

Quick Start - PostgreSQL iDataAgent

TABLE OF CONTENTS

OVERVIEW

Introduction

Key Features

- Wide Range of Backup and Restore Options
- Efficient Job Management and Reporting
- Block Level Deduplication

Terminology

SYSTEM REQUIREMENTS - POSTGRESQL /DATAAGENT

INSTALLATION

Install Software from the CommCell Console

Install the PostgreSQL iDataAgent - Unix

Silent Install - Unix Agents

CONFIGURATION - - POSTGRESQL /DATAAGENT - BASIC

Overview

Creating an Instance

BACKUP - POSTGRESQL /DATAAGENT - BASIC

Overview

Your First Full Backup

- Running a Full Backup Immediately
- Scheduling a Full Backup

Viewing Backup Job History

RESTORE DATA - POSTGRESQL /DATAAGENT - BASIC

Overview

Your First Restore

- Basic Restore
- Browse and Restore

Viewing Restore Job History

- View Restored Items Associated With a Restore Job

Overview - PostgreSQL iDataAgent

TABLE OF CONTENTS

Introduction

Key Features

- Wide Range of Backup and Restore Operations
- Efficient Job Management and Reporting
- Block Level Deduplication

Terminology

INTRODUCTION

The PostgreSQL iDataAgent is designed to efficiently manage and protect business critical data in PostgreSQL databases. You can use this iDataAgent to backup and restore the entire PostgreSQL server or individual databases. If necessary, you can also restore individual tables.

KEY FEATURES

The PostgreSQL iDataAgent offers the following key features:

WIDE RANGE OF BACKUP AND RESTORE OPERATIONS

PostgreSQL iDataAgent provides the flexibility to backup the PostgreSQL databases in different modes and restore them in minimal time. You can perform a full or log backups of the entire PostgreSQL server or individual databases or archive logs at any point of time.

The following section describes the backups and restores that can be performed in different modes:

BACKUP AND RESTORE ENTIRE POSTGRESQL SERVER

This iDataAgent provides the ability to restore an entire PostgreSQL Server. All the databases that reside on a source server can be restored on to a destination server.

In addition to performing backup and restores of databases, this iDataAgent performs database and log backups like any other database iDataAgent and effortlessly restores all the databases in a PostgreSQL server.

BACKUP AND RESTORE SELECTIVE INDIVIDUAL DATABASES

You can use the CommCell Console to create user-defined subclients and distribute some of the database content for dump based backups. Distributing the client data using subclients in this way can help improve backup performance by organizing the workload on the client into logical groupings. While creating user defined subclients for dump based backup data, you can also choose selected object list. This option allows you to save time while restoring this subclient's data.

You can either define an individual database or a group of databases as subclient data and perform backup and restores.

BACKUP LOGS


You can perform backups of only the logs in the PostgreSQL server. Log files are generated for each transaction made to the data files in the database. These log files can be used to recover the database transactions that have been lost due to an operating system or disk failure. You can apply them to the File System backup to recover the entire PostgreSQL server to a specific point-in-time.

RESTORE DATABASE TABLES

This iDataAgent allows you to restore and recover individual database tables to the same database or to another database. While restoring a dump based backup, you can select the Table View from the browse options to view all the tables in a database. During the restore operation, you can select individual tables from the list of database tables in a tree view from browse window and restore them to a destination location.

RESTORE DATABASES TO A POINT-IN-TIME

You can restore and recover an entire PostgreSQL server to a point in time for a File System Based Backup. A point in time restore job allows you to restore your PostgreSQL server to a specific point in time, and, typically, this point in time would be just before an undesired transaction, corruption, or update to your database.



EFFICIENT JOB MANAGEMENT AND REPORTING

You can view and verify the status of backup and recovery operations from the Job Controller and Event Viewer windows within the CommCell Console. You can also track the status of the jobs using Reports, which can be saved and easily distributed. Reports can be generated for different aspects of data management. You also have the flexibility to customize the reports to display only the required data and save them to any specified location in different formats. For example, you can create a backup job summary report to view at-a-glance the completed backup jobs.

In addition, you can also schedule these reports to be generated and send them on email without user intervention.

BLOCK LEVEL DEDUPLICATION

Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

Deduplication at the data block level compares blocks of data against each other. If virtual machines contains blocks of data that are identical to each other, block level deduplication eliminates storing the redundant data and reduces the size of the data in storage. This dramatically reduces the virtual machine backup data copies on both the disk and tapes.

TERMINOLOGY

The following terminology is used by this product:

CLIENT	The computer in which the iDataAgent is installed and contains the data to be secured.
INSTANCE	The PostgreSQL database to be used for the backup and restore operations.
BACKUPSET	A backupset is a logical grouping of subclients. This iDataAgent creates backupsets by default while creating an instance.
DUMP BASED BACKUPSET	A dump based backupset contains selected individual or a group of databases herein referred as dumps.
FILE SYSTEM BASED BACKUPSET	A file system based backupset contains an entire PostgreSQL server which includes all the databases and logs.
SUBCLIENT	The PostgreSQL data to be backed up.

[Back to Top](#)

System Requirements - PostgreSQL iDataAgent

[System Requirements](#) | [Supported Features](#)

The following requirements are for the PostgreSQL iDataAgent:

PostgreSQL version 9.x is supported with SP5A and the Enterprise DB version 9.0 is supported with SP5B.

APPLICATION/OPERATING SYSTEM		PROCESSOR
POSTGRESQL VERSION 8.2 (OR HIGHER) AND ENTERPRISE DB VERSION 8.3 (OR HIGHER) ON:		
LINUX	RED HAT ENTERPRISE LINUX/CENTOS	
	Red Hat Enterprise Linux/CentOS 6.x with glibc 2.12.x	Intel Pentium, x64 or compatible processors
	Red Hat Enterprise Linux/CentOS 5.x with glibc 2.5.x	Intel Pentium, x64 or compatible processors
	Red Hat Enterprise Linux/CentOS 4.x with a minimum of glibc 2.3.4	Intel Pentium, x64 or compatible processors
	SOURCE MAGE LINUX	
	Source Mage Linux	x64
	SUSE LINUX (SLES)	
	SuSE Linux 11.x with glibc 2.9.x and above	Itanium, x64 or compatible processors
	SuSE Linux 11.x with glibc 2.6.x	Intel Pentium, x64 or compatible processors
	SuSE Linux 10.x with glibc 2.4.x	Intel Pentium, x64 or compatible processors

HARD DRIVE

220 MB minimum of hard disk space for software

50 MB of additional hard disk space for log file growth

725 MB of temp space required for install (where the temp folder resides)

MEMORY

64 MB RAM per stream/drive minimum required beyond the requirements of the operating system and running applications

Swap space = 2*RAM size

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

The File System iDataAgent will be automatically installed during installation of this software, if it is not already installed. For System Requirements and install information specific to the File System iDataAgents, refer to System Requirements - Linux File System iDataAgent.

NETWORK

TCP/IP Services configured on the computer.

DRIVERS

Drivers for SCSI Adapters (Supplied by your SCSI Adapter provider)

Drivers for media drives (Supplied by your media drive provider)

DISCLAIMER

Minor revisions and/or service packs that are released by application and operating system vendors are supported by our software but may not be individually listed in our System

Requirements. We will provide information on any known caveat for the revisions and/or service packs. In some cases, these revisions and/or service packs affect the working of our software. Changes to the behavior of our software resulting from an application or operating system revision/service pack may be beyond our control. The older releases of our software may not support the platforms supported in the current release. However, we will make every effort to correct the behavior in the current or future releases when necessary. Please contact your Software Provider for any problem with a specific application or operating system.

Additional considerations regarding minimum requirements and End of Life policies from application and operating system vendors are also applicable

Install the PostgreSQL iDataAgent - Unix

TABLE OF CONTENTS

Install Requirements

Before You Begin

Install Procedure

- Getting Started
- Select Components for Installation
- Base Software Installation
- Kernel Parameters
- Enable Global Filters
- Client Group Selection
- Storage Policy Selection

Post-Install Considerations

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing the PostgreSQL iDataAgent and the Unix File System iDataAgent. The PostgreSQL iDataAgent is installed on the computer on which the PostgreSQL application resides. (This computer is referred to as the *Client* computer in this install procedure.)

Verify that the computer in which you wish to install the software satisfies the minimum requirements specified in System Requirements - PostgreSQL iDataAgent and System Requirements - Linux File System iDataAgent.

Review the following Install Requirements before installing the software:

GENERAL

- Review Install Considerations before installing the software.
- Agents should be installed only after the CommServe and at least one MediaAgent have been installed in the CommCell. Also, keep in mind that the CommServe and MediaAgent must be installed and running (but not necessarily on the same computer), before you can install the Agent.
- Ensure there is an available license on the CommServe for the Agent.
- Verify that you have the Software Installation Disc that is appropriate to the destination computer's operating system.

PACKAGE MANAGEMENT SYSTEMS

You can use the Red Hat Package Manager (RPM) package management system to install all Linux versions of the agent. Also, you can use the Advanced Packaging Tool (APT) package management system to install all Debian Linux versions of the agent. For step-by-step instructions, see Base Software Install for RPM Package Manager and Base Software Install for APT Package Manager.

RED HAT LINUX

Red Hat Linux will create an entry in the `/etc/hosts` file when it is first installed, in the following format:

```
<ip_address> <host name> localhost
```

For example, if the host name of your computer is `bluesky`, the entry will look something like this:

```
192.168.1.111 bluesky localhost
```

If you have not already done so, edit the `/etc/hosts` file. The edited entry should look like this:

```
127.0.0.1 localhost
```

Depending upon your environment, and using the above example again, you may also need an entry similar to this:

```
192.168.1.111 bluesky
```

BEFORE YOU BEGIN

- The PostgreSQL application has been installed on the client.
- Log on to the client as `root`.
- The install package requires `super-user` permissions to execute.

INSTALL PROCEDURE

GETTING STARTED

- Place the software installation disc for the Unix platform into the disc drive.
You can also install the product using a disc drive mounted on another computer on the network.

- On Solaris, double-click the **cvpkgadd** program from the File Manager window.
- On other Unix platforms, open the Terminal window, navigate to the software installation disc and then enter **./cvpkgadd**.

- Enter the number corresponding to the setup task you want to perform.

NOTES

- For Install data protection agents on this computer option, follow the steps described in this procedure.
- Advance options provide additional setup features such as record and play setup, creating a custom package and External Data Connector Agent software.

To create a custom package and for record and play setup, follow the steps described in Custom Package - Unix.

To install the External Data Connector Agent, follow the steps described in External Data Connector - Unix.

- The product banner and other information is displayed.

Press **Enter** to continue.

- Read the license agreement. Type **y** and press **Enter** to continue.

- If your computer is 32-bit, press **Enter**.

If your computer is 64-bit, see Install Unix Agents on 64-bit Platform for step-by-step procedure.

- This prompt is displayed only when you are installing on HP-UX, Linux, or Solaris computers.

Press **Enter** to continue

NOTES

- When you install on non-clustered computer, you must select the number associated with the option **Install on a physical machine**.

- If you have only one network interface, press **Enter** to accept the default network interface name and continue.

If you have multiple network interfaces, enter the number corresponding to the network interface that you wish to use as default, and then press **Enter** to continue.

NOTES

- The interface name and IP addresses depend on the computer in which the software is installed and may be different from the example shown.

If you have multiple network interfaces, enter the number corresponding to the network interface that you wish to use as default, and then press **Enter** to continue.

NOTES

- The interface name and IP addresses depend on the computer in which the software is installed and may be different from the example shown.

Verify the default network interface name.

Press **Enter** to accept the default network interface name and continue, or Type the default network interface name, and then press **Enter** to continue.

NOTES

Please select a setup task you want to perform from the list below:

Advance options provide extra setup features such as creating custom package, recording/replaying user selections and installing External Data Connector software.

- 1) Install data protection agents on this computer
- 2) Advance options
- 3) Exit this menu

Your choice: [1]

This machine supports both 32 bit and 64 bit binaries. By default, we will install 32 bit binary set that has full support for all the modules included in this package. Please note that 64 bit binary set currently only support limited modules.

- 1) All platforms (32 bit)
- 2) FS and MA only (64 bit)

Your choice: [1]

Certain Calypso packages can be associated with a virtual IP, or in other words, installed on a "virtual machine" belonging to some cluster. At any given time the virtual machine's services and IP address are active on only one of the cluster's servers. The virtual machine can "fail-over" from one server to another, which includes stopping services and deactivating IP address on the first server and activating the IP address/services on the other server.

You now have a choice of performing a regular Calypso install on the physical host or installing Calypso on a virtual machine for operation within a cluster.

Most users should select "Install on a physical machine" here.

- 1) Install on a physical machine
- 2) Install on a virtual machine
- 3) Exit

Your choice: [1]

Network interface with the following IPs have been found available on your system. One of these interfaces should be associated with the physical machine being installed. It will also be used by the CommServe to initiate connections to the physical machine. Note that you will be able to additionally customize Datapipe Interface Pairs used for the backup data traffic later in the Calypso Java GUI.

Please select the correct network interface below.

- 1) client (201.42.33.598)
- 2) hk97:489:9glg:hk8d:9490
- 3) client.company.com (hr90:8842:2:78:013:8ghh:hg8k:9x54)

Interface number: [1]

Network interfaces with the following IPs have been found configured on your system. Please select one to be used by the CommServe to contact the client.

- 1) dunk.company.com (480.19.65.333)
- 2) 480.19.66.326

Interface number: [1]

Please verify the physical machine interface name below. Make it as complete (with fully qualified domain name) as possible:

Physical Machine Host Name: [angel.company.com]

- This prompt will be displayed only when you have multiple network interfaces for the computer.

8. Specify the client name for the computer.

Press **Enter** to accept the default name and continue, or
Enter a new client name for the computer and then press **Enter** to continue.

Please specify the client name for this machine.

It does not have to be the network host name: you can enter any word here without spaces. The only requirement is that it must be unique on the CommServe.

Physical Machine Client name: [angel]

SELECT COMPONENTS FOR INSTALLATION

9. Enter the number corresponding to the **CVGxPostGres** module and press **Enter**.

```
Install Calypso on physical machine client.company.com
Select the Calypso module that you would like to install
[ ] 1) Media Agent [1301] [CVGxMA]
[ ] 2) PostGresiDataAgent [1209] [CVGxPostGres]
>) >>>> NEXT PAGE >>>>>>
```

```
[a=all n=none r=reverse q=quit d=done >=next <=previous ?
=help]
```

Enter number(s)/one of "a,n,r,q,d,>,<," here: 2

10. A confirmation screen will mark your choice with an "X". Type "d" for **Done**, and press **Enter** to continue.

NOTES

- To select multiple component, enter the number by adding a space.
- Your screen may look different from the example shown.
- Components that either have already been installed, or which cannot be installed, will not be shown.
- In addition, the list of modules that appear depends on the specific Unix File System in which the package is installed. (e.g., **CVGxWA** will appear only when the installation package is run on a Solaris computer.)

```
Install Calypso on physical machine client.company.com
Select the Calypso module that you would like to install
[ ] 1) Media Agent [1301] [CVGxMA]
[X] 2) PostGresiDataAgent [1209] [CVGxPostGres]
>) >>>> NEXT PAGE >>>>>>
```

```
[a=all n=none r=reverse q=quit d=done >=next <=previous ?
=help]
```

Enter number(s)/one of "a,n,r,q,d,>,<," here: d

BASE SOFTWARE INSTALLATION

11. If you wish to install the agent software for restore only, enter **Yes** and press **Enter** to continue. See *Installing Restore Only Agents* for more information.

Do you want to use the agents for restore only without consuming licenses? [no]

Otherwise, accept **no**, press **Enter** to continue.

12. Type the appropriate number to install the latest software scripts and press **Enter** to continue.

Installation Scripts Pack provides extra functions and latest support and fix performed during setup time. Please specify how you want to get this pack.

NOTES

- Select **Download from the software provider website** to download the latest software scripts from your software provider website.
Make sure you have internet connectivity when you are using this option.
- Select **Use the one in the installation media**, to install the software scripts from the disc or share from which the installation is performed.
- Select **Use the copy I already have by entering its unix path**, to specify the path if you have the software script in an alternate location.

If you choose to download it from the website now, please make sure you have internet connectivity at this time. This process may take some time depending on the internet connectivity.

```
1) Download from the software provider website.
2) Use the one in the installation media
3) Use the copy I already have by entering its unix path
Your choice: [1] 2
```

13. Enter **Yes** to download and install the latest service packs and post packs from the software provider.

Keep Your Install Up to Date - Latest Service Pack

NOTES

- Internet connectivity is required to download updates.
- This step is applicable for multi instancing.

Latest Service Pack provides extra functions and latest support and fix for the packages you are going to install. You can download the latest service pack from software provider website.

If you decide to download it from the website now, please make sure you have internet connectivity at this time. This process may take some time depending on the internet connectivity.

Do you want to download the latest service pack now? [no]
Press <ENTER> to continue ...

14. Specify the location where you want to install the software.

Please specify where you want us to install Calypso binaries.

NOTES

- The amount of free space required depends on the components selected for install, and may look different from the example shown.

It must be a local directory and there should be at least 209MB of free space available. All files will be installed in a "calypso" subdirectory, so if you enter "/opt", the files will actually be placed into "/opt/calypso".

Press **Enter** to accept the default path and continue, or
Enter a path and then press **Enter** to continue.

Installation Directory: [/opt]

Press **Enter** again to confirm the path.

```
..
Calypso will be installed in /opt/calypso.
Press ENTER to continue ...
```


15. Specify the location for the log files.

NOTES

- All the modules installed on the computer will store the log files in this directory.
- The amount of free space required depends on the components selected for install, and may look different from the example shown.

Press **Enter** to accept the default path and continue, or
Enter a path and then press **Enter** to continue.

Press **Enter** again to confirm the path.

16. Indicate whether you would like to launch processes with inherent database access rights.

Press **Enter** to assign a new group, or
Type **No** and then press **Enter** to continue.

17. If you indicated **Yes** in the previous step, you will be prompted for the group name that must be used to launch processes.

Enter the group name and then press **Enter** to continue.

Press **Enter** again to continue.

For installs on a Solaris computer, proceed to the next step. Otherwise, skip to Storage Policy Selection.

18. Type a network TCP port number for the Communications Service (CVD) and press **Enter**.

Type a network TCP port number for the Client Event Manager Service (EvMgrC) and press **Enter**.

NOTES

- For more information about Network TCP Ports, see Network TCP Port Requirements.
- For more information about these services, see Services.
- If the port number you entered already exists, a message will be displayed `Port #### is already reserved in /etc/services`. To work around this issue, enter different port number.

19. If this computer and the CommServe is separated by a firewall, type **Yes** and then press **Enter** to continue.

For firewall options and configuration instructions, see Firewall Configuration and continue with the installation.

If you do not wish to configure the firewall services, type **No** and then press **Enter** to continue.

20. Type the name of the CommServe computer and press **Enter** to continue.

NOTES

- Ensure that the CommServe is accessible before typing the name; otherwise the installation will fail.
- If you enter a short name which resolves to the same IP address as the fully qualified CommServe name, you will be asked if you would prefer to use the fully qualified name.

21. Enter the **username** and **password** information for an external domain user account or a CommCell user account. This authorizes the installation of an agent on the CommCell.

NOTES

- This is only displayed when the **Authentication for Agent** feature is enabled in

Please specify where you want to keep Calypso log files.
It must be a local directory and there should be at least 100MB of free space available. All log files will be created in a "calypso/Log_Files" subdirectory, so if you enter "/var/log", the logs will actually be placed into "/var/log/calypso/Log_Files".

Log Directory: [/var/log]

..

Calypso log files will be created
in /var/log/calypso/Log_Files.
Press ENTER to continue ...

Most of Calypso processes run with root privileges, but some are launched by databases and inherit database access rights. To make sure that registry and log files can be written to by both kinds of processes we can either make such files world-writeable or we can grant write access only to processes belonging to a particular group, e.g. a "calypso" or a "dba" group.

We highly recommend now that you create a new user group and enter its name in the next setup screen. If you choose not to assign a dedicated group to Calypso processes, all temporary and configuration files will be created with -rw-rw-rw permissions.

If you're planning to backup Oracle DB you should use "dba" group.

Would you like to assign a specific group to Calypso?
[yes]

Please enter the name of the group which will be assigned to all Calypso files and on behalf of which all Calypso processes will run.

In most of the cases it's a good idea to create a dedicated "calypso" group. However, if you're planning to use Oracle iDataAgent or SAP Agent, you should enter Oracle's "dba" group here.

Group name: dba

REMINDER

If you are planning to install Calypso Informix, DB2, PostgreSQL, Sybase or Lotus Notes iDataAgent, please make sure to include Informix, DB2, etc. users into group "dba".

Press <ENTER> to continue ...

Every instance of Calypso should use a unique set of network ports to avoid interfering with other instances running on the same machine. The port numbers selected must be from the reserved port number range and have not been registered by another application on this machine.

Please enter the port numbers.

Port Number for CVD : [8600]

Port Number for EvMgrC: [8602]

Is there a firewall between this client and the CommServe?
[no]

Please specify hostname of the CommServe below. Make sure the hostname is fully qualified, resolvable by the name services configured on this machine.

CommServe Host Name:

Enter your CommCell user name and password:

User Name :

Password :

Press <ENTER> to continue ...

the CommCell Properties. Users must belong to a User Group with Agent Management capabilities to enable this feature. For more information, see Authentication for Agent Installs.

Click **Enter** to continue.

KERNEL PARAMETERS

22. Enter the appropriate number of streams, and then press **Enter** to continue, or Press **Enter** to accept the default number of streams and continue.

NOTES

- The number of streams specified ensures that concurrent backup/restore streams would have sufficient system resources. For more information on the subject, see Configuring Kernel Parameters for Macintosh and Configuring Kernel Parameters for Solaris.

This prompt is relevant only when you install/upgrade on a Macintosh or Solaris computer as appropriate.

23. Indicate whether you would like modifications to be made to the `/etc/system` configuration file.

Type **Yes**, and then press **Enter** to automatically update the file and continue, or Press **Enter** to accept the default **No** and continue (if you do not want to automatically update the file).

This prompt is displayed only when you install/upgrade on a Solaris (8 or 9) or Macintosh computer.

Please enter the total number of streams that you plan to run at the same time. We need to make sure that you have enough semaphores and shared memory segments configured in `/etc/system`.

Number of streams: [10]

We now need to modify the `/etc/system` configuration file on this computer. It is done to make sure that there will be enough shared memory and semaphores available for Calypso programs.

Please review the changes below and answer "yes" if you want us to apply them to the `/etc/system` file. Otherwise, the installation will proceed, the changes will be saved to some other file, and you will have to apply them manually.

```
set shmsys:shminfo_shmmni=8570 (was 7930)
set shmsys:shminfo_shmseg=8420 (was 7780)
set semsys:seminfo_semms=10320 (was 9680)
set semsys:seminfo_shmmni=8570 (was 7930)
set semsys:seminfo_semmsl=8570 (was 7930)
```

Do you want us to apply these changes now? [no]

Changes saved into `/etc/system.gal.1744`

Press <ENTER> to continue.

Although a 'no' answer can be selected to this question during install, the user should make sure the min requirements (below) for shared memory are met, otherwise the backups may fail (the message in logs is 'could not start the pipeline').

```
set shmsys:shminfo_shmmax=4199304
set shmsys:shminfo_shmmni=1
set semsys:shminfo_shmmni=640
set semsys:shminfo_shmseg=640
set semsys:seminfo_semms=640
set semsys:seminfo_shmmni=640
set semsys:seminfo_semmsl=640
set maxusers=256
```

24. If you indicated **No** in the previous step, the file to which the changes have been saved is displayed. Make sure that these values are established later to ensure that all the requirements for this setup is satisfied.

NOTES

- The settings that are displayed are the maximum or minimum required settings. Value '640', which is provided for various shared memory segment or semaphore requirements, is a maximum value based on 10 streams.

Press **Enter** to continue.

This prompt is displayed only when you install/upgrade on a Solaris (8 or 9) computer, in cases where the install detects that the computer does not have the maximum or minimum required shared memory settings.

ENABLE GLOBAL FILTERS

25. Type the appropriate number for configuring Global Filters for the default subclient and press Enter to continue.

NOTES

- Select **Use Cell level Policy** to inherit the global filter policy configuration set for the CommCell, i.e., if the **Use Global Filters on All Subclients** option is selected in the **Global Filters** dialog box (from the CommCell Console's Control Panel), then this policy will be applied to the default subclient as well. If is not selected, then the global filters will not be applied to the default subclient.
- Select **Always use Global filters** to always apply the global filters policy to the default subclient regardless of the policy set for the CommCell.
- Select **Do not use Global filters** to disregard applying the global filters to the default subclient regardless of the policy set for the CommCell.

Commcell Level Global Filters are set through Calypso GUI's Control Panel in order to filter out certain directories or files from backup Commcell-widely. If you turn on the Global filters, they will be effective to the default subclient. There are three options you can choose to set the filters.

- 1) Use Cell level policy
- 2) Always use Global filters
- 3) Do not use Global filters

Please select how to set the Global Filters for the default subclient? [1]

CLIENT GROUP SELECTION

26. Type the number of a Client Group and press **Enter**.

A confirmation screen will mark your choice with an "X". Type **d** for done with the selection, and press **Enter** to continue.

NOTES

Client Group(s) is currently configured on CommServe cs.company.com. Please choose the group(s) that you want to add this client client.company.com to. The selected group(s) will be marked (X) and can be deselected if you enter the same number again. After you are finished with the selection, select "Done with the Selection".

- This screen will be displayed only if Client Groups are configured for the CommCell. For more information, see Client Computer Groups.

```
[ ] 1) Unix
[ ] 2) DR
[ ] 3) DKS

[a=all n=none r=reverse q=quit d=done >=next <=previous ?
=help]

Enter number(s)/one of "a,n,r,q,d,>,<,>?" here: 2

+-----+
IMPORTANT:

In addition to installing Calypso on this computer, you
will also need to create a Job Schedule for each
iDataAgent that has been installed on this client
computer.

Job Schedules are required to allow the Calypso
iDataAgents to perform automated backup and restore
operations.

Job Schedule creation can be accomplished through the
Calypso CommCell Console administration interface.

+-----+
```

27. Press **Enter** to continue.

NOTES

- Schedules help ensure that the data protection operations for the Agent are automatically performed on a regular basis without user intervention. For more information, see Scheduling.

STORAGE POLICY SELECTION

28. Enter the number corresponding to the storage policy through which you want to back up the File System iDataAgent and then press **Enter** to continue.

NOTES

- A storage policy directs backup data to a media library. Each library has a default storage policy.
- When you install an Agent, the install program creates a default subclient for most Agents.
- If desired, you can change your storage policy selection at any time after you have installed the client software.
- If this screen appears more than once, it is because you have selected multiple agents for installation and are configuring storage policy association for each of the installed agents.

```
Please select one storage policy for this IDA from the
list below:

1) SP_StandAloneLibrary2_2
2) SP_Library3_3
3) SP_MagLibrary4_4
4) fornax_fornax_HWCmp
5) ranger_ranger_HWCmp
6) fornax_fornax_CIntCmp
7) fornax_fornax_MACmp
8) fornax_fornax_NoCmp

Storage Policy: [3]
```

29. If you have multiple instances of the database, ensure that the items listed are satisfied. Press **Enter**.

```
IMPORTANT

If you are planning to use Calypso <package> with several
instances of the database, each running on behalf of
different UNIX users, you must make sure that:

1) All database users have been added to each of the
database groups involved. For example, if there is a
"db2inst1" user from group "dba1" and "db2inst2" user from
group "dba2", then both "db2inst1" and "db2inst2" should
be included into "dba1" and "dba2" groups. On most UNIXes
this can be accomplished by editing the /etc/group file.

2) The .profile of each database user must contains a
"umask 002" command at the end, which will guarantee that
all files created by one database user can be open for
writing by other database users.

3) All database users should be added to the "dba" group
that you have chosen for Calypso at the time of CVGxBase
install.

If all of the above requirements are not fulfilled,
Calypso will run into various "Permission denied" errors,
and backups will not work.

Press <ENTER> to continue ...
```

30. The install program now starts copying the software to the computer. The progress of the operation is displayed.

```
.....
.....
.....
.....

Successfully copied xx files
```

31. Press **Enter** to continue.

```
.....
.....

Successfully installed <package_name>.

Press ENTER to continue ...
```

32. This prompt is displayed only when you are installing on HP-UX, Linux, or Solaris computers. Enter the number corresponding to the **Exit** option and then press **Enter** to continue.
The installation is now complete.

```
Certain Calypso packages can be associated with a virtual
IP, or in other words, installed on a "virtual machine"
belonging to some cluster. At any given time the virtual
machine's services and IP address are active on only one
of the cluster's servers. The virtual machine can "fail-
over" from one server to another, which includes stopping
services and deactivating IP address on the first server
and activating the IP address/services on the other
server.
```

```
Currently you have Calypso installed on physical node  
stone.company.com.
```

```
Now you have a choice of either adding another package to  
the existing installation or configure Calypso on a  
virtual machine for use in a cluster.
```

- 1) Add another package to stone.company.com
- 2) Install Calypso on a virtual machine
- 3) Exit

```
Your choice: [1]
```

POST-INSTALL CONSIDERATIONS

GENERAL

- Review Install Considerations after installing the software.
- Install post-release updates or Service Packs that may have been released after the release of the software. When you are installing a Service Pack, ensure that it is the same version as the one installed in the CommServe Server. Alternatively, you can enable Automatic Updates for quick and easy installation of updates in the CommCell component.

Configuration - PostgreSQL iDataAgent

Basic | **Advanced**

TABLE OF CONTENTS

Overview

Creating an Instance

OVERVIEW

Once you install PostgreSQL iDataAgent, you need to create an instance in order to perform your first backup. This automatically creates the following default backupsets, each with a default subclient:


- DumpBasedBackupSet
- FSBasedBackupSet

CREATING AN INSTANCE

This section describes how to create an instance for this iDataAgent in order to perform backup operations.

1. From the CommCell Browser, locate the correct client computer node.

Right-click the agent node for which you want to create a new instance, click **All Tasks**, and then click **New PostgreSQL Server**.



2. Add the following information in the **PostgreSQL Create Instance** dialog:

- Enter the name of PostgreSQL Server in **Server Name**.
- Type the user name to access the Postgres application on a Unix client in the **Postgres User Name** box.

In order to create a PostgreSQL instance, the PostgreSQL user should have Administrator (DBA) and System Administrator (SYSDBA) privileges. Make sure that you have database administrator privileges to access the application.

- Type the Postgres User Account Password and Confirm the Password in **Postgres Account Password** and **Postgres Confirm Password** boxes.
- Type the name of a system database which is used as **Maintenance DB**. A Maintenance database is used for connecting PostgreSQL server. Any database that exists on PostgreSQL can be selected as a Maintenance database. This database should exist in the PostgreSQL server. The Backups and Restores will fail, if the Maintenance database is not existing or dropped.
- Type the path to the Postgres application files in the **Binary Directory** box or click **Browse** to locate the path. For example, `/opt/PostgreSQL/8.4/bin`. Make sure to avoid using special characters such as slash(/), underscore(_), and spaces.
- Type the path to the library directory for the instance in the **Lib Directory** box or click **Browse** to locate the path. For example, `/opt/PostgreSQL/8.4/lib`

Postgres Lib directory path must contain libpq.so file and other related files. This library is used to connect to the database and execute queries. Execute the following command from the bin directory to locate required library path:

```
pg_config --libdir
```

For example:

```
bash-3.2$ cd /opt/PostgreSQL/8.4/bin
```

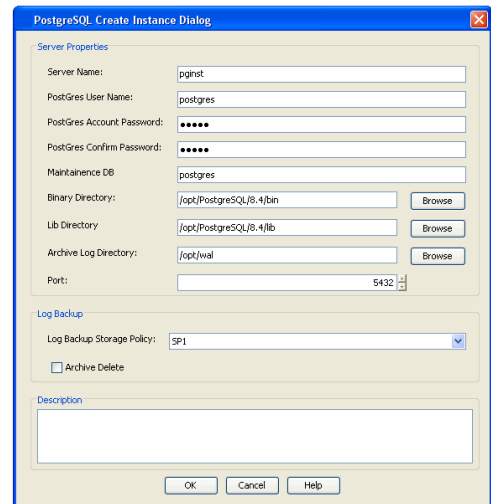
```
bash-3.2$ ./pg_config --libdir
```

```
bash-3.2$ /opt/PostgreSQL/8.4/lib
```

- A separate directory should be created for archive log directory. Type the path to this archive log directory for the instance in the **Archive Log Directory** box or click **Browse** to locate the path. Make sure that this path does not point to the **pg_log** and **pg_xlog** directories. You should have the database administrator privileges (r-w-x) to access this directory.

For example, `/opt/wal`

From PostgreSQL 8.3 version onwards, use the following command to turn on the `archive_mode`. This feature is not supported for PostgreSQL 8.2 and earlier



versions. Also, specify this path in the **postgresql.conf** file prior to performing the PostgreSQL FS backup.

```
archive_mode = on  
archive_command = 'cp %p /opt/wal/%f'
```

From PostgreSQL 9.x.x version onwards, use the following configuration:


```
Set wal_level = archive instead of default wal_level = minimal
```

- Select the **Port** to open the communication between PostgreSQL Server and Clients.
- Select the **Log Backup Storage Policy** to be used for the transaction log backup jobs under Log Backup section.
- Select **Archive Delete**, if you need to delete the archived log backup files.

Refer Troubleshooting for any issues during the creation of an instance.

3. Click **OK** to save the instance.

The backupsets and subclients necessary for your backup operation are automatically created by default, once you create an instance.



[Back to Top](#)

Backup - PostgreSQL iDataAgent

Basic Advanced

TABLE OF CONTENTS

Overview

Your First Full Backup

- Running a Full Backup Immediately
- Scheduling a Full Backup

Viewing Backup Job History

OVERVIEW

Once you have created an instance, you can run your first backup operation on the default subclient. The first backup on the default subclient will be a full backup and will include the entire PostgreSQL server.

The PostgreSQL iDataAgent backs up the following types of data:

For Dump Based Backup:

- PostgreSQL system databases
- PostgreSQL user databases

For File System Based Backup:

- PostgreSQL databases (data and logs)
- Log files

Backup operations do not support database names with special characters.

Backups can be run on demand or can be scheduled to be run at a later point of time. The following section explains in detail the procedure to run your first backup operation.

YOUR FIRST FULL BACKUP


Backups for any client start with a full backup. The full backup becomes a baseline to which subsequent backup types are applied.

A full backup contains all the data that comprises a subclient's contents. If a client computer has multiple agents installed, then the subclients of each agent require a full backup in order to secure all of the data on that client.

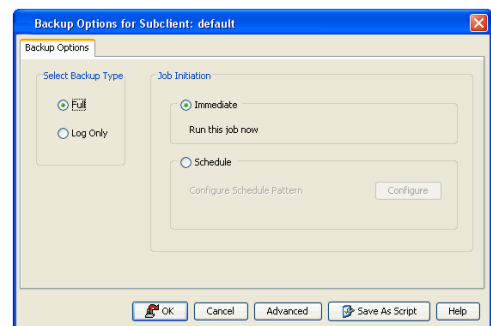
The following sections provide step-by-step instructions for running your first backup, either immediately or at a scheduled time.

RUNNING A FULL BACKUP IMMEDIATELY

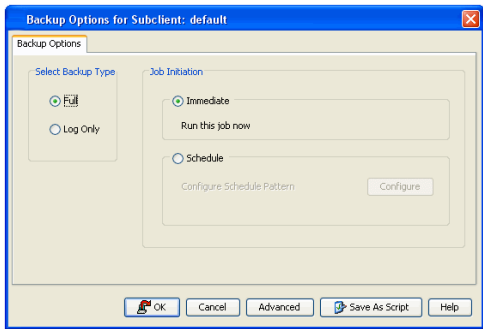
- From the CommCell Browser, right-click the default subclient, and click **Backup**.



- From the **Backup Options** dialog box, select **Full** backup type.



- Select **Immediate** to run the backup operation immediately.




4. Click **OK**. You can track the progress of the backup job from the **Job Controller** window. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically. Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

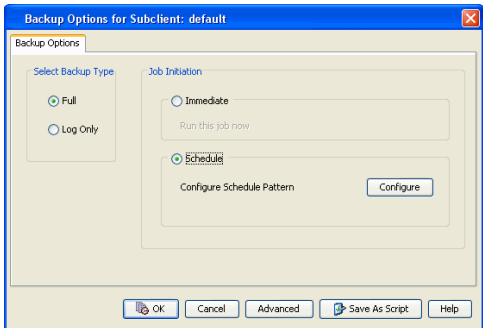
Once completed, the details of the job can be viewed in the **Job History** window.

SCHEDULING A FULL BACKUP

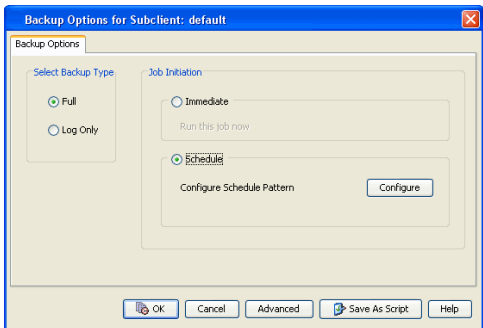
1. From the CommCell Browser, right-click the subclient, and click **Backup**.



2. From the **Backup Options** dialog box, select **Full** backup type



3. Select **Schedule** to schedule the backup for a specific time.



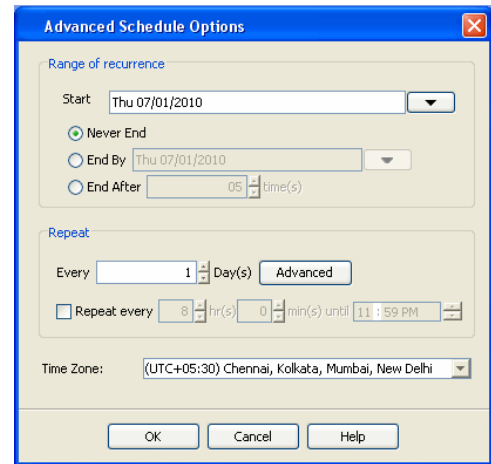

4. Click **Configure** to view the **Schedule details** dialog box. Specify the Schedule Name, Select Time, On these Days and Recurrence options. Click **OK** once all the desired options have been configured.

You can also confirm and/or edit (where applicable) your choices from the **Job Summary** tab. For a monthly or yearly schedule, you can select either a standard calendar or a custom calendar.

5. Click **Options** to view the **Advanced Schedule details** dialog box.

Specify the following options:

- **Range of recurrence:** Select the range of recurrence for the backup job that is scheduled to run.
- **Repeat:** Select the value for which you want to run the job repeatedly on the day in which the job is scheduled to run.
- **Time Zone:** Select a specific time zone from which the job schedule time will be based.



6. Click **OK**.

Your backup operation will execute according to the specified schedule. Once initiated, you can track the progress of the backup job from the **Job Controller** window.

Once completed, the details of the job can be viewed in the **Job History** window.

VIEWING BACKUP JOB HISTORY

Once a backup job has completed, it may be useful to view specific historical information about the job, such as the following:

- Items that failed during the job
- Items that succeeded during the job
- Details of the job
- Media associated with the job
- Events of the job
- Log files of the job.

This information can be viewed in the Job History window, which is accessible at the client computer, iDataAgent, instance, or subclient level in the CommCell Console.

The history provided depends on the entity from which the job history is viewed. For example, viewing job history from the subclient level will yield information for the jobs performed for that subclient. Conversely, viewing job history from the instance level will yield information for jobs run from all subclients within that instance.

To view the backup job history:

1. From the CommCell Browser, right-click the entity (client computer, iDataAgent, instance, or subclient) whose backup history you want to view, click **View**, and then click **View Backup History**.
2. From the Backup History filter window select the filter options, if any, that you want to apply, and then click OK. The system displays the Backup Job History window.
3. Once you have chosen your filter options, they are displayed in the **Backup Job History** window.

To view the additional options discussed above, right-click the desired job choose the appropriate option.

4. Click **OK**.

Back to Top

Restore - PostgreSQL iDataAgent

Basic | Advanced

TABLE OF CONTENTS

Overview

Your First Restore

- Basic Restore
- Browse and Restore

Viewing Restore Job History

View Restored Items Associated With a Restore Job

OVERVIEW

As restoring your backup data is very crucial, it is recommended that you perform a simple restore operation immediately after your first full backup to understand the process. The following sections explain the steps for performing a simple restore.

Restores can be run on-demand or can be scheduled to be run at a later point of time. The following sections explain the procedure to perform your first restore operation.

YOUR FIRST RESTORE

The restore operation can be done using one of the following methods:

- Basic Restore, which is useful when you want to restore and recover an entire database with logs or just the logs or control file.
- Browse and Restore, which is useful when you need to locate and restore specific data files or table spaces.

For this first simple restore, let us perform an out-of place restore of your entire PostgreSQL server. If desired, you can perform subsequent test restores as desired. Refer to the Advanced documentation for information on other types of restoring data.

BASIC RESTORE



To perform a basic restore of the entire database, do the following:

- From the CommCell Browser, right-click the instance, select **All Tasks**, and then click **Restore PostgreSQL Server**.
- From the **Restore Options (General)** tab, select the name of the client computer with the database you want to restore in the **Destination Server** from the list.
- If you need to run the restore operation immediately, select **Immediate** from the **Job Initiation** tab.
Go to step 8.

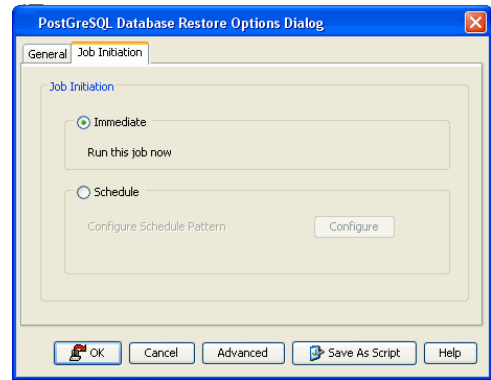
RELATED TOPICS

Capabilities and Required Actions

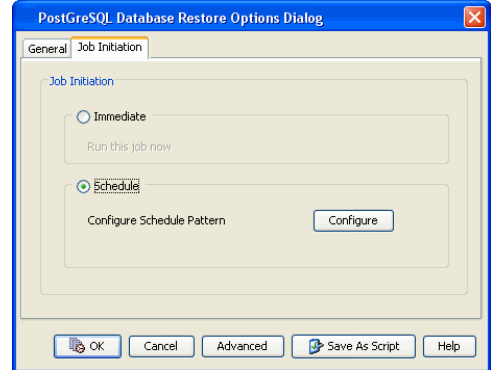
Provides comprehensive information on licensing.



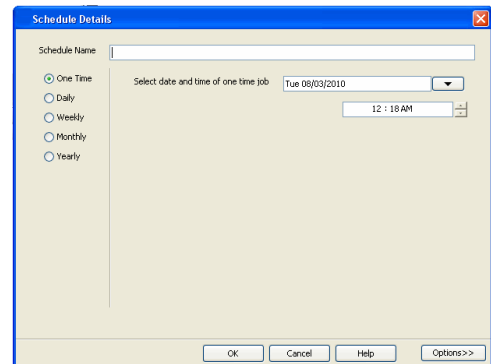
- If you need to schedule the restore operation, select **Schedule** from the Job Initiation tab and click **Configure**.




- From the **Schedule Details** dialog box that appears, select the appropriate scheduling options.
Click **Options** to view the **Advanced Schedule Options** dialog box.



- From the **Advanced Schedule Options** dialog box, select a specific time zone from which the job schedule time will be based.
Click **OK** to close the **Advanced Schedule Options** dialog box.



- Click **OK**.
- Click **OK** to perform the **Restore** operation.
You can track the progress of the restore job from the **Job Controller** window.
You can view the details of the job in the **Job History** window, once the restore is completed.




BROWSE AND RESTORE

To browse and restore the entire database:

- From the CommCell Browser, right-click the level whose data you want to browse, click **All Tasks**, and then click the **Browse Backup Data** option from that level.

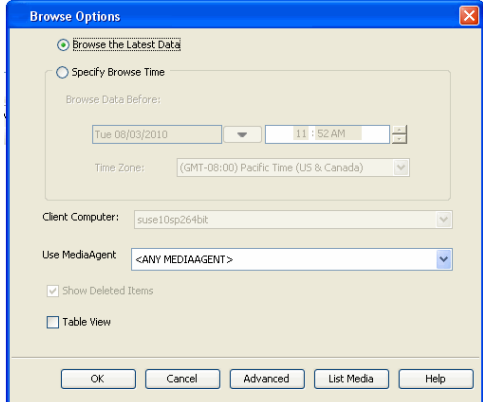
2. From the **Browse Options** dialog box, click **OK** to execute the browse using the **Browse the Latest Data** option.



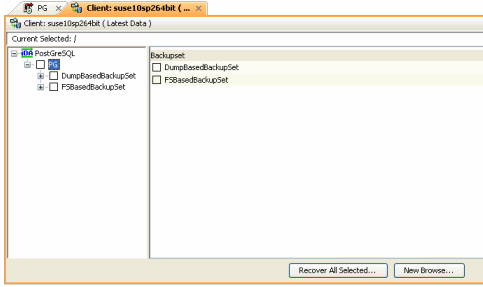
3. The **Browse** window is displayed, containing data objects such as files, folders, databases, etc. and consists of two parts:

- The left pane displays the object tree at the selected level.
- The right pane displays the contents of the selected object.

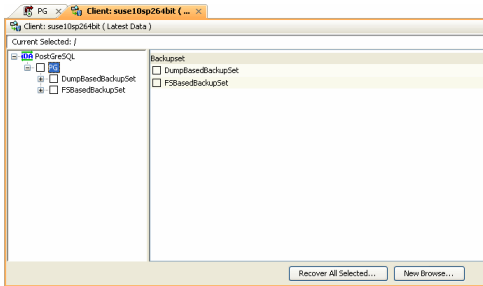
Note that the window displays only the data that was obtained by a data protection operation. Data that is excluded by a filter, or data which did not exist as of the specified browse time does not appear in the window.




4. Under an instance node, either select the DumpBasedBackupSet or the FSBasedBackupSet in the left pane. You can also browse to select any backupset of your choice.
Click **Recover All Selected** at the bottom of the window.



5. For Dump based backupset, select the name of the client server with the database to restore in the **Destination Server** from the list available in the **Restore Options (General)** tab.

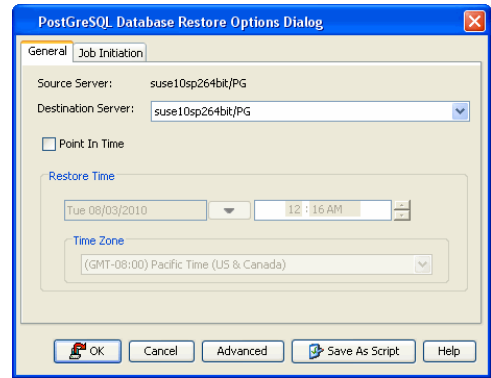


6. For File System based backupset, select the name of the client server with the database to restore in the **Destination Server** from the list available in the **Restore Options (General)** tab.

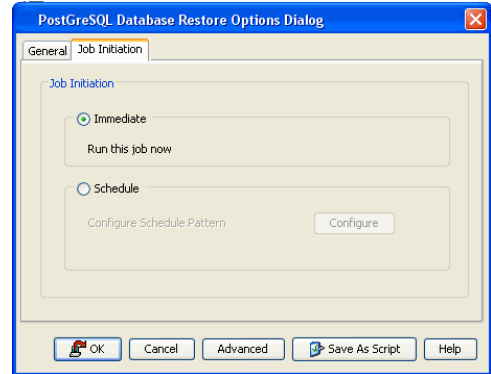


- If you need to run the backup operation immediately, select **Immediate** from the **Job Initiation** tab.

Go to step 12.

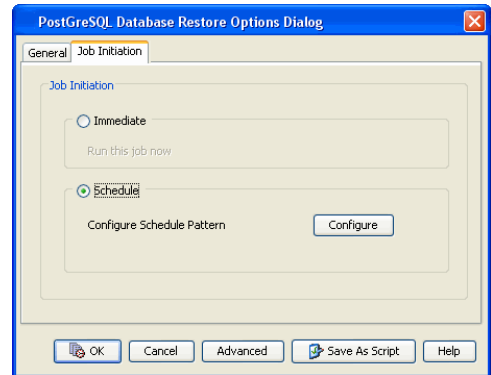


- If you need to schedule the restore operation, select **Schedule** from the Job Initiation tab and click **Configure**.




- From the **Schedule Details** dialog box that appears, select the appropriate scheduling options.

Click **Options** to view the **Advanced Schedule Options** dialog box.




- From the **Advanced Schedule Options** dialog box, select a specific time zone at which the job schedule time will be based.

Click **OK** to close the **Advanced Schedule Options** dialog box.



- Click **OK** to close the **Schedule Details** window.



12. Click **OK** to close the **Restore Options** window.

You can track the progress of the restore job from the **Job Controller** window.

You can view the details of the job in the **Job History** window, once it is completed.

VIEWING RESTORE JOB HISTORY

Once a restore job has completed, it may be useful to view specific historical information about the job, such as the following:

- View Restore Items; items in the job that were **Successful**, **Failed**, **Skipped** or **All**.
- View Job Details
- View Events of the restore job.
- View Log files of the restore job.

This information can be viewed in the Job History window, which is accessible at the client computer, iDataAgent, instance, or subclient level in the CommCell Console.

The history provided depends on the entity from which the job history is viewed. For example, viewing job history from the instance level will yield information for the jobs performed for that instance. Conversely, viewing job history from the agent level will yield information for jobs run from all subclients within that instance.

VIEWING RESTORED ITEMS ASSOCIATED WITH A RESTORE JOB

1. From the CommCell Browser, right-click the entity (e.g., instance, client computer, etc.) whose job restore history you want to view, click **View**, and then click **Restore History**.
2. From the **Job History** filter window, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the **Job History** window, right-click the job whose restored items you want to view; click **View Restore Items**, and select from the type of items to view: **Successful**, **Failed**, **Skipped** or **All**.
4. The **Restored Files** window will display the selected type of restored items for the job.
5. Click **OK**.

[Back to Top](#)