Agent 6.73 for Solaris and Oracle Plug-In Guide

July 2011

This document describes how to install and use the Agent software Version 6.73 for Solaris and the Oracle Plug-In for backups and restores.

Contents

Revision: This manual is updated for Version 6.73
Software Version: 6.73 (July, 2011)

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 Agent, Agent Console, and Vault applications have the added encryption option of 128/256 bit 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).
See:**http://csrc.nist.gov/encryption/aes/round2/r2report.pdf** for details.

The Agent and Vault applications have the added security feature of an over the wire encryption method.

1. Introduction

This Guide is intended for the administrator responsible for ensuring that users' servers (computers) 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.

Different servers may require different files and directories backed up depending on what data needs to be secured. Some may require more frequent backups or different schedules. This can depend on your backup requirements, how frequently the data changes (its volatility), as well as your bandwidth capabilities.

Backups and restores with Agent Console are described in the *Windows Agent Console Operations Guide* and the online help for Web Agent Console from the point of view of "how to operate" the program.

1.1 Features/Changes in this Release

- Support for x86 processors on Solaris 10.
- Support for Global and Non-Global Zones on Solaris 10.
- NFSv4 ACL Model supported on ZFS File System on Solaris 10.
- Extended Attributes are protected on Solaris 9 and Solaris 10.
- Separate Agent installation kits for Solaris 9 and 10.
- Separate Oracle Plug-In kits for Solaris 9 and 10
- For Agent versions 6.72 and above, the Preserve folder structure checkbox has been removed from the Restore Wizard. If you have selected your restore data using the recursive option, the child directory structure is restored recursively from the selection point. This includes all sub-folders and the data they contain. The Parent folder is not restored. Parent folders can be created manually within the restore location if desired. The restore selection can then be directed to the re-created Parent directory. Alternatively, from within the restore wizard, the user can specify an alternative path adding the Parent folder and the Agent will create it automatically. Files and data that are not desired for the restore can be manually excluded from the restore Job.
- 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.

1.2 How the Software Works

The Agent Console, Vault, and Agent software comprise a data protection software suite that securely backs up and restores data from servers across a network to a remote data

Vault. The applications provide an automated lights-out method for protecting your valuable computer data without the need for tape devices or other backup media.

This Agent for Solaris Guide & Oracle Plug-In Guide explains how to install, configure and manage the Solaris Agent and Oracle Plug-In on individual computers. The Agent Console Operations Guide and online help explain how to use the Windows Agent Console or Web Agent Console application to configure backups and restores from the computers that run the Agent.

1.2.1 Agent Software

The Agent software runs on the individual computers to be backed up. Backups and restores on the Agent computers are configured and scheduled by the Agent Console application. The Agent communicates its backup data directly to the Vault.

The Agent software consists of these components:

- The backup process (VV), which performs backup and restore functions.
- The Agent (VVAgent or BUAgent), which runs as a *Solaris* daemon.

1.2.2 Agent Console Software

Windows Agent Console and Web Agent Console provide a centralized point of control for managing all computers running the Agent software on a large computer network. The Agent Console applications can centrally control Windows and Solaris Agents.

Generally within an organization, the configuration and scheduling of backup and restore Jobs is done through the computer running the Agent Console software. Agent Console connects to an organization's computers running the Agent software activating that computer's backup Job.

1.2.3 Vault Software

The Vault controls and manages the pooling and storage of data at a remote secure location (Vault). This data is communicated to the Vault from the Agent computers over a WAN, LAN, the Internet, or imported from alternate media.

1.2.4 Oracle Plug-In Software

The Oracle Plug-In is an add-on to the Solaris Agent. It allows a user to perform a database backup on an Oracle database. The Plug-In is installed on top of the Agent on the database host to perform the backups, either on demand or scheduled.

1.2 Agent vs. Agent Console

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.

The setup of Agents, Jobs, scheduling, and monitoring is done from the Agent Console application. The actual backup is done from the server with the Agent, to the system with the Vault. No user data goes through Agent Console. The Vault has to be previously set up with a valid account to receive your Agent's commands and data.

1.3 Agent / Agent Console Configuration Overview

The Agent program runs as a service on the server (computer) that will be backed up. The way to configure and control it is with the Agent Console program. One Agent Console program controls many Agents on many different servers on a network.

You need to provide names, passwords, and permissions to allow the Agent to connect with Agent Console:

1. Each server (computer to be backed up) needs an Agent.
2. Need to connect (from Agent Console) to an Agent (when you create a new Agent).
3. Supply Name, IP address, and user/password.
4. Then, register the computer on the Vault.

You must register a computer on a Vault in order to log on to the Vault and establish a connection. The Vault must be informed that this Agent is valid and is authorized to perform its functions.

You need to “re-register” a computer if you restore from another computer, or perform a disaster recovery (described in chapter 4 of this guide.).

When you create a backup Job, you register information about it. For example:

1. Which profile (i.e., which Vault) applies?
2. What data does it need to back up?
3. What types of logs will it produce?
4. What type of encryption (if any) will it use?
5. When will it run?

Note: The first backup is a “seed” (complete backup). The next and subsequent ones are deltas (i.e., changes only), but they are still considered to be full backups.

Depending on the configuration of your system:

- There might be more than one Vault that you can use.
- One Agent Console usually controls all of the Agents on your network.
- You can also back up to your own local disk.

2. Agent Configuration

2.1 Introduction

For a newly installed application, you can use the following steps to quickly perform your first backup. Regarding the Agent Console program, the "Agent Console Operations Guide" describes all of the features, options, and further details.

An Agent configures, manages, runs, and monitors backup Jobs. You can manage and control many Agents through one Agent Console application (GUI). An Agent can have multiple Jobs.

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.

A 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.

Steps for a QUICK START:

1. Create an Agent profile.
This is the local name (used by Agent Console) 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. Agent Console 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 Agent Console 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.

2.2 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.

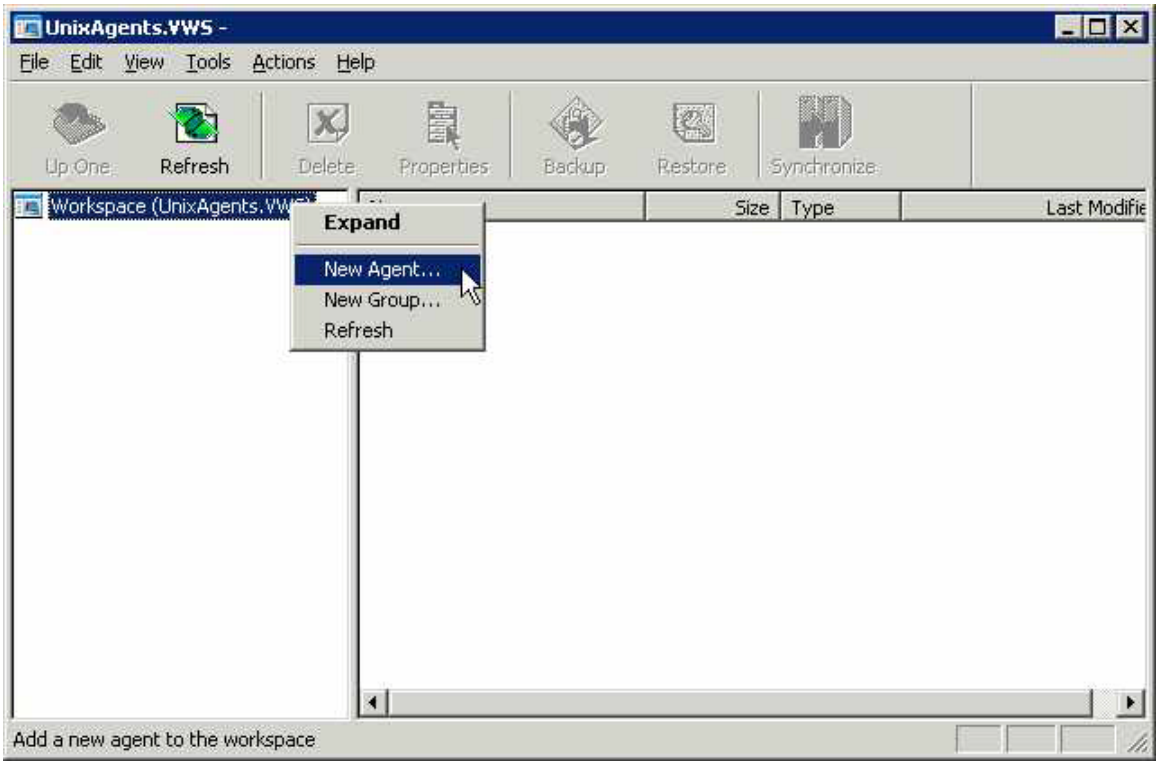

Fig 1. Create an Agent profile

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.
- Right-click on the workspace, and then click on New Agent (see Figure 1).

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 Agent Console.
- Domain: Windows domain (if applicable).

Click the Check 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.

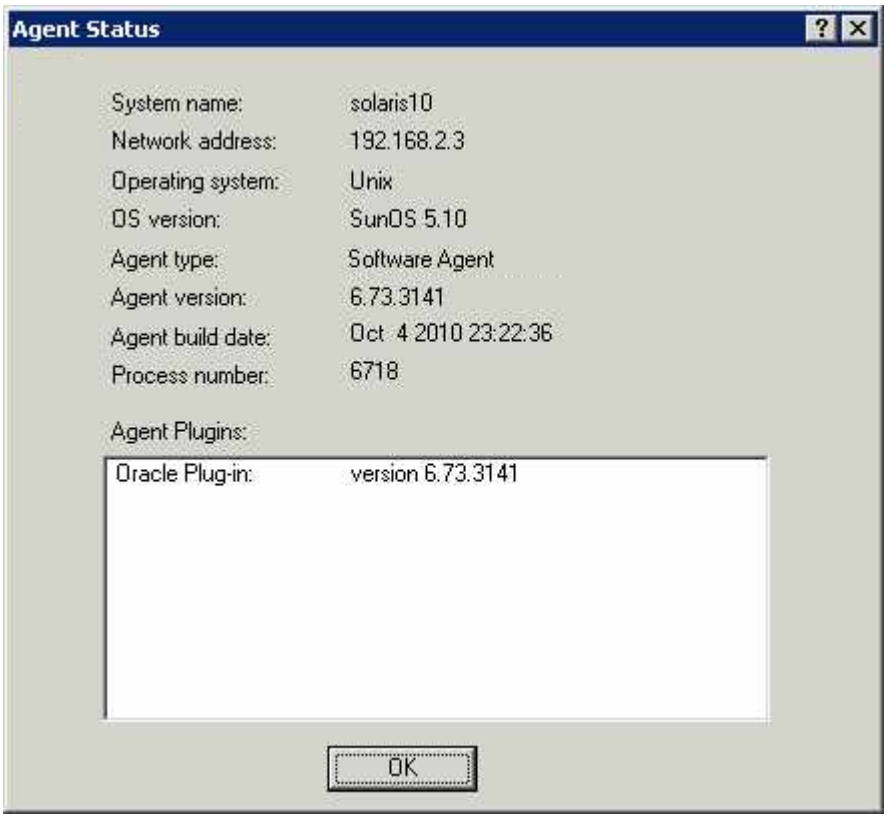

Fig 2. Check Agent Status

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 Agent Console GUI.

Note: In this screen and others, you can use the "What's This" help (the '?' in the upper right corner) for further information about the fields. You can also use the main Help (F1) for information and assistance.

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.

2.3 Save the Workspace

Save the default workspace as a named workspace. To save all your configurations (Agents, Jobs and options) you must save your workspace with a name of your choice. Agent Console will prompt you to save any changes before you exit the application. You can save more than one workspace, but you can open only one at a time.

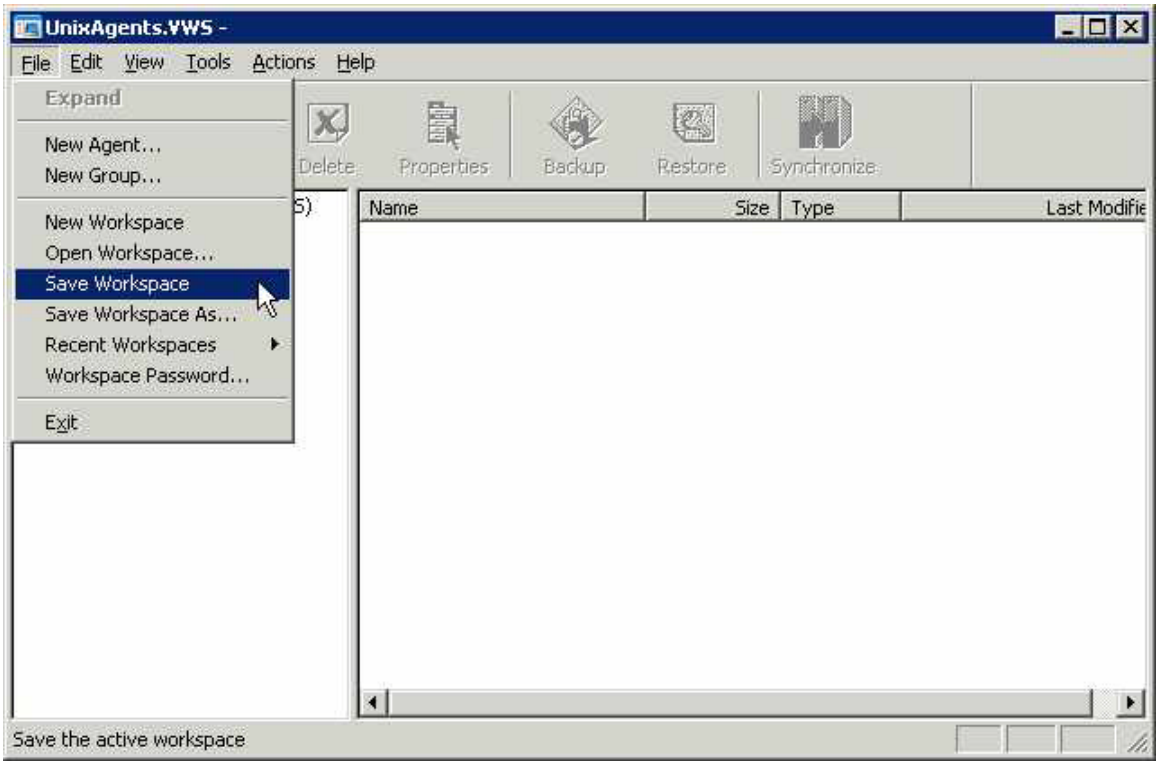

Fig 3. Save the Workspace

Choose a meaningful name for your workspace. As well as saving the current workspace as a new name, you may create new ones, open existing ones, save the current one, and see the recent ones.

Because a workspace contains important user names and passwords necessary for access to do backups, it is advisable to optionally encrypt these workspaces so that unauthorized users cannot gain access to them.

The “Workspace Password” option allows you to add or change a password, as well as choose from encryption types with different cipher strengths.

If this is the first time that you are using a password here, there will not be an “Old Password”, so leave that field blank. Select an encryption type, and create a password (case sensitive). If you lose this password you will need to recreate the workspace.

2.4 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.

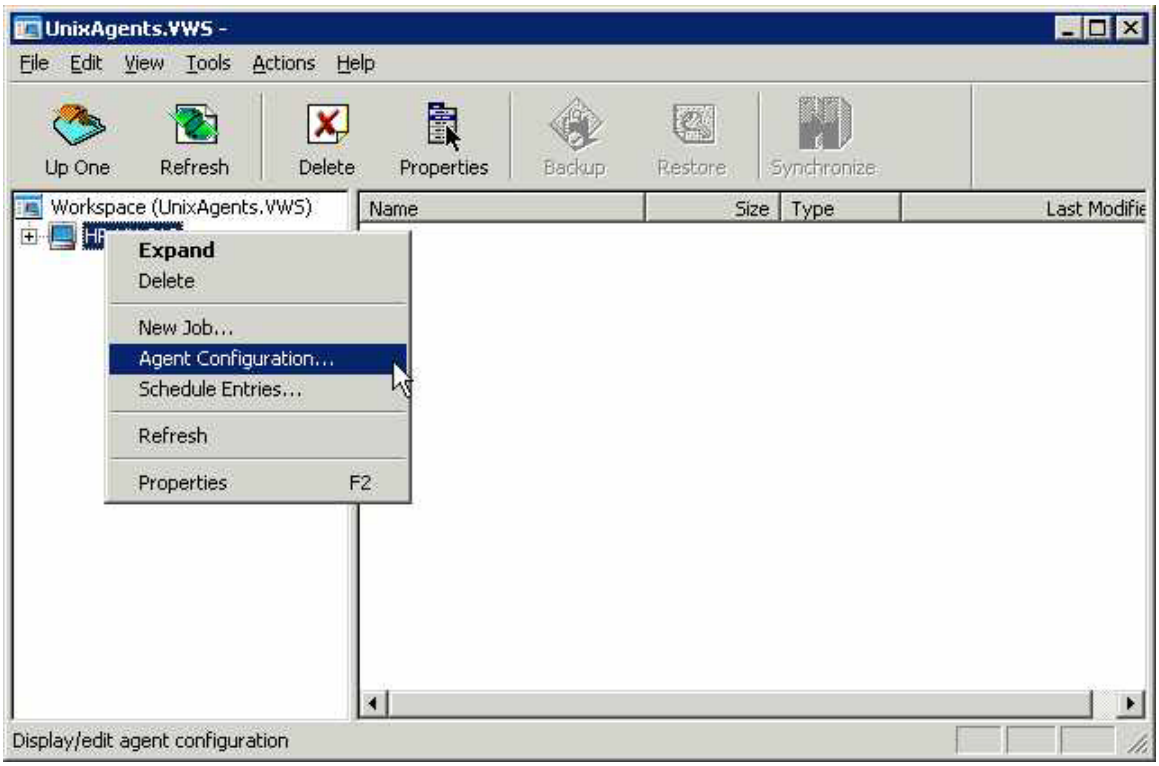

Fig 4. Agent Configuration

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 Figure 5). 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. See the Plug-In manuals.

2.5 Create a Job

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

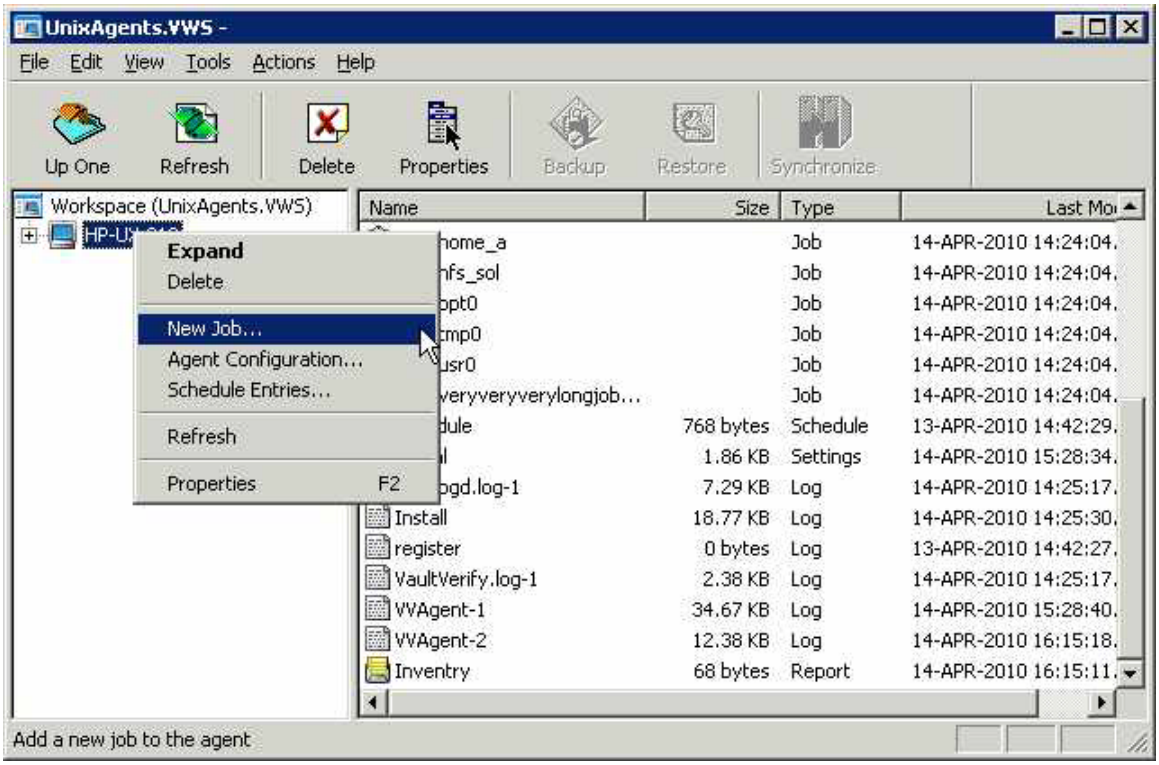

Fig 5. Create a Job

Select New Job to start the New Job Wizard (a program that asks you questions and prompts for details regarding your 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 simply 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 Section 3.1.1 for more information on 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.

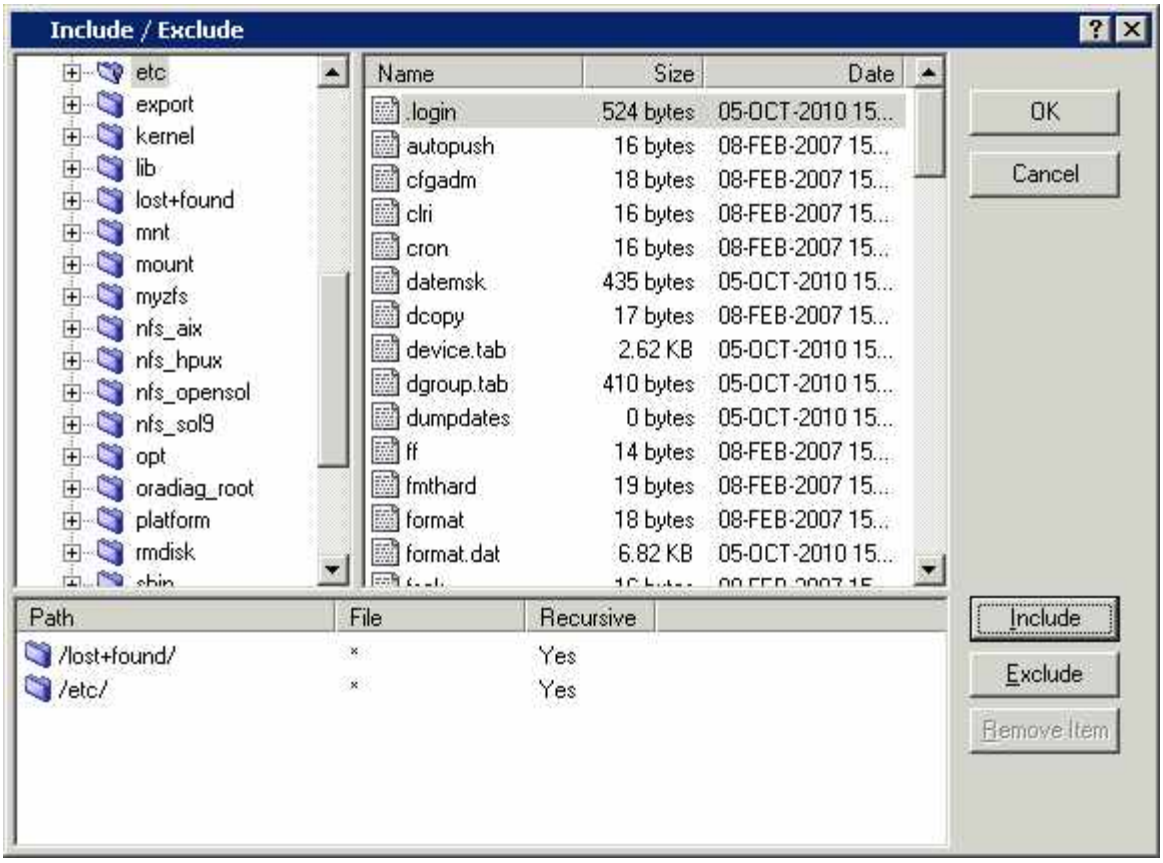

Fig 6. Include/Exclude Directories and Files

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.

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 the section about "Adding a File or Directory to your new Backup Job". Later you may want to add or remove files or directories from the Backup Job.

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

Select the "Source" tab in the Job Properties window.

Fig 7. Source tab in Job Properties

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 selected entry (or entries)?”

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

Click **OK** when you finish.

2.6 Schedule the Job

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

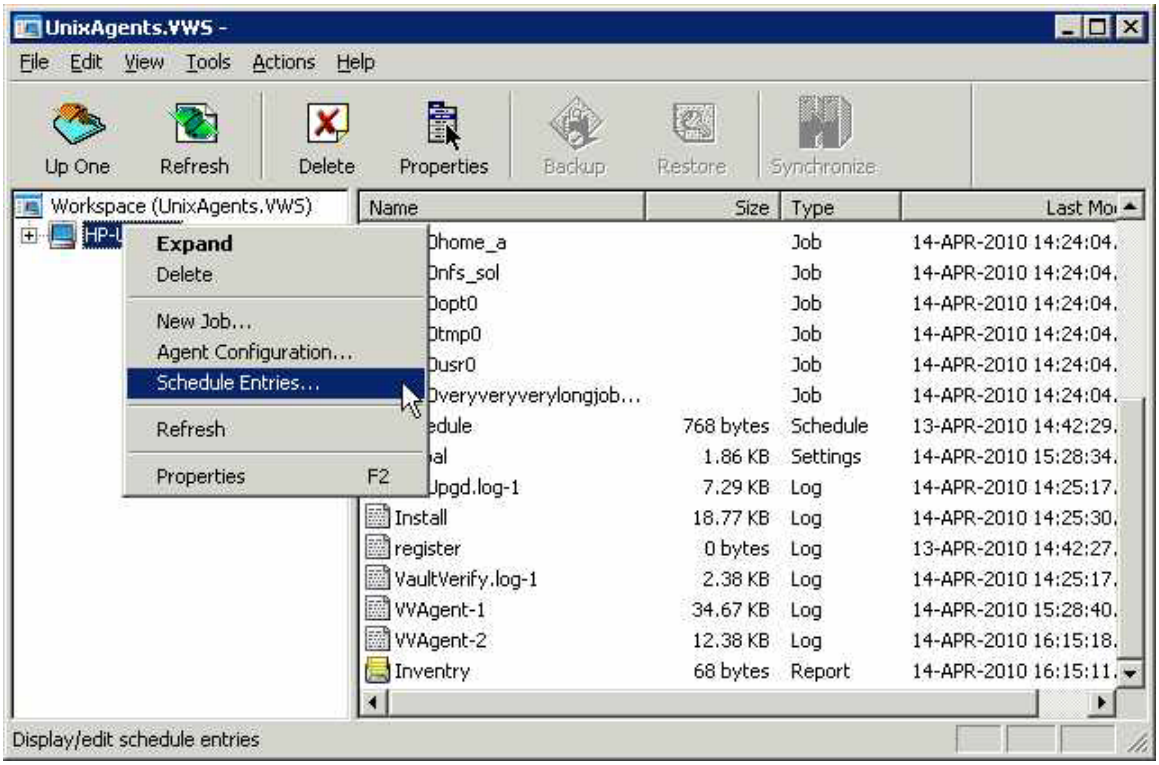

Fig 8. Schedule the Job

Start the scheduling from Tools → Schedule Entries, or right-click on a selected Agent in the left pane (see Figure 8). 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 you finish.

3. Performing Backups

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 a new password, the next backup will be a re-seed. In the case of a re-seed, your backup will take longer to complete. If a reseed occurs, the full backup “delta” and “pre-delta” bytes will be the same (the all-stream value). This is shown in the log summary.

3.1 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

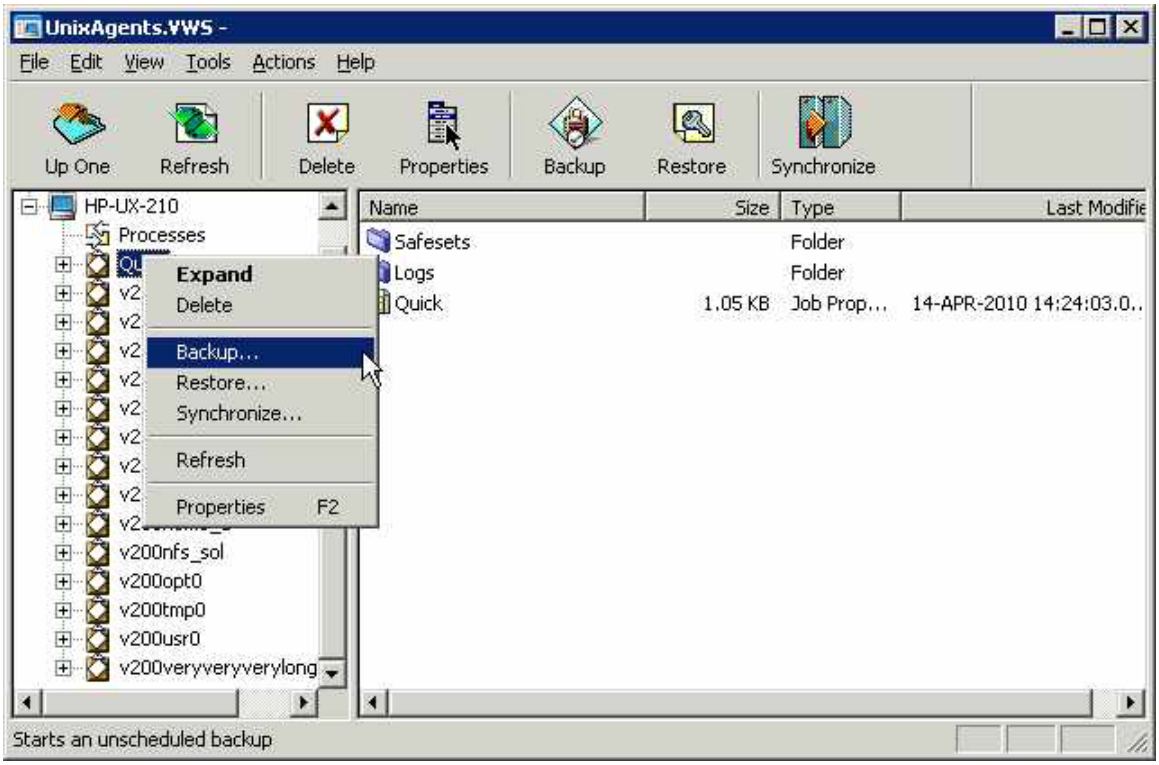

Fig 9. Ad-hoc backups

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.**
- Retention type. Select a retention scheme: Daily, Weekly, Monthly or Custom. 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.

3.1.1 File Backup Options for Solaris

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 (unchecked), 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 the original backup.

3.2 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.

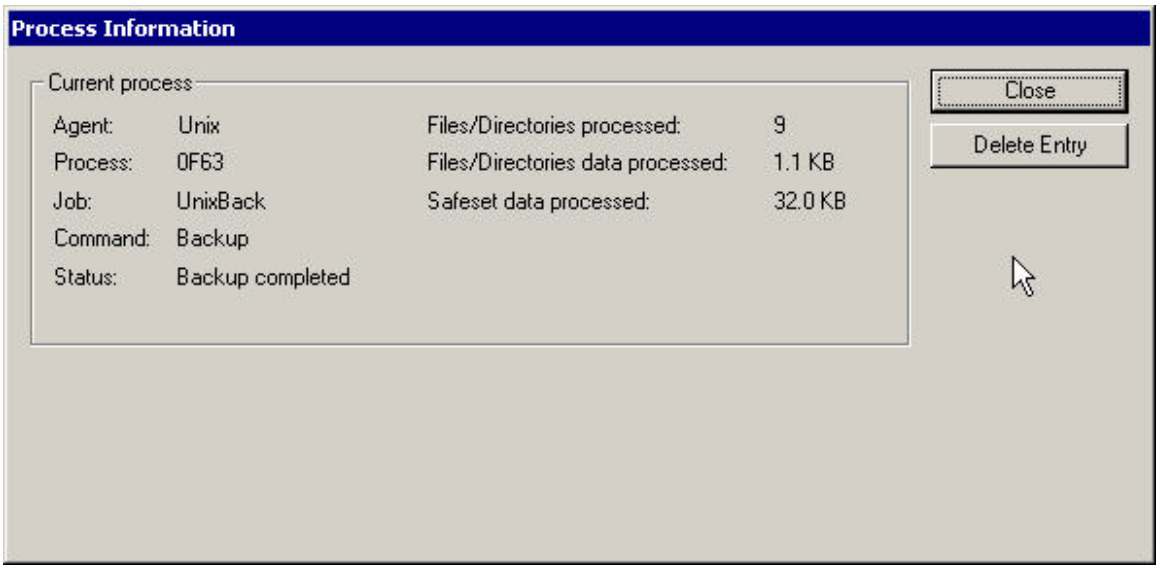

Fig 10. Checking backup results

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.

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.

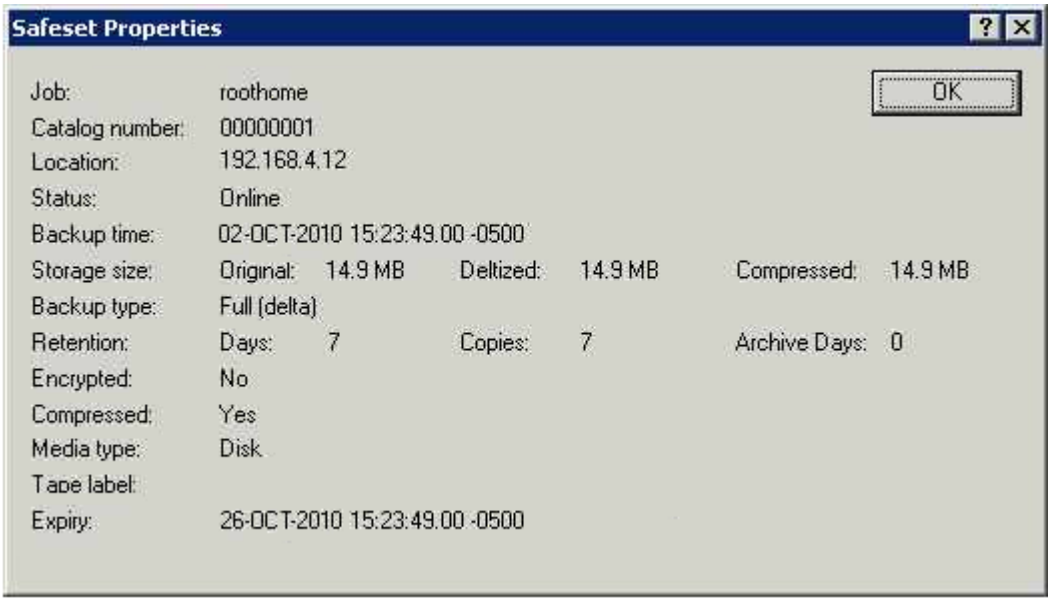

Fig 11. Safeset Properties

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.

3.3 Zones – Global and Non-Global

The 6.73 Solaris Agent is Zone-aware for Solaris 10. You can choose to protect non-Global Zone containers from the Global zone, or you can install the Agent inside each non-Global Zone independently.
There are some limitations when protecting non-Global Zones:

- The Agent does not back up loopback filesystem (LOFS) mounts. This means that if `/usr` is mounted as a loopback filesystem from the Global Zone, it will not be protected by the Agent inside the non-Global zone.

Note: If you clone a zone that contains an existing Solaris backup Agent, you must uninstall the newly cloned Agent and then perform a new Agent installation in the newly cloned zone. A complete uninstall must be done, not just the program files. The Agent must be separately registered both to Web Agent Console and Vault. The Agent on the cloned zone should be uninstalled and then a new Agent reinstalled with new registration. If the Agent is not reinstalled with a new registration, the cloned zone will have the same registration with the Vault. This will cause conflicts on backups because the Vault will interpret the two zones as the same computer.

3.3.1 Protecting Non-Global Zones from a Global Zone

You can back up a non-global zone directly by installing a Solaris Agent within the non-global zone itself. In a disaster recovery situation, you can protect the entire zone. Other than installing the agent within the non-global zone, you can protect a non-global zone from within the global zone by backing up the zonepath location of the zone(s). (see section 4.5 for Zone Restore Steps)

4. Performing Restores

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 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.

4.1 Restoring 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 Job, 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

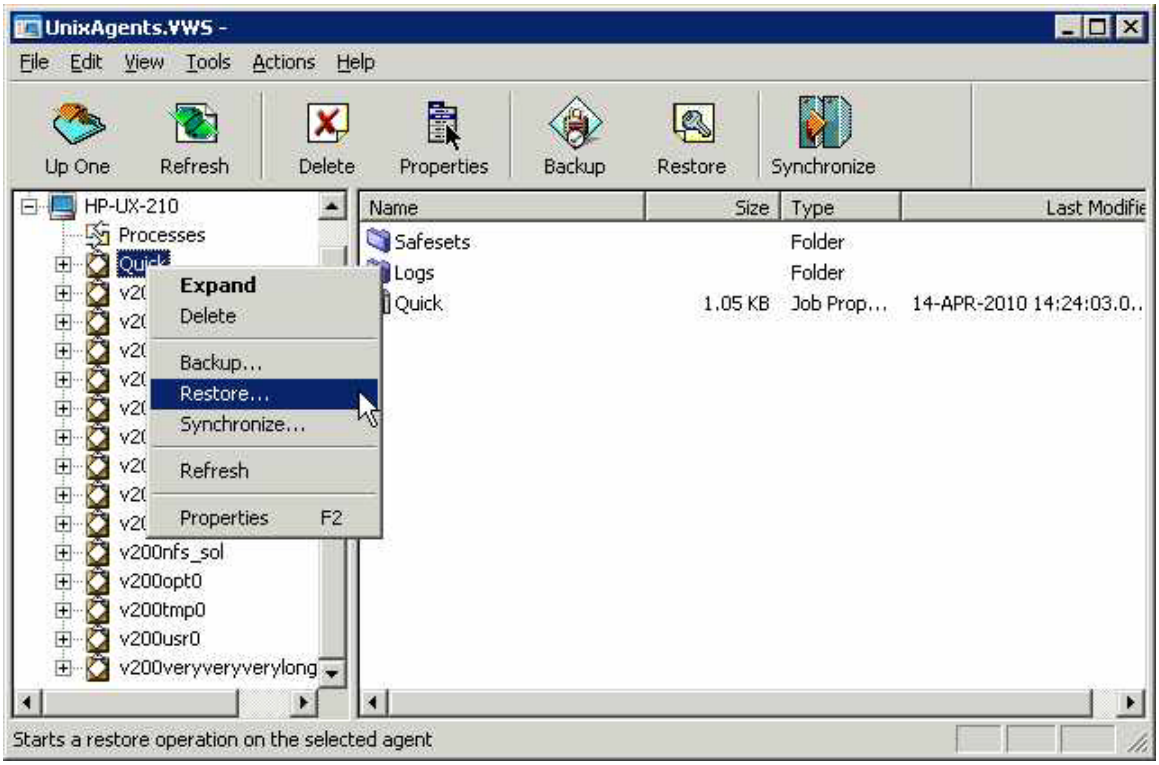

Fig 12. Restore

The Restore Wizard starts allowing you to:

- Select a Vault or backup safeset. 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). You can select or deselect files to include and exclude on 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 only 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.

4.1.1 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).

4.1.2 NFS – Network File System

The Agent has the ability to back up NFS mount points to a remote NFS Server using NFS v3 and v4.

In the Agent Console 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.

4.2 Cross-Computer Restores

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

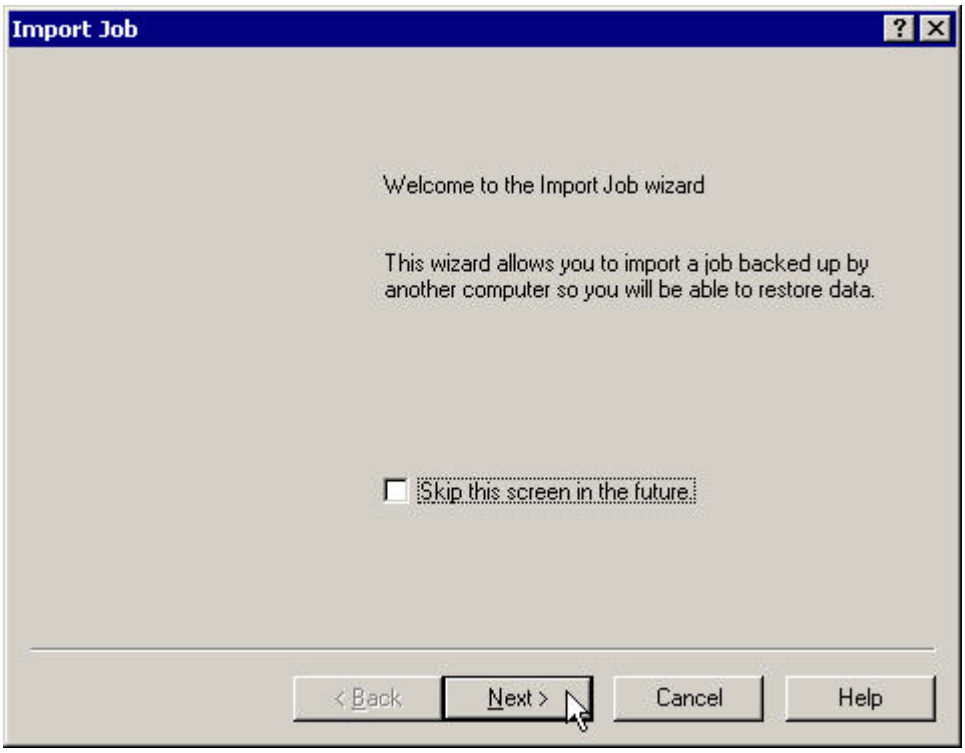

Fig 13. Restoring from another computer

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).

4.3 Disaster Recovery

“Disaster Recovery” is not a menu choice in the Agent Console 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. See section 6 for details.

4.4 Restoring ACLs

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

There are two types of ACL (Access Control List) models in Solaris: The POSIX ACL model and the NFSV4 type ACL Model used by ZFS. ACLs in general, allow the enhanced security on a file by providing access-control granularity beyond the traditional owner/group/other permissions on files and directories. ACLs may not be enabled by default on your file system.

The Solaris Agent supports both POSIX and NFSv4 ACL Models. Which ACL model is used depends on which type of file system is being used. Solaris UFS File System supports the POSIX ACL Model. The Solaris ZFS file system and NFSv4 to a remote ZFS file system supports the NFSv4 ACL model.

POSIX ACLs are set using setfacl and getfacl and use a single entry to define both allowed or denied permissions on a specific user or group. NFSv4 ACLs are set using chmod and displayed using ls -v. NFSv4 ACLs use two types of ACE (Access Control Entries). One is for ALLOW and the other is for DENY.

The two ACL models are not compatible. If you attempt to set an NFSv4 ACL on a UFS file system, you may receive an error message "ACL types are different". If you attempt to set a POSIX ACL on a ZFS files ystem, you may receive "File system doesn't support aclent_t style ACLs.". Similarly, if you attempt to restore ACLs backed up from a UFS file system to ZFS, the ACLs will not be restored. A warning message will appear in the backup log.

If you back up a non-trivial ACL on a ZFS file system (which does not contain additional user/group information other than the default owner, group and everyone permissions), you may see different results on the restore of these ACLs when "Data Steams only" is selected. That is, datastream-only restores will still inherit parent directory permissions and unmask settings.

4.5 Zone Restore Steps

Note: Protecting Non-Global Zones from a Global Zone.
You can back up a non-global zone directly by installing a Solaris Agent within the non-global zone itself. In a disaster recovery situation, you can protect the entire zone.
Other than installing the Agent within the non-global zone, you can protect a non-global zone from within the global zone by backing up the zonepath location of the zone(s).

The steps below assume that you are protecting a non-global zone from within a Global zone by recursively backing up the entire zone's root directory. It also assumes that the non-global zone is in a RUNNING state.

1. Bring down the Zone (detach).
zlogin -S [zone name] init 0

2. Verify that the zone is in INSTALLED state.
zoneadm list -civ

3. Uninstall the zone.
zoneadm -z [zone name] uninstall

4. Verify that the zone is in CONFIGURED state.
zoneadm list -civ

5. Delete the zone from disk.
zonecfg -z [zone name] delete [-F]

6. Create a new zone using your backup configuration file. If you don't have one, you will have to manually reconfigure your zone.
zonecfg -z [zone name] -f [zonename.cfg]

7. Verify that zone is in a CONFIGURED state.
zoneadm list -civ

8. Install the zone from Global.
zoneadm -z [zone name] install

9. Verify that the zone is in INSTALLED state.
zoneadm list -civ

10. Detach the installed zone.
zoneadm -z [zone name] detach

11. Verify that the zone is in CONFIGURED state.
zoneadm list -civ

12. RESTORE the zone using Original location or new zonepath using OVERWRITE.

13. Attach the configured zone.
zoneadm -z [zone name] attach

Note: The system administrator is notified of required actions to be taken if either or both of the following conditions are present:

- Required packages and patches are not present on the new machine.
- The software levels are different between machines.

14. Boot the restored zone.
zoneadm -z [zonename] boot

15. Login to the zone.
zlogin -C [zonename]

5. Installation

The software versions noted here are current for this manual. Refer to the Release Notes for any recent changes. Only version 6.7x of Solaris is covered here. Other (previous) versions are covered in earlier manuals.

There are separate Agent installation kits as well as separate Oracle Plug-In kits for the Solaris operating system that you are using as well as the processor type that you are using (SPARC or x86). Only the appropriate Agent installation kit can be installed on the associated operating system or processor. For example:

Agent installation kits:

- On Solaris 9 with SPARC: Agent-6.73-Solaris-9.tar.gz
- On Solaris 10 with SPARC: Agent-6.73-Solaris-10.tar.gz
- On Solaris 10 with x86: Agent-6.73-Solaris-x86.tar.gz

Oracle Plug-In kits:

- On Solaris 9 with SPARC: Oracle-Plug-In-6.73-Solaris-9.tar.gz
- On Solaris 10 with SPARC: Oracle-Plug-In-6.73-Solaris-10.tar.gz

Note: See the latest release notes for specific installation requirements.

The installation kit typically comes as a zipped tar file. This must be unzipped only on the machine it is intended for (the target machine). That is, do not unzip it on a non-Solaris machine, or even on another type of Solaris. 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.)

5.1 Installation – Install.sh Options

For versions 6 and up, the installation directory has a shell file called “install.sh”. The “-help” option shows all of the commands 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.
```

5.1.1 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 Agent Console) and buagent (for Agents controlled by Web Agent Console). This single script affects both vvagent and buagent.

For example:

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

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

5.1.2 Web Agent Console Registration

During the Agent installation you are prompted to register the Agent to Web Agent Console. You are also asked to choose a default language for email and command-line log viewing.

After installation, you can change the registration (new or re-register) and/or the default language. The Agent must be stopped before re-registration. Then, the Agent must be restarted for these changes to take effect.

Web Agent Console Registration:

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

If you are already registered to a Web Agent Console server, you will see:

```
Do you wish to register as a new computer?
This will invalidate your previous registration. (Y/[N])
```

Whether New or Re-Register, you will be prompted for:

```
What is the Web-based Agent Console address?  ("-" to cancel)
What is the Web-based Agent Console connection port?  ("-" to cancel)
What is your Web-based Agent Console username?  ("-" to cancel)
What is your Web-based Agent Console password?  ("-" to cancel)
```

The address is the name or IP address of Web Agent Console.
The port number is defined by the Web Agent Console Administrator.
Your user name/password authentication is set by the Web Agent Console Administrator.

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

```
wccAddress=ADDRESS_OF_AMP_SERVER
wccPort=PORT_OF_AMP_SERVER # Defaults to 8086
wccLogin=WEBCC_USER_LOGIN
wccPassword=WEBCC_USER_PASSWORD
```

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

Note: This command only applies during installation of the Agent. That is, it does not work with the “register” script, only the install.sh script.

5.1.3 Web Agent Console Language Selection

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

After installation, you can change the registration (new or re-register) and/or 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 the information about the `-help` option for more information about the language selection.

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

5.2 Agent for Solaris Installation

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

Windows Agent Console or Web Agent Console is needed to communicate with, configure, and manage the Agent for Solaris.

Note: See the latest release notes for specific installation requirements.

5.2.1 System Requirements

Hardware

- ☑ CPU – SPARC (32 and 64 bit are supported)
  - – x86 processors are supported
- ☑ RAM - 2 GB (minimum)
- ☑ Note: If your operating system suggests different minimum hardware requirements, use the most advanced requirements.

Software

- ☑ OS – Solaris 9 or 10 with SPARC processors
- ☑ OS – Solaris 10 with x86 processors

Privilege Requirements

- ☑ Installation

  To extract the installation files for the Agent for Solaris, 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 Agent Console 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 Solaris Agent from “fresh” or upgrade it from version 5.x.

5.2.2 Installation Procedures

There are separate Agent installation kits as well as separate Oracle Plug-In kits for the Solaris operating system that you are using. Only the appropriate Agent installation kit can be installed on the associated operating system. For example:

Agent installation kits:

- On Solaris 9 with SPARC: Agent-6.73-Solaris-9.tar.gz
- On Solaris 10 with SPARC: Agent-6.73-Solaris-10.tar.gz
- On Solaris 10 with x86: Agent-6.73-Solaris-x86.tar.gz

Oracle Plug-In installation kits:

- On Solaris 9: Oracle-Plug-In-6.73-Solaris-9.tar.gz
- On Solaris 10: Oracle-Plug-In-6.73-Solaris-10.tar.gz

The installation kit typically comes as a zipped tar file. This must be unzipped only on the machine it is intended for (the target machine). That is, do not unzip it on a non-Solaris machine, or even on another type of Solaris. 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.)

5.2.2.1 Requirements

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

The Agent for Solaris Installation kit.
A target system running a supported version of Solaris.
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 –F` (but this only applies to the initial check for actual installation space, rather than extra temporary space for rolling back).

5.2.2.2 Running the Installation Kit

Depending on your service provider, you may have the option of downloading the file from the Web, from diskette and/or from a CD.

To install the Agent for Solaris:

1. Download the appropriate Agent installation kit for Solaris 9, 10 with the appropriate processor type (SPARC or x86).

   For SPARC:

   `Solaris 9: Agent-6.73-Solaris-9.tar.gz`,

   `Solaris 10: Agent-6.73-Solaris-10.tar.gz`

   For x86:

   `Solaris 10: Agent-6.73-Solaris-x86.tar.gz`

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

2. Extract the files from the package. To do so, use this command (in which "PACKAGENAME" could be "Agent- Solaris -6.73"):
   `gunzip -c PACKAGENAME.tar.gz | tar xvf -`
3. Run the installation script.

   `# ./install.sh`

Note: The installation script will interactively prompt you for configuration information such as web registration (address, port number and authentication), log file name, and language selection for logs and command lines. See section 5.1.2 of this guide.

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

The `Install.log` is in the installation directory, if successful. For example:

`<Agent installation directory>/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>.`

5.2.3 Uninstall Procedures

To uninstall the Agent for Solaris:

1. Log onto the target system.
2. Go to the installation directory (by default `/opt/BUAgent/`).

3. Run "`uninstall.sh`". A message will appear, asking if you want to remove the Agent. (The uninstall script will Stop 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 will appear.

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

5.2.4 Upgrading

This Agent version supports upgrades from Agent 5.x and above.

Upgrading an Agent includes the following tasks:

1. Meeting System and Software Requirements
2. Preparing the Computer
3. Upgrading Program Files and Configuration Files
4. Running the upgrade (installation) Job

When you run the Installation Kit, the Agent is upgraded to Version 6.7x and the following tasks are automatically performed:

- The new BuildVersion and VVCVersion are created in Global.vvc and Job vvc files:
  `BuildVersion = 6.7X.XXXX`

  `VVCVersion = XXX`
- Server profiles in Global.vvc are updated with the computer registration information, such as computer GUID (Global Unique ID), computer name and Vault GUID. These must match the information stored in the Vault's database.
- Job vvc files are updated with Job registration information. This includes computer GUID (must be the same as in Global.vvc), Job GUID (must match the record of the Vault's database) and Vault GUID (must be the same as in Global.vvc).
- All Delta files are upgraded to Version 6.7x 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.

5.2.4.1 Meeting System and Software Requirements

To upgrade the Agent, your system must meet the minimum requirements mentioned in the User Guide.

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

5.2.4.2 Preparing the Computer

To prepare your existing machine for upgrading the Agent, complete these tasks:

a. Back up the previous Agent Files

   We strongly recommend that you make at least one backup of your previous Agent files, including all files and subdirectories under the Agent installation directory. Do not attempt an upgrade without a backup.

b. Clean up Server Profiles in Global.vvc

   From Agent Console, open up the Agent Configuration of the Agent that you want to upgrade. Go to the Vaults section, look for a server configuration that is no longer in use, and delete it. Also, highlight every server configuration and click Edit to double-check that the information in this server profile is valid. Then click OK to save your changes.

c. Clean up Jobs

   After the Global.vvc has been cleaned up, check all backup Jobs to see if there is any Job backing up to a Vault that has been deleted from Agent Configuration. If so, delete that Job.

   If you have Jobs that are backing up to Directory on Disk, they are local Jobs and leave them unchanged. During the upgrade, they will be registered to the first Vault indicated in Agent Configuration.

d. Synchronize all backup Jobs

   After cleaning up the Jobs, check the backup logs of each Job to see if any errors show “Validation failed: “. If so, you need to verify the validation information with your Vault operator to make sure it is valid. If the latest backup log shows no errors, do a Synchronize with the Vault and check the Synch log.

e. Verify eligible Vault version

   For every Vault that you are backing up to, make sure it is running Vault Version 5.53 or higher. Otherwise, you need to upgrade the Vault before you upgrade the Agent.

5.2.4.3 Upgrading Program Files and Configuration Files

We recommend starting the installation when the Directors in Agent Configuration are not busy with other Jobs.

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

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

5.2.4.4 Upgrade Steps

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.*

1. Log onto the target system.
2. Go to the installation kit directory.
3. Use the VVAgent script to stop the Agent (see section 5.1.1).
4. Download the appropriate `Agent-6.73-Solaris-9.tar.gz` or `Agent-6.73-Solaris-10.tar.gz` package for Solaris 9 or 10.

   Note: You must do this locally (i.e., on the target machine).
5. Extract the files from the package. To do so, use this command (where "PACKAGENAME" could be "Agent-Solaris-6.73"):

   ```
   gunzip -c PACKAGENAME.tar.gz | tar xvf –
   ```
6. 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 an upgrade fails, the Global.vvc, Job vvc and Delta files will roll back to their old versions. The Global, Job and Deltas will not work with new executables, but an upgrade failure will roll back the executables, too.

You may try to run the upgrade program again. If it still fails, contact your service provider for support. To completely roll back to the old version, you need to manually copy back the previous backups.

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

5.2.5 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.

6. Solaris System Recovery

The purpose of this chapter is to illustrate techniques for recovering a Solaris file system. The procedures provided describe the minimum resources and information required to rebuild the Solaris 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 for a Solaris system 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 with a Solaris system, ensure that your hardware configuration is at least sufficient to hold the programs, data, and system state previously installed on the system.

Note: When performing a complete system restore (DR), you need to ensure there is ample disk space for the creation of large restore 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.

6.1 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.

6.2 Software Requirements

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

Solaris 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 Solaris Restoration Steps

This section describes the steps to perform a system restore.

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 appropriate for your operating system.
2. Configure the Agent according to the instructions in section 2 of this manual. It is important to re-register 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 Solaris 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 restoration of files to critical locations, but not all critical locations are necessarily detected. Also, the locations can vary between varieties of Solaris.

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

6.3.4 Perform post-restore 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 restore

Once the restore procedure is complete, determine if the restoration 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 restoration 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 Solaris 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 restoration?
- 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 supports Windows and Solaris Agents. It is an add-on that allows you to perform database backups on Oracle databases.

The Plug-In installs on top of the Agent on the database host to perform backups.

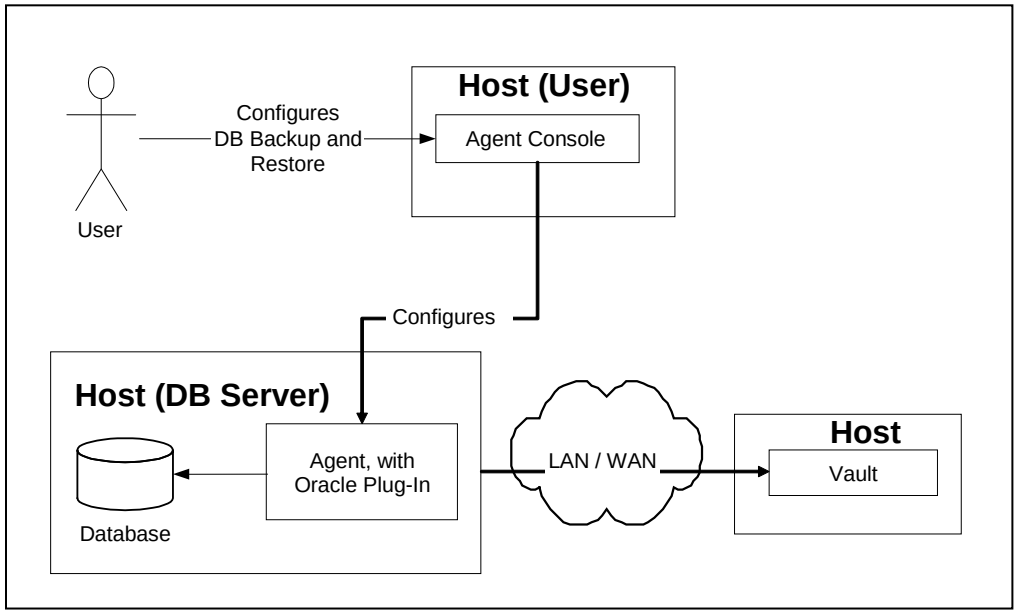

Fig 14. Oracle Plug-In Diagram

This diagram illustrates the basic product implementation. A user, typically a DBA, configures the backup via the Web Agent Console or Windows Agent Console application. Agent Console configures the Agent, which typically resides on a different host from Agent Console.

A user can schedule a backup of the database, at which time the Agent (with the help of the Oracle Plug-In) will send the database information to the Vault.

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.
- Agent versions 5.6 and above specify databases using Oracle Service Names. They do not require script-level or backup-level ORACLE_HOME customization.
- Database passwords are encrypted for enhanced security over script-based methods.

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.

7.1.3 Release Notes and Help

Release Notes

Release notes provide “up to the minute” information about the released product. They also contain an overview of new features, known defect (bug) fixes incorporated since the last release, a description of any known issues, and a section about product support. Release notes are available from your service provider.

Online Help

Agent Console (Windows Agent Console/Web Agent Console) provides online help, which contains information similar to the contents of this User Guide.

There is also context-sensitive “What’s This” help on each Windows Agent Console GUI screen. You can access the context-sensitive “What’s This” help by clicking the Help icon (question mark) in the Agent Console application. Note: If the Windows Agent Console 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.

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

7.2 Installing the Oracle Plug-In on Solaris

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

Installing the Oracle Plug-In for Solaris requires that you have previously installed the Solaris Agent application.

The Oracle Plug-In program typically arrives in a tar file. There are separate Oracle Plug-In kits for the Solaris operating system that you are using. Only the appropriate Oracle Plug-In installation kit can be installed on the associated operating system. For example:

Oracle Plug-In installation kits:

- On Solaris 9: Oracle-Plug-In-6.73-Solaris-9.tar.gz
- On Solaris 10: Oracle-Plug-In-6.73-Solaris-10.tar.gz

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

- Solaris 9: Oracle 10g R2 (10.2.0.0) and Oracle 11g R1 (11.1.0.0) (32-bit and 64-bit Solaris SPARC processors)
- Solaris 10: Oracle 10g R2 (10.2.0.0), Oracle 11g R1 (11.1.0.0), and Oracle 11g R2 (11.2.0.0) (32-bit and 64-bit Solaris SPARC processors)

7.2.3 Before Installing or Upgrading

- The Solaris Agent and the Oracle Plug-In must be installed on the system that has the Oracle database server.
- The Solaris Agent and the Oracle Plug-In should always have the same version number.
- 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 has been tested in homogeneous and heterogeneous installations of Oracle 10g and 11g with a single listener. The Oracle Plug-In can *only* find the TNS name list (`tnsnames.ora`) in the global location `/var/opt/oracle`. 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 appropriate `Oracle-Plug-In-6.73-Solaris-9.tar.gz` or `Oracle-Plug-In-6.73-Solaris-10.tar.gz` package.

2. Extract the files from the package. To do so, type the following (note: `Solaris-9` is used in this example):

```
# cd /tmp
# tar xvf Oracle-Plug-In-6.73-Solaris-9.tar
```

3. Next, type the following:

```
# cd Oracle-Plug-In-6.73-Solaris-9
```

4. Run the installation script.

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

5. Follow the installation instructions on the screens.

7.2.5 Upgrading the Plug-In

An upgrade of the Plug-In is similar to a new installation. The installation script will stop the Agent prior to the upgrade. It will offer to start the Agent after a successful installation, and it will not start the Agent after a failed installation.

Upgrade the Oracle Plug-In as a **root** user**.**

1. Download the appropriate `Oracle-Plug-In-6.73-Solaris-9.tar.gz` or `Oracle-Plug-In-6.73-Solaris-10.tar.gz` package.

2. Extract the files from the package. To do so, type the following (note: `Solaris-9` is used in this example):

```
# cd /tmp
# tar xvf Oracle-Plug-In-6.73-Solaris-9.tar
```

3. Next, type the following:

```
# cd Oracle-Plug-In-6.73-Solaris-9
```

4. Run the installation script.

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

5. Follow the installation instructions on the screens.

7.2.6 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 is in the installation directory (typically `/opt/BUAgent`). After you run the uninstall script, stop and start the Agent.

7.2.7 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 Solaris Oracle Plug-In to create backup Jobs, the license must be available on the Vault. A version 6.04 or greater Vault will supply the license for the Plug-In. 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 back up, and database account information)	Database Service Name = User Name = Password =	Validates the fields, and allows connection to the database
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	Immediate = 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)

Fig 15. Table of Backup information

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. From the Web Agent Console GUI program (which communicates with the Agent and the Plug-In), create a new Job.

2. Select "Oracle" for the Backup Source Type. The Oracle Options will appear on the page.

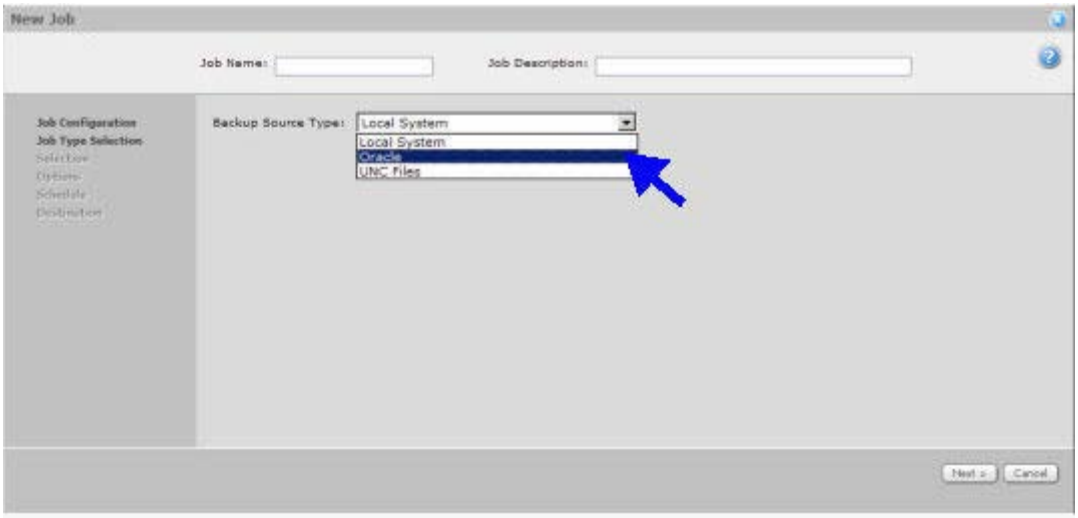

Fig 16. Oracle Backup Source

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

   For Oracle 11g in Web Agent Console, set the Database Service Name to the *Database Instance* from Oracle (rather than the *Instance Name* from Oracle).

   For Oracle 11g in Windows Agent Console, 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 are connecting to a database that is listening on a port other than the default, the format for the database service name is:

Service Name:port #
i.e. 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.

Log files are created on the Web Agent Console machine, under the installation directory. Their names match the Job names. You normally view them from Web Agent Console.

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

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

1. Shut down the database.
2. Restore the files using the Restore to an Alternate Location option.
   (**Note:** You can restore to the original location by directing the restore to the original location)
3. If the file names have been renamed, they must be changed 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 on Solaris

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.