

EVault Software
Agent 6.0 for AIX
User Guide



Revision: This manual has been updated for Version 6.01.

Software Version: 6.01 (May 2016)

© 2016 Carbonite, Inc.

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.evault.com

Carbonite, EVault Software, EVault SaaS, and EVault DeltaPro, are registered 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 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	5
2	Installation	7
2.1	Ports.....	7
2.2	Install the Agent.....	7
2.2.1	Installation Script.....	8
2.3	Uninstall the Agent.....	9
2.4	Upgrade the Agent.....	10
2.5	Starting and Stopping the Agent.....	11
2.6	Web CentralControl Registration.....	11
2.7	Web CentralControl Language Selection.....	12
3	Agent Configuration	14
3.1	Create an Agent Profile.....	14
3.2	Configure the Vault – Agent Configuration.....	16
3.3	Create a job.....	17
3.3.1	Adding a File or Directory to a New Backup Job.....	18
3.3.2	Existing Backup Job – Adding or Removing a File or Directory.....	20
3.4	Schedule the job.....	21
4	Backing up data	22
4.1	Running an Ad-Hoc Backup.....	22
4.1.1	File backup options (“Unix options” checkbox).....	23
4.2	Check the Backup Results.....	24
5	Restoring data	26
5.1.1	Symbolic Links.....	27
5.1.2	NFS (Network File System).....	27
5.2	Cross-Computer Restores.....	28
5.3	Disaster Recovery.....	29
5.4	Restoring ACLs.....	29

6	System Recovery	31
6.1	Hardware Requirements.....	31
6.2	Software Requirements	31
6.3	Recovery Steps.....	31
6.3.1	Install the minimal operating system	32
6.3.2	Install and configure the Agent	32
6.3.3	Restore the backed up system	32
6.3.4	Perform post-recovery maintenance	32
6.3.5	Verify the recovery	33
6.3.6	Recovery problems	33

1 Introduction

Agent for AIX securely backs up and restores data from AIX servers across a network to a remote data vault.

The Agent software runs as a background service on the individual computers to be backed up. Backups and restores on the Agent computers are configured, managed and scheduled using Portal or a CentralControl application. The Agent sends its backup data directly to the vault.

Different systems require different file and directory backups on different schedules, depending on what data needs to be secured. Some may require more frequent backing up, depending on how the data changes (its volatility).

This guide explains how to install, configure and manage the Agent on individual computers. This guide is intended for the administrator responsible for ensuring that servers are configured properly for backups.

The Agent includes:

- Support for Workload Partitioning (WPAR).
- Trusted Execution (TE) and enhanced Role-Based Access Control (RBAC).
- Support for passwords that are more than eight (8) characters long.
To take advantage of this AIX feature, select one of the new algorithms for password encryption. AIX 5.2 and 6.1 currently support SHA, MDS, and Blowfish.
- Support for multi-CPU awareness, Agent-wide bandwidth throttling, Delta re-creation, LVR, authenticated SMTP, and stronger AES 256-bit encryption.
- Support for Advanced Filtering, longer path names, restoring from another computer, and cross-catalog searches.
- User-configurable backups and restore process priority with 10 levels of granularity.

Workload Partitioning

You can install an Agent in a Workload Partition (WPAR), performing backup and restore operations in the WPAR as you would in a normal physical system. Also, an Agent instance running in a Global Environment can perform backup and restore operations in WPARs.

If you restore a WPAR within an ordinary directory (i.e., not at the WPAR root), the WPAR will restore as a directory.

Enhanced Role-Based Access Control (RBAC)

The AIX Agent can back up and restore the RBAC database.

Before you restore, make sure that the associated files (i.e., command files) and users exist on the system.

Trusted Execution (TE)

The AIX Agent can back up and restore the Trusted Signature Database (TSD). It can also back up trusted command files that are managed by the Trusted Execution environment.

To control trusted command files, use the `TSD_FILES_LOCK` option of the `trustchk` command:

- If `TSD_FILES_LOCK` is off, you can restore trusted command files. If the contents of a trusted command file have changed, the command will not run. This is because the integrity check that the TE environment performs at run time will have failed.
- If `TSD_FILES_LOCK` is on, all trusted command files are locked. You cannot modify them.

2 Installation

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

2.1 Ports

The Agent uses ports 807 and 2546.

2.2 Install the Agent

You must have root privileges to run the installation script. No special privileges are required to extract the installation files for the Agent.

Domain Name Resolution is disabled on most AIX Systems. You may have to use IP addresses.

To install the Agent:

1. Use your system utilities to extract the Agent package from the gzipped tarball.
2. Using your console, change the directory to the extracted directory.
3. As root, run `./install.sh`

This starts the installation process. You are prompted for the following configuration information:

- *Installation directory? [/opt/BUAgent]. Press **Enter**.*
 - */opt/BUAgent doesn't exist. Create it? ([Y]/n). Enter **Y**.*
 - *Select language: [en-US]. Press **Enter**.*
 - *Do you wish to register to a Web-based Agent Console server? ([Y]/n). Enter **Y**.*
 - *What is the Web-based Agent Console address? ("- " to cancel). Enter **<WebCC Server>** (e.g., protect.evault.com).*
 - *What is the Web-based Agent Console connection port? [8086] ("- " to cancel). Press **Enter**.*
 - *What is your Web-based Agent Console username? ("- " to cancel). Enter **<Username>** (e.g., john.doe@company.com).*
 - *What is your Web-based Agent Console password? ("- " to cancel). Enter **<Password>**.*
4. Exit.

If the installation succeeds, a completion message appears, and the Agent daemon will be running. The installation log is in the installation directory. For example:

```
/usr/local/BUAgent/Install.log
```

If the installation fails and rolls back, the installation log will be in the <Installation Failure directory>. If it fails and does not roll back, the installation log will be in the <Installation Kit directory>.

To configure, manage, and communicate with this Agent, you must register it with Portal or CentralControl.

If registration to Web CentralControl fails, verify /opt/BUAgent/buagent daemon is running and Username and Password are entered correctly. If the issue remains, run ping and telnet tests to Web CentralControl from port 8086.

```
ps au | grep buagent
```

Register to Web CentralControl from the EVault Software Agent install directory in AIX Shell

```
cd /opt/BUAgent
./register
```

2.2.1 Installation Script

After you download the installation kit and extract from it, the installation directory has a shell file called `install.sh`.

The `-help` option shows all of the commands available for `install.sh`.

```
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 automatically supply responses. For example:

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

wccLogin=WEBCC_USER_LOGIN
wccPassword=WEBCC_USER_PASSWORD
```

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

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

2.3 Uninstall the Agent

To uninstall the Agent:

1. Log onto the target system.

2. Go to the installation directory (by default `/usr/local/BUAgent/`).
3. Use the `vvagent` script to stop the Agent.
4. Run `"uninstall.sh"`. A message appears, asking if you want to remove the Agent.

Select **Yes** to completely remove the Agent, including all Job files and settings.

Select **No** to remove the VVAgent service entry, executables and scripts. This choice leaves your directory, Job files and settings intact for future use.

If you choose to completely uninstall the Agent, a confirmation prompt appears.

The log will be in the `/tmp/Agent-Uninstall-<timestamp>.log` file.

2.4 Upgrade the Agent

This Agent supports upgrades from Agent Versions 5 and later.

Before upgrading the Agent, we recommend that you make at least one backup of your previous Agent files, including all files and subdirectories under the Agent installation directory.

When you run the installation kit, all executables are replaced with new versions and a log file is created in the Agent installation directory.

Note: Available free space on the volume on which the Agent is installed should be bigger than the size of all Delta files + the size of the largest Delta file + a reasonable cushion (at least 100 MB).

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

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

To upgrade the Agent:

1. Log on to the target system.
2. Go to the installation directory.
3. Stop the Agent.
4. Download the `Agent-AIX-6.xx.xxxx.tar.gz` package.
5. Use your system utilities to extract the Agent package from the gzipped tarball.
6. Using your console, change the directory to the extracted directory.
7. As root, run `./install.sh`

Always check the log file after an upgrade process. The log file can be useful for troubleshooting in the case of failure.

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

2.5 Starting and Stopping the Agent

Commands for the Agent (e.g., `start` and `stop`) are actually “rc” (run control) scripts. You can determine the location of these scripts by viewing the `Install.log` file.

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 Web CentralControl).

```
/etc/rc.d/vvagent {start/stop/restart/status}
```

This single script affects both VVAgent and BUAgent. Its parameters are `stop`, `start`, `restart` and `status`.

Here are examples that show how to use the `vvagent` parameters:

```
[root@sunny]$ /etc/rc.d/vvagent start
```

```
Starting VVAgent...
```

```
Starting buagent...
```

```
[root@sunny]$ /etc/rc.d/vvagent status
```

```
VVAgent is running (PID: 11652).
```

```
buagent is running (PID: 11653).
```

```
[root@sunny]$ /etc/rc.d/vvagent stop
```

```
Stopping VVAgent... stopped.
```

```
Stopping buagent... stopped.
```

2.6 Web CentralControl Registration

During Agent installation, you are prompted to register the Agent with Web CentralControl. You are also asked to choose a default language for email, command lines, and log viewing.

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

Registration Script

Run `<Installation Directory>/register` to register the Agent with Web CentralControl.

If you are already registered with a Web CentralControl 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 Web-based CentralControl address? ("- " to cancel)  
What is the Web-based CentralControl connection port? ("- " to cancel)  
What is your Web-based CentralControl username? ("- " to cancel)  
What is your Web-based CentralControl password? ("- " to cancel)
```

The address is the name or IP address of Web CentralControl.

The port number is defined by the Web CentralControl Administrator.

Your user name/password authentication is set by the Web CentralControl Administrator.

2.7 Web CentralControl Language Selection

During Agent installation, you are prompted to register the Agent with Web CentralControl. You are also asked to choose a default language for email, command lines, and log viewing.

After installation, you can change the registration (reregister) and/or the default language. The Agent must be restarted for these changes to take effect.

Language Selection

Run `<Installation Directory>/set_language` to specify the language that the local Agent will use by default for email, command lines, and log viewing.

The Agent supports these languages:

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

Select a language:

- `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 the information about the `-help` option for more details about the language selection.

(Note that Web CentralControl will use its own language selection – which might differ from yours – to display its log files.)

3 Agent Configuration

To configure a newly installed Agent, you can perform the following steps in Windows CentralControl:

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 the vault.

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.

Each Agent on 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.

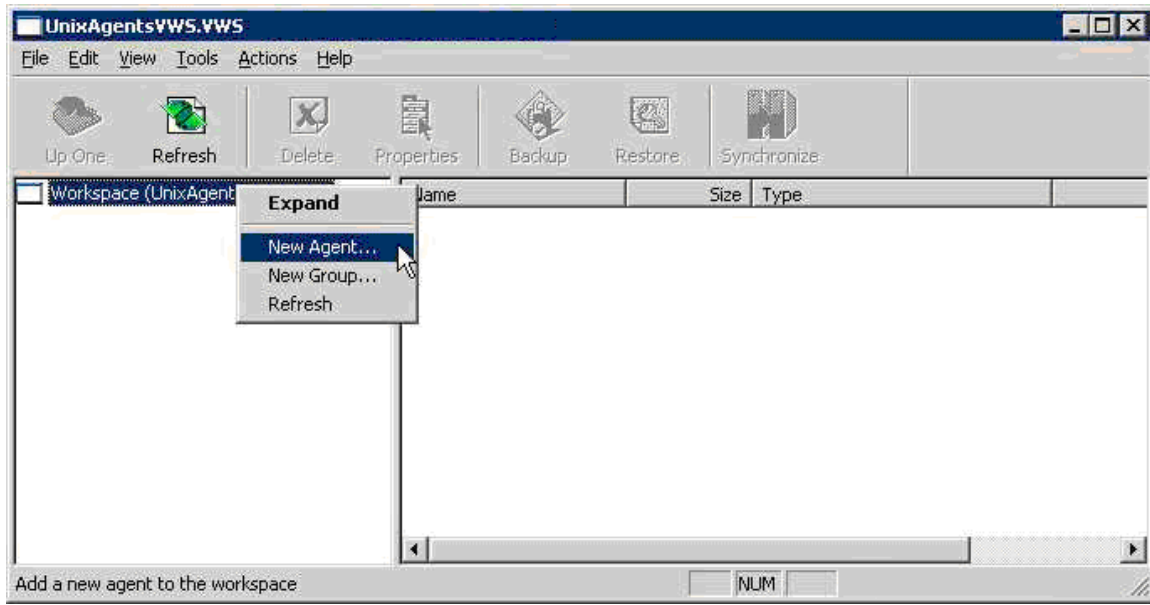
You can run your job at predetermined times. You can also run it manually (“ad-hoc”) whenever you want.

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.

3.1 Create an Agent Profile

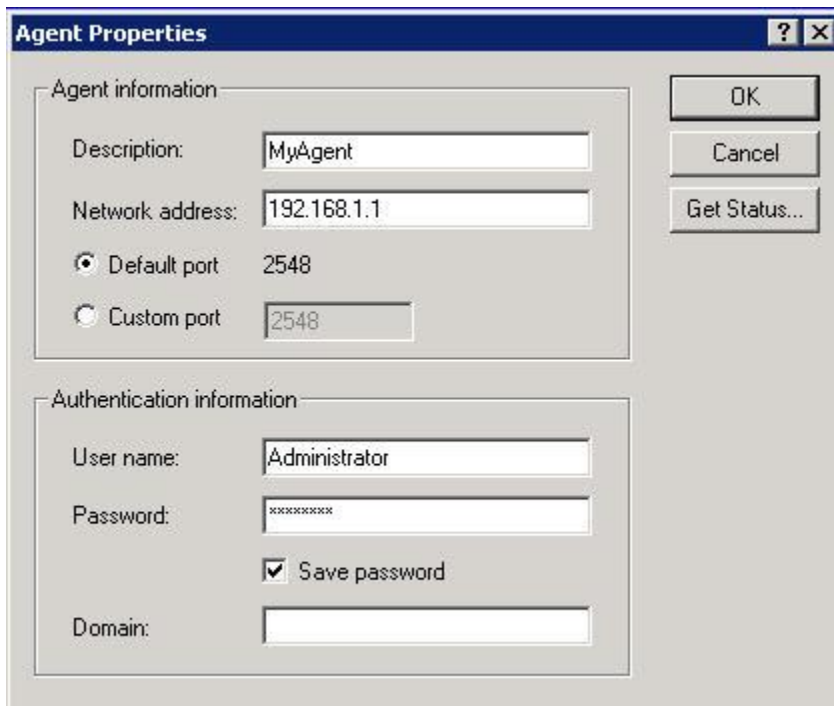
This is the named function that will initiate the backups. You may (at this stage), when you create the Agent, continue right through to creating a job, configuring the vault, and running the backup. This chapter, though, will describe the steps for configuration only, as outlined here, with the backup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



Click the Get Status button to ensure the communication is valid and you can talk to 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.

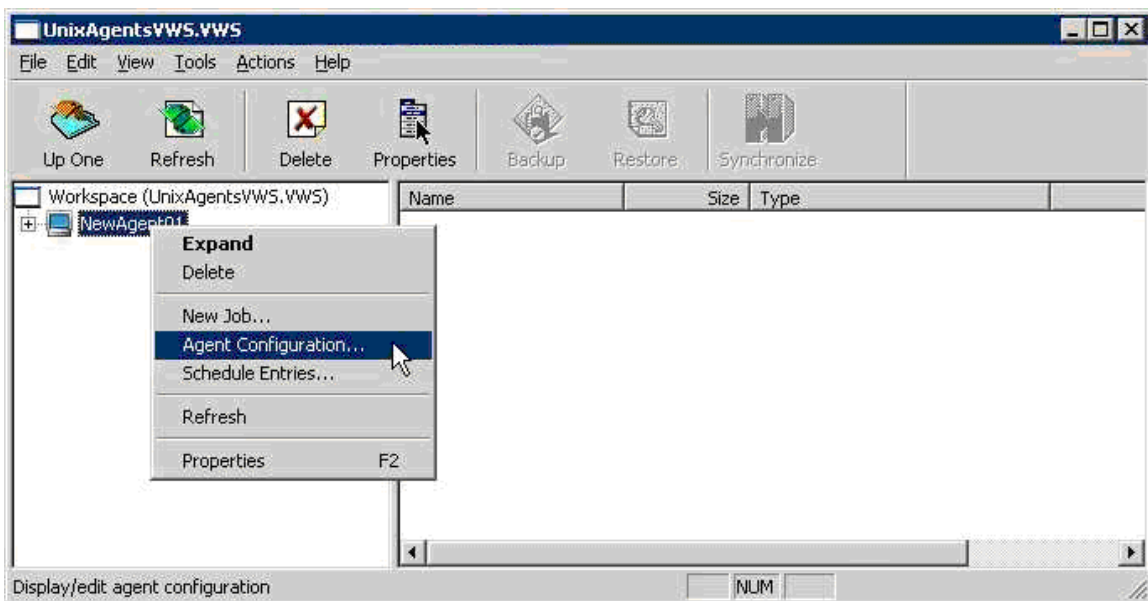
Your new Agent's name will appear in the left pane of the CentralControl GUI.

Note: For this screen and others, use the "What's This" help (the '?' in the upper right corner) for further information about the fields, as well as the main Help menu (F1) for general help.

If the F1 Help screen is open (even minimized), the "What's This" help will not be active. The F1 help must be closed for the "What's This" help to function properly.

3.2 Configure the Vault – Agent Configuration

Configure the vault through 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 (see the figure). The Agent Configuration screen has several tabs available. Some, like 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.

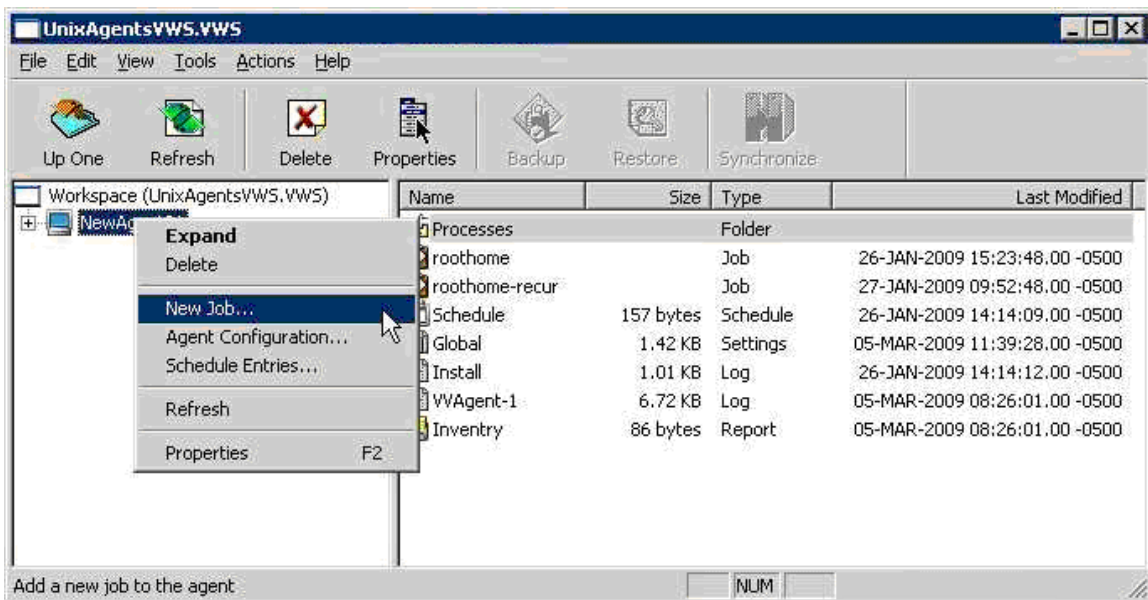
Retention: Decide on the number of days online, copies online and number of days archived for your backups. This may affect the cost of your backups.

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. See the Plug-In manuals.

3.3 Create a job

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



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.

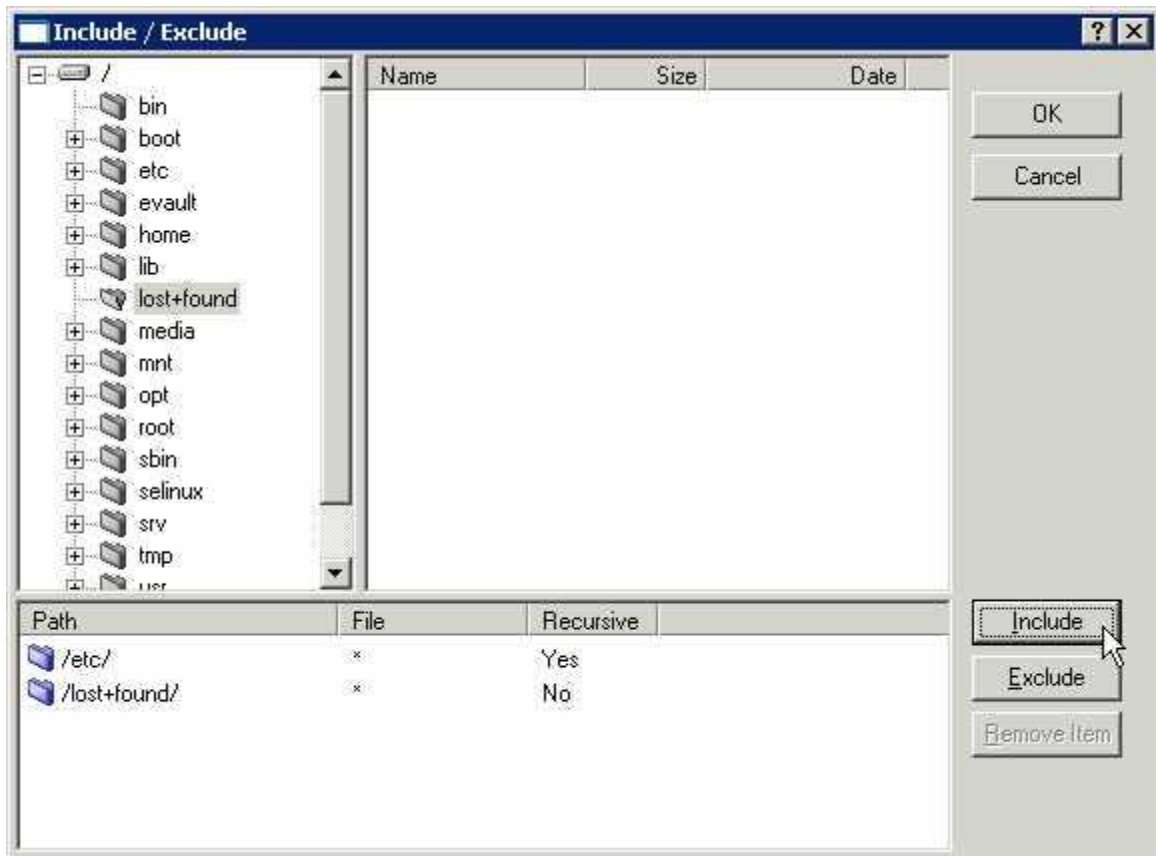
3.3.1 Adding a File or Directory to a New Backup Job

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 prescan pass through the file selection. (See Section 3.1.1 for more information about these 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.



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 only the ones that 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.

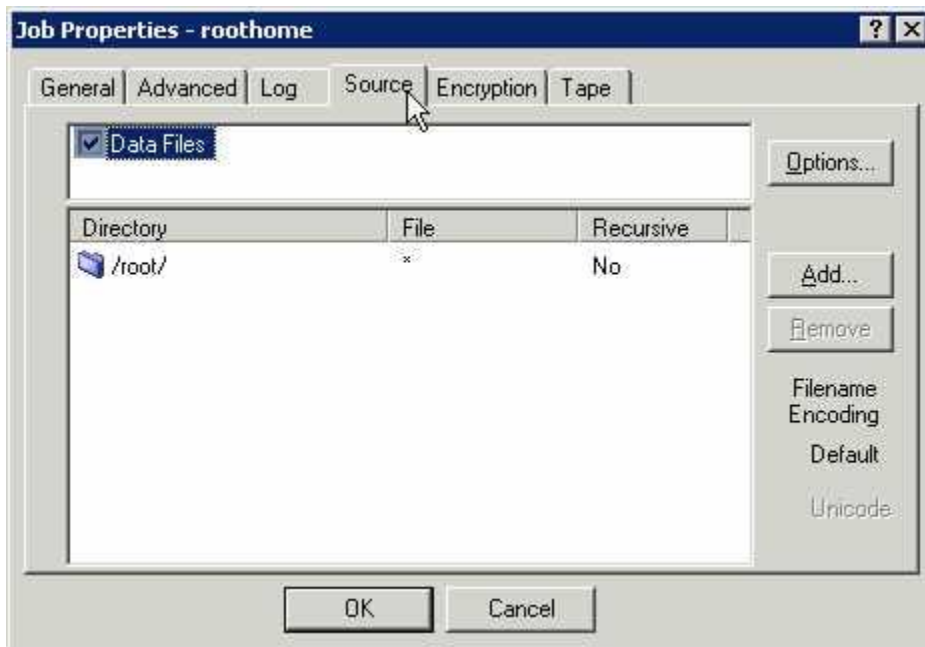
When you have finished selecting (and including) all the files and directories you want to be in this backup Job, click **OK** and you will be back at the Source screen, where you can click **Next** to continue the next step of the New Job Wizard. See the information in the preceding section about creating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

3.3.2 Existing Backup Job – Adding or Removing a File or Directory

When you first create a backup job, you must include one or more files, or directories (folders). See the section about “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.

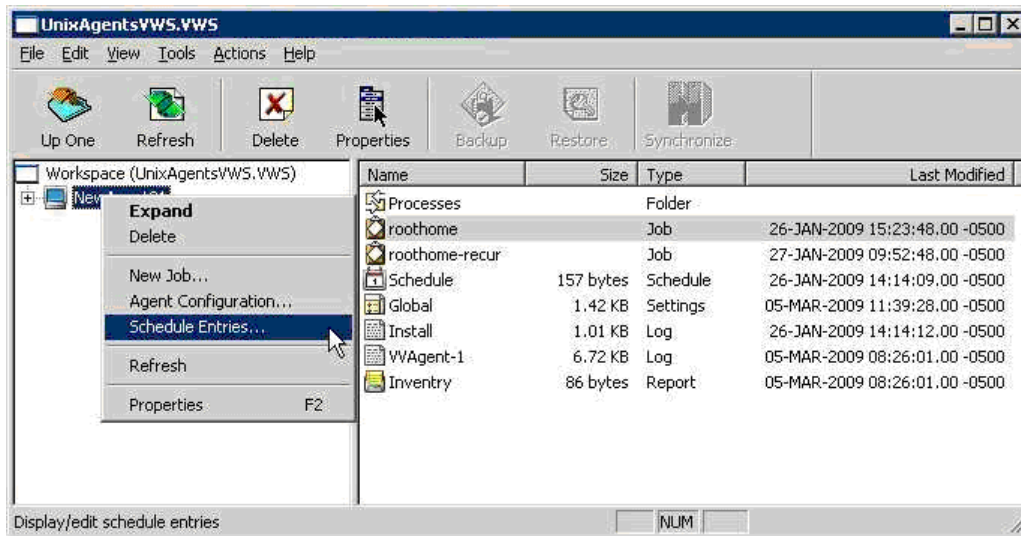


This displays the existing list of files and directories for this backup job. You may select (highlight) one or more in the lower window, and click **Remove**. You will be prompted with a message “Are you sure you wish to delete the scheduled entry (or entries)?”

The **Add** and **Options** buttons work as described in preceding sections. Click **OK** when you finish.

3.4 Schedule the job

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



Start the scheduling from Tools → Schedule Entries, or right-click on a selected Agent in the left pane (see the figure). 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. Choose Daily, Weekly, or Monthly from the list. 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 Backing up data

Once all the Agent configuration information has been entered, and a schedule set up, as in the previous chapter, the 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 one 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 deltas (changes in file), which are applied to the first full backup to create subsequent backups. This way a current full backup is always available.

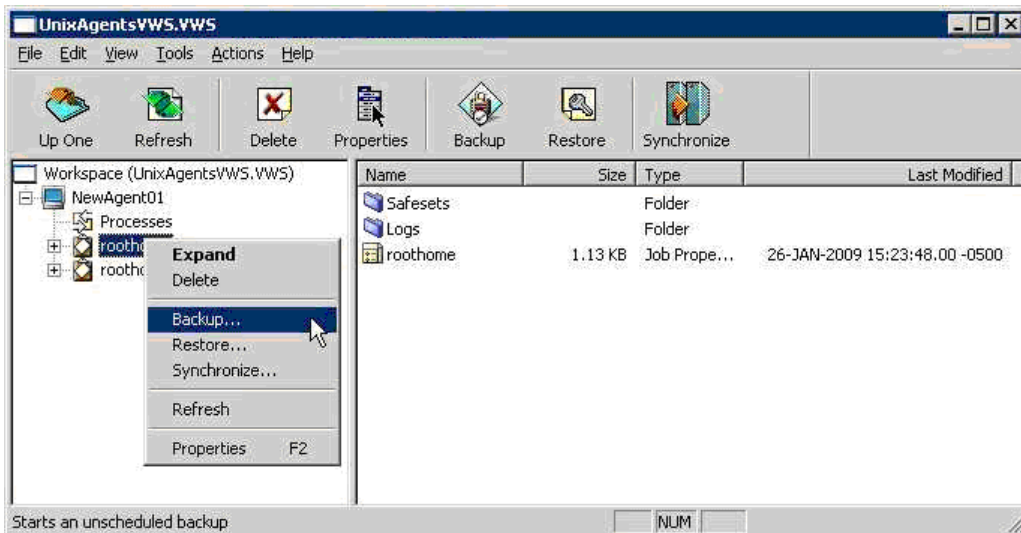
If the Agent detects changes, such as the encryption type or password changing, the next backup will be a re-seed.

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 this screen.
- Retention type. Select a retention scheme: daily, weekly or monthly. 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 (“Unix options” checkbox)

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 the checkbox is selected (this is the default), the backup is slower, as a second pass of the file selection (prescan) 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.

The prescanning 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 the checkbox is not selected, it makes the backup faster, but the total backup size is larger because each link (occurrence) gets backed up separately.

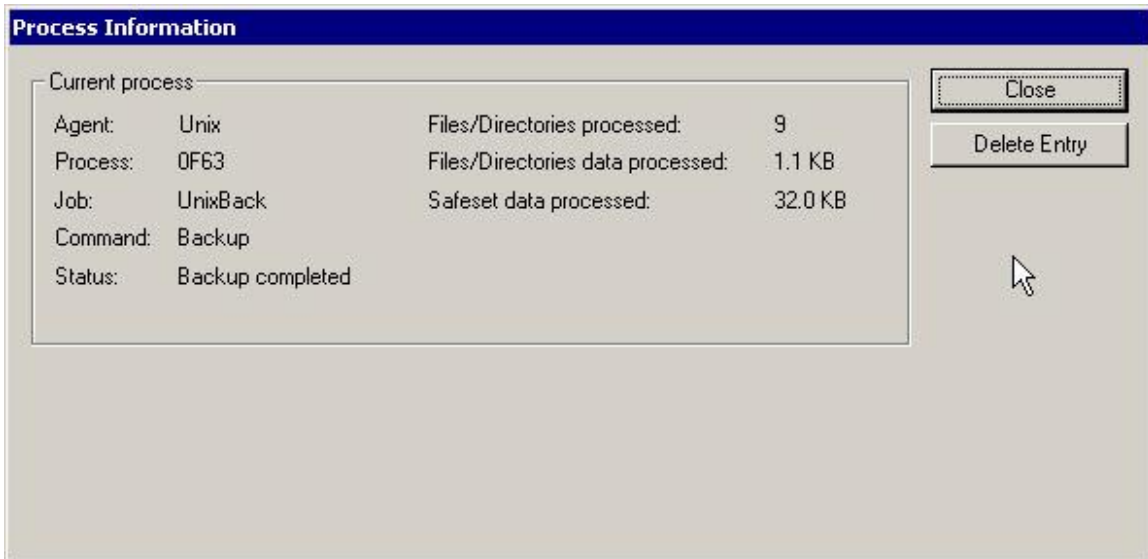
Disabling hard link pre-scanning means that if there are hard links in the file selection list, they will be backed up more than once.

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.

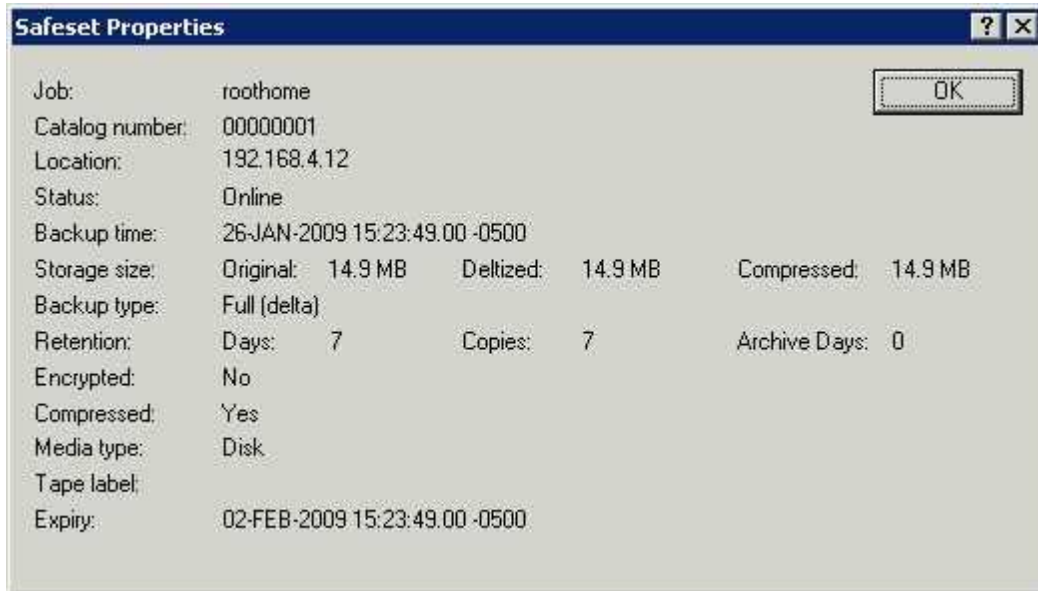
4.2 Check the Backup Results

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.



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 Synchronize operation.

Below each job in the left pane are Safesets and Logs. Safesets are "sets" of backup data (sequentially numbered) on the vault. They remain until their retention date (configured by you) expires. Double-click a backup (Safeset) to see its properties.



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

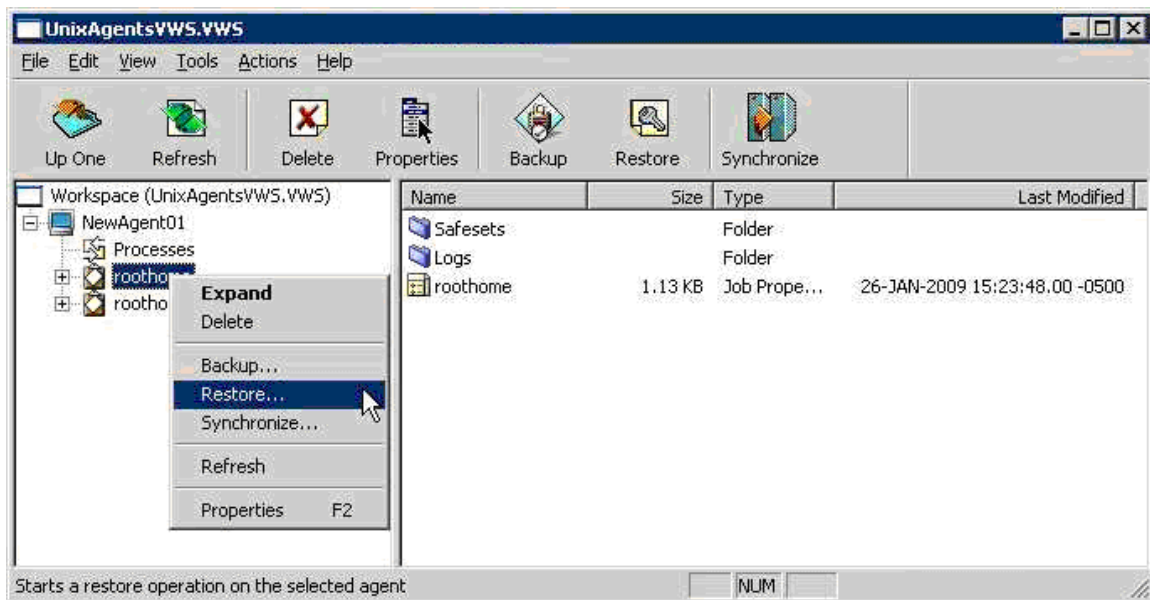
There are several reasons for which you might want to do restores:

- 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.

Note: For successful AIX system restores, the Technology Level (TL) of the target system must equal or surpass the TL of the source system.

To start a restore using Windows CentralControl, 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



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. Restore 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 particular 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.

5.2 Cross-Computer Restores

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



What the “Restore from another computer” option does is allow the user to redirect the (original) restore job to a different client (location). It re-registers where the configuration file was originally pointing, so that the restore job can be redirected to another location. It does this by getting, authenticating and copying 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 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. If a job already exists with that name, you will receive a prompt regarding an overwrite.

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

5.3 Disaster Recovery

“Disaster Recovery” is not a menu choice in the CentralControl program. 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.

Note: For successful AIX system restores, the Technology Level (TL) of the target system must equal or surpass the TL of the source system.

5.4 Restoring ACLs

You can back up and restore Access Control Lists (ACLs). The following behaviors can occur when you restore ACLs on an AIX 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 differ by variety of Unix, 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 (i.e., a file system that does not support ACLs, or where you are using another variety of Unix), the ACLs will not be restored. An error message will appear in the backup log.

If you restore to a compatible system (i.e., the original system, or a different system with the same variety of Unix), 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 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-recovery maintenance.
5. Verify the recovery.

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 recovery of programs, system state, and data. Otherwise the recovery 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.

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.
- Agent installation media identical to that installed on the original system.

6.3 Recovery Steps

For successful AIX system recovery, the Technology Level (TL) of the target system must equal or surpass the TL of the source system.

6.3.1 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.

6.3.2 Install and configure the Agent

1. Install the Agent according to the instructions in this manual.
2. Configure the Agent according to the instructions in this manual. It is important to reregister with 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 recovery according to the instructions in 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 can 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 is complete, the process of verifying the integrity of the restore can commence.

6.3.4 Perform post-recovery maintenance

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

6.3.5 Verify the recovery

When the recovery procedure finishes, determine whether or not it 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.

6.3.6 Recovery problems

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

- Was the system restored using the same version of AIX?
- 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 AIX patches applied?
- Was there sufficient disk space to handle all of the restored data?
- Was the Technology Level (TL) of the target system lower than the TL of the source system? For successful AIX system restores, the TL of the target system must equal or surpass the TL of the source system.