EVault Software

Agent for Linux and Oracle Plug-in 8.1

User Guide

Revision: This manual has been updated for version 8.1.
Software version: 8.10 (July 2016)

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.

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

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.

The EVault Software Agent, EVault Software CentralControl, and EVault Software 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 EVault Software Agents and EVault Software Director applications also have the added security feature of an over the wire encryption method.

Contents

1 Introduction to the Linux Agent 6

2 Installing, upgrading, and uninstalling the Agent 7

2.1 System requirements 7

2.1.1 Privilege requirements 7

2.2 Installing the Agent 8

2.2.1 Installation script – install.sh commands 11

2.2.2 Starting and stopping the Agent 12

2.2.3 Portal registration 12

2.2.4 Language selection 13

2.2.5 Kernel configuration parameters 14

2.3 Upgrading the Agent 14

2.4 Uninstalling the Agent 15

3 Configuring the Agent 16

3.1 Quick Start steps 16

3.2 Create an Agent profile 17

3.3 Configure the vault – Agent configuration 18

3.4 Create a backup job 19

2.5.1 Adding a file or directory to a new backup job 20

2.5.2 Adding/removing a file or directory with an existing backup job 21

3.5 Schedule the job 22

4 Running backups 24

4.1 Running an ad-hoc backup 24

4.1.1 File backup options 25

4.2 Check the backup results 26

5 Restoring data 28

5.1 Restoring from a backup 28

5.1.1 Symbolic links 29
5.1.2 NFS – Network file system 30
5.1.3 NFS support for ACLs and extended attributes 30
5.2 Cross-computer restores 30
5.3 Disaster recovery 31
5.4 Restoring ACLs 31
6 System recovery 33
6.1 Hardware requirements 33
6.2 Software requirements 33
6.3 Recovery steps 34
6.2.1 Install the minimal operating system 34
6.2.2 Install and configure the Agent 34
6.3.3 Restore the backed up system 34
6.2.3 Perform post-recovery maintenance 35
6.2.4 Verify the recovery 35
6.2.5 Recovery problems 35
7 Oracle Plug-in 36
7.1 Overview 36
7.1.1 Features 36
7.1.2 Limitations 36
7.2 Installing the Oracle Plug-in 36
7.2.1 System requirements 37
7.2.2 Supported platform combinations 37
7.2.3 Before installing or upgrading 37
7.2.4 Installing the Plug-in 37
7.2.5 Uninstalling the Plug-in 38
7.2.6 Before you run the Plug-in 38
7.3 Backups 40
7.3.1 Table of backup information 40
7.3.2 Oracle instance protection 41
7.3.3 How the backup works 42

7.4 Restores 42

7.4.1 Guidelines for restoring 43

1 Introduction to the Linux Agent

This guide is intended for the administrator responsible for ensuring that Linux systems are configured properly for backups. This guide will show the administrator how to configure the Agents, select the data to be backed up, and schedule when the backups will be run. Those who use the servers do not necessarily need to be aware that their systems are being backed up.

Linux Agent is available as a 32-bit application and a 64-bit application. You do not need to install 32-bit compatibility libraries on 64-bit Linux distributions to run the 64-bit Linux Agent.

This guide describes how to install, configure and manage the Linux Agent and Oracle Plug-in. The guide includes procedures for running backups and restores using the legacy Windows CentralControl. For Portal procedures, see the Portal online help.

Each computer that needs to be backed up must have the Agent software installed, running, and connected to the network. The Agent runs on the computer as a background service, and starts automatically when the system starts.

Backups and restores on the Agent computers are configured and scheduled by Portal or Windows CentralControl. The Agent communicates its backup data directly to the vault.

The Oracle Plug-in is an add-on to the Agent. It allows a user to back up an Oracle database. The Plug-in is installed with the Agent on the database host to perform backups.

2 Installing, upgrading, and uninstalling the Agent

This section describes how to install the Agent for Linux. The installation requires that you have the Agent for Linux installation kit and a system running Linux.

There is a separate installation kit for the Oracle Plug-in for Linux. For more information, see Installing the Oracle Plug-in.

2.1 System requirements

The Linux Agent is available as a 64-bit application and a 32-bit application.

Note: The 64-bit version of the Linux Agent can only be installed on 64-bit systems. The 32-bit version of the Linux Agent can only be installed on 32-bit systems.

The usrlibacl.so.1 shared library is required and must be installed on the protected system.

2.1.1 Privilege requirements

☑ Installation

To extract the installation files for the Agent, no special privileges are required. However, to run the installation script, you must have root privileges.

☑ Functional

To communicate with the Agent remotely, the user specified must have full root privileges.

Note: Enhanced privileges are required for the User ID that you use for Windows CentralControl Agent authentication.

To log in to the account:

* The User ID must be enabled.

* The User ID must not be suspended or locked out due to invalid password attempts.

* The password must not have expired.

* The User ID should not have time of day limits for when you can log in.

* The User ID must belong to the "root" group.

For additional security, rather than disabling the account, you can set the shell to be `/bin/false`. This can usually be done with the following command:

```
usermod -s /bin/false buagent
```

If you are running an ftp server, for additional security, it may be necessary to review the ftp server configuration to deny logins for this User ID. It is often sufficient to ensure /bin/false is not listed as a valid shell in /etc/shells. Your server may vary.

Note: You can install the Agent from "fresh" or upgrade it.

2.2 Installing the Agent

The installation kit is provided as a tar.gz file. This must be unzipped only on the machine it is intended for (the target machine). That is, do not unzip it on another type of machine. This may cause unpredictable results.

The amount of disk space needed for the installation varies from system to system. In all cases, the installation program will determine if there is enough disk space for the installation to continue. (This determination also includes any temporary space required for an upgrade.)

Before beginning the installation, make certain that the following requirements and materials are available:

- The Agent for Linux installation kit.
- A target system running a supported version of Linux.
- Root privileges on the target system to install the product.
- Sufficient disk space for the new installation, and later job activities. Note that if the available disk space is insufficient for a complete installation, the installation directory will roll back to its original state. You can override the space checking with `./install.sh` (but this only applies to the initial check for actual installation space, rather than extra temporary space for rolling back).

To install the Agent:

1. Download the 64-bit or 32-bit Linux Agent tar.gz installation kit.

   Note: You must do this locally (i.e., on the target machine).

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

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

   ```
   >  tar -zxf packageName.tar.gz
   ```

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

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

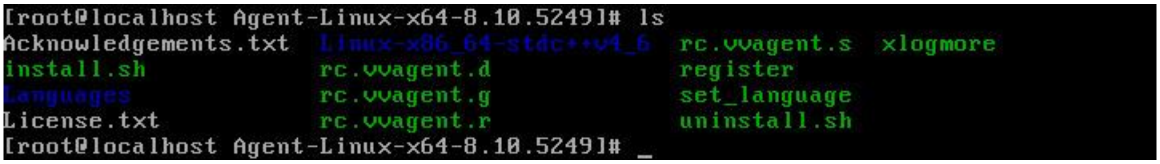

```
[root@localhost Agent-Linux-x64-8.10.5249]# ls
Acknowledgements.txt  Linux-x86_64-stdc++v4_6  rc.vvagent.s  xlogmore
install.sh            rc.vvagent.d             register
Languages             rc.vvagent.g             set_language
License.txt           rc.vvagent.r             uninstall.sh
[root@localhost Agent-Linux-x64-8.10.5249]# _
```

3. Run the following command to start the installation:

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

4. Press **Enter** for the installation directory.

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

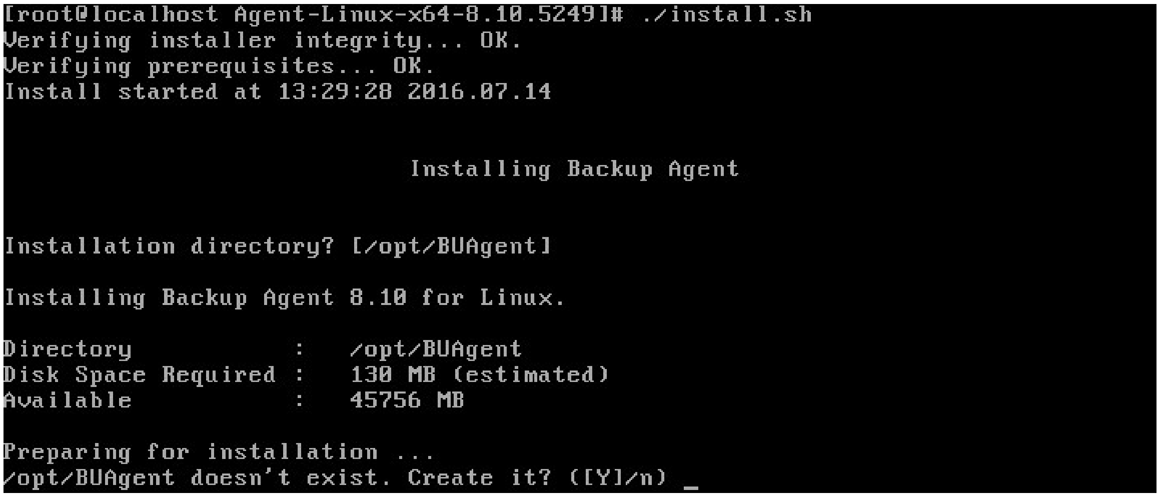

```
[root@localhost Agent-Linux-x64-8.10.5249]# ./install.sh
Verifying installer integrity... OK.
Verifying prerequisites... OK.
Install started at 13:29:28 2016.07.14


                            Installing Backup Agent


Installation directory? [/opt/BUAgent]

Installing Backup Agent 8.10 for Linux.

Directory           :   /opt/BUAgent
Disk Space Required :   130 MB (estimated)
Available           :   45756 MB

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

5. Enter **Y** to create the BUAgent directory.

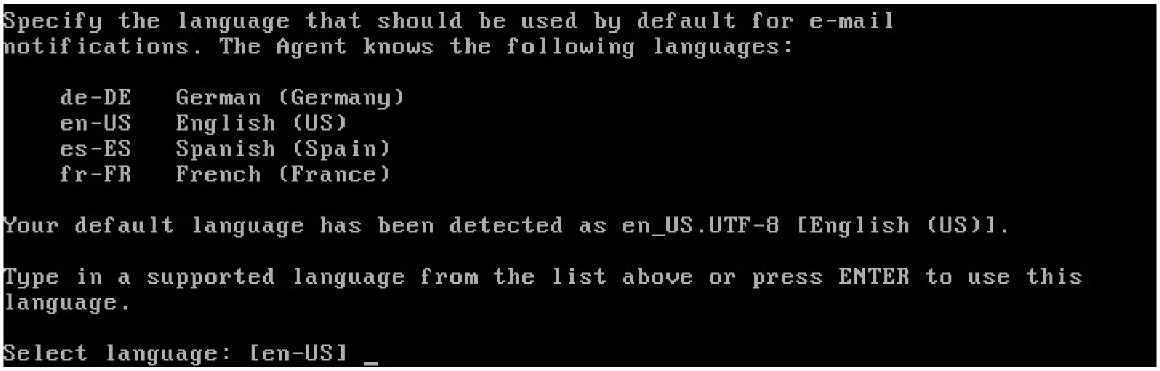

```
Specify the language that should be used by default for e-mail
notifications. The Agent knows the following languages:

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

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

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

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

6. Enter the language. The default language is English [en-US].

7. At the Register to Portal prompt, enter **Y**.

8. At the Portal address prompt, enter the Portal address.

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

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

11. At the Portal password prompt, enter the password for the user specified in Step 10.

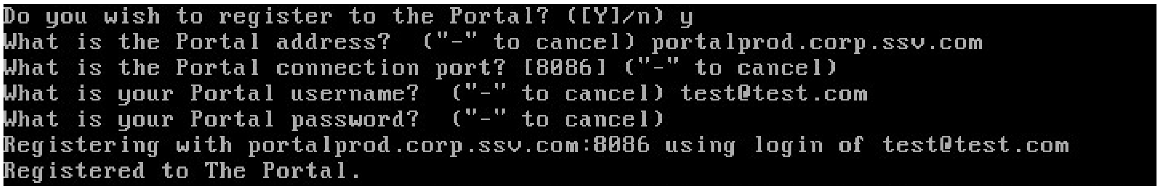

```
Do you wish to register to the Portal? ([Y]/n) y
What is the Portal address?  ("-" to cancel) portalprod.corp.ssv.com
What is the Portal connection port? [8086] ("-" to cancel)
What is your Portal username?  ("-" to cancel) test@test.com
What is your Portal password?  ("-" to cancel)
Registering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

The next step in the installation is to choose the encryption method. By default, the Agent encrypts data using an encryption method that is integrated in the Agent. For audit purposes, some organizations require the Agent to use an external encyption library that is provided with the Agent. Using the external encryption library can degrade Agent performance.

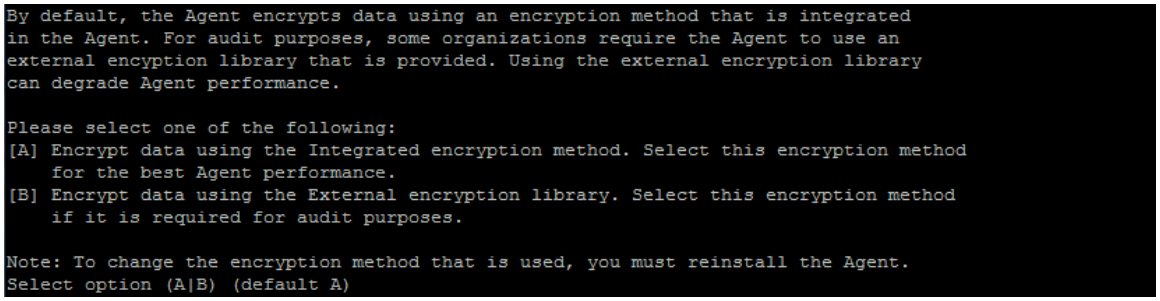

```
By default, the Agent encrypts data using an encryption method that is integrated
in the Agent. For audit purposes, some organizations require the Agent to use an
external encyption library that is provided. Using the external encryption library
can degrade Agent performance.

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

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

12. Do one of the following:

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

    The installation proceeds. When complete, a completion message will appear, and the Agent daemon will be running.

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

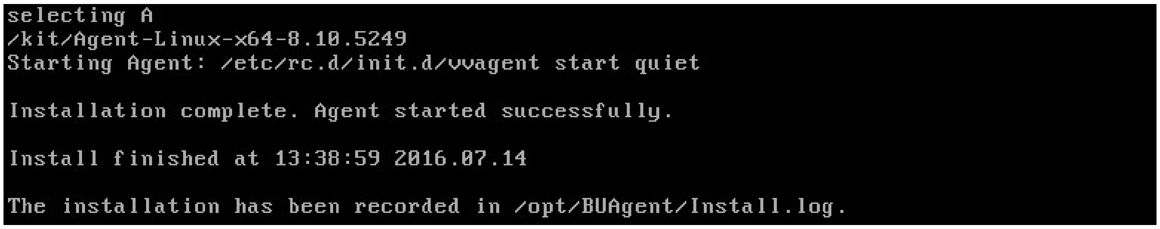

```
selecting A
/kit/Agent-Linux-x64-8.10.5249
Starting Agent: /etc/rc.d/init.d/vvagent start quiet

Installation complete. Agent started successfully.

Install finished at 13:38:59 2016.07.14

The installation has been recorded in /opt/BUAgent/Install.log.
```

2.2.1 Installation script – install.sh commands

The installation directory has a shell file called “install.sh”. The “- help” option shows the commands that are available for installation.

```
Usage: install.sh [options]

-shutdown | -s          Force the Agent shutdown, if running.
-force | -F             Force the installation; skip the initial free
                        space check.
-defaults | -D          Use the default values for installation.
-force-defaults         Force the installation using the defaults
                        (assumes -s and -F).
-web-registration=off   Turns off web console registration.
    -W-
-web-registration=FILE  Attempts to register to the web console with
    -W=FILE             the values in FILE.
-quiet | -Q             Quiet install; does not echo output to the
                        screen. If user interaction is required in
                        quiet mode, the install will fail unless
                        -force-defaults is specified.
-log=NAME | -L=NAME     Writes the installation log to the specified
                        FILE.
-lang=NAME | -l=NAME    Selects NAME as the language. Must begin with
                        an ISO language code. May optionally be
                        followed by a dash or underscore and an ISO
                        country code (e.g., fr, fr-FR, and fr_FR are
                        acceptable). Character set markers (e.g.,
                        UTF-8) are ignored. Languages that cannot be
                        matched will report an error and the language
                        will be defaulted to en-US [English (US)]. If
                        not specified, the language will be guessed
                        from your system value of "en_US.UTF-8".
-backup=DIR | -B=DIR    Backs up the current installation of the Agent
                        to the specified directory.
-verify | -V            Verifies the integrity of the installation kit.
-help                   Shows this text.
```

Registration options

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

```
wccAddress=PortalAddress

wccPort=PortalConnectionPort # Defaults to 8086

wccLogin=PortalUsername

wccPassword=PortalPassword
```

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

2.2.2 Starting and stopping the Agent

Stop and Start commands for the Agent are determined by the specific OS version that you use. They are actually “rc” (run control) scripts. You can determine the location of these scripts by viewing the “Install.log”.

The “vvagent” script is used to start, stop or check the status of both vvagent (for Agents controlled by Windows CentralControl) and buagent (for Agents controlled by Portal). This single script affects both vvagent and buagent.

For example:

```
Linux:   /etc/init.d/vvagent {start/stop/restart/status}
```

In this example, “stop”, “start”, “restart” and “status” are the parameters.

2.2.3 Portal registration

During Agent installation, you are prompted to register the Agent with Portal.

After installation, you can change the registration (reregister). You must stop the Agent before reregistering, and you must restart it for the changes to take effect.

Registration Script:

Run `<Agent installation directory>/register` to register the Agent with Portal.

If you are already registered to a Portal server, you will see:

```
Do you wish to register as a new computer?

This will invalidate your previous registration. (y/[N])
```

For a new registration or a reregistration, you will be prompted as follows:

```
What is the Portal address?  ("-" to cancel)
What is the Portal connection port?  ("-" to cancel)

What is your Portal username?  ("-" to cancel)

What is your Portal password?  ("-" to cancel)
```

The address is the name or IP address of Portal.

The port number is defined by the Portal Administrator.

Your user name/password authentication is set by the Portal administrator.

2.2.4 Language selection

During the Agent installation, you are asked to choose a default language for email and command-line log viewing.

After installation, you can change the default language. The Agent must be restarted for these changes to take effect.

Language Selection:

Run `<Agent installation directory>/set_language` to specify the language that the local Agent will use for e-mails and command-line log viewing.

```
Specify the language that should be used by default for e-mails
and command-line log viewing. The Agent knows the following
languages:

    de-DE, en-US, es-ES, fr-FR

Which language do you want? [en-US]
```

de-DE is German (Germany)

en-US is English (USA)

es-ES is Spanish (Spain)

fr-FR is French (France)

Refer to information in the `-help` option for more information about the language selection.

Note: Portal will use its own language selection (which may be different from this) to display its log files.

2.2.5 Kernel configuration parameters

You may see core dumps that are related to the limit of semaphores on the system. Semaphore limits can be increased in the kernel configuration parameters.

Note: Please refer to the latest Agent release notes for the recommended minimum semaphore values. If another program requires a larger value than specified in the release notes, use the larger value.

2.3 Upgrading the Agent

Before you upgrade the Agent, ensure that your system meets the minimum requirements described in the release notes.

When the installation kit is launched, it detects the previously installed versions of the Agent and starts to upgrade it.

The following tasks are automatically performed during an upgrade:

- Configuration files are upgraded.
- All Delta files are upgraded to the new version format.
- A backup of the old Global.vvc, Job vvc and Delta files is saved under a subdirectory of Agent installation directory.
- A log file will be created in Agent installation directory.
- All executables and documents are replaced by new versions.

IMPORTANT: *When the upgrade process starts, wait until it finishes. Do not run more than one upgrade at the same time.*

Note: To upgrade the Agent properly, you must select the same installation directory that was used for the previous Agent. *Otherwise, the upgrade will proceed as if it were a new installation.*

To upgrade the Agent:

1. Log on to the target system.
2. Go to the installation kit directory.
3. Download the Linux Agent tar.gz installation kit.

   Note: You must do this locally (i.e., on the target machine).
4. Extract the files from the package. To do so, use this command (where "*PACKAGENAME*" is the name of the Agent installation kit):

   ```
   tar -zxf PACKAGENAME.tar.gz
   ```

5. Run the installation script.

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

Always check the log file after an upgrade. The log file will be used when troubleshooting in the case of failure. If the upgrade fails, the configuration files and executables will roll back to the previous version.

You may try to run the upgrade program again. If it still fails, contact your service provider for support.

Recommendation: Do at least one backup for each job after upgrading successfully. This allows the Agent to upload new configuration files to the vault.

2.4 Uninstalling the Agent

To uninstall the Agent:

1. Log on to the target system.
2. Go to the installation directory (by default `/opt/BUAgent/`).
3. Run the following command

   ```
   ./uninstall.sh
   ```

   A message asks whether you want to stop the Agent.
4. Enter Y.

   A message asks whether you want to remove the installation directory.
5. Enter Y to completely remove the Agent, including all job files and settings.

   Enter N to remove the Agent service entry, executables and scripts. This choice leaves your directory, job files and settings intact for future use.

   The log is saved in `/tmp/Agent-Uninstall-<timestamp>.log`

3 Configuring the Agent

You can use Portal or the legacy Windows CentralControl to manage and configure the Linux Agent.

Note: This guide provides instructions for performing tasks using the legacy Windows CentralControl. For instructions using Portal, see the Portal online help.

3.1 Quick Start steps

For a newly installed Agent, you can use the following steps to configure the Agent and perform your first backup.

Note: The CentralControl Operations Guide describes all available features, options, and further details.

1. Create an Agent profile.

This is the local name (used by CentralControl) of the Agent program that will initiate the backups. You need an Agent profile name for each computer that you back up.

2. Save the default workspace as a named workspace.

To save your configurations (for Agents, jobs, and options), you need to assign a workspace name. CentralControl will prompt you to save any changes. You can create more than one workspace, but you can open only one workspace at a time.

3. Configure a vault profile.

A vault profile_defines the vault configuration that your Agent will use. It matches a job to an account on a vault. The job uses the profile to validate the backup to the vault, and to know where to put the data. A profile can apply to more than one job.

To connect with your account on the vault, create a profile with the properties of this Agent. Some users may have only one profile to service one account (i.e., all jobs back up to a single account). Others may have multiple profiles (and accounts) on one or more Vaults.

4. Create a job.

A job defines the parameters associated with backups, restores, and other processes. Parameters can include: file selections and filters; compression; and encryption settings. A job can belong to only one Agent. Job names are unique on that Agent.

Each Agent in CentralControl has jobs with names that are unique to that Agent. Other Agents may have similar or different job names, even if they perform similar functions. A named job can be one of many for different types of backups, in different ways, at different times. When you create a job, specify a profile that you have created. This allows you to access the vault (i.e., your account).

5. Schedule the job.

You can run your job at predetermined times. You can also run it manually (“ad-hoc”) whenever you want.

When you have completed these steps, you are ready to run a backup.

The remainder of this chapter describes the steps in more detail. Backups are described in the next chapter.

3.2 Create an Agent profile

This is the named function that will initiate the backups. You may (at this stage), when you create the Agent, continue right through to creating a job, configuring the vault, and running the backup. This chapter, though, will describe the steps for configuration only, as outlined here, with the backup being run as described in the next chapter.

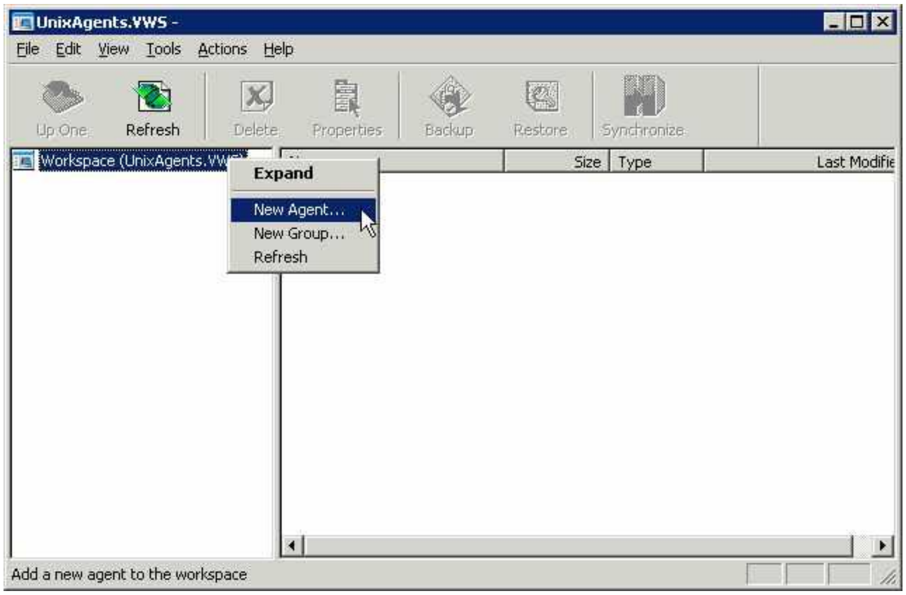

To create an Agent profile, you must have the Workspace selected (highlighted). From here you may either:

- From the pull down menus, use File → New Agent, or
- Right-click on the workspace, and then click on New Agent.

This brings up an Agent Properties screen.

- Description: a description meaningful to you.

- Network Address: either the IP or name of the server the Agent software is on.
- Port: the communications port number reserved for this service (the default is 808).
- User name: authentication to communicate with the Agent service.
- Password: password assigned to the user above.
- (Check to save the password): saves the password on this machine with CentralControl.
- Domain: Windows domain (if applicable).

Click the Check Status button to ensure the communication is valid and you can access the remote Agent. If not, check with your support or vault service provider. Click OK to exit the Status window, and OK again to finish and exit the New Agent window.

Your new Agent's name will show up in the left pane of the CentralControl GUI.

After creating the Agent Profile, save the CentralControl workspace.

3.3 Configure the vault – Agent configuration

Configure the vault with Agent Configuration (i.e., Agent Properties). These are the properties that the Agent will use to connect to this vault. The settings are specific to the Agent, and affect all jobs run under that Agent.

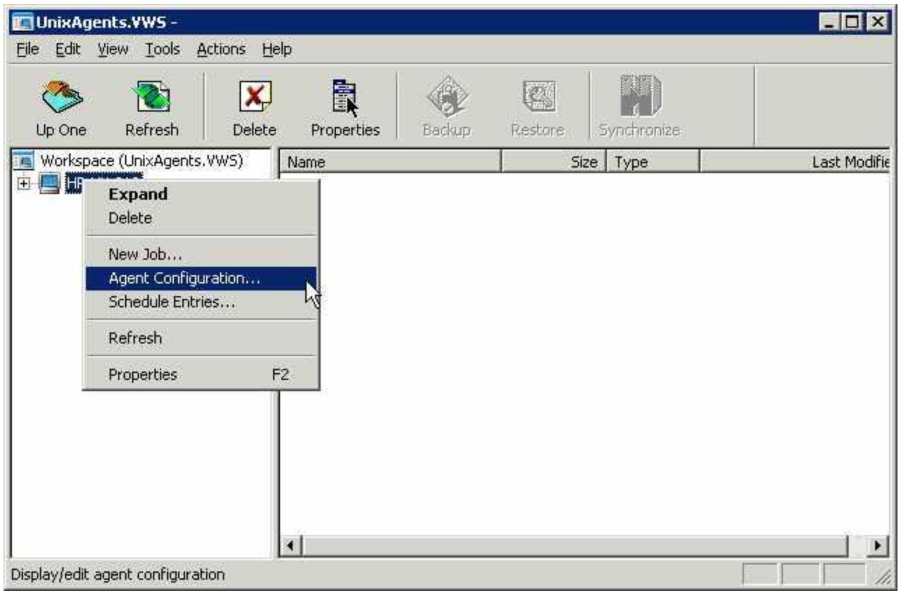

You can start the Agent Configuration from either the Tools → Agent Configuration pull-down menus, or by right-clicking on a selected Agent in the left pane. The Agent Configuration screen has several tabs available. Some, such as Notification or Plug-in, you might not use here, depending on your system and company/organization policies.

Vaults: Adds new vaults, and edits and/or deletes existing ones

- New: You want to select a new (existing) vault, and enter the following information, supplied by the vault service provider.
- Registration: The first time is always New. (Re-Registration is used for changes to the profile.)
- Profile Name: A meaningful name that points to your account on the vault.
- Network Address: vault machine address (IP or server name).
- Ports: Use a communication port.
- Reconnection: How to reconnect if there are communication problems.
- Authentication: Account, user name, and password to access your vault account.

Retention: Decide on the number of days online, copies online and number of days archived for your backups. This may affect the cost of your backups.

Notification: Do you want to be alerted by emails, to successful or failed backups?

Plug-ins: Allows you to set and use optional Plug-in software.

3.4 Create a backup job

This named job can be one of many used to do different types of backups, in different ways, at different times.

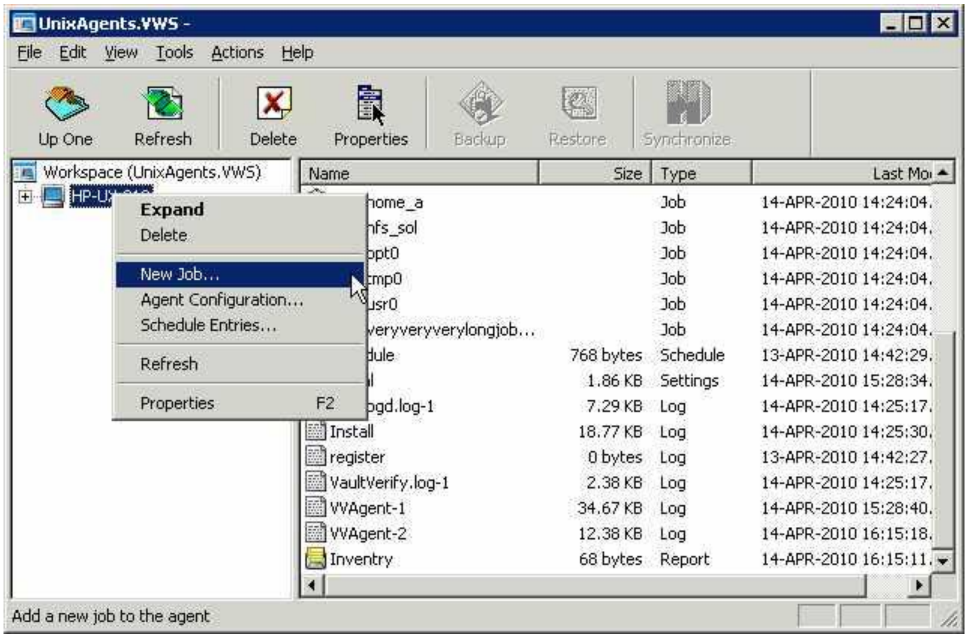

Name	Size	Type	Last Mo
home_a		Job	14-APR-2010 14:24:04.
nfs_sol		Job	14-APR-2010 14:24:04.
opt0		Job	14-APR-2010 14:24:04.
tmp0		Job	14-APR-2010 14:24:04.
usr0		Job	14-APR-2010 14:24:04.
veryveryverylongjob...		Job	14-APR-2010 14:24:04.
dule	768 bytes	Schedule	13-APR-2010 14:42:29.
l	1.86 KB	Settings	14-APR-2010 15:28:34.
ogd.log-1	7.29 KB	Log	14-APR-2010 14:25:17.
Install	18.77 KB	Log	14-APR-2010 14:25:30.
register	0 bytes	Log	13-APR-2010 14:42:27.
VaultVerify.log-1	2.38 KB	Log	14-APR-2010 14:25:17.
VVAgent-1	34.67 KB	Log	14-APR-2010 15:28:40.
VVAgent-2	12.38 KB	Log	14-APR-2010 16:15:18.
Inventry	68 bytes	Report	14-APR-2010 16:15:11.

Add a new job to the agent

Select New Job to start the New Job Wizard (a program that asks you questions and prompts for details regarding the new job).

- Backup source type – choose a local drive or mapped network drives.
- Vault profile – choose an existing one created earlier, or “branch out” from this Wizard and create a new one here.
- Job name – choose a unique, meaningful job name.
- File list backup source - Select Data files. You can include/exclude files and subdirectories.
- Set the options – Quick File Scanning (on/off) and Backup time Options. (These are also accessible in the Schedule Job Wizard.)
- Select an encryption type – choose one from the list, or none. You must supply a password if you choose to encrypt your data on the vault. The data cannot be recovered if you lose the password.
- Configure the logs – set log options and log copies. Choices here depend on your backup activity, and your need for detailed logs and their length of retention. Changes here only affect the logs that will be created, not those already created.
- Finish – Run immediately, schedule a backup, or just exit.

To do an “ad-hoc” backup, we could choose to run this job immediately. For this chapter, we are going to schedule the job to run later. Choose either “Schedule a Backup” and go to the next section, or “Exit” and start the schedule in the next section.

2.5.1 Adding a file or directory to a new backup job

When you first create a backup job, you must include one or more files, or directories (folders). You may modify this list of files and directories afterwards.

In the New Job Wizard (described in the previous section), the Source screen asks you to select files and/or directories to include in the Backup.

If you are selecting Data Files, the **Options** button allows you to select Backup files opened for write (that is, shared read, not opened exclusive), or back up a single instance of all selected hard linked files. This requires a pre-scan pass through the file selection. See File Backup Options.

Click **Add** to start adding files/directories to the list to be backed up. This brings up the Include/Exclude screen, which displays a hierarchy of the disks and directories that you may select from for the backup.

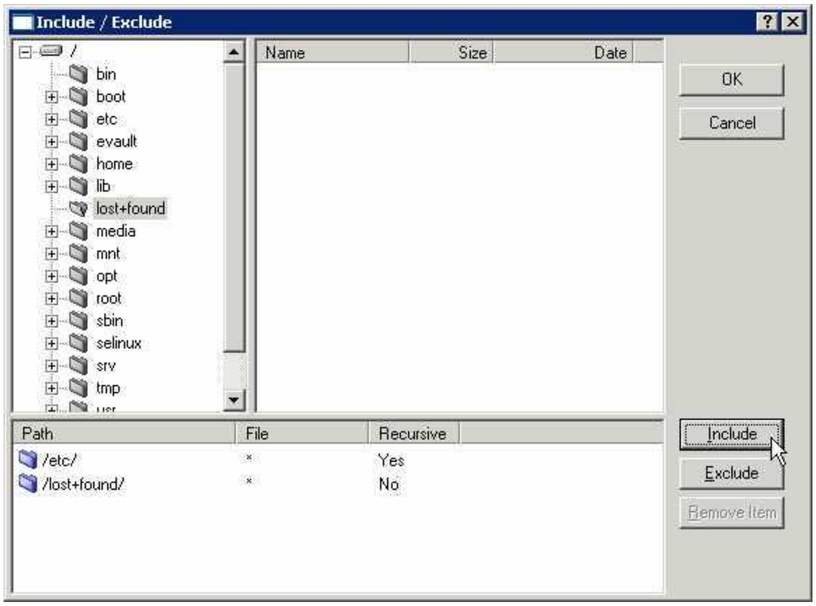

Path	File	Recursive
/etc/	*	Yes
/lost+found/	*	No

You can "open" the tree in the left pane by clicking on the + signs. The files in that directory are displayed in the right pane, where you can select one or more files. Use the CTRL key and the mouse to select multiple files in that directory. Click **Include**. The file/directory names are moved to the lower part of the screen. The **Remove Item** button allows you to un-select names from this lower list, if you change your mind, before you click the **OK** button.

If you have a directory with a large number of files, and you want to select most of them, it might be easier to **Include** them all, and then **Exclude** (from the list) the ones you don't want.

You may also select one directory (folder) at a time to be backed up. When you click **Include**, you will get a message asking if you want to include all files, or just some of them which match your selection criteria (filter).

"Recursive" means to include all files and directories below this directory. Otherwise you may choose to select certain files, depending on their names and extensions. The asterisk (*) means all files with any name or extension.

When you have finished selecting (and including) all the files and directories you want to be in this backup job, click **OK** and you will be back at the Source screen, where you can click **Next** to continue the next step of the New Job Wizard. See the information in the preceding section about creating a job.

2.5.2 Adding/removing a file or directory with an existing backup job

When you first create a backup job, you must include one or more files, or directories (folders). See Adding a File or Directory to a new Backup Job. Later you may want to add or remove files or directories from the backup job.

Select a job in the CentralControl window, and select "Properties" for that job, either from the icons, or by right-clicking or by using F2.

Select the "Source" tab in the Properties window.

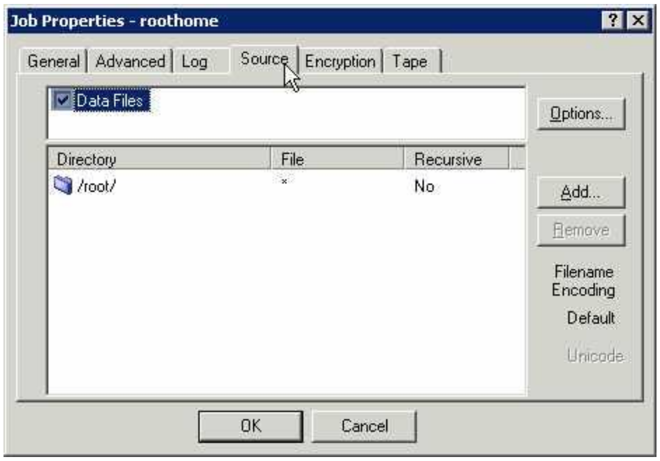

Directory	File	Recursive
/root/	*	No

This displays the existing list of files and directories for this backup job. You may select (highlight) one or more in the lower window, and click **Remove**. You will be prompted with a message "Are you sure you wish to delete the scheduled entry (or entries)?"

The **Add** and **Options** buttons work as described in preceding sections.

Click **OK** when you finish.

3.5 Schedule the job

This job can be run at predetermined times. All jobs can also be run "manually" (ad-hoc) when desired.

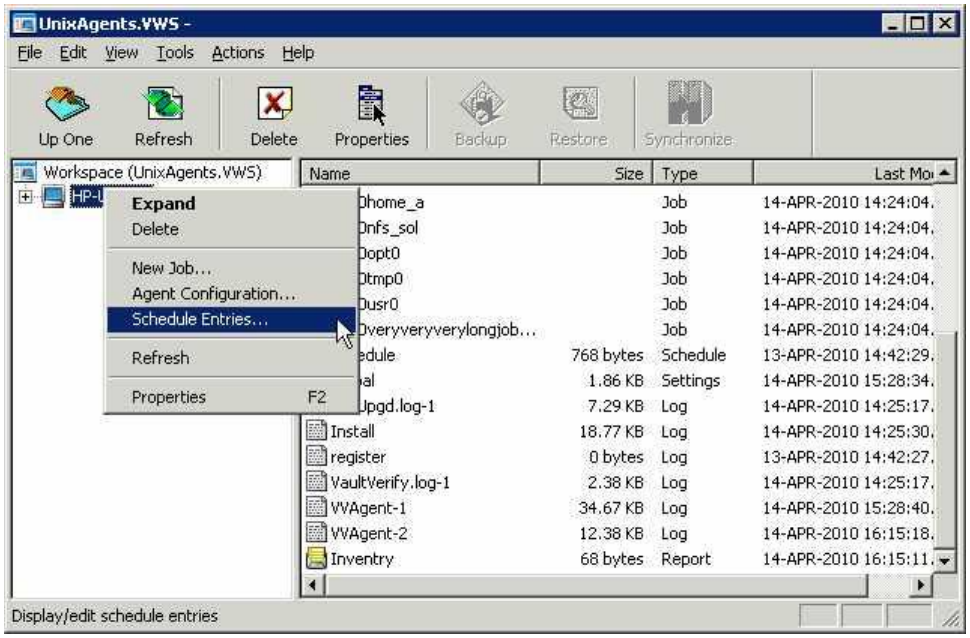

Name	Size	Type	Last Mo
home_a		Job	14-APR-2010 14:24:04.
nfs_sol		Job	14-APR-2010 14:24:04.
opt0		Job	14-APR-2010 14:24:04.
tmp0		Job	14-APR-2010 14:24:04.
usr0		Job	14-APR-2010 14:24:04.
veryveryverylongjob...		Job	14-APR-2010 14:24:04.
edule	768 bytes	Schedule	13-APR-2010 14:42:29.
al	1.86 KB	Settings	14-APR-2010 15:28:34.
Upgd.log-1	7.29 KB	Log	14-APR-2010 14:25:17.
Install	18.77 KB	Log	14-APR-2010 14:25:30.
register	0 bytes	Log	13-APR-2010 14:42:27.
VaultVerify.log-1	2.38 KB	Log	14-APR-2010 14:25:17.
VVAgent-1	34.67 KB	Log	14-APR-2010 15:28:40.
VVAgent-2	12.38 KB	Log	14-APR-2010 16:15:18.
Inventry	68 bytes	Report	14-APR-2010 16:15:11.

Display/edit schedule entries

Start the scheduling from Tools → Schedule Entries, or right-click on a selected Agent in the left pane. This brings up the Schedule List screen. For a new installation, this will be empty. Click New to add a new schedule. This will start the Schedule Wizard, which will take you through the steps to configure a schedule.

- Select a Command to schedule. You may choose: Backup, Synchronize, or Custom command. For now choose "Backup".
- Select a job from the list. It shows the Target and Destination for each.
- Select a Backup type. (Note: This screen will not display for a vault backup.) Specify a Backup type and Processing Options for local disk.
- Select a Retention. This determines how long your backup will be kept online.
- Set the Options. Choose Quick File Scanning (on/off), and Backup Time Options. (These are also accessible in the Create a Job Wizard.)
- Select a Command Cycle. Choose Weekly, Monthly or a Custom Cycle for backups. When you have selected one, and defined the days and times, the Wizard will finish. The command you have just created will now show in the Schedule List. You may Edit, Remove or Disable it. If you have more than one schedule in the list, you may move them up or down in position (priority), so that any conflicts are resolved by taking the parameters in the first (highest) one, and overriding any others. Click **OK** when done.

4 Running backups

Once all the Agent Configuration information has been entered, and a schedule has been set up, as described in the previous chapter, backups will take place automatically.

On occasion you may need to run a "one-time" backup for a special reason. You can either use an existing Agent and job (and modify it), or create a job specifically for this backup.

Seeding and Re-seeding:

When you run your first backup, a full backup is created on the vault. This first backup contains all the data selected for backup and is called a "seed". Subsequent backups are much smaller and deltas only (changes in file), which are applied to the first full backup to create subsequent backups. This way, a current full backup is always available.

If the Agent detects changes, such as the encryption type or password changing, the next backup will be a re-seed.

In this case of a re-seed, your backup will take longer to complete and a message about re-seeding is created in the log file.

4.1 Running an ad-hoc backup

To start an unscheduled (ad-hoc) backup job, select (highlight) a job, and then perform one of these actions:

- Choose Actions → Backup
- Click the backup icon (or use CTRL+B)
- Right-click a job in the left pane

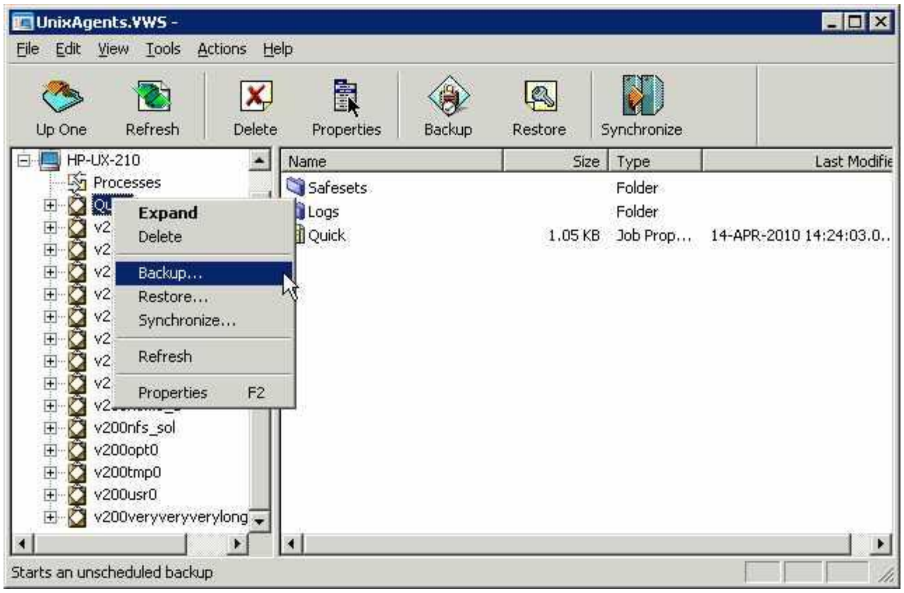

This starts the Backup Wizard, which asks you for:

- A destination (vault or directory on disk). You may choose "Skip further configuration and **Backup Now**", or click **Next.**
- Backup type and options. Depending on your choice of vault or disk, make selections here for type and options. Note that a vault backup will skip over this screen.
- Retention type. Select a retention type. This is the same as in the scheduling of jobs.
- Other options. Quick file scanning, and backup time options. This is the same as in the scheduling of jobs.
- Click **Finish** to complete the configuration and start the backup.

4.1.1 File backup options

The "File Backup Options" will have "Unix Options" enabled for Linux Agents.

Note: A **hard link** is a reference, or pointer, to physical data on a storage volume. The name associated with the file is simply a label that refers the operating system to the actual data. As such, more than one name can be associated with the same data.

Prescanning reads through the file system, gets each inode, and stores it in a map. The larger the file system, the more memory this map requires, and the more time it takes to process. Prescanning only makes a difference on hard-linked files. These share the same initial inode and are therefore the same file. Hard-linked files can only exist on the same disk. They cannot cross disk boundaries.

Backup single instance – option is selected:

If this option is set (this is the default), the backup is slower, as a second pass of the file selection (pre-scan) is required to follow all the links. Some files may have many hard links, and the process of searching them all may take considerable time. The backup size is smaller, as only one "copy" (inode) of the data is backed up, as well as all the links.

Unix Options: *"Backup a single instance of all selected hard linked files. This requires a pre-scan pass through the file selection"*

The pre-scan process can take a significant amount of time and memory depending on the number of files in the file selection (hard links may not cross physical file system boundaries).

On a restore (to original or alternate location), the data (with a new inode) and its hard links are restored.

Backup single instance – option is not selected:

If this option is not set (not checked), it makes the backup faster, but the total backup size is larger, as each link (occurrence) gets backed up separately.

Disabling hard link pre-scanning means that if there are hard links in the file selection list, they will be backed up more than once.

On restore, the hard-link relationship will not be re-established. Each file will be restored individually and applications depending on hard links may not be automatically restored.

Additionally, the restore may require more space than the size of original backup.

4.2 Check the backup results

After a backup (scheduled or ad-hoc) you can check the results for success, or any possible errors. Note that you may have chosen, in Agent Configuration, to be notified by email of successful or failed backups.

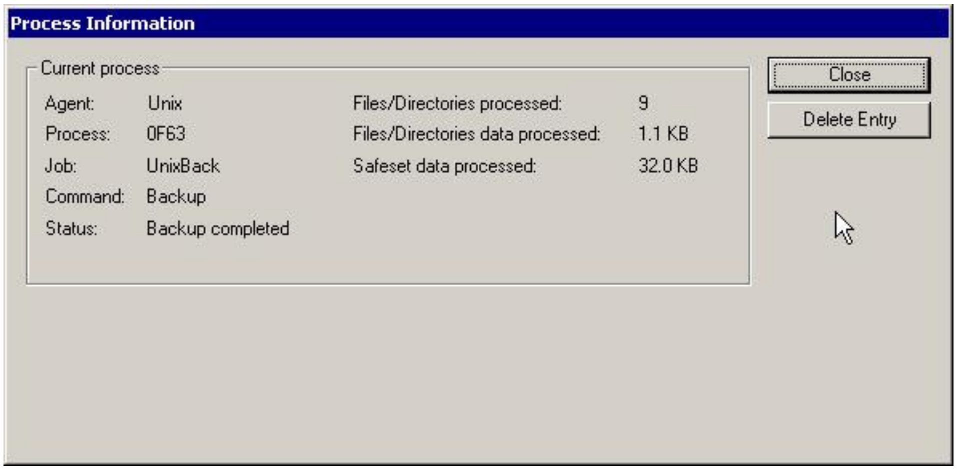

Agent:	Unix	Files/Directories processed:	9
Process:	0F63	Files/Directories data processed:	1.1 KB
Job:	UnixBack	Safeset data processed:	32.0 KB
Command:	Backup		
Status:	Backup completed		

Processes are the “jobs” that the system has performed, such as backups, synchs, and restores. If you select “Processes” in the left pane, you can see a list of processes. Double clicking on one will show you the details. These processes will normally be deleted after approximately one hour in this list. To ensure an accurate (current) picture of the processes, you must perform a Refresh operation by clicking the Refresh button or F5.

Below each job in the left pane are Safesets and Logs. Safesets are ”sets” of backup data (sequentially numbered) on the vault. Double-click a backup (Safeset) to see its properties.

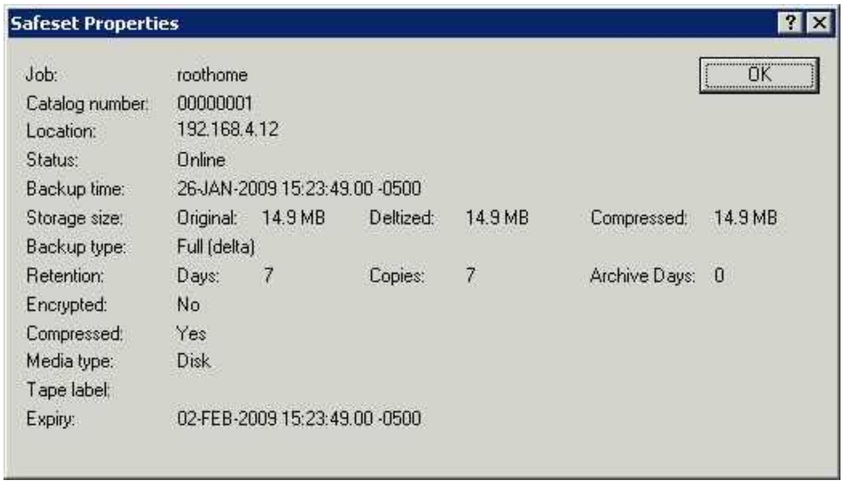

Job:	roothome					
Catalog number:	00000001					
Location:	192.168.4.12					
Status:	Online					
Backup time:	26-JAN-2009 15:23:49.00 -0500					
Storage size:	Original:	14.9 MB	Deltized:	14.9 MB	Compressed:	14.9 MB
Backup type:	Full (delta)					
Retention:	Days:	7	Copies:	7	Archive Days:	0
Encrypted:	No					
Compressed:	Yes					
Media type:	Disk					
Tape label:						
Expiry:	02-FEB-2009 15:23:49.00 -0500					

Log files are the system transcripts of what happened while the backup, synch, or restore function proceeded. Double-clicking on a log will display the contents, which you can also print.

5 Restoring data

There are several reasons for which you might want to restore:

- To recover one or more data files or directories. You can restore them to their original location, overwriting any that are there, or restore them to a different location on that disk, so that you can then decide on which files you want to copy (restore).
- To recover data that was backed up from one computer, to be restored on another (similar) computer system.
- To recover a complete system (i.e., perform a disaster recovery) when the original system has been lost.

5.1 Restoring from a backup

Restoring a backup is the most common usage, allowing you to recover anything from a single file, a directory structure, to a complete system.

To start a restore, select (highlight) a job, and then perform one of these actions:

- Choose Actions → Restore
- Click the Restore icon (or use CTRL+R)
- Right-click a job in the left pane

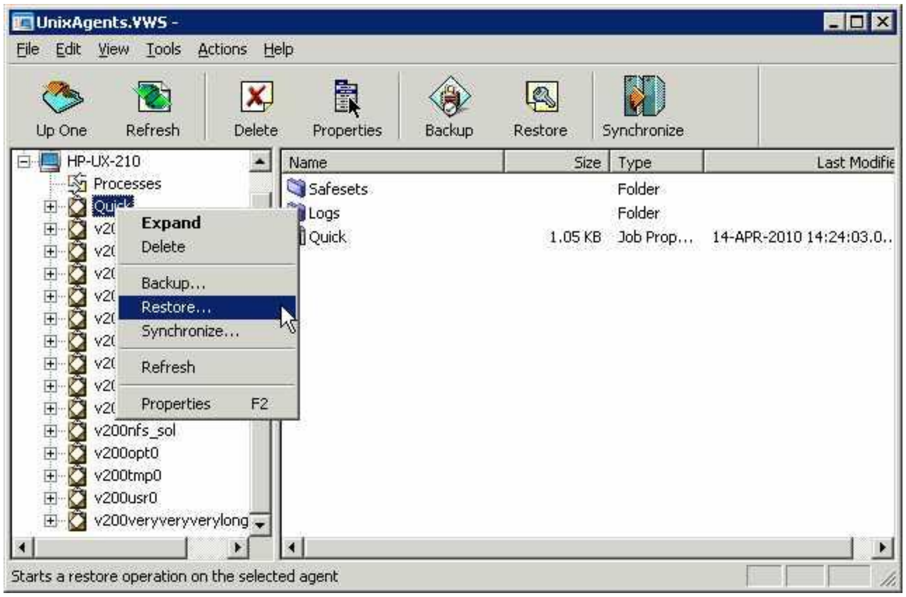

Name	Size	Type	Last Modifie
Safesets		Folder	
Logs		Folder	
Quick	1.05 KB	Job Prop...	14-APR-2010 14:24:03.0..

Starts a restore operation on the selected agent

The Restore Wizard starts allowing you to:

- Select a type of source device, vault or directory. Depending on what you choose here, you may also select a vault and a backup. You can also choose to restore from a particular safeset, or from a range of Safesets.
- Enter the password if the backup is encrypted. You may not see this screen if the backup was not encrypted. If you have lost the password, you cannot access the backup data.
- Select the restore objects (files or directories). You can expand the directories (if available) and select or deselect files to include in the restore.
- Enter the restore destination options. You may choose to restore files to their original locations, or to alternate locations; create sub-directories; overwrite already existing files.
- Select the other restore options. You may overwrite files that are locked; choose all streams or just data streams. You may choose to create a log file with different levels of detail.

Press the **Finish** button to start the restore process. The restore proceeds, and the process information is displayed. You may wish to review the log file afterwards. Recovery logs are prefixed with “RST” in the log listings.

5.1.1 Symbolic links

A symbolic link (also called a symlink or soft link) consists of a special type of file that serves as a reference to another file or directory. A symbolic link contains a path that identifies the target of the symbolic link. The term “orphan” refers to a symbolic link whose target has moved or been deleted.

During a backup, a symbolic link gets backed up with the timestamp of the link. Restoring a symbolic link sets its modification date and time to the date and time of the restore (rather than the date and time of the symbolic link when it was backed up).

5.1.2 NFS – Network file system

To back up data (at a local or remote mount point) to a vault, you can use NFS.

In the CentralControl application, create a new job using "New Job Wizard - Backup Source Type", and then select "Mapped Network Drive Only" from the drop-down list.

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

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

If the local disk does not have sufficient space, this may cause a problem.

If you do not realize that a restore is local, and overwrite is enabled, you will overwrite the local data. You will think, however, that you are overwriting the mount-point data.

5.1.3 NFS support for ACLs and extended attributes

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

5.2 Cross-computer restores

From the menus, select Options → Restore from another computer. This starts the Job Import Wizard.

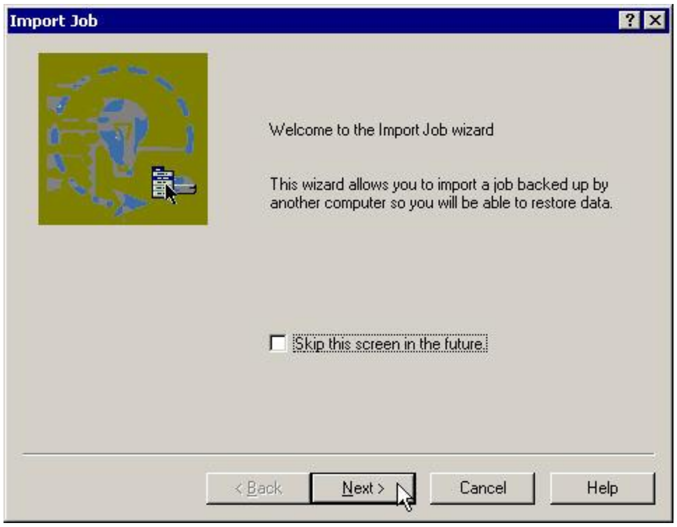
Import Job

Welcome to the Import Job wizard

This wizard allows you to import a job backed up by another computer so you will be able to restore data.

☐ Skip this screen in the future.

The “Restore from another computer” option allows the user to redirect the (original) job to a different client (location) for restore. It does this by getting configuration information - vault name, computer name, and job name - from the original configuration, and adding it to your location so that the restore can be accomplished there. The different client also should be registered to the same vault using the same credentials.

The steps that the Wizard takes you through to do this are:

- Select an existing vault profile.
- Select the computer that has backed up the job that you wish to import.
- Select the job you want to restore.

The Wizard will now copy the job to your local workspace.

From here, the restore proceeds normally (as outlined in the previous section).

5.3 Disaster recovery

“Disaster Recovery” is not a menu choice in CentralControl. Rather, it is a way of restoring a complete backup to a new system. You would want to do this, for example, if a system has crashed, and the disk has been replaced. This is one of the times at which you would want to recover all system and user data back to that disk.

Reinstalling the O/S, applications, and 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. A successful disaster recovery brings your new system to the state of the original system after its last full-drive backup. For more information, see System Recovery.

5.4 Restoring ACLs

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

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

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

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

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

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

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

6 System recovery

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

The basic recovery procedure is:

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

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

6.1 Hardware requirements

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

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

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

6.2 Software requirements

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

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

6.3 Recovery steps

This section describes the steps to perform a system recovery.

6.2.1 Install the minimal operating system

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

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

6.2.2 Install and configure the Agent

1. Install the Agent according to the instructions in this manual appropriate for your operating system.

2. Configure the Agent according to the instructions in section 2 of this manual. It is important to reregister to the vault where the data was backed up.

3. Synchronize the job to ensure that local copies of job catalogs are created.

6.3.3 Restore the backed up system

1. Start a restore according to the instructions in section 4 of this manual.

2. Select the files you wish to restore. The Agent will restore most files to their original locations and protect against many known restore problems (for file systems mounted in their default locations), but some files may cause unpredictable results if restored. These files vary by OS to OS and may generally be restored to alternative locations without problems.

3. Ensure that the files are not being restored to a file system that is mounted read-only.

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

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

6.2.3 Perform post-recovery maintenance

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

6.2.4 Verify the recovery

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

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

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

6.2.5 Recovery problems

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

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

7 Oracle Plug-in

7.1 Overview

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

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

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

7.1.1 Features

- The Oracle Plug-in provides ARCHIVELOG-based, non-RMAN backups of whole online database instances. All non-temporary tablespaces and instance parameter files are automatically backed up.[1]
- Full and partial databases are restored through normal user-managed Oracle recovery mechanisms.
- Agents specify databases using Oracle Service Names (see Oracle instance protection). They do not require script-level or backup-level ORACLE_HOME customization.
- Database passwords are encrypted for enhanced security over script-based methods.

7.1.2 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.

7.2 Installing the Oracle Plug-in

The Oracle Plug-in integrates into existing architecture and allows you to protect Oracle databases. The Agent also performs recovery processes, providing data that you can use to recover Oracle databases.

[1] Oracle Corporation recommends that backups take place in periods of low database activity.

Installing the Oracle Plug-in for Linux requires that you have previously installed the Linux Agent application.

The Oracle Plug-in installation kit is provided in a tar.gz file.

Note: For information about creating a new Agent, creating a backup job, scheduling backups, and disaster recovery, refer to the operations guides and user guides.

7.2.1 System requirements

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

```
SELECT banner
  FROM v$version

SELECT version
  FROM v$instance
```

7.2.2 Supported platform combinations

- See the latest release notes

7.2.3 Before installing or upgrading

- The Agent and the Oracle Plug-in must be installed on the system that has the Oracle database server.
- The Agent must be installed before the Plug-in.
- The Plug-in requires a separate license (usually obtained from a vault).

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

7.2.4 Installing the Plug-in

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

1. Download the Oracle Plug-in for Linux tar.gz installation package.

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

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

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

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

4. Run the installation script:

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

5. Follow the installation instructions on the screens.

7.2.5 Uninstalling the Plug-in

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

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

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

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

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

7.2.6 Before you run the Plug-in

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

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

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

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

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

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

```
SELECT d.file_name, b.status

  FROM dba_data_files d, v$backup b
```

```
        WHERE b.file# = d.file_id;
```

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

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

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

7.3 Backups

7.3.1 Table of backup information

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

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

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

7.3.2 Oracle instance protection

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. In Portal, create a new job.

2. Select "Oracle" for the Backup Source Type.

The Oracle Options will appear on the page.

3. Supply the Database Service Name, User Name, and Password.

In Portal, set the Database Service Name to the *Database Instance* from Oracle (rather than the *Instance Name* from Oracle).

In Windows CentralControl, set the Oracle Service Name to the *Database Instance* from Oracle (rather than the *Instance Name* from Oracle).

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).

Note: If you are connecting to a database where a domain name has been included in the tnsnames.ora, you may have to use the following connection string: \\IPaddress\servicename.

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

4. Select or confirm the databases that you want to back up.

5. If you wish, select an encryption type, and supply an encryption password. 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.

You can start the backup immediately, or let it run on a schedule.

7.3.3 How the backup works

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

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

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

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

7.4 Restores

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

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

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

5. Start and recover the database.

6. Re-open the database for use.

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

7.4.1 Guidelines for restoring

The Oracle Plug-in has been tested in several data recovery scenarios.

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

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

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

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

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

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

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

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

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

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