.

Before you can restore from an alternate location, you need to mount the disk containing the safeset. You can do this through the vSphere Agent command console.

Use the **mount** line command to mount the remote disk. You can then locate the disk by browsing. Beside the **Select folder** field, click **Browse**.

Browse through the mount points, and select the folder that contains the safeset. Next, select the safeset (backup version) that is located on the remote disk.

Provide the encryption password (specified during job configuration). If you have lost this password, you will not be able to restore.

Click **Next**. A selection screen will open. This is where you will choose which items to restore.

8.2.3 Items to Restore

For file/folder recoveries, see [Items to Share](#).

Select the items that you want to restore. Click **Include**.

The item names that you choose will appear in the right-hand pane. You must include at least one item to continue.

Note that you cannot select incomplete VMs from lists.

For virtual disks:

You can choose individual disks to restore by browsing into the VMs. To display individual disk names, click the small icons to the left of the VM list.

You can restore all virtual disks for a particular VM by selecting the VM itself. Selecting a VM only restores its virtual disks (rather than fully restoring the VM).

8.2.4 Items to Share

The **Select a virtual disk to share** screen shows the names of VMs whose files and folders can be restored from the safeset.

From the **Virtual Machines** list, choose the VM from which you want to restore files or folders. The names of the virtual disks associated with the VM will appear in the right-hand pane.

Select the disk from which you want to restore files, and then click **Next**.

8.2.5 Options for VM Restores

8.2.5.1 Where to Restore

This screen allows you to choose where to restore your VMs. Open the menu, and select a datastore.

Now choose one of the options below the menu. These options define how to use the selected datastore:

Restore all selected Virtual Machines to the selected datastore only

This option forces the whole restore to the selected datastore.

Restore to the selected datastore only when a Virtual Machine's original datastore is not available

By default, a VM will only be restored to the selected datastore if the VM's original datastore no longer exists.

Regardless of which option you choose, if a VM spans more than one datastore, and one (or more) of those datastores no longer exists, the entire VM will be restored to the datastore selected from the menu.

For your convenience, the screen provides a link that allows you to review which VMs have been selected for the restore. You can do this without returning to previous screens. For each selected VM, you can also see the original host registration and datastore location.

Click **Next** to continue.

8.2.5.2 Where to Register

This screen allows you to choose where to register your VMs. Open the menu, and select a host. Only hosts that have access to the datastore selected on the preceding page will appear in the menu.

Now choose one of the registration options below the menu. These options define the conditions for registering your VMs with the selected host:

Register all selected Virtual Machines with the selected host only

This option forces the restore to register all of the selected VMs with the selected host.

Register with the selected host only when a Virtual Machine's original host is not available

By default, a VM is registered with the selected host only if the host that held the registration at backup time no longer exists.

Also, if you wish, enable the **Power ON VMs after restoring** option. This automatically powers on all restored VMs when the restore finishes.

For your convenience, the screen provides a link that allows you to review which VMs have been selected for the restore. You can do this without returning to previous screens. For each selected VM, you can also see the original host registration and datastore location.

Click **Next** to continue.

8.2.6 Options for Virtual Disk Restores

8.2.6.1 Datastore Destination

When you restore virtual disks, you can only direct them to a datastore. This dialog allows you to browse the available datastores (and folders within) to choose a target location.

You can specify a new folder to create within the selected datastore. Its name cannot exceed 80 characters. Click **Apply** to use the new folder name. Your selected datastore and folder will display.

You must select a destination in order to continue.

8.2.7 Share Options

Share Options

In the **Idle time** field, enter the number of minutes of inactivity after which the share should automatically stop. This value can range from **2** to **180**.

Bandwidth Option

By default, starting the share uses all available bandwidth. If you do not want this, clear the **Use all available bandwidth** option. This allows the *bandwidth settings that were configured for the Agent* to apply.

8.2.8 Lists of vSphere Selections

(For information about virtual disk shares for vSphere Agents, see [Shared Disk Path](#).)

To see a list of the VMs that you have chosen to restore, click the **View selected Virtual Machines** link. This link is available on some of the Restore screens.

The **Selection Summary** screen opens, displaying information about the selected items. For example:

- **Name** – VM name (applies to Virtual Machine restores)
- **Virtual Machine** – VM name (applies to Virtual Disk restores)
- **Virtual Disk** – In this column, * means that all virtual disks for the selected virtual machine will be restored.
- **Size** – Disk size required to restore the selected virtual machines or virtual disks
- **Host** – Host with which the VM was registered at backup time
- **Datastores** – Location of the VM at backup time. If a VM spanned multiple datastores, they will be listed here. To see the whole list, pause your mouse over a portion of it, or resize the dialog.

An icon on the final screen of the workflow links you to the same sort of information (**Virtual Machines to restore**, or **Virtual Disks to restore**, as applicable).

8.2.9 vSphere Restore Options

You can choose specific options for logging and performance.

By default, the restore uses all available bandwidth so that it runs as fast as possible. If you do not want this, clear the **Use all available bandwidth** option. This allows the *bandwidth settings that were configured for the Agent* to apply to the restore.

Click **Next** to continue.

8.2.10 vSphere Restore Summary

At the end of the **Restore from a backup** workflow, a summary screen will appear. For your convenience, the screen provides an icon that links you to a [list of the selected items](#).

If you are satisfied with your settings, click **Run Restore**.

Note: For information about virtual disk shares and share summaries, see [Shared Disk Path](#).

8.2.11 Shared Disk Path

For shared virtual disks, a summary screen opens after you provide the share details. If you are satisfied with the settings on this screen, click **Share**.

The **Process Information** screen opens, showing the process status. When the disk is shared, a UNC path appears in the **Path** field.

You can use this path to access the UNC share. Right-click the path, and click **Copy**.

You can also send the path to other users so that they can restore the files.

The disk remains shared while files are being copied, until there is no activity for the specified amount of idle time, or until you click the **Unshare** button.

Note: For virtual disks that include Windows dynamic disks, you require the Dynamic Disk tool to access all of the available data. For more information, see the *vSphere Agent User Guide*.

8.2.12 Restoring from Another Computer

8.2.12.1 Start and Select Resources

To restore from another computer:

1. Select a vSphere Agent.
2. Click **Actions > Restore from another computer**. The **Restore From Another Computer** screen will open.
3. Select the **vault**, **computer**, and **job** from which to restore.
4. Click **Next**.

8.2.12.2 Downloading Job Information

The software will attempt to download information about the job that you have selected. If the job is not a vSphere Agent job, or it has not produced a usable backup, the download will fail.

8.2.12.3 Remaining Steps

After the job information downloads, the remaining steps are the same as the steps for regular [vSphere Agent restores](#).

At one point, the software will attempt to download catalog information from the vault. If the vault cannot be reached, or the catalog file cannot be retrieved, the download will fail. If this happens, select another job, or try again later.

9 Resolving Common Errors

The topics in this section cover many general problems and solutions.

Some situations that you might encounter are described here. Error messages (if applicable) are listed, probable reasons for situations are given, and solutions to these problems are provided.

9.1 Email Notification Not Received

Sometimes notification email is not received properly.

The Agent software generates a unique **E-mail from address** for its notification messages. These addresses are based on the name of the program (i.e., Agent), the name of the computer running the software, and the domain of the network upon which the machine resides.

The address takes the following form:

Agent_{machine name}@{domain}

For example:

VVAgent_BILLINGPC@acmecorp.com

Solution

Use the **from** override to customize the email "from address" of these messages and override the default action. (If the **from** override is not present, the email address is composed as above.) See [Notification Tab](#) for specific instructions.

9.2 Failed to Authorize – Insufficient Privileges

Sometimes the CentralControl application does not connect to the Agent because access is denied.

Message example: Failed to authorize as <username> or insufficient privileges.

Possible causes:

- Incorrect password or misspelled user name. The password is case-sensitive.
- Insufficient privileges assigned to user.
- Agent Service was previously disabled.

Solutions:

- Re-enter your password.
- Assign adequate privileges (contact your administrator). Restart Agent Service.

9.3 Failed to Connect

Sometimes the CentralControl application does not connect to the Agent.

Verification:

Try to "ping" or perform a "tracert" to the Agent system using the IP or DNS address.

Also, perform a DNS lookup with your DNS server to determine if the DNS name you are using is valid.

Possible causes:

The Agent system is not running a TCP/IP stack. Check the system's network setup.

The Agent system is not running the Agent (Windows Service). Check the Windows Services.

Your system is not running a TCP/IP stack. Check the system's network setup.

The network card is not properly configured on the CentralControl application or Agent. Check the system's network setup.

The network is down between the CentralControl application system and the Agent system. Check routing.

The network subnet mask is not set up properly. Check the system's network setup.

Solution:

Contact your LAN or WAN Administrator to determine which of the previous components is not functioning.

9.4 I/O Device Error

If you receive this message, it usually means that a controller, cable, termination or drive failure occurred while backing up. Other possible reasons are driver problems, operating system problems or low memory problems.

Make sure that:

- Each device is working properly.
- You are using the proper drivers for devices, and that your operating system is not corrupted.
- You have enough RAM to meet the demands of the system and applications that you are running.

9.5 Limit on Number of Shared Memory Segments

```
Error opening file /usr/local/Software/Test/SYNCH.MIR for output
```

```
SSET-E-0111 error committing synchronization data
```

```
Error opening file /usr/local/Software/Test/00000004.cat for output
```

Problem

In Solaris, the VVAgent may report the error message that follows.

```
PROG-E-0003 O/S error: Too many open files.
```

This problem occurs because of the limit on the number of shared memory segments per process. There are tunable parameters in the Solaris kernel that you can change.

Solution

Increase the default setting.

1. Add these lines to `/etc/system`: (if a setting already exists, increase it accordingly)

```
set shmsys:shminfo_shmmni=200
set shmsys:shminfo_shmseg=200
set semsys:seminfo_semmni=200
set semsys:seminfo_semmns=5000
set semsys:seminfo_semmap=5000
set rlim_fd_cur=256
```

2. Verify that the kernel parameter has changed using:

```
grep shmsys /etc/system
grep semsys /etc/system
```

3. Restart the system.

9.6 VVAgent Unexpected Shutdown on Unix

Sometimes the VVAgent shuts down unexpectedly on Unix systems.

Solution

Wait for 1 to 4 minutes. Then restart the VVAgent.

If you try to restart the VVAgent sooner, it will terminate with the message: 'Stop with Error 42.'

9.7 System Trouble after Active Directory Restore

Sometimes after you perform a System State recovery, "my system" does not start properly.

Suggestions

Restart the computer. If the computer does not restart after recovery because of Hardware Abstraction Layer (HAL) mismatches, you can use the Windows installation disk to perform an in-place installation or repair. This type of repair occurs after you accept the licensing agreement, and Setup searches for previous versions to repair. When the installation that is damaged or needs repair is found, press R to repair the selected installation. Setup re-enumerates your computer's hardware (including HAL) and performs an in-place upgrade while maintaining your programs and user settings. This also refreshes the `%SystemRoot%\Repair` folder with accurate information that you can use for normal repairs.

If the computer does restart after the recovery, log on as Administrator and initiate an in-place upgrade by running Winnt32.exe. This refreshes the Setup.log and registry files in the %SystemRoot%\Repair folder, and ensures the proper HAL is in use.

For additional information about user profiles, refer to the following article in the Microsoft Knowledge Base:

Q214653 How to Set the Path for the All Users Profile

9.8 Configuration File Missing Info on Reregistration

If you delete an Agent from a vault, you are deleting the actual profile on the Agent computer. If you then add that same Agent (it uses the Agent's computer name) to the vault, the vault recognizes it and prompts you for a reregistration. This will also happen for **Restore from another computer**.

The original profile is downloaded from the vault back to the Agent, but in this case, it is missing several fields:

- Encrypted password
- Domain, user name and password of the account used to perform a MAPI backup (Windows)
- Domain, user name and password of the account used to back up SQL Server (Windows)
- Domain, user name and password of the account used to back up a mapped drive (Windows)

You will receive messages similar to this example in the error log when a backup or restore fails because of a reregistration, or a problem restoring from another computer.

PARS-W-0002 Due to a computer registration, configuration file "weekend" is missing the following information:

PARS-W-0002 Enc_Password (Encryption Password)

Please use CentralControl to re-enter the missing information.

Solution

The Agent reregistration process creates a **register** log that reports any missing job settings. The log file can be viewed via CentralControl once the reregistration has completed.

Any attempt to perform a backup or restore using one of the affected job files will fail until the job has been reconfigured. Should this failure happen, the backup or restore log will contain information similar to that of the **register** log, indicating which job settings are missing.

10 Additional Information

10.1 Recursive Backups

If you select **Recursive** when you create a backup job, the backup will include the selected file/directory and all subordinate branches.

If you do not select **Recursive**, only the first branch of the selected directory is backed up. Folders within the selected directory appear in the backup, but the files contained within those folders are not backed up.

10.2 User and System State Data

System State data mainly consists of the operating system used on a specific computer, including service packs, upgrades, application settings, security information, and user profiles. This data usually differs from user data – user data is often application output.

10.2.1 Backing up System State Data

All operating systems contain system-specific data that can be backed up and restored. In Windows, these system-specific files have been grouped together under the "System State" heading. Backups of System State data are critical in the event of an operating system failure. The more up-to-date your System State backup is, the better your chances for a fast and complete system recovery.

By selecting **System State** when you create a new job, you automatically back up all system-specific data when your job runs.

System State data is backed up at the beginning of a job.

1. Highlight an Agent, and click **New Job** on the File menu. The New Job Wizard opens.
2. When you reach the New Job Wizard – Source screen, enable the **System State** checkbox.
3. If you are running Windows, you can click **Options** to open the System State Backup Options screen. Enable **Backup system files** to back up all files that are part of your operating system. This option is selected by default. If you are backing up to a tape device with slow throughput, you might prefer to disable this option. Click **OK** to return to the wizard.
4. Enable the **Data Files** checkbox to include other files in the job, or click **Next** to proceed with job creation.

You can add System State files to an existing job by clicking **Properties** and selecting the **Source** tab.

The System State files backed up for each operating system will differ depending on the OS and how it has been set up and maintained.

Note: If you have a customized Windows server set up with certain services disabled, a System State or BMR backup may fail. These VSS writers are required on all “healthy” Windows computers: **COM+ REGDB Writer**, **Registry Writer**, and **System Writer**.

10.2.2 Restoring System State Data

CentralControl allows you to recover System State data through the Restore wizard. The recovery procedure differs depending on the operating system, but the following guidelines apply to all platforms:

- You must recover all of the System State data that was backed up.
 - System State recovery should not be done on a computer that has a system root location different from the root location where the backup was done.
 - Ideally, System State recovery should be performed on the same computer on which the backup occurred.
 - To restore System State on a different computer, the new computer must have at least the same number of disk drives, and similar hardware configuration as the original computer.
1. If you are running Certificate Services, stop the service.
 2. If you are running Active Directory, be sure to restart in Directory Services Restore mode.
 3. Start the Restore workflow.
 4. On the Select objects to restore screen, enable the **System State** checkbox.
 5. Click **Next**, and continue the recovery.
 6. When the recovery is complete, check the log for errors.
 7. Restart after the recovery.

10.2.3 Windows 2008 System State Recovery

In Windows 2008, System State information and cluster quorum data can be restored. If necessary, cluster quorum data can be restored independently of System State. Before restoring Active Directory on Windows 2008, you must restart in Directory Services Restore mode.

1. Shut down all applications and third-party products before starting.
2. Restart the computer.
3. During the phase of startup in which the operating system is normally selected, press **F8**.
4. On the Windows Advanced Options Menu, select **Directory Services Restore Mode** and press **Enter**. This ensures that the Active Directory is offline.
5. At the “Please Select the Operating System To Start” prompt, select **Microsoft Windows Server** and press **Enter**.
6. Log on as a Local Administrator.
7. In the message box that warns you that Windows is running in Safe Mode, click **OK**.

Note: When you restart the computer in Directory Services Restore Mode, you must log on as an Administrator using a valid Security Accounts Manager (SAM) account name and password, rather than the Active Directory Administrator's name and password. This is because Active Directory is offline, and account verification cannot occur. Rather, the SAM accounts database is used to control

access to Active Directory while it is offline. This password was specified during the setup of Active Directory.

8. Restore using the command-line interface or the CentralControl application interface.
9. Restart.

10.3 Active Directory Restores

In Directory Services Restore mode, the Active Directory can be restored using the Restore wizard. You must recover the Active Directory in one of three ways:

1. Primary Restore
2. NON-AUTHORITATIVE Restore
3. AUTHORITATIVE Restore

These are described in the following sections. But to illustrate, here are three examples that show when they can be used:

Scenario 1

You want to recover the Active Directory when all Domain Controllers are lost. Use Primary Restore.

Scenario 2

You have at least one working Domain Controller. You want to recover the Active Directory on one of the Domain Controllers. Use NON-AUTHORITATIVE Restore.

If the AD already exists on the controller:

- You must be in Directory Services Restore mode.
- Restore the System State only.
- Restart.

If the AD does not exist on the controller (for example if the entire machine was rebuilt, in a Disaster Recovery scenario):

- Restore the data files first.
- Restore the System State.
- Restart.

Scenario 3

You have at least one working Domain Controller. You want to recover an object to the AD. Use AUTHORITATIVE Restore.

- Restore System State only. At the end of the recovery, do NOT restart as prompted.
- Use Microsoft's NTDSUTIL.exe program.
- Restart.

To handle the SYSVOL part of Active Directory restore:

SYSVOL is a replicated dataset that contains the policies and scripts that are used by Active Directory. SYSVOL uses Windows file replication for distribution throughout the network. The three options for SYSVOL restore are identical to the options for file replication: primary, non-authoritative (the default), and authoritative restores.

Note: Although you typically restore SYSVOL and Active Directory together, they are explained separately in order to clarify the issues involved for each process.

10.3.1 Primary Restore

Perform a primary restore when all domain controllers in the domain are lost, and you want to rebuild the domain from the backup. (Do not perform a primary restore if any other working domain controller in this domain is available.) Use primary restore for the first domain controller. Later, use non-authoritative restore for all other domain controllers.

A primary restore builds a new File Replication Service (FRS) database by loading the data present under SYSVOL onto the local domain controller.

To perform a primary restore, start using the Agent to restore the System State. Click **Options** on the Select Restore Objects screen, and then mark the restored data as **Primary** in the Replicated data section.

Important: If this domain controller is a member of an FRS replica set other than SYSVOL, those other replica sets will also be restored as **Primary**. If you only want to restore the SYSVOL replica set, select **Primary**, but exclude other replica sets.

10.3.2 Non-Authoritative Restore

Perform a non-authoritative (normal) restore when at least one other domain controller in the domain is available and working. (Do not perform a non-authoritative restore when this domain controller is the only one in the domain.) You use a non-authoritative restore when you want this domain controller to receive SYSVOL data from a controller that has not failed.

A non-authoritative restore ignores all SYSVOL data that is restored locally. After a restart, the File Replication Service (FRS) receives SYSVOL data from its inbound partner domain controllers. After the non-authoritative restore completes, the SYSVOL tree on the local machine is the same as the SYSVOL tree on the inbound partners.

Performing a non-authoritative restore

1. Start ("boot") the machine in **Directory Services Restore Mode**.
2. While you are in **Directory Services Restore Mode**, open Windows CentralControl, and select the job that has run the System State backup.
3. Click the **Restore** icon.
4. Select the source from which to restore. Click **Next**.
5. On the **Select objects to restore screen**, select **System State**. Now click **Options**.

6. The System State Restore Options screen opens.
7. Under **Replicated data**, choose to mark the restored data as **Non-Authoritative**. Click **OK**. Click **Next** to proceed.
8. Restoring to the original location is enabled by default. Accept this **original location** setting. Click **Next** to proceed.
9. Under **Advanced restore options**, make selections or accept the default settings. Click **Next** to proceed.
10. Click **Finish** to launch the restore process.
11. Click **OK** when you are prompted to restart the computer.

In this mode, active directory data is restored to your computer from a backup, and then updated with the most current version from other domain controllers. For example, if the last backup was performed a week ago, and the active directory is restored non-authoritatively, changes made after the backup are replicated from the other domain controllers. The active directory replication system updates the restored data with newer data from your other servers.

10.3.3 Authoritative Restore

Perform an authoritative restore when you have accidentally deleted critical SYSVOL data from the local domain controller, and the deletion has propagated to other controllers. (Do not perform an authoritative restore if the local domain controller is not working, or it is the only controller in the domain.) You can perform an authoritative restore of SYSVOL only on a working domain controller (that is, changes to SYSVOL are replicating from this controller to other controllers).

An authoritative restore replicates any changes made to the current SYSVOL tree to its outbound replication partners.

When you perform an authoritative SYSVOL restore:

- Start using the Agent to restore the System State. On the System State Restore Options screen, specify an alternate location for the SYSVOL data in the Replicated data section. You will need the SYSVOL copies later.
- Do not click **OK** when you are prompted to restart the computer. Use NTDSUTIL to authoritatively restore Active Directory. This step is required because it is always advisable to restore Active Directory along with SYSVOL (so that they are not out of synch).
- Restart the system in normal mode and allow the SYSVOL to be published (this may take several minutes).
- Copy the old SYSVOL (from the alternative location) over the existing one.

Always authoritatively restore the SYSVOL whenever you authoritatively restore Active Directory, and vice-versa. This ensures that the SYSVOL and Active Directory are in synchrony.

1. Follow steps 1 through 9 as described in [Non-Authoritative Restore](#).
2. DO NOT CLICK **OK** when you are prompted to restart the computer. Perform the following steps instead.
3. From the Start menu, select **All Programs > Accessories > Command Prompt**.
4. At the command prompt, type `ntdsutil`.
5. At the **NTDSUTIL** prompt, type `authoritative restore`, and follow the documentation for NTDSUTIL regarding Active Directory objects that need to be restored authoritatively.
6. Restart the computer.

In this mode, Active Directory information is restored to your computer, and then replicated to all other machines on your domain. Use this recovery method if, for example, you accidentally delete users or groups from Active Directory, and you want to recover the system so that the deleted objects are recovered and replicated.

To authoritatively recover Active Directory data, you must run NTDSUTIL after you have performed a non-authoritative restore of the System State data, but before you restart the server in normal mode.

NTDSUTIL allows you to mark objects as authoritative. Marking objects as authoritative changes the update sequence number of an object so it is higher than any other update sequence number in the Active Directory replication system. This ensures that any replicated or distributed data that you have restored is properly replicated or distributed throughout your organization.

The NTDSUTIL utility can be found in the `systemroot\system32` directory.

If you experience problems starting your system after restoring the Active Directory, refer to [Resolving Common Errors](#).

10.4 Backing Up/Restoring Event Log Databases

Event log databases store events that are viewed through the Windows Event Viewer program. To back up all event logs currently available on the system, enable the **Event Logs** checkbox in the backup interface.

To restore event logs, enable the **Event Logs** checkbox in the restore interface.

10.5 Backing Up/Restoring Terminal Service License Databases

Terminal Service licenses are stored in a database that needs to be properly handled during backup/restore. To back up the database, select the **Terminal Service licensing** checkbox in the backup interface.

To restore the database, select the **Terminal Service licensing** checkbox in the restore interface.

10.6 Using the Exchange Plug-In

This section provides information about the workflows for backing up and restoring Microsoft Exchange databases and mailboxes.

This section also provides information about setting up a [DR safeset share](#).

For detailed information about this plug-in, see the *Exchange Plug-In Guide*. For information about supported platforms, see the Windows Agent release notes.

10.6.1 Working with Backups – Exchange 2010/2013/2016 DR

10.6.1.1 Creating a Job

1. Right-click an Exchange Agent, and select **New Job**. The New Job Wizard opens.
2. Select a **Backup source type** of **Exchange 2010/2013/2016 DR** from the list. The Exchange plug-in must be installed in order for Exchange options to appear here.
3. The **Encoding** type is **Unicode**, so some database names might not display properly in selection screens. You can back these databases up, though.
4. Click **Next** to continue.
5. Select a **Destination** for the backup. Click **Next**.
6. Enter a name for your job. Enter a job description (optional). Click **Next**.

10.6.1.2 Source Screen

1. On the **Source** screen, click **Add** to open the **Include/Exclude** dialog.
2. Information about the Exchange server will appear. Expand the entry to display its contents.
3. Select databases or components to back up, and click **Include**. The items that you choose will appear in the lower pane of the screen.
4. Click **OK** to continue.

10.6.1.3 Backup Options

1. After you have selected items to back up, click **Options**. The **Backup Options** screen will open.
2. Select a **Backup type** of **Full** or **Incremental**.
The first backup is always a full “seed” of the Exchange database, regardless of the **Incremental** or **Full** selection. After that, only the transaction logs are backed up when you select **Incremental**.
3. Enable **Validate Exchange database** to check the integrity of the Exchange data during backup.
Note: You should run a full DR backup at least once per week.

10.6.2 Exchange 2007 DR

10.6.2.1 Creating a Job

1. Right-click an Exchange Agent, and select **New Job**. The New Job Wizard opens.
2. Select a **Backup source type** of **Exchange 2007 DR** from the list. The Exchange plug-in must be installed in order for Exchange options to appear here.

3. The **Encoding** type is **Unicode**, so some database names might not display properly in selection screens. You can back these databases up, though.
4. Click **Next** to continue.
5. Select a **Destination** for the backup. Click **Next**.
6. Enter a name for your job. Enter a job description (optional). Click **Next**.

10.6.2.2 Source Screen

1. On the **Source** screen, click **Add** to open the **Include/Exclude** dialog.
2. Information about the Exchange server will appear. Expand this entry to display its contents.
3. Select components to back up, and click **Include**. The items that you choose will appear in the lower pane of the screen.

You cannot select specific databases to back up, but you can back up their Storage Groups. Then, from a Storage Group backup, you can restore a specific database.

4. Click **OK** to continue.

10.6.2.3 Backup Options

1. After you have selected items to back up, click **Options**. The **Backup Options** screen will open.
2. Select a **Backup type** of **Full** or **Incremental**.

The first backup is always a full “seed” of the Exchange database, regardless of the **Incremental** or **Full** selection. After that, only the transaction logs are backed up when you select **Incremental**.
3. Enable **Validate Exchange database** to check the integrity of the Exchange data during backup.
4. For Exchange 2007 configurations that support local replication, you can enable **Only back up active instance**. This causes the plug-in to use only active copies of databases during backup.

Note: You should run a full DR backup at least once per week.

10.6.3 Remaining Steps for Job Creation

1. **Options** screen
 - **Quick file scanning:** Enabling this option (where available) reduces the amount of data read during the backup process. Any file streams that are deemed unchanged since the last backup are skipped over. If this setting is disabled, files are read in their entirety.
 - **Disable deferring:** This option allows you to run the job without stopping, even if it means extending the run beyond the **Backup time window**.
 - **Backup type:** The first time you back up an Exchange database, it will be a backup type of **Full**. If you choose **Incremental** for your first backup, a **Full** backup will run instead. Subsequent backups can run as **Incremental**. (Remember that you should run a full DR backup at least once per week.)
2. **Encryption** screen: For new backup jobs, AES (256-bit) is the **Encryption type** for data on the vault. Enter a password and password hint.

Note: If an existing job uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES, None), 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.

Important: To recover data, you must remember the encryption password. If you forget the encryption password for a job, you cannot restore data from the job.

3. Choose from the **Log Options** that follow. Log files are created on the server machine (with the Agent) in directories using the job names.
 - **Create log file:** Enable this option to generate log files for each job executed. These printable log files report start-connect-completion and disconnect times, file names (i.e., the name of each file that was copied during a backup process), and any processing errors.
 - **Log detail level:** You can select a detail level of **None**, **Summary**, **Directories**, or **Files**. Detailed logging creates large log files, but this is useful for troubleshooting problems.
 - Changing the **Log detail level** only affects log files that are created from that point on. It does not affect any previously created log files.
 - **Automatically purge expired log files only:** You can automatically purge expired log files, or keep a selected number of them before they get deleted. The oldest file is deleted first.
 - **Keep the last <number of> log files:** You can specify how many log files to keep. When that number is reached, the oldest log file will be deleted to make space for the new one.
4. **Run the job immediately, Schedule the job, or Just exit from this wizard:** You can run your job immediately, or schedule it for later. If you click **Finish** and simply exit, the job will still be available.

10.6.4 Running Exchange DR Backups

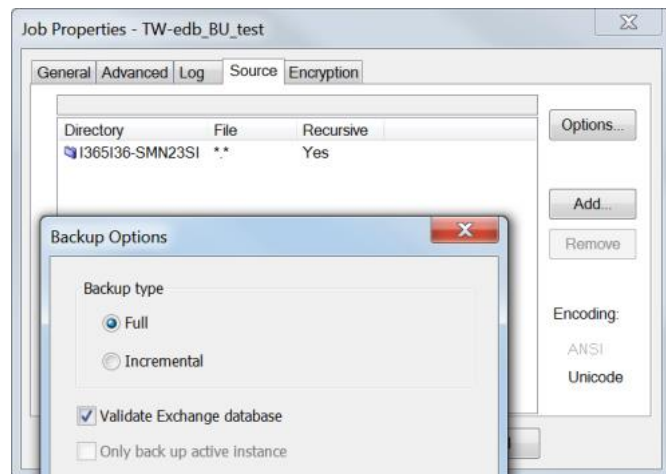
Exchange DR plug-ins support backup types of **Full** and **Incremental**. The first backup is always a full “seed” of the Exchange database, regardless of the **Incremental** or **Full** selection. After that, only the transaction logs are backed up when you select **Incremental**.

The following options can also apply:

- **Validate Exchange database:** Enable this option to check the integrity of the Exchange data during backup (Exchange 2007 DR and Exchange 2010/2013/2016 DR only).
- **Only back up active instance:** For configurations that support local replication, enable this option if you want to use only active copies of databases during backup (Exchange 2007 DR only).

10.6.5 Editing a Backup Job

Use the **Source** tab within Job Properties to change settings that are specific to Exchange plug-in jobs:



10.6.6 Working with Restores

Restore Mode (Exchange 2007 DR and Exchange 2010/2013/2016 DR Only)

Select what you want to restore.

Exchange Databases: Restore entire Exchange databases. This is the usual selection for disaster recovery (DR).

Mailboxes, messages and other Exchange objects: Choose this option to set up a DR VSS safeset as a shared resource for the **Granular Restore for Microsoft Exchange** application. (Depending on your configuration, this option might be disabled, or a warning about a related tool might appear.)

For information about the Granular Restore application, see the *Granular Restore for Microsoft Exchange* manual.

10.6.7 Exchange DR (Disaster Recovery)

Include/Exclude for Exchange 2007 DR and 2010/2013/2016 DR

On the Source or Select objects to restore screen, click the **Add** button to open the Include/Exclude screen.

From there, browse for items and select them. To include an item, highlight it and click **Include** (or double-click it). To exclude an item, highlight it and click **Exclude**. You must include at least one item in your backup or recovery.

If you choose an entire server, a confirmation screen will open. You can include all of the data items, or you can filter them according to their names. To see examples of the filtering scheme, click [here](#).


Note that wildcard searches do not work for item names that contain semicolons.

After you finish making your Include/Exclude selections, click **OK**.

10.6.8 Exchange 2007 DR and 2010/2013/2016 DR Options

Start hard recovery: Applies the database and replays the log files. Also rolls the logs forward.

A "roll forward" means to restore the information from the backup, keeping any log files created since the last backup.



For detailed information, see the *Exchange Plug-In Guide*.

10.6.9 Granular Restore for Microsoft Exchange

You can expose an Exchange 2007 DR or 2010/2013/2016 DR safeset as a shared resource for use with the Granular Restore for Microsoft Exchange and SQL application.

When you expose a DR safeset as a shared resource (a "share"), you can use Granular Restore to restore mailboxes, messages, and other Exchange objects to a Personal Storage Table (PST) file. The Granular Restore application is available as a separate installation kit. You can download this application and run it on your Exchange server.

The application allows you to:

1. Browse within the Exchange database (from an Exchange 2007 DR or 2010/2013/3016 DR backup)
2. Select individual or multiple items (such as messages or mailboxes) to export to a PST file
3. Import the PST file back into your Exchange system

For more information about the Granular Restore application, see the *Granular Restore for Microsoft Exchange and SQL* guide.

10.6.10 Setting Up a DR Safeset Share

(On the source selection screen, choose a safeset from which to restore.)

In order to use the Granular Restore for Microsoft Exchange and SQL application, you need to set up an Exchange safeset as a shared resource (a "share").

To adjust the options for this setup, click Advanced Share Options. The Advanced Share Options screen will open.

Share options

In the **Idle time** field, enter the number of minutes of inactivity after which the share should automatically stop. This value can range from **2** to **180**.

Bandwidth option

By default, starting the share uses all available bandwidth. If you do not want this, clear the **Use all available bandwidth** option. This allows the *bandwidth settings that were configured for the Agent* to apply.

Click **OK** to save your settings. Click **Next** to continue through the **Encryption options** and **Share summary** screens.

10.6.10.1 Share Summary and Path

Review the settings on the **Share summary** screen.

Click **Share** to continue. Otherwise, click **Cancel** or **Back**.

When you click **Share**, the Process Information screen opens. This screen reports the status of the share process. It also returns the share path, which you can copy to use elsewhere. (Depending on your configuration, the share process might not succeed, or a warning about a related tool might appear.)

When the share is ready to use, copy its path to your Windows clipboard. You can now paste the path into the Granular Restore application.

If you no longer need to share the safeset, click **Unshare**. When you click **Unshare** (or the share period expires), the **Process Information** screen indicates that the share is no longer available. The original path for the share also displays.

To see the **Process Information** screen again, open the **Processes** folder for the Agent, and double-click the process.

10.7 Using the SQL Server Plug-In


The SQL Server Plug-in integrates into the existing Agent architecture, allowing you to back up/restore SQL Server databases to/from the vault.

Windows CentralControl communicates with the Agent and the plug-in.

Backups and restores require Windows or SQL Administrator rights.

10.7.1 Creating a SQL Server Backup Job

1. Right-click a SQL Server Agent, and select **New Job**. The New Job Wizard starts.
2. Select a **Backup source type** (i.e., **SQL Server**) from the list. (Your list might look different from the one in the diagram.) The SQL Server Plug-in must be installed in order for **SQL Server** to appear here.



3. Click **Next** to continue.
4. Select a **Destination** for the backup. Click **Next**.
5. Enter a name for your job. Enter a job description (optional). Click **Next**.

10.7.1.1 Selecting a SQL Server Instance

In the list of available SQL Server instances, select a SQL Server instance.

The list begins with an entry named **<Default>**, designating the default instance. Other instances will appear if they can be retrieved. If there is only one entry, you must select **<Default>**.

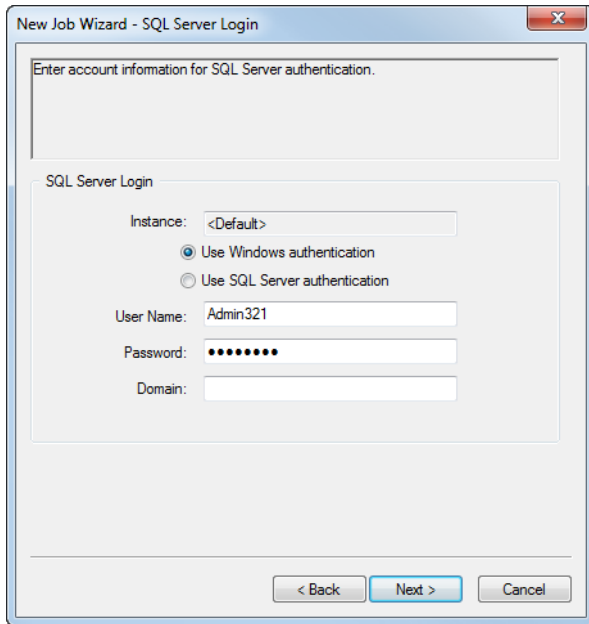
10.7.1.2 SQL Server Login

The **SQL Server Login** screen allows you to select the authentication type (Windows or SQL Server).

- For SQL Server authentication, access is controlled through the User Name and Password on this screen. Enter a user name and password to connect to the selected instance. The maximum length of the password is 31 characters.
- For Windows authentication, access is controlled by default through the Windows login.

The **Instance** name appears here in read-only format, for information purposes.

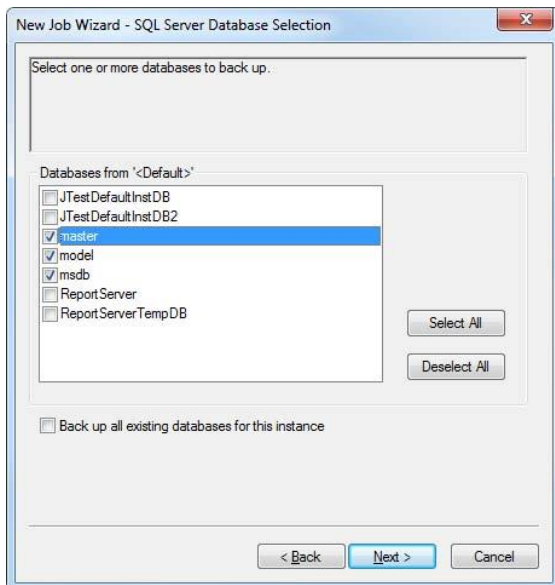
After you have supplied valid credentials, click **Next**.



10.7.1.3 SQL Database Selection

Select one or more databases to back up.

You can enable **Back up all existing databases for this instance** (so that if databases are added to or removed from the instance, the job will automatically include all of them at backup time).



It is strongly recommended that you back up the system databases (**master**, **model** and **msdb**) in one job. Then back up any other (user) databases in one or more other jobs.

For a disaster recovery scenario, you must restore the **master** database first, by itself. Then you can restore the other system databases.

10.7.1.4 Remaining Steps for Job Creation

1. Options screen

Quick file scanning: Enabling this option reduces the amount of data read during the backup process. Any file streams that are deemed unchanged since the last backup are skipped over. Without this setting, files are read in their entirety.

Disable deferring: This option allows you to run the job without stopping, even if it means extending the run beyond the **Backup time window**.

Backup type: The first time you back up a SQL Server database, it will be a backup type of **Full** or **Full with Include Transaction Logs**. If you run a backup type of only **Transaction Logs** as a first backup, the backup will be **Full** with **Include Transaction Logs**. Subsequent backups will be **Transaction Logs** only.

For subsequently scheduled backups, you can choose between **Full**, **Full with Include Transaction Logs**, and **Transaction Logs**. **Full** backs up the entire database. **Full with Include Transaction Logs** backs up the entire database with the latest transaction logs. **Transaction Logs** (on its own) backs up the transaction logs.

Note: Transaction logs are not backed up if the database recovery model is set to **Simple**.

2. Encryption: For new backup jobs, AES (256-bit) is the **Encryption type** for data on the vault. Enter a password and password hint.

Note: If an existing job uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES, None), 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.

Important: To recover data, you must remember the encryption password. If you forget the encryption password for a job, you cannot restore data from the job.

3. Choose from the **Log Options** that follow. Log files are created on the server machine (with the Agent) in directories using the job names.

- **Create log file:** Enable this option to generate log files for each job executed. These printable log files report start-connect-completion and disconnect times, file names (i.e., the name of each file that was copied during a backup process), and any processing errors.
- **Log detail level:** You can select a detail level of **None**, **Summary**, **Directories**, or **Files**. Detailed logging creates large log files, but this is useful for troubleshooting problems.

Changing the **Log detail level** only affects log files that are created from that point on. It does not affect any previously created log files.

- **Automatically purge expired log files only:** You can automatically purge expired log files, or keep a selected number of them before they get deleted. The oldest file is deleted first.
- **Keep the last <number of> log files:** You can specify how many log files to keep. When that number is reached, the oldest log file will be deleted to make space for the new one.

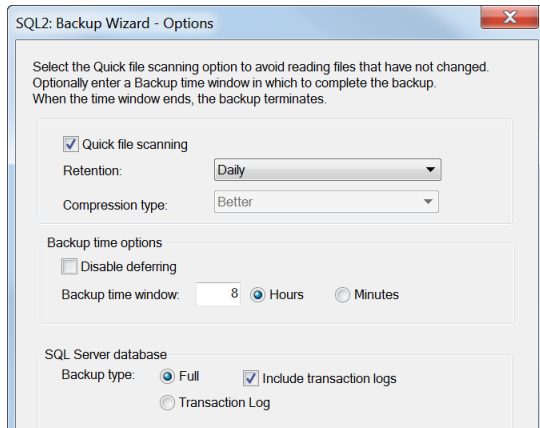
4. **Run the job immediately, Schedule the job, or Just exit from this wizard:** You can run your job immediately, or schedule it for later. If you click **Finish** and simply exit, the job will still be available.

10.7.2 Running a Backup

Backups and restores for the SQL Server plug-in require Windows or SQL Administrator privileges.


The plug-in supports backup types of **Full**, **Full with Include Transaction Logs**, and **Transaction Logs**. **Disable deferring** is only available for VSS jobs.

Settings that are specific to this plug-in appear on the **Options** screen of the Backup Wizard:



10.7.3 Editing a Backup Job

Use the **Source** tab within Job Properties to change settings that are specific to SQL Server plug-in jobs:



10.7.4 Restoring SQL Data

For a SQL Server database or transaction-log restore, you must select a backup (safeset) from which you would like to restore. Restores can be to the original location, an alternate location, or a file on disk.

Here are some restore scenarios:

- Restoring the full database, with any incremental backups, overwriting the existing database.
- With no system backup, restoring the system from the ground up (“bare metal restore”) – installing the OS, applications, and then the full database (plus any incremental backups) onto a new system.

- If there is a SQL Server backup and a full system backup, install the OS, and then restore the System State and SQL Server system.

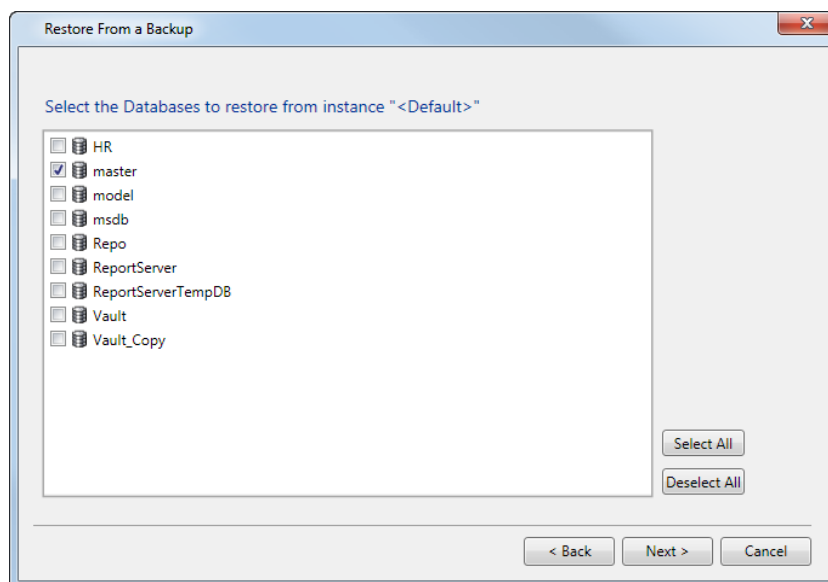
10.7.5 Restoring from a SQL Server Backup Job

1. Right-click a SQL Server Plug-in backup job, and select **Restore**. The Restore from a Backup wizard starts.
2. On the Choose How to Restore page, select one of the following, and then click **Next**:
 - **Restore to SQL Server** (Select databases to restore directly into a SQL Server instance.)
 - **Restore to flat files** (Select databases to restore as flat files.)
 - **Granular restore of SharePoint data**
3. On the Select a Source page, do the following, and then click **Next**:
 - Choose a source from which to restore: a vault or **Alternate safeset location**.
 - Choose a safeset.
 - Enter the data encryption password for the backup data.
4. On the Select the Databases to Restore page, select one or more databases to restore. You must select at least one database in order to move to the next screen.

If you use **Select All**, and databases are added to/removed from the instance, the restore will automatically include all of them when it runs.

For an [Alternate database name](#) recovery, select one database only. If you need to restore more than one database, you must restore each database on its own (i.e., one database per restore operation). An **Alternate database name** recovery always restores the data as a file.

For a disaster recovery scenario, you must restore the **master** database first, by itself. Then you can restore the other system databases.



10.7.5.1 Restore to SQL Server: Server Instance

Select the SQL Server instance where databases should be restored, and provide credentials. You can use Windows or SQL Server authentication.

For Windows authentication, access is controlled by default using the Windows login credentials.

If you wish, click **View selected databases**. The **Selection Summary** screen will open, listing the databases that you have selected within the instance.

The screenshot shows the 'Restore From a Backup' dialog box. The title bar reads 'Restore From a Backup'. The main content area has the following elements:

- A heading: 'Select the SQL Server instance where the database should be restored'.
- A dropdown menu currently showing '<Default>'. A mouse cursor is pointing at the dropdown arrow.
- A heading: 'Provide the credentials for the instance'.
- Two radio buttons:
 - Use Windows authentication
 - Use SQL Server authentication
- Text labels and input fields:
 - User Name: Admin321
 - Password: [masked with dots]
 - Domain: [empty field]
- A blue link: 'View selected databases'.
- At the bottom, three buttons: '< Back', 'Next >', and 'Cancel'.

10.7.5.2 Restore to SQL Server: Where to Recover Databases

The screenshot shows the 'Restore From a Backup' dialog box. The title bar reads 'Restore From a Backup'. The main content area has the following elements:

- A heading: 'Select where to recover the database(s)'.
- Two radio buttons:
 - Original database names
 - Alternate database name: [empty text box]
- A heading: 'Recovery options'.
- Two checkboxes:
 - Overwrite existing databases
 - Restore using No Recovery option
- A heading: 'Fallback file location'.
- Text: 'If the original drive for the database files is not available, recover the files to:'.
- Two radio buttons:
 - Default location of the selected instance
 - Alternate path: [empty text box] [Browse button]

The following recovery scenarios can target the same SQL instance or a different SQL instance:

- **Original database names**
- **Original database names, with Overwrite existing databases**
- **Original database names, with Restore using No Recovery option**

- **Original database names**, with **Overwrite existing databases** and **Restore using No Recovery option**
- **Alternate database name**
- **Alternate database name**, with **Restore using No Recovery option**

About the **Alternate database name** option:

- Only available for single-database recoveries. If you have multiple databases, you must restore them one at a time.
- Always restores the database as a file.
- Recovers over an existing database. The application will create a new database if the database does not exist within the SQL instance.

About the **Restore using No Recovery option**:

After restoring, you must apply transaction logs manually (e.g., through SQL Management Studio).


About the **Fallback file location**:

The **Default location of the selected instance** is defined during database installation.

10.7.5.3 Restore to flat files: Destination

On the restore options screen, enter a **Destination** path, or click **Browse** to navigate to a location.

For information about the other options on this screen, see [Remaining Restore Steps](#).

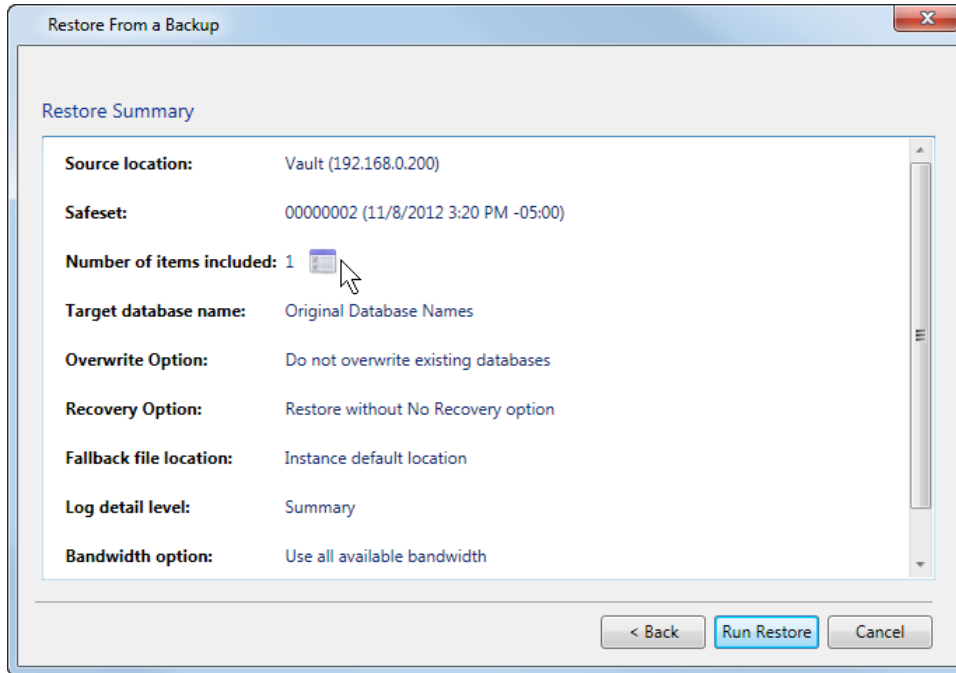


The screenshot shows a dialog box titled "Restore From a Backup". Under the heading "Select the restore options", there are three sections: "Destination" with a text box and a "Browse" button; "Log file detail level" with a dropdown menu set to "Summary"; and "Bandwidth option" with a checked checkbox for "Use all available bandwidth".

10.7.5.4 Remaining Restore Steps

1. Select a Log file detail level.
2. Enable or disable the **Use all available bandwidth** option, and click **Next**. A screen showing the restore settings will open.
3. Review the list of restore settings.

4. If you wish, click the icon on the line for **Number of items included**. The **Selection Summary** screen will open, listing the selected databases.
5. To revise your settings, click the **Back** button.
6. When you are ready, click **Run Restore**.



10.8 Using the Cluster Plug-In

A cluster consists of two or more computers that work together to provide higher availability, reliability, and scalability than you can obtain from a single computer.

The Cluster plug-in allows you to configure an Agent and jobs on a virtual node of a cluster. This allows jobs to redirect workflows to the active node when a failover occurs in the cluster.

Use the Virtual Cluster Agent to create jobs to protect the applications that run on the virtual node of the cluster.

Note: If you use Windows authentication for jobs in a clustering environment, you must supply a domain name.

Jobs created with this Agent will fail over with the virtual node, so they can continue to protect applications without reseeding (regardless of which physical node the virtual node runs on). The Agent can still access its configuration (on a shared drive), and scheduled backups can occur as usual without appearing to be "different" backups (which cause reseeding). However, if a failover occurs when a backup is in progress, the backup will fail, and you will need to run it again. To be notified of failures, adjust the notification settings.

If there are multiple virtual nodes in a cluster where each virtual node is set up to protect a specific application, create the job to protect the application through the Virtual Cluster Agent for the virtual node.

To use the Agent and the Cluster plug-in in a Windows Cluster environment, set up these conditions:

- Install the Agent and Cluster plug-in (plus any other plug-ins required to protect the data on the cluster) on each physical node of the cluster. All of the Agents must be installed with the same set of plug-ins on all of the servers in the cluster.
- Register each physical node with the same vault.
- To create jobs to protect the applications and files that are associated with the virtual nodes of the cluster, [configure a Virtual Cluster Agent](#). The Virtual Cluster Agent is associated with the virtual node. Its configuration files and jobs are stored together in the cluster. They will fail over between the physical nodes as the virtual node fails over.
- To help protect the physical systems, create a Bare Metal Restore job for each physical system.

10.8.1 Configuring the Cluster Plug-In

The Virtual Cluster Agent is for creating jobs to protect the applications that run on the virtual node of the cluster.

When you connect to a virtual Agent, the Agent icon displays a white screen:



You must configure your virtual Agent before you can create jobs on the virtual node. In the right-hand pane, double-click the **Global** file. The **Agent Configuration** dialog will open.

Note: Register the virtual node with the same vault that the physical nodes use.

Jobs created with this Agent will fail over with the virtual node, so they can continue to protect these applications without reseeding (regardless of which physical node the virtual node runs on). However, if a failover occurs when a backup is in progress, the backup will fail, and you will need to run it again. To be notified of failures, adjust the notification settings.

If there are multiple virtual nodes in a cluster where each virtual node is set up to protect a specific application, create the job to protect the application through the Virtual Cluster Agent for the virtual node.

10.8.2 Microsoft Server Clustering Services (MSCS)

Vault profiles, schedules and jobs can be associated with a physical node in a cluster in order to back up the System Volume and System State.

Vault profiles, schedules and jobs can be associated with a cluster virtual server. On failover, this information will “travel” with the cluster virtual server to the new physical node in the cluster.

From a user’s perspective (in the GUI), one server is shown for each physical node, and one server for each cluster virtual server. Virtual servers display a specific icon.

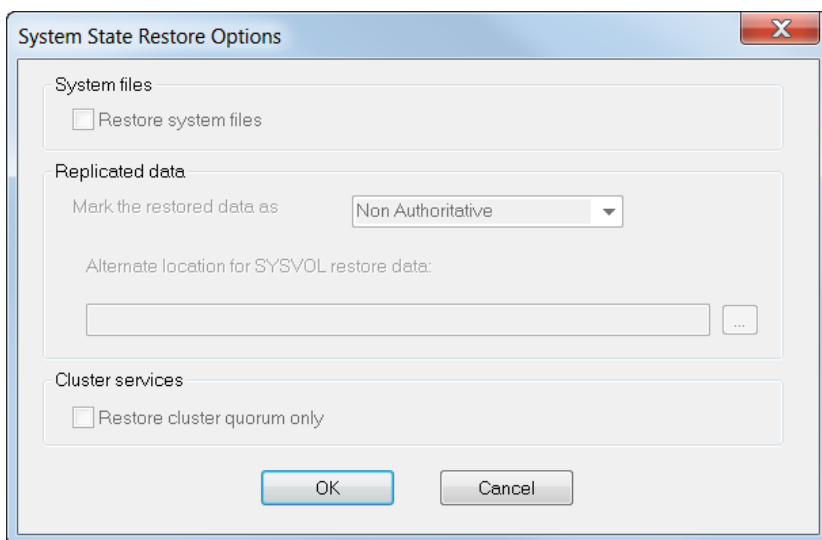
The Process Information screen for the virtual node only shows processes related to virtual node jobs. The Process Information screen for the physical node only shows processes for physical node jobs. After failover, any existing processes from the virtual node will cease to be visible.

10.8.3 Cluster Quorum Restore Operations

CentralControl offers the option to restore only Windows cluster quorum data. If you select **Restore cluster quorum only**, only cluster quorum data is restored. If **Restore cluster quorum only** is not selected (default), both the cluster quorum and System State data are restored.

To perform a cluster quorum only recovery:

1. On the Select objects to restore screen, enable the **System State** checkbox.
2. Click **Options**. The System State Restore Options screen opens.
3. Select **Restore cluster quorum only** to restore cluster data only.
4. Click **OK**.
5. Continue to follow the Restore wizard.
6. When you are prompted to restart, choose **Respond to request later** (if you are planning to recover the quorum disk). Use the **clusrest.exe** utility, and then restart the computer.



10.9 Using the Oracle Plug-in

To protect Oracle databases, install the Oracle Plug-in with the Windows Agent on the Oracle database server. You can then add and run backup jobs that specify which databases to back up, and where to save the backup data.

The Plug-in provides ARCHIVELOG-based, non-RMAN backups of whole online database instances. All non-temporary tablespaces and instance parameter files are automatically backed up. Full and partial databases are restored through normal user-managed Oracle recovery mechanisms.

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

For installation and configuration information, see the Windows Agent guide or Portal online help. For supported platform information, see the Windows Agent release notes.

10.9.1 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.

10.9.2 Creating an Oracle Backup Job

To back up an Oracle database, install the Agent on the same system as the Oracle database server. Create a new Job using "Oracle" as the Backup Source Type. The New Job wizard will direct you through the process. Briefly, the steps are:

1. Open Windows CentralControl and create a new job.
2. Select Oracle in the Backup Source Type list. The Oracle Options will appear on the page.
3. Enter the Database Service Name, User Name, and Password.

For Oracle 11g in Windows CentralControl, set the Oracle Service Name to the Database Instance from Oracle (rather than the Instance Name from Oracle).

Jobs back up only one database at a time. There can be more than one Job doing backups on different databases (but you cannot run multiple Jobs at the same time on the same database).

4. Select or confirm the databases that you want to back up.
5. For new backup jobs, AES (256-bit) is the **Encryption type** for data on the vault. Enter a password and password hint. Also, select any advanced options (e.g., compression and logging levels) that you want.
6. Specify a schedule if you wish. Oracle Corporation recommends that backups take place in periods of low database activity.
7. Choose a destination (i.e., vault) for the backup data.
8. Start the backup immediately, or let it run on a schedule.

10.9.3 How Backups Work

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.

Note: Configuration files that are not instance-specific (such as 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.

10.9.4 Control File Name and Format

The control file is backed up with a different name. The control file's format is:

```
BACKUP_ORACLE_SID_CONTROLEFILE_safesetname.CTL
```

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.

10.9.5 Restoring Oracle Databases

Restores might be necessary in a variety of situations:

- A requirement to restore the full database.
- With no system backup, restoring the system from the ground up (“bare metal”) – installing the OS, applications, and then the full database (plus any transaction logs) onto a new system.

If there is an Oracle backup and a full-system backup, 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 then 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:

- Shut down the database.
- Restore the files.
- If necessary, reset the control information for the database.
- Start and recover the database.
- Re-open the database for use.

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

10.10 Encryption, Compression and OTW

During a backup, you can:

- Encrypt the backup data so that it is secure when it is stored on the vault. For new backup jobs, AES (256-bit) is the Encryption type for data on the vault. If an existing job uses another encryption type (e.g., AES 128-bit, Blowfish, DES, Triple DES, None), 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. To recover data, you

must remember the encryption password. If you forget the encryption password for a job, you cannot restore data from the job.

- Compress the data, so that it can be transmitted and stored on the vault efficiently. Compressing the data may allow for more efficient transmission and storage. This depends on whether or not you have sufficient CPU, communication bandwidth, and storage space.
- Use Over-The-Wire encryption (OTW). OTW compresses the backup data, as well as any other communications between Agents and vaults (such as commands). This is important if there are concerns about unauthorized people seeing this communication.

11 Carbonite Server Backup Support

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

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

What can we help you with?


Popular Searches

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

11.1 Contacting Carbonite

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

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



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.