

User Guide - Oracle RAC iDataAgent

TABLE OF CONTENTS

OVERVIEW

Introduction

Key Features

Add-On Components

Terminology

NEW FEATURES

SYSTEM REQUIREMENTS

SUPPORTED FEATURES

GETTING STARTED

DEPLOYMENT ON WINDOWS

DEPLOYMENT ON UNIX

CONFIGURATION

BACKUP

RESTORE

ADVANCED

CONFIGURATION

Understanding the CommCell Console

Managing Instances

Managing Subclients

Managing Control Files

Configuring Table Restores

Configuring Archive Log Destinations

Disabling Log Switch

Enabling Log Deletion after Backup

Changing the Oracle Home Directory

Changing the TNS ADMIN Directory

Configuring Lights Out Script for Offline Backups

Including the Server Parameter (SP) File during Backups

Enabling Backups of Flash Recovery Area

Configuring User Accounts for Backups

Using Recovery Catalog for Backups

Configuring Streams for Backups

Enhancing Backup Performance

Assigning Unique Identification Tags for Backups

Excluding Data During Backups

Simulating a Backup

Disabling the RMAN Crosscheck

Enabling Multiple Backup Copies

Modifying an Agent, Instance, or Subclient

Deleting an Agent, Instance or Subclient

BACKUP

Full Backups

Incremental Backups

Archive Log Backups

Control File Backups

Command Line Backups

Scheduling Backups

Marking Backups with a Unique Identification Tag

Enhancing Backup Performance

Validating Backups For Restore

Managing Jobs

Additional Options

RESTORE

DATA AGING

DISASTER RECOVERY

ADDITIONAL OPERATIONS

BEST PRACTICES

FREQUENTLY ASKED QUESTIONS

BACKUP TROUBLESHOOTING

RESTORE TROUBLESHOOTING

Overview - Oracle RAC iDataAgent

TABLE OF CONTENTS

Introduction

Key Features

- Full Range of Backup Options
- Selective Online Full Backup
- Effective Management of Backups and Restores
- Resource Enhancement on Nodes for accelerated Backups and Restores
- Backup and Recovery Failovers
- Efficient Job Management and Reporting
- Block Level Deduplication

Terminology

INTRODUCTION

The Oracle RAC iDataAgent works in conjunction with the Oracle RAC database to provide simplified end-to-end backup and recovery solutions for Oracle databases in your enterprise without using multiple subclients and storage policies. This Product allows you to load-balance Oracle backups and restores across multiple database nodes. The product can be used to perform both full system rebuilds and granular recovery of the data and logs.

KEY FEATURES

FULL RANGE OF BACKUP AND RECOVERY OPTIONS

The Oracle RAC iDataAgent provides the flexibility to backup the Oracle database in different environments. This is very essential since the Oracle database is always subject to constant changes.

You can perform a full or incremental backup of the entire database or individual datafiles/ tablespaces, or archive logs at any point of time. The following section describes the backups that can be performed in different environments.

OFFLINE BACKUP

When the database is shutdown and not available for use, you can perform a full backup of the database without the logs. This is especially used when the data is consistent and there are no transactions in the database.

ONLINE BACKUP

In cases, when you cannot bring down the database to perform an offline backup, you can use the online backup method. Here, you can perform full or incremental backups when the database is online and in ARCHIVELOG mode. This is very useful when you want to perform a point-in-time restore of the database.

You also have the facility the backup only the archive logs when the database is online. These logs can be applied to an online backup to recover the database to the current point-in-time.

You can also protect the non-database files and profiles using the appropriate File System iDataAgent.

SELECTIVE ONLINE FULL BACKUP

This iDataAgent allows you to backup and store copies of valid data from a source copy of a specific storage policy to all or one active secondary copy within a storage policy providing for a better tape rotation. An online full backup job is copied to a selective copy, if the full backup job cycle completes successfully thereby allowing you to select, store and protect your valuable data on a secondary copy for future restores in a more viable and economic mode.

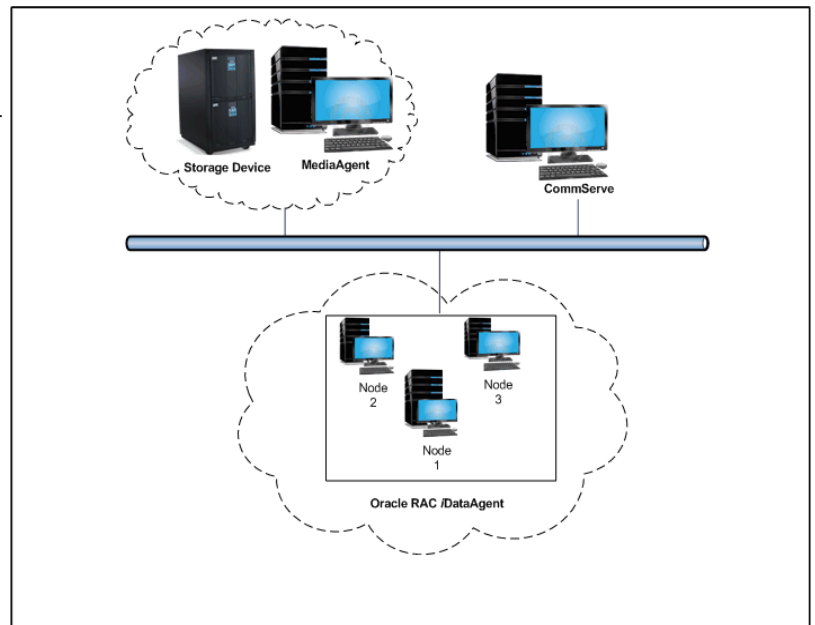
EFFECTIVE MANAGEMENT OF BACKUPS AND RESTORES

This iDataAgent allows you to group any desired number of Oracle iDataAgent instances under one or more Oracle RAC database logical entities. As such, Oracle backups and restores as well as other job types and functions (including Data Aging, Scheduling, Job Management) are all consolidated and easy to manage. This allows you to maintain your data irrespective of whether you add or remove Oracle iDataAgent instances from the RAC database.

WHERE TO GO NEXT

Install the Oracle RAC iDataAgent

Walks you through the process of installing the Oracle RAC iDataAgent.



RESOURCE ENHANCEMENT ON NODES FOR ACCELERATED BACKUPS AND RESTORES

You can configure various resources on your RAC nodes to enable specific node strengths. For example, you can configure a specific RAC node to involve multiple processors and therefore use more streams for a backup or restore.

BACKUP AND RECOVERY FAILOVERS

In the event that a MediaAgent used for the backup or recovery operation fails, it is automatically resumed on alternate MediaAgents. In such cases, the backup or restore job will not restart from the beginning, but will resume from the point of failure. This is especially useful for backups and restores of large amount of file system data.

In the event, that a network goes down, the backup and recovery jobs are resumed on alternate data paths. Similarly, in the event of a device failure, the jobs are automatically switched to alternate disk and tape drives.

Also, this iDataAgent automatically checks the status of each Oracle instance during a backup or restore and allocates RMAN channels only for the instances that are active. Therefore, even if a specific instance fails, the backup or restore will continue.

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 an object (file, database, etc.) contains blocks of data that are identical to each other, then block level deduplication eliminates storing the redundant data and reduces the size of the object in storage. This way dramatically reduces the backup data copies on both the disk and tapes.

TERMINOLOGY

The Oracle documentation uses the following terminology:

RAC	A Real Application Cluster environment which allows multiple nodes in a clustered system to access a single database that resides on a shared storage location.
CLIENT	The computer in which the iDataAgent is installed and contains the data to be secured.
INSTANCE	The Oracle database to be used for the backup and restore operations.
SUBCLIENT	The Oracle data to be backed up.



New Features - Oracle RAC iDataAgent

NEW FEATURES FOR ORACLE RAC iDATAAGENT

COMMAND LINE INTERFACE

Oracle third party command line operations can now be executed from the Oracle Enterprise Manager application.

DATA PROTECTION OPERATIONS

- For Oracle and Oracle RAC iDataAgents, you can now assign tags to a specific backup operation in order to facilitate restore based on the assigned tag. For more information, [Learn more...](#)
- For archive log backups, you now have the facility to disable the log switch of current log files.
[Learn more...](#)
- Multiple copies of data or log backups are now supported from the CommCell Console.
[Learn more...](#)

DATA RECOVERY OPERATIONS

- You can now restore the control file and SP file from a specific backup piece from the CommCell Console. [Learn more...](#)
- Oracle archive log files can be restored based on tags from the CommCell Console. [Learn more...](#)
- Oracle RMAN scripts for backup and restore operations can now be viewed and edited from the CommCell Console. [Learn more...](#)
- When performing table level restore, you now have the facility to select the dependent and referenced tables of the selected tables for the restore operation. You can also choose to include triggers, constraints, indexes, grants, stored procedures, or other external parameters (such as COMPRESSION or PARALLEL) on the selected tables for the restore operation. [Learn more...](#)
- You can resubmit restore jobs from the Job History window without the need to reconfigure the job with the same options. If required, you can also modify the existing configuration and resubmit the job. [Learn more...](#)
- During Oracle restore operations, the restore and recover processes are run as separate RMAN run blocks and hence when resumed, the job is restarted from the last failed RMAN run block. [Learn more...](#)

NEW COMMCELL FEATURES SUPPORTED FOR ORACLE E RAC iDATAAGENT

DEPLOYMENT

INSTALL FROM THE COMMCELL CONSOLE

The software installation for this component can be initiated and managed from the CommCell Console, which facilitates the building of your CommCell and eliminates the need to manually install the software. Additionally, the installation of this component can be scheduled to occur at a time suitable for your environment. For more information, see [Install Software from the CommCell Console](#).

UNINSTALL FROM THE COMMCELL CONSOLE

This component can be uninstalled using the CommCell Console. The Uninstall Software utility allows you to quickly see a list of the software packages installed on the selected computer, from which you can then select the components to uninstall. With this, you can easily manage removing software components from client computers and MediaAgents in your CommCell without having to directly access each computer. For more information, see [Uninstall Components using the CommCell Console](#).

SCHEDULE POLICY

A data protection schedule policy can now be created for this specific agent. For more information, see [Agent-Specific Data Protection Schedule Policy](#).

[Back to Top](#)

System Requirements - Oracle RAC iDataAgent

System Requirements | Supported Features

The following requirements are for the Oracle RAC iDataAgent:

APPLICATION/OPERATING SYSTEM		PROCESSOR
ORACLE 10G/11G (R1, R2 OR HIGHER) RAC DATABASES ON:		
AIX	AIX 7.1	Power PC (Includes IBM System p)
	AIX 6.1	Power PC (Includes IBM System p)
	AIX 5.3	Power PC (Includes IBM System p)
HP-UX	HP-UX 11i v3 (11.31)	PA-RISC
	HP-UX 11i v3 (11.31)	Itanium
	HP-UX 11i v2 (11.23)	Itanium
	HP-UX 11i v2 (11.23)	PA-RISC
LINUX	ORACLE LINUX	
	Oracle Linux 6.x with glibc 2.12.x	Intel Pentium, x64 or compatible processors
	Oracle Linux 5.x with glibc 2.5.x	Intel Pentium, x64 or compatible processors
	Oracle Linux 4.x with a minimum of glibc 2.3.4	Intel Pentium, x64 or compatible processors
	RED HAT ENTERPRISE LINUX/CENTOS	
	Red Hat Enterprise Linux/CentOS 6.x with glibc 2.12.x	Intel Pentium, Itanium, x64, Power PC (Includes IBM System p) or compatible processors
	Red Hat Enterprise Linux/CentOS 5.x with glibc 2.5.x	Intel Pentium, Itanium, x64, Power PC (Includes IBM System p) or compatible processors
	Red Hat Enterprise Linux/CentOS 4.x with a minimum of glibc 2.3.4	Intel Pentium, Itanium, x64, Power PC (Includes IBM System p) or compatible processors
	SUSE LINUX (SLES)	
	SuSE Linux 11.x with glibc 2.9.x and above	Intel Pentium, Itanium, x64, Power PC (Includes IBM System p) or compatible processors
	SuSE Linux 10.x with glibc 2.4.x	Intel Pentium, Itanium, x64, Power PC (Includes IBM System p) or compatible processors
SOLARIS	Solaris 11.x	x64, Sparc5 (or higher recommended)
	Solaris 10.x	x64, Sparc5 (or higher recommended)
WINDOWS	WINDOWS 2008	
	Microsoft Windows Server 2008 Editions with a minimum of Service Pack 1* *Core Editions not supported	All Windows-compatible processors supported
	WINDOWS 2003	
	Microsoft Windows Server 2003 Editions with a minimum of Service Pack 1	All Windows-compatible processors supported

HARD DRIVE

See System Requirements - Oracle iDataAgent

MEMORY

See System Requirements - Oracle iDataAgent

ORACLE EXADATA DATABASE MACHINE SUPPORT

Oracle Exadata Database Machine is supported on the following Oracle Database versions:

- Oracle 11g (11.1.0.7)
- Oracle 11g R2 (11.2.0.1 or higher)

SOLARIS ZONES/CONTAINERS SUPPORT

Data Protection of data residing on global and non-global zones is supported.

For a comprehensive list of supported components, see Unix Virtualization.

AIX LPAR/WPAR SUPPORT

Data protection on Logical Partitioning (LPAR) and Workload Partitioning (WPAR) is supported.

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

NETWORK

TCP/IP Services configured on the computer.

The Oracle iDataAgent must be installed and configured prior to configuring the Oracle RAC iDataAgent. For System Requirements and install information specific to the Oracle iDataAgent, refer to System Requirements - Oracle iDataAgent.

SELINUX

If you have SELinux enabled on the client computer, create the SELinux policy module as a root user before performing a backup. The SELinux Development package must be installed on the client.

To create an SELinux policy module, perform the following steps as user "root":

1. Create the following files in the /usr/share/selinux/devel directory:

File Name	Content of the File
<p><directory>/<file_name>.te</p> <p>where:</p> <p><directory> is /usr/share/selinux/devel</p> <p><file_name> is the name of the Unix file, created to save the policy module statement. It is a good idea to use the same name for policy module and the file.</p> <p>For example: When you are creating a policy module for backup_IDA application, you can use the following file name: backup_IDA.te</p>	<p>The content of the file should be as follows:</p> <pre>policy_module(<name>,<version>) ##### where: <name> is the name of the policy module. You can give any unique name to the policy module, such as a process or application name. <version> is the version of the policy module. It can be any number, such as 1.0.0. For Example: While creating a policy module for the backup_IDA application, you can use the following content. policy_module(backup_IDA,1.0.0)</pre>
<p><directory>/<file_name>.fc</p> <p>where:</p> <p><directory> is /usr/share/selinux/devel</p> <p><file_name> is the name of the Unix file, created to save the policy module statement. It is a good idea to use the same name for policy module and the file.</p> <p>For example: When you are creating a policy module for backup_IDA application, you can use the following file name: backup_IDA.fc</p>	<p>The content of the file should be as follows:</p> <p>Note that the following list of files is not exhaustive. If the process fails to launch, check /var/log/messages. Also, if required, add it to the following list of files.</p> <pre>/opt/<software installation directory>/Base/libCTreeWrapper.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libCVMAGuiImplgso -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2locale.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2osse.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libDb2Sbt.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2trcapi.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0)</pre>

```

/opt/<software installation directory>/Base/libDrDatabase.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libIndexing.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libSnooper.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)

```

2. Create the policy file from command line. Use the following command. Ensure that you give the following commands in the `/usr/share/selinux/devel` directory.

```

[root]# make backup_IDA.pp
Compiling targeted backup_IDA module
/usr/bin/checkmodule: loading policy configuration from tmp/backup_IDA.tmp
/usr/bin/checkmodule: policy configuration loaded
/usr/bin/checkmodule: writing binary representation (version 6) to tmp/backup_IDA.mod
Creating targeted backup_IDA.pp policy package
rm tmp/backup_IDA.mod tmp/backup_IDA.mod.fc
[root]# semodule -i backup_IDA.pp
[root]#

```

3. Execute the policy module. Use the following command:

```

[root]# restorecon -R /opt/<software installation directory>

```

SELinux is now configured to work with this application.

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

Supported Features - Oracle RAC iDataAgent

System Requirements | **Supported Features**

The following table lists the features that are supported by this Agent.

FEATURE	SUB-FEATURE	SUPPORT	COMMENTS
ADVANCED BACKUP/ARCHIVE OPTIONS	Data tab - Catalog		
	Data tab - Create New Index		
	Data tab - Verify Synthetic Full		
	Job Retry tab	✓	
	Media tab - Allow other Schedule to use Media Set	✓	
	Media tab - Mark Media Full on Success	✓	
	Media tab - Reserve Resources Before Scan		
	Media tab - Start New Media	✓	
	Startup tab	✓	
	VaultTracking tab	✓	
	Comments	✓	Includes several additional options in the Backup Archive Logs, Delete Archive Logs, Custom RMAN Script, and Oracle Options tabs.
ADVANCED FILE SYSTEM IDATAAGENT OPTIONS	Automatic File System Multi-Streaming		
	On Demand Data Protection Operation		
	Restore by Jobs		
	Restore Data Using a Map File		
	Comments		
ALERTS AND MONITORING	Global Alerts	✓	
	Job-Based Alerts*	✓	
	Comments		
AUTOMATIC UPDATES	Automatic Updates	✓	
	Comments		
BACKUP/ARCHIVE OPTIONS	Differential Backup		
	Full Backup	✓	
	Incremental Backup	✓	
	Other Backup Types		
	Synthetic Full Backup		
	Comments		
BACKWARD COMPATIBILITY	Version 7	✓	
	Version 8	✓	
	Version 9		
	Comments	✓	Automatically upgraded with the Oracle iDataAgent. See: Upgrade Oracle RAC iDataAgent for more information.
BROWSE	Browse from Copy Precedence	✓	
	Browse the Latest Data	✓	
	Exclude Data Before	✓	
	Find		
	Full Backup Transparent Browse		
	Image Browse		
	No Image Browse		
	Page Size		
	Specify Browse Path		
	Specify Browse Time	✓	
	Subclient Browse	✓	

	Use MediaAgent		
	View All Versions		
	Comments		
CLUSTERING	Netware cluster		
	Unix Cluster		
	Windows - Microsoft Cluster (MSCS)		
	Windows - Non-Microsoft Cluster		
	Comments	✓	Oracle RAC uses its own clustering technology; therefore, it does not depend on operating system-level clusters.
COMMAND LINE INTERFACE	Command Line Interface	✓	
	Comments		
COMMAND LINE INTERFACE - SPECIFIC COMMANDS	Qcreate - Backup set/SubClient		
	Qcreate - Instance		
	Qdelete - Backup set/Subclient		
	Qdelete - Client/Agent	✓	
	Qlist globalfilter		
	Qmodify - instance		
	Qoperation - Backup		
	Qoperation - move		
	Qoperation - Restore		
	Comments		
COMMCELL MIGRATION	CommCell Migration		
	Comments		
CONTENT INDEXING	Offline Content Indexing		
	Comments		
DATA AGING	Basic Retention Rules	✓	
	Extended Retention Rules	✓	
	Unique Data Aging Rules	✓	
	Comments	✓	Extended Retention Rules are supported for offline/selective online fulls.
DATA CLASSIFICATION ENABLER	Data Classification Enabler		
	Comments		
DATA COMPRESSION	Client Compression	✓	
	Hardware Compression	✓	
	MediaAgent Compression	✓	
	Comments		
DATA ENCRYPTION	Data Encryption Support	✓	
	Third-party Command Line Encryption Support	✓	
	Comments		
DATA MULTIPLEXING	Multiplexing	✓	
	Comments		
DEDUPLICATION	MediaAgent Deduplication	✓	
	Source Deduplication	✓	
	Comments		
ERASE BACKUP/ARCHIVED DATA	Erase Data by Browsing		
	Erase Stubs		
	Comments		
GLOBAL FILTERS	Global Filters		
	Comments		
INSTALLATION	Custom Package		
	Decoupled Install		
	Remote Install		
	Restore Only Agents		
	Silent Install	✓	

	Comments		No software installation is required. See Deployment - Oracle RAC iDataAgent for more information.
INSTALLING 32-BIT COMPONENTS ON A MICROSOFT WINDOWS X64 PLATFORM	Install 32-bit On x64		
	Comments		
JOB RESTART - DATA PROTECTION	Not Restartable		
	Restarts from the Beginning		
	Restarts from the Beginning of the Database		
	Restarts from the Point-of-Failure	✓	
	Comments	✓	Third-party command line operations and selective online full backup jobs are Not Restartable . Offline backup jobs Restart from the Beginning .
JOB RESTART - DATA RECOVERY	Not Restartable		
	Restarts from the Beginning		
	Restarts from the Beginning of the Database		
	Restarts from the Point-of-Failure	✓	
	Comments		
LIST MEDIA	List Media Associated with a Specific Backup Set or Instance	✓	
	List Media Associated with Index		
	List Media Associated with Specific Files and/or Folders		
	List Media Associated with Specific Jobs		
	Comments		
MULTI INSTANCING	Multi Instance		
	Comments		
PRE/POST PROCESSES	Pre/Post Process with Data Protection and Recovery		
	Comments		
RESTORE/RECOVER/RETRIEVE DESTINATIONS	Cross-Application Restores (Different Application version)		
	Cross-Platform Restores - Different Operating System		
	Cross-Platform Restores - Same Operating System - Different Version	✓	
	In-place Restore - Same path/ destination - Same Client	✓	
	Out-of-place Restore - Different path/ destination	✓	
	Out-of-place Restore - Same path/ destination - Different Client	✓	
	Restore Data Using a Map File		
	Restore to Network Drive /NFS-Mounted File System		
	Comments		See Advanced Restore - Oracle RAC iDataAgent for this iDataAgent.
RESTORE/RECOVER/RETRIEVE OPTIONS	Automatic Detection of Regular Expressions		
	Filter Data From Recover Operations		
	Rename/ Redirect Files on Restore	✓	
	Restore Data Using Wildcard Expressions		
	Restore Data with Pre/Post Processes	✓	
	Restore from Copies	✓	
	Skip Errors and Continue		
	Use Exact Index		
	Use MediaAgent	✓	
	Comments		
RESTORE/RECOVER/RETRIEVE OVERWRITE OPTIONS	Overwrite Files		
	Overwrite if file on media is newer		

	Restore only if target exists		
	Unconditional Overwrite		
	Unconditionally overwrite only if target is a DataArchiver stub		
	Comments		
SCHEDULE POLICY	Agent Specific Data Protection Schedule Policy	✓	
	All Agent Types Schedule Policy	✓	
	Comments		
STORAGE POLICIES	Incremental Storage Policy*	✓	
	Standard Storage Policies	✓	
	Comments	✓	Incremental Storage Policy does not support Transaction Log backups for this agent.
STORAGE POLICY COPIES	Data Verification	✓	
	Job Based Pruning	✓	
	Manual Retention	✓	
	Mark Job Disabled	✓	
	Selective Copy	✓	
	Comments	✓	Selective Copy is supported for Selective Online and Offline Fulls.
SUBCLIENT POLICIES	SubClient Policy		
	Comments		
UPGRADE	Netware - Local		
	Unix - Remote (Push)		
	Unix/Linux/Macintosh - Local		
	Unix/Linux/Macintosh - Silent		
	Upgrade from CommCell Console		
	Windows - Local		
	Windows - Remote (Push)		
	Windows - Silent		
	Comments	✓	See Upgrade Oracle RAC iDataAgent
USER ADMINISTRATION AND SECURITY	Backup Set/Archive Set		
	Subclient		
	Comments		

Additional features are listed below:

Activity Control	Auxiliary Copy
CommCell Console	Deconfiguring Agents
GridStor	Languages
Log Files	MediaAgent
Operation window	QR Volume Creation Options
Robust Network Layer	Scheduling
SnapProtect Backup	Snapshot Engines
VaultTracker Enterprise	VaultTracker
Report Output Options	Restore/Recover/Retrieve - Other Options
Cloud Storage	Job Restart - Data Collection

Getting Started Deployment on Windows - Oracle RAC iDataAgent



WHERE TO INSTALL

Install the Oracle iDataAgent software on all the nodes of Oracle RAC which satisfy the minimum requirements specified in the System Requirements.

INSTALLATION

1. Log on to the client computer as Administrator or as a member of the Administrator group on that computer.

2. Run **Setup.exe** from the **Software Installation Package**.

If you are installing on Windows Server Core editions, navigate to Software Installation Package through command line, and then run **Setup.exe**.

3. Select the required language.

Click **Next**.

4. Select the option to install software on this computer.

The options that appear on this screen depend on the computer in which the software is being installed.

5. Select **I accept the terms in the license agreement**.

Click **Next**.

6. Expand **Client Modules** | **Backup and Recovery** | **Database** and then click **Oracle iDataAgent**.

Click **Next**.

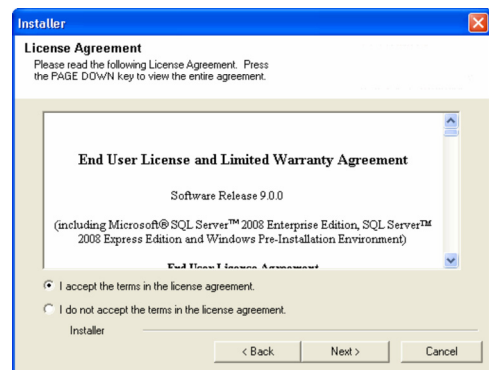


BEFORE YOU BEGIN

Download Software Packages

Download the latest software package to perform the install.

Verify System Requirements

Make sure that the computer in which you wish to install the software satisfies the System Requirements.



- If this computer and the CommServe is separated by a firewall, select the **Configure firewall services** option and then click **Next**.

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

If firewall configuration is not required, click **Next**.

- Enter the fully qualified domain name of the **CommServe Host Name**.

Click **Next**.

Do not use space and the following characters when specifying a new name for the CommServe Host Name:

`\|`~!@#$%^&*()+=<>/?,[\]{}:;'"`





- Click **Next**.

- Select **Add programs to the Windows Firewall Exclusion List**, to add CommCell programs and services to the Windows Firewall Exclusion List.

Click **Next**.

This option enables CommCell operations across Windows firewall by adding CommCell programs and services to Windows firewall exclusion list.

It is recommended to select this option even if Windows firewall is disabled. This will allow the CommCell programs and services to function if the Windows firewall is enabled at a later time.



11. Verify the default location for software installation.

Click **Browse** to change the default location.

Click **Next**.

- Do not install the software to a mapped network drive.
- Do not use the following characters when specifying the destination path:

/ : * ? " < > | #

It is recommended that you use alphanumeric characters only.





12. Select a Client Group from the list.

Click **Next**.

This screen will be displayed if Client Groups are configured in the CommCell Console.

13. Click **Next**.

14. Click **Next**.



15. Select a **Storage Policy** from the drop-down list.
Click **Next**.

If you do not have Storage Policy created, this message will be displayed.

Click **OK**.

You can create the Storage Policy later in step 22.

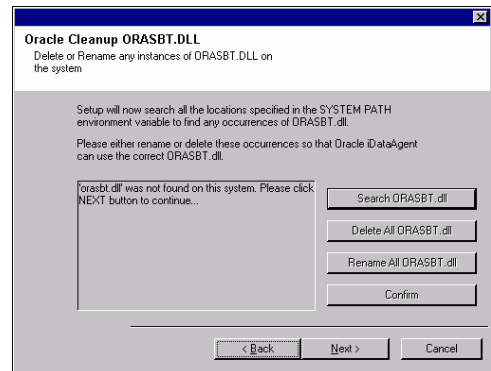
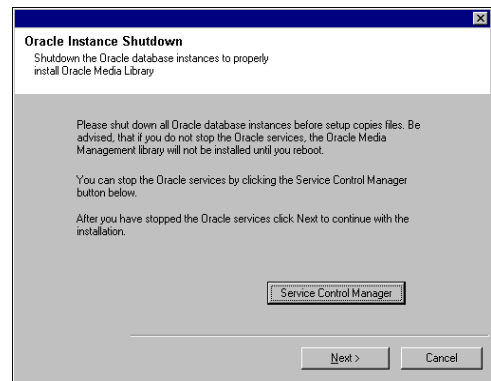



16. Click **Next**.

Stop the Oracle services by clicking the **Service Control Manager** button.

17. Click **Next**.

Click **Search ORASBT.dll** button to search for any existing ORASBT.dll file. If found, rename or delete the ORASBT.dll file.





18. Click **Next**.



19. Click **OK**.

20. Click **Next**.

21. Click **Finish**.



If you already have a storage policy selected in step 15, proceed to the **Configuration** section.


If you do not have Storage Policy created, continue with the following step.

22. Create a Storage Policy:

1. From the CommCell Browser, navigate to **Policies**.
2. Right-click the **Storage Policies** and then click **New Storage Policy**.
3. Follow the prompts displayed in the Storage Policy Wizard. The required options are mentioned below:
 - o Select the Storage Policy type as **Data Protection and Archiving** and click **Next**.
 - o Enter the name in the **Storage Policy Name** box and click **Next**.
 - o From the **Library** list, click the name of a disk library to which the primary copy should be associated and then click **Next**.

Ensure that you select a library attached to a MediaAgent operating in the current release.

- From the **MediaAgent** list, click the name of a MediaAgent that will be used to create the primary copy and then click **Next**.
- For the device streams and the retention criteria information, click **Next** to accept default values.
- Select **Yes** to enable deduplication for the primary copy.
- From the **MediaAgent** list, click the name of the MediaAgent that will be used to store the Deduplication store.
Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to select the folder and then click **Next**.
- Review the details and click **Finish** to create the Storage Policy.



Getting Started Deployment on a UNIX Computer - Oracle RAC iDataAgent

◀ Previous Next ▶

SKIP THIS PAGE IF YOU ARE NOT INSTALLING THIS AGENT ON A UNIX COMPUTER

WHERE TO INSTALL

Install the Oracle iDataAgent software on all the nodes of Oracle RAC which satisfy the minimum requirements specified in the System Requirements.

INSTALLATION

- Logon to the client computer as **root** or as a sudo user.
If you are installing the software using a sudo user account, make sure that sudo user account is configured on this computer. For more information, see FAQ - Install.
- On Unix machine, run the following command to mount the cd:
`mount -t iso9660,udf /dev/cdrom /mnt/cdrom`
Run the following command from the Software Installation Package or mount point:
`./cvpkgadd`
- The product banner and other information is displayed.
Press **Enter**.
- Read the license agreement. Type **y** and press **Enter**.
- Press **Enter**.
- Press **Enter**.
- 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 interface name that you wish to use as default, and then press **Enter**.

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

BEFORE YOU BEGIN

Download Software Packages

Download the latest software package to perform the install.

Verify System Requirements

Make sure that the computer in which you wish to install the software satisfies the System Requirements.

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]

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]

We found one network interface available on your machine. We will associate it with the physical machine being installed, and it will also be used by the CommServe to connect 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 check the interface name below, and make connections if necessary:

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

8. Press **Enter**.
- 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]
9. Type the number associated with the **Oracle iDataAgent** and press **Enter**.
- Install Calypso on physical machine 172.19.99.62
- Please select the Calypso module(s) that you would like to install.
- [] 1) MediaAgent [1301] [CVGxMA]
 [] 2) UNIX File System iDataAgent [1101] [CVGxIDA]
 [] 3) Oracle iDataAgent [1204] [CVGxOrIDA]
- [a=all n=none r=reverse q=quit d=done >=next <=previous ? =help]
- Enter number(s)/one of "a,n,r,q,d,>,<,>?" here:3
- Install Calypso on physical machine 172.19.99.62
- Please select the Calypso module(s) that you would like to install.
- [] 1) MediaAgent [1301] [CVGxMA]
 [] 2) UNIX File System iDataAgent [1101] [CVGxIDA]
 [X] 3) Oracle iDataAgent [1204] [CVGxOrIDA]
- [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
10. A confirmation screen will mark your choice with an "X".
Type **d** for **Done**, and press **Enter**.
- Do you want to use the agents for restore only without consuming licenses? [no]
11. Press **Enter**.
12. Type the appropriate number to install the latest software scripts and press **Enter**.
- Select **Download from the software provider website** to download the latest software scripts. Make sure you have internet access.
 - Select **Use the one in the installation media** to install the software scripts from the package or share from which the installation is currently being 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.
- Installation Scripts Pack provides extra functions and latest support and fix performed during setup time. Please specify how you want to get this pack.
- 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. Press **Enter**.
- Keep Your Install Up to Date - Latest Service Pack
- 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]
14. Press **Enter** to accept the default path.
- If you want to specify a different path, type the path and then press **Enter**.
 - If you want to install the software binaries to an NFS shared drive, specify the directory on which you have mounted the NFS file system and then press **Enter**.
- In order to make sure that the client computer has `read/write` access to NFS shared drive, review the steps described in *Installing Software Binaries to an NFS Shared Drive*.
- Do not use the following characters when specifying the path:
- !@#%&*():/?\
- Please specify where you want us to install Calypso binaries.
- It must be a local directory and there should be at least 176MB 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".
- Installation Directory: [/opt]
15. Press **Enter** to accept the default location.
- Enter a path to modify the default location and press **Enter**.
 - All the modules installed on the computer will store the log files in this directory.
- 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]
16. Type **Yes** and press **Enter**.
- Most of Software 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

If you do not want to assign a group to software, type **no**, press **Enter** and proceed to step 18.

17. Type the **Group name** and then press **Enter**.

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 "oinstall" 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 Software processes, you will need to specify the access permissions later.

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

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

Please enter the name of the group which will be assigned to all Software files and on behalf of which all Software 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 "oinstall" group here.

Group name: oinstall

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 "oinstall".

18. Type **d** for done with the selection and press **Enter**.

This screen is displayed if you do not assign any group to the software and type **no** in step 16.

Access Permissions for Other Users. Installer will assign full access rights to root user and its belonging group for all installed Calypso files and its processes. For any other users, you can specify the access permissions now. However, since you chose not to assign a dedicated group in previous step, make sure you specify sufficient access rights for other users if you are also planning to install Calypso agents involving third party software protection.

[X] 1) Allow read permission to other users
[X] 2) Allow write permission to other users
[X] 3) Allow execute permission to other users

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

19. This prompt is relevant only when you install on Solaris. Press **Enter** to accept the default value for **Number of Streams**.

You can type the **Number of Streams** that you plan to run at the same time and then press **Enter**.

Number of Streams

IMPORTANT : Please read install document "Configure Kernel Parameters - Unix/Macintosh" from "Books Online" before you start configuring kernel parameters. 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]

20. Press **Enter** if you do not want the changes to be updated automatically.

- If you want the changes to be made automatically, type **Yes** and then press **Enter**.
- You will come across this prompt when you install the software on the earlier versions of Solaris.

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_semmns=10320 (was 9680)
set semsys:seminfo_semmni=8570 (was 7930)
set semsys:seminfo_semmnl=8570 (was 7930)
```

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

21. Press **Enter**.

Changes saved into /etc/system.gal.1744

Press <ENTER> to continue.

22. Press **Enter**.

You will see this prompt if you have accepted the default **no** and pressed **Enter** in step 20.

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_shmmmax=4199304
set shmsys:shminfo_shmmmin=1
set semsys:shminfo_shmmni=640
set semsys:shminfo_shmseg=640
set semsys:seminfo_semmns=640
set semsys:seminfo_semmni=640
set semsys:seminfo_semmnl=640
set maxusers=256
```

Press <ENTER> to continue.

Every instance of Calypso should use a unique set of

23. 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**.
24. If you do not wish to configure the firewall services, press **Enter**.

If this computer is separated from the CommServe by firewall(s), type **Yes** and then press **Enter**.

For firewall options and configuration instructions, see Firewall Configuration and continue with the installation.
25. Type the fully qualified CommServe host name and press **Enter**.

Ensure that the CommServe is accessible before typing the name; otherwise the installation will fail.
26. Press **Enter**.
27. Type the number associated with the Client Group and press **Enter**.

This screen will be displayed only if Client Groups are configured for the CommCell.
28. A confirmation screen will mark your choice with an "**X**".
Type **d** for done with the selection, and press **Enter** to continue.
29. Enter the number associated with the storage policy you want use and press **Enter**.

If you do not have Storage Policy created, this message will be displayed.
You may not be prompted for user input.

You can create the Storage Policy later in step 31.
30. Type **3** to the **Exit** option and press **Enter**.
The installation is now complete.
- 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 : [8400]
Port Number for EvMgrC: [8402]

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: mycommservice.company.com

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(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.
[] 1) Unix
[] 2) DR
[a=all n=none r=reverse q=quit d=done >=next <=previous ? =help]
Enter number(s)/one of "a,n,r,q,d,>,<,>?" here: 1

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.
[X] 1) Unix
[] 2) DR
[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

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

1) SP_StandAloneLibrary2_2
2) SP_Library3_3
3) SP_MagLibrary4_4

Storage Policy: [1]

There seem to be no Storage Policies configured on the CommServe. Before you can run any backups of this IDA, you will need to install a MediaAgent, create a Storage Policy and assign it to all subclients..

Adjusting modes and permissions of files
Successfully installed Calypso

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 angel.company.com.



If you already have a storage policy selected in step 29, proceed to the Configuration section.

If you do not have Storage Policy created, follow the procedure given below.

31.

1. From the CommCell Browser, navigate to **Policies**.
2. Right-click the **Storage Policies** and then click **New Storage Policy**.
3. Follow the prompts displayed in the Storage Policy Wizard. The required options are mentioned below:
 - Select the Storage Policy type as **Data Protection and Archiving** and click **Next**.
 - Enter the name in the **Storage Policy Name** box and click **Next**.
 - From the **Library** list, click the name of a disk library to which the primary copy should be associated and then click **Next**.
Ensure that you select a library attached to a MediaAgent operating in the current release.
 - From the **MediaAgent** list, click the name of a MediaAgent that will be used to create the primary copy and then click **Next**.
 - For the device streams and the retention criteria information, click **Next** to accept default values.
 - Select **Yes** to enable deduplication for the primary copy.
 - From the **MediaAgent** list, click the name of the MediaAgent that will be used to store the Deduplication store.

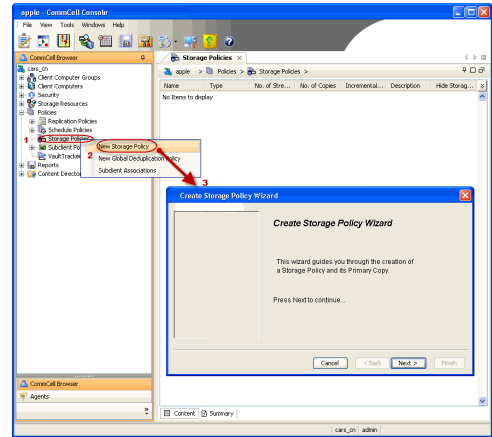
Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to select the folder and then click **Next**.

 - Review the details and click **Finish** to create the Storage Policy.

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 angel.company.com
- 2) Install Calypso on a virtual machine
- 3) Exit

Your choice: [1] 3




Getting Started Configuration - Oracle RAC iDataAgent

◀ Previous Next ▶

CREATING A RAC CLIENT

After installing the Oracle iDataAgent on the nodes of the RAC, configure the Oracle RAC client using the following steps:

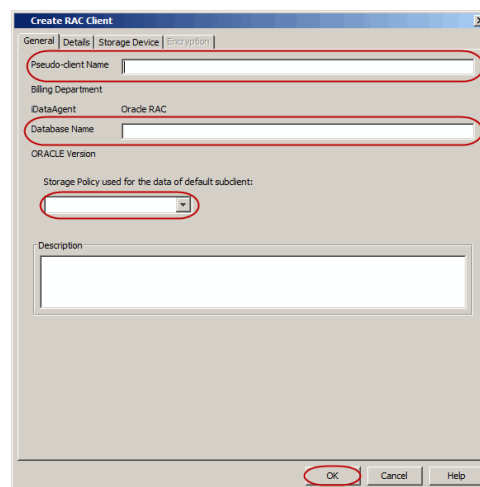
- From the CommCell Console, right-click the **Client Computers** and then click **New Client**.
- Select **RAC Client** from the client list and then click **OK**.



- In the **Pseudo-client Name** box, type a name for the RAC client.
 - In the **Database Name** box, type the name of the Oracle RAC database to which you are assigning this client.
 - From the **Storage Policy** list, select the storage policy used for the data of the default subclient.
 - Click **OK**.

To create a RAC client for Windows, go to step 6.

To create a RAC client for Unix, go to the next step.

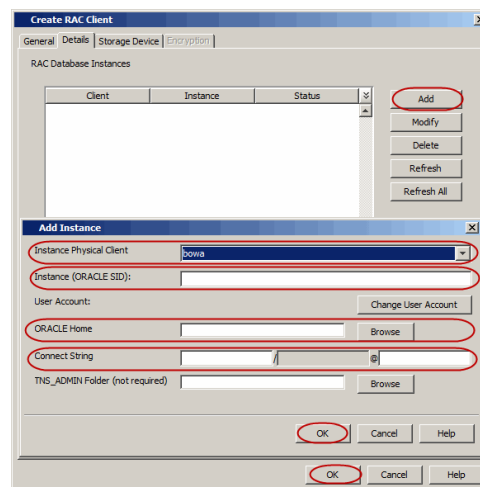


- On Unix:
 - Click the **Details** tab.
 - Click **Add** to add an instance for RAC client.
 - From the **Instance Physical Client** list, select an Unix instance client containing the Oracle iDataAgent instances that you want to include in the Oracle RAC node.
 - In the **Instance (ORACLE SID)** box, type the name of the instance that you want to include in the Oracle RAC node.
 - In the **ORACLE HOME** box, type or click **Browse** to locate the Oracle application install path.
 - In the **Connect String** box, type the credentials to access the Oracle RAC database. Use SYS login i.e., sys/password@Service for Oracle RAC client. For example, sys/pwd12@orc14.
 - In the **TNS_ADMIN Folder**, type or click **Browse** to specify the path for TNS_Admin directory. Not that the TNS_admin path is automatically appended to the path of the \$ORACLE_HOME directory even if you do not specify the path.
 - Click **OK**.

Repeat this step for each Oracle iDataAgent instance that you want to include.

Once you have added the instances, go to step 6.

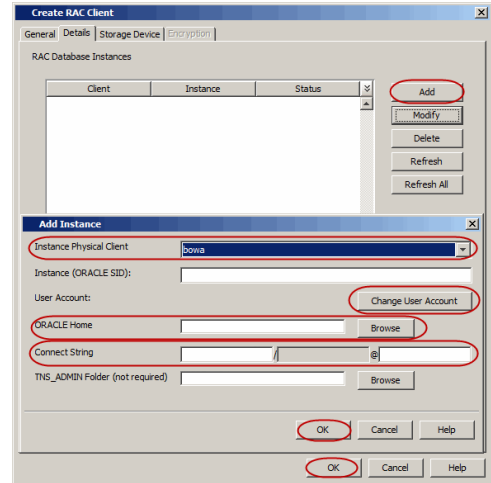
- On Windows:
 - Click the **Details** tab.
 - Click **Add** to add an instance for RAC client.
 - From the **Instance Physical Client** list, select an Windows instance client containing the Oracle iDataAgent instances that you want to include in the Oracle RAC node.
 - Click **Change User Account** and type the Oracle user name and password credentials.
 - In the **ORACLE HOME** box, type or click Browse to locate the Oracle application



install path.

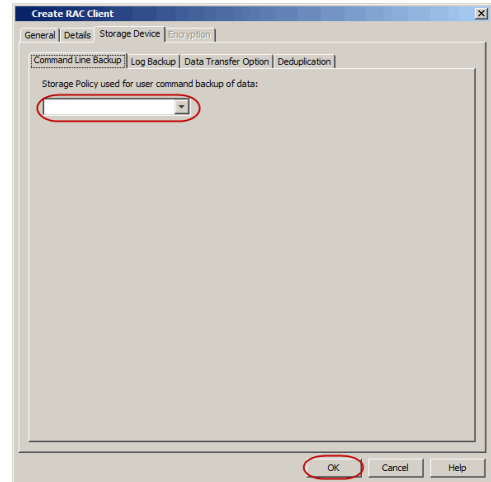
- In the **Connect String** box, type the credentials to access the Oracle RAC database. It is recommended to use SYS login i.e., `sys/password@Service` for Oracle RAC client. For example, `sys/pwd12@orc14`.
- In the **TNS_ADMIN** Folder, type or click Browse to specify the path for TNS_Admin directory. Not that the TNS_admin path is automatically appended to the path of the \$ORACLE_HOME directory even if you do not specify the path.
- Click **OK**.

Repeat this step for each Oracle iDataAgent instance that you want to include.



6.
 - Click the **Storage Device** tab.
 - In the **Storage Policy used for user command backup of data** box, select a storage policy name.
 - Click the **Logs Backup** tab.
 - In the **Storage Policy used for all Archive Log backups** box, select a storage policy name.
 - Click **OK**.

If you do not have a storage policy created, go to step 9 to create a storage policy.



7.
 - Click **Encryption** tab, if the encryption is enabled.
 - Type the Passphrase value.
 - Click **OK**.
8. After you configure all your instances as appropriate, click **OK**.

9. Create a Storage Policy:


1. Click **Create Storage Policy**.
2. Follow the prompts displayed in the Storage Policy Wizard. The required options are mentioned below:
 - Select the Storage Policy type as **Data Protection and Archiving** and click **Next**.
 - Enter the name in the **Storage Policy Name** box and click **Next**.
 - From the **Library** list, click the name of a disk library to which the primary copy should be associated and then click **Next**.

Ensure that you select a library attached to a MediaAgent operating in the current release.

- From the **MediaAgent** list, click the name of a MediaAgent that will be used to create the primary copy and then click **Next**.
- For the device streams and the retention criteria information, click **Next** to accept default values.
- Select **Yes** to enable deduplication for the primary copy.
- From the **MediaAgent** list, click the name of the MediaAgent that will be used to store the Deduplication store.

Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to select the folder and then click **Next**.





- Review the details and click **Finish** to create the Storage Policy.



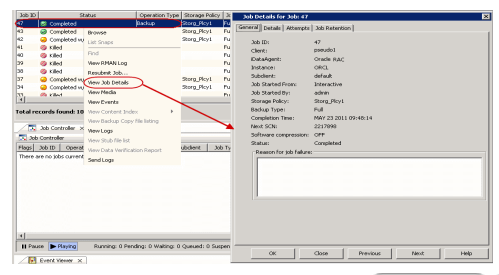
Getting Started Backup - Oracle RAC iDataAgent

WHAT GETS BACKED UP	WHAT DOES NOT GET BACKED UP
Oracle database files that include datafiles (*.dbf) and control files (*.ctl) Archived redo logs Oracle Managed Files (OMF) Parameter files (SP File)	Oracle application files associated with the Oracle installation. Use the File System iDataAgent to back up the above mentioned components.

- From the CommCell Browser, navigate to **Client Computers** | **<RAC Client>** | **<Instance>**.
 - Right-click the default subclient and click **Backup**.
- Click **Full** as backup type and then click **Immediate**.
 - Click **OK**.
- You can track the progress of the job from the **Job Controller** window of the CommCell console.
- Once the job is complete, view the job details from the **Backup History**. Right-click the **Subclient** and select **Backup History**.
- Click **OK**.




- Right-click the job to:
 - Browse the database that was backed up.
 - View RMAN Logs.
 - Resubmit the job.
 - View job details.
 - View media associated with the job.
 - View events associated with the job.
 - View or send the log file that is associated with the job.



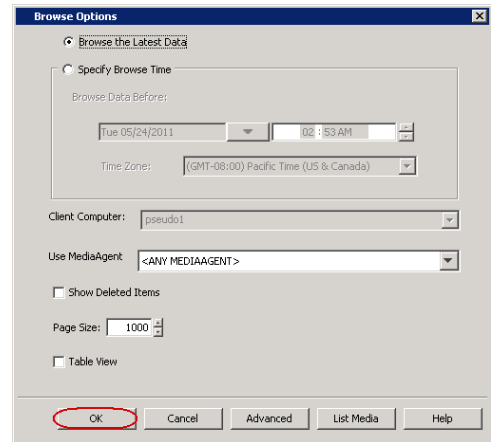
Getting Started Restore - Oracle RAC iDataAgent

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


- From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
 - Right-click the **<Instance>**, point to **All Tasks**, and then click **Browse Backup Data**.




- Click **OK**.



- In the right pane of the Browse window, click the **<Instance>** and select all the entities.
 - Click **Recover All Selected**.

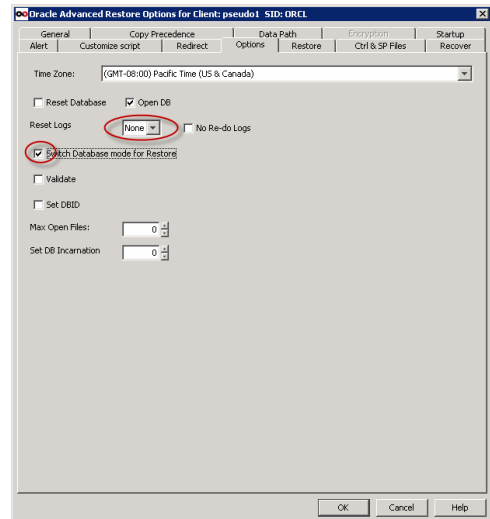


- Click **Advanced**.

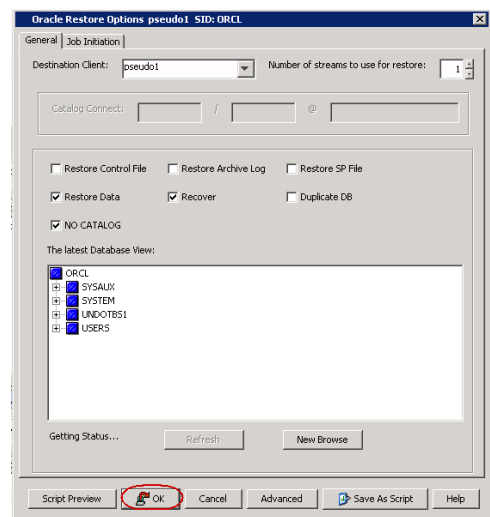


- Click the **Options** tab.
 - In the **Reset Logs** box, select **None**.
 - Select the **Switch Database mode for Restore** checkbox.
 - Click **OK**.

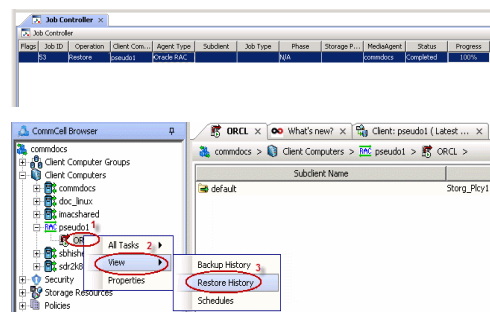
6. Click **OK**.



7. You can monitor the progress of the restore job in the **Job Controller**.



8. Once the restore job has completed, right-click the **<Instance>** and click **View | Restore History**.






9. Click **OK**.

10. You can view the following details about the job by right-clicking the job:

- View Restore Items
You can view them as **Successful, Failed, Skipped** or **All**.
- View Job Details
- View Events of the restore job.
- View Log files of the restore job
- View Paths
- View RMAN Logs

11. Once the database is restored, verify that the restored database and log files are available in the original location.



CONGRATULATIONS - YOU HAVE SUCCESSFULLY COMPLETED YOUR FIRST BACKUP AND RESTORE.

If you want to further explore this Agent's features read the **Advanced** sections of this documentation.

Advanced Configuration - Oracle RAC iDataAgent

TABLE OF CONTENTS

<p>Understanding the CommCell Console</p> <p>Managing Instances Manually Adding an Instance to the RAC Client</p> <p>Configuring RAC One Node</p> <p>Managing Subclients Creating a Subclient for Offline Backups Creating a Subclient for Online Backups Creating a Subclient to Backup Individual Datafiles/Tablespaces Creating a Subclient for Log Backups Creating a Subclient for Selective Online Full Backups</p> <p>Managing Control Files Enabling Automatic Backup of Control Files for All Subclients Enable/Disable Control File Backups for a Specific Subclient Enabling Automatic Backup of the Control Files from RMAN</p> <p>Configuring Table Restores Enabling Table Browse for Restores Setting Up the Auxiliary Instance</p> <p>Configuring Archive Log Destinations Setting Up the Log Destination for Deletion Deleting the Logs after a Backup</p> <p>Disabling Log Switch</p> <p>Enabling Log Deletion after Backup</p> <p>Changing the Oracle Home Directory</p> <p>Changing the TNS ADMIN Directory</p> <p>Configuring Lights Out Script for Offline Backups</p> <p>Including the Server Parameter (SP) File during Backups</p> <p>Enabling Backups of Flash Recovery Area</p> <p>Configuring User Accounts for Backups Configuring User Account to Access the Oracle Application Configuring User Account to Access the Oracle Database</p> <p>Using Recovery Catalog for Backups Synchronizing the Control File with the Recovery Catalog</p> <p>Configuring Streams for Backups Data Backups Log Backups</p> <p>Enhancing Backup Performance</p> <p>Assigning Unique Identification Tags for Backups</p> <p>Excluding Data During Backups</p> <p>Simulating a Backup</p> <p>Disabling the RMAN Crosscheck</p> <p>Enabling Multiple Backup Copies</p> <p>Modifying an Agent, Instance, or Subclient</p> <p>Deleting an Agent, Instance or Subclient Deleting an Agent Deleting an Instance Deleting a Subclient</p>	<p>Command Line Operations Log on to the CommServe Available Parameters Creating a RAC Pseudo Client Adding First RAC Node Adding Second RAC Node Getting RAC Instance Properties Modifying RAC Instance properties Deleting RAC Instance</p>
--	--

UNDERSTANDING THE COMMCELL CONSOLE

The Oracle RAC iDataAgent uses the following logical entities to manage backup and restore operations from the CommCell Console.

AGENT


Facilitates Oracle instance discovery.

INSTANCE

Defines the Oracle database to be backed up.

SUBCLIENT

Defines the Oracle database objects to be backed up.



MANAGING INSTANCES

Oracle RAC instances are created as part of the procedure to create a RAC pseudo-client. Multiple Oracle iDataAgent instances can be configured using a RAC pseudo-client. You can create one Oracle RAC instance per Oracle RAC node. See Creating a RAC Client for more information.

At least one Oracle iDataAgent instance must be configured within each Oracle RAC instance, and you cannot remove an Oracle iDataAgent instance from an Oracle RAC instance unless more than one Oracle iDataAgent instance is configured therein. Removing an Oracle iDataAgent instance from an Oracle RAC instance does not delete the backup data; this data is still available for restore by another eligible client. The Oracle User should have DBA and SYSDBA privileges for creating an Oracle RAC instance.

MANUALLY ADDING AN INSTANCE TO THE RAC CLIENT

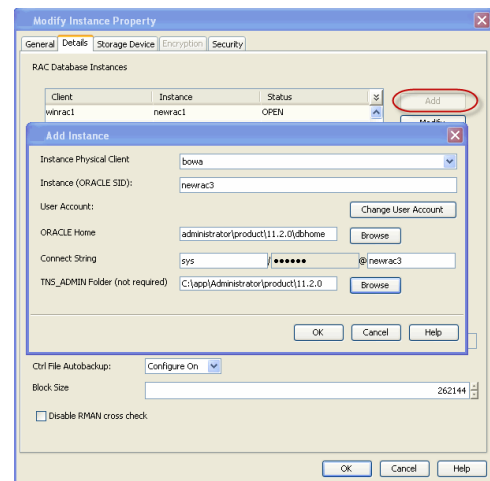
You can add a new instance to the existing RAC client. While adding a new instance make sure to use SYS login as connect string to connect to the database. Make sure that the Connect String value does not include "/" when used with the Use SQL Connect option as RAC will have multiple nodes.

If you want to use a SCAN connect string for an instance, use connect string as: `sys/sys@nodename:1521/service_name` of the database where, service name is the database name.

For Example: `sys/sys@devrac11gr2-1.xyzcompany.com:1521/DEVVRAC11.xyzcompany.com`

Use the following steps to add an instance to the existing RAC client.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<instance>**, and click the **Details** tab, and then click **Add**.
3. In the **Instance Physical Client** list, select the Instance Physical client name.
4. In the **Instance (ORACLE SID)** box, type the Instance name.
5. Click the **Change User Account**.
6. In the **User Account** box, type the user account name.
7. In the **Enter Password** box, type the password for the user account.
8. In the **Confirm Password** box, type re-type the password for the user account.
9. Click **OK**.
10. In the Oracle Home box, type the Oracle installed path. Alternatively, you can click **Browse** to select the location.
11. In the **Connect String** box, type the Connect String (SYS login):
`<sys>/<syspassword>@<Oracle service>`
 Example: `sys/password1@neworac1`
12. In the **TNS_ADMIN Folder** box, type the TNX ADMIN folder name. Alternatively, click **Browse** to select the location.
13. Click the **Storage Device** tab, in the **Storage Policy used for user command backup of data** box, select a storage policy.
14. Click the **Log Backup** tab, in the **Storage Policy used for all Archive Log backups** box, select a storage policy name.
15. Click **OK**.



CONFIGURING RAC ONE NODE

In RAC one node configuration, a single instance of Oracle RAC database runs on one node of the cluster while the second instance is in a cold standby mode. If there is a failure or fault in the first node and the instance cannot be restarted on the same node, then RAC One Node detects it and relocates the instance to the second node. The benefit of this feature is that it automates the instance relocation without any downtime and manual intervention. In this configuration, you must use unique instance names while configuring instances on nodes. See Creating a RAC Client for more information. For example, if you have two RAC nodes and a database that is not relocated to node2, configure instance on node1 and node2. If instance is running on node1, instances on node1 must be instance_1. If a database is relocated to node2, then the instances on node2 will be instance_2.

For example:

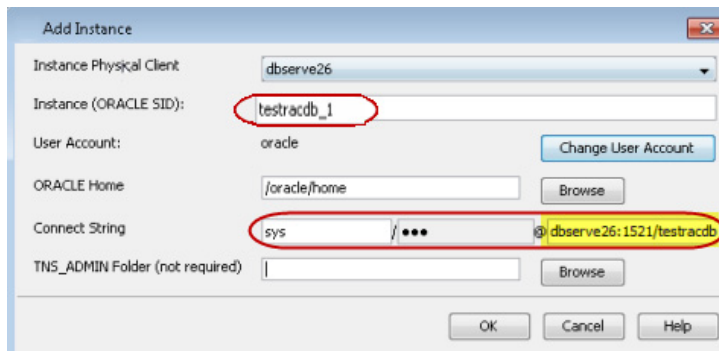
If the RAC database name is testracdb.

Instances are testracdb_1 and testracdb_2.

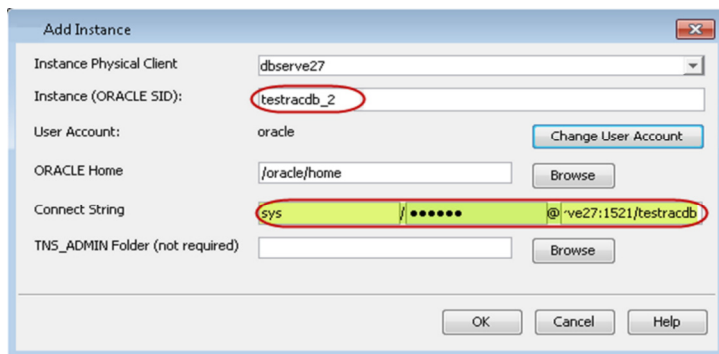
If the node1 is active, the instances will be testracdb_1 on both nodes. If the database is relocated to node2, then the instances will be testracdb_2 on both nodes.

Use the following steps to configure the RAC on one node:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<instance>**, and click the **Details** tab, and then click **Add**.
3. In the **Instance Physical Client** list, select the Instance Physical client name.
4. In the **Instance (ORACLE SID)** box, type the Instance name as testracdb_1 for Instance_1.



Type the Instance name as testracdb_2 for Instance_2



5. You will see the following warning message. Click **OK**.
6. In the **Database Name** box, type the name of the Oracle RAC database (testracd) to which you are assigning this client.
7. In the **Connect String** box, type the Scan Connect String (SYS login):
username/password@physical-nodename:1521/service name of the database where, service name is the database name.

It is important to use the physical node name in the scan connect because the backup process requires pre-identified node for backup resource reservation.

Examples:

Node 1 Connect String: sys/password1@dbserve26:1521/testracdb

Node 2 Connect String: sys/password1@dbserve27:1521/testracdb

You must use local scan connect to configure RAC on one node in a multiple node environment.

8. In the **Oracle Home** box, type the Oracle installed path. Alternatively, you can click **Browse** to select the location.
9. In the **TNS_ADMIN Folder** box, type the TNS ADMIN folder name. Alternatively, click **Browse** to select the location.
10. Click the **Storage Device** tab, in the **Storage Policy used for user command backup of data** box, select a storage policy.
11. Click the **Log Backup** tab, in the **Storage Policy used for all Archive Log backups** box, select a storage policy name.
12. Click **OK**.



See Can you restore from a RAC One Node in a multiple node configuration for important considerations that you need to follow while restoring from a RAC One Node.

MANAGING SUBCLIENTS

The Oracle RAC iDataAgent is designed to back up and restore an Oracle RAC database (*.dbf) along with control files (*.ctl) and archived redo log files. After you create an Oracle RAC pseudo-client, a default subclient is created by the system only after you first configure an Oracle RAC database instance from the CommCell Console. Since an Oracle RAC instance references a RAC database, the default subclient includes the entire database associated with that instance. However, you can use the CommCell Console to create user-defined subclients to include only some of the database content by configuring the Subclient Properties (Content) tab. Any content that is not included in the user-defined subclients will remain in the default subclient.

You can create user-defined subclients to backup the specific components or conditions, such as the following:

- Whether a backup will be offline or online.
- Whether the entire database will be backed up, or only a subset of objects within the database should be backed up.
- Whether the archive logs need to be backed up separately
- Manage your archive log and control file backups.

If a new database object is added to the database, and if none of the subclients contain that database object, then that object is automatically assigned to the content of the default subclient.

CREATING A SUBCLIENT FOR OFFLINE BACKUPS

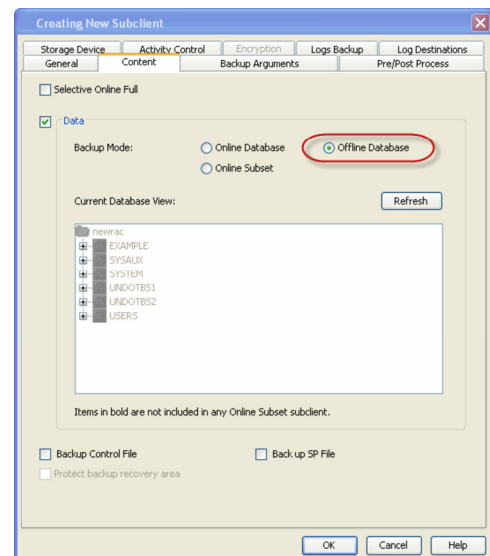
During an offline backup, the database is shutdown and is not available for use. Since incremental backups require access to various tablespaces and datafiles, it is always recommended that you perform a full backup of the database when it is offline. The full backup includes all the datafiles, tablespaces, and control file of the Oracle database. Note that, offline backups do not include the archived log files.

Offline backups can be performed when the database is in offline or online mode. If the database is online, it shuts down the database, performs the backup and then brings up the database back.

In order to backup the Oracle database when it is offline, you need to create a separate user-defined subclient for offline backup. Make sure that the database is in the MOUNT mode during the backup.

Use the following steps to create a subclient for offline backups:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, point to **All Tasks**, and then click **New Subclient**.
3. In the **Subclient name** box, type the subclient name.
4. Click the **Content** tab.
5. Click **Offline Database**.
6. Click the **Storage Device** tab.
7. In the **Data Storage Policy** box, select the storage policy name.
8. Click **OK**.



CREATING A SUBCLIENT FOR ONLINE BACKUPS

In some environments, it may not be possible to bring down the entire database to perform an offline backup. In such situations, you can choose online backups, where the database will be up and running during the backup.

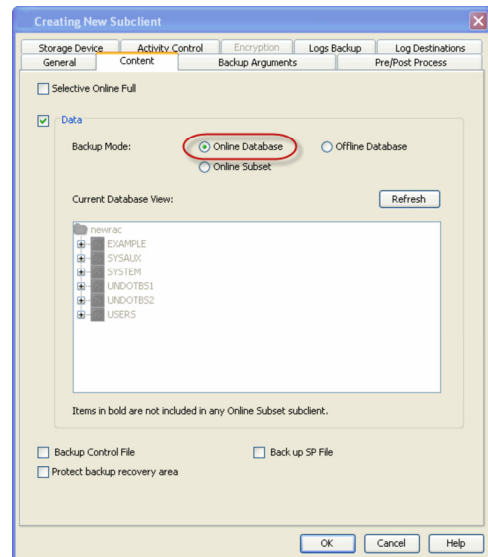
You can perform either a full backup or incremental backup of the Oracle database when it is online. Since full backups includes all the datafiles, tablespaces and logs, it is very time consuming, hence you can plan for a full online backup less frequently (say, on weekly basis). On the other hand, incremental backups includes the data and logs that have been changed after the last full backup, they are faster and can be performed more frequently (say, on a daily basis).

If a new database object is added to the database, and if none of the subclients contain that database object, then that object is assigned to the content of the default subclient

Use the following steps to create a subclient for online backups:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**

- Right-click the **<Instance>**, point to **All Tasks**, and then click **New Subclient**.
- In the **Subclient name** box, type the subclient name.
- Click the **Content** tab.
- Click **Online Database**.
- Click the **Storage Device** tab.
- In the **Data Storage Policy** box, select the Storage policy.
- Click **OK**.

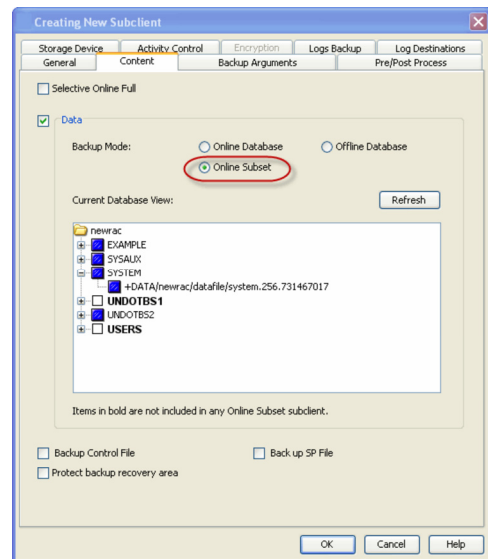


CREATING A SUBCLIENT TO BACKUP INDIVIDUAL DATAFILES/TABLESPACES

You can also create subclients to backup specific datafiles and tablespaces, which undergo frequent changes in the Oracle database.

Use the following steps to create a subclient to backup individual datafiles and tablespaces:

- From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
- Right-click the **<Instance>**, point to **All Tasks**, and then click **New Subclient**.
- In the **Subclient name** box, type the subclient name.
- Click the **Content** tab.
- Click **Online Subset**.
- Select the database objects to be included in the backup.
- Click the **Storage Device** tab.
- In the **Data Storage Policy** box, select the storage policy name.
- Click **OK**.



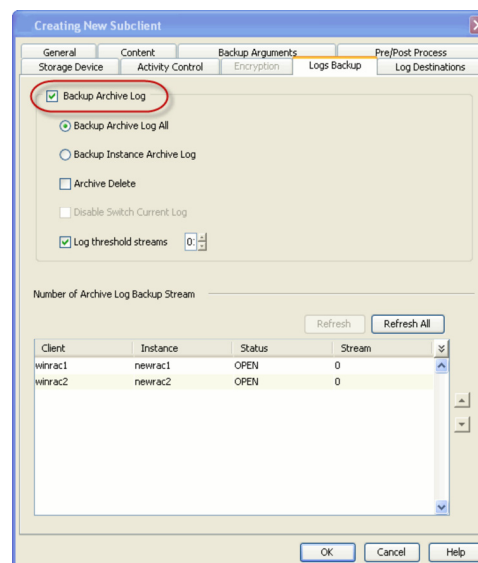
CREATING A SUBCLIENT FOR LOG BACKUPS

Archive logs are required to recover database transactions that have been lost due to an operating system or disk failure. You can apply these archive logs to an online backup in order to recover a database.

Though online full backups can include both data and logs, because of their importance in recovering data, it is recommended that you create separate subclients to backup archive log files.

- From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
- Right-click the **<Instance>**, point to **All Tasks**, and then click **New Subclient**.
- In the **Subclient name** box, type the subclient name.
- Click the **Logs Backup** tab.
- Select the **Backup Archive Log** check box.
- Click the **Storage Device** tab.
- In the **Data Storage Policy** box, select the storage policy name.

- Click **OK**.



CREATING A SUBCLIENT FOR SELECTIVE ONLINE FULL BACKUPS

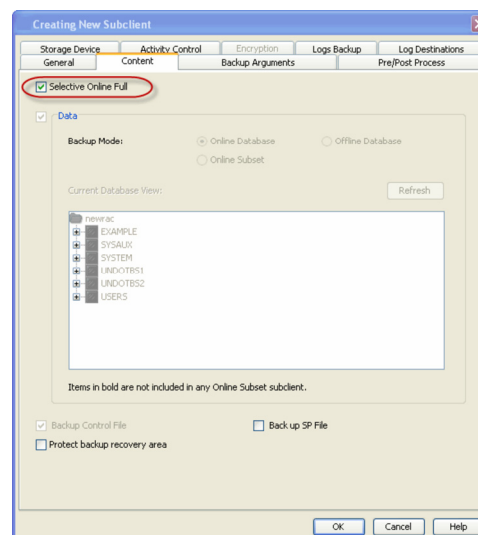
You can selectively copy full backup to a selective copy during an auxiliary copy operation. The advantage of this type of backup is that both data and logs use the same storage policy, which means that they reside together on the same media. They are completely self-contained for restore and long term archiving purposes. Also, the data aging rules for selective online full backups are different from regular full backups, as both data and logs are aged together under the same storage policy.

It is recommended to create a different storage policy during the creation of this subclient.

When performing the selective online full backup, note the following:

- A selective online full backup job will wait for other Oracle backup jobs currently running on the same instance to complete before it begins.
- For this type of backup, both data and archive logs will use the same storage policy as defined for data in the subclient, and will ignore the storage policy setting for archive logs (which is defined at the instance level).
- During selective online full backups, if the data streams (defined at the subclient level) is less than the archive log streams (defined at the instance level), then both the data phase and the archive log phase will use the same number of streams defined for the data in the subclient, and will ignore the number of streams set for the archive logs.
- While the data backup phase of a selective online full backup is running, the only other Oracle backup jobs that are allowed to run on the same instance are archive log backups. During the log backup phase of a selective online full backup, no other Oracle backups are allowed to run (neither logs nor data) on the same instance.
- If other Oracle archive log backup jobs are running at the same time as the selective online full backup, they will be forced to use the same storage policy used by the selective online full backup during the time-frame in which the selective online full backup job is running, and will be included in the same auxiliary copy operation.

- From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
- Right-click the **<Instance>**, point to **All Tasks**, and then click **New Subclient**.
- In the **Subclient name** box, type the subclient name.
- Click the **Content** tab.
- Select the **Selective Online Full** check box. Click **OK** on the warning message dialog box.
- Click the **Storage Device** tab.
- In the **Data Storage Policy** box, select the storage policy name.
- Click **OK**.



MANAGING CONTROL FILES

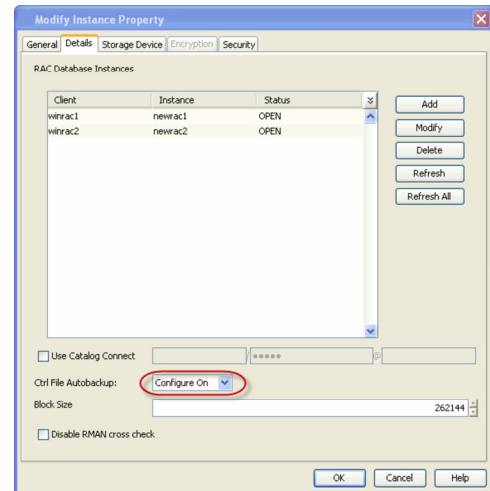
As the Control file stores the status of the physical structure of the database. It is required when you want to recover the database to the current state. Hence it is important to include control files in your backups. By default, control files are included in the subclient for selective online full backups. You can also include/exclude control files for online or offline backups. The following sections describe the methods by which you can backup control files.

ENABLING AUTOMATIC BACKUP OF CONTROL FILES FOR ALL SUBCLIENTS

You can configure the instance to automatically backup Control Files whenever you a backup the subclient in that instance.

Use the following steps to enable automatic backup of control files for all the subclients in the instance.

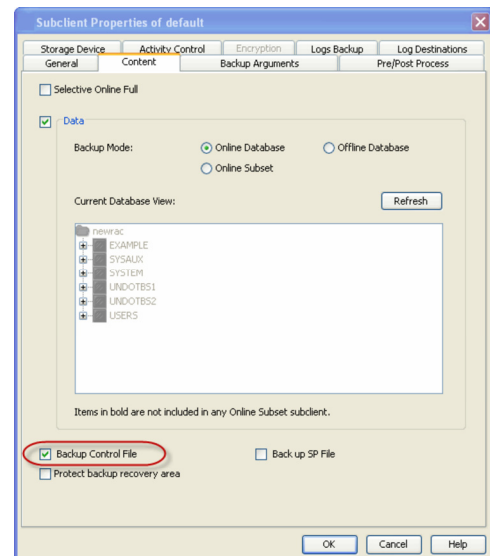
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**.
3. Click the **Details** tab, in the **Ctrl File Autobackup** box,
 - o Select **Configure ON**. This will backup the control file separately.
 - o Alternatively, select **Configure Off** to backup the control file along with the datafile.
4. Click **OK**.



ENABLE/DISABLE CONTROL FILE BACKUPS FOR A SPECIFIC SUBCLIENT

If the control file backup is enabled at the instance level, you can also include/exclude control files for backups from a specific subclient. Use the following steps to enable control file backups for a specific subclient.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Content** tab.
4. Select the **Backup Control File** check box.
5. Click **OK**.



ENABLING AUTOMATIC BACKUP OF THE CONTROL FILES FROM RMAN

You can enable autobackup from RMAN command line to automatically backup the control file in every backup job operation.

Enter the following commands/parameters at the RMAN prompt (the catalog connect string is optional):

```
$rman target <connect string for target> catalog <connect string for catalog>RMAN>configure controlfile autobackup on;
```

For Example: \$rman target systemuser/systempass@rman92 RMAN>configure controlfile autobackup on;

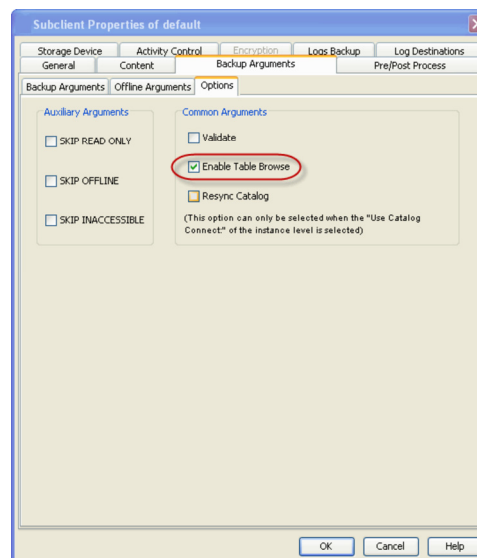
CONFIGURING TABLE RESTORES

By default, all data is backed up which includes tablespaces. If you want to restore individual tables, you need to enable table level backup.

ENABLING TABLE BROWSE FOR RESTORES

In order to backup and restore database tables, you need to enable table level backups for the subclient. Use the following steps to enable backups at table level.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab, and then click the **Options** tab.
4. Select the **Enable Table Browse** check box.
5. Click **OK**.



SETTING UP THE AUXILIARY INSTANCE

By default, when you restore database tables to a target instance, the system automatically duplicates the source database to an auxiliary instance in a temporary staging location specified during the restore operation. The database will be automatically imported from this auxiliary instance after the restore.

Use the following steps to set up a specific database as an auxiliary instance. This is useful when you want to restore a table to a specific failure point.

1. Copy the `init<SID>.ora` file from the source database to the auxiliary database instance.
2. Update the database name and the database file locations in the `init<SID>.ora` file for the auxiliary database instance.
3. Add the `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT` parameters in the `init<SID>.ora` file. These parameters will redirect the datafiles, temp files, and log files to the auxiliary instance.
4. Add the `log_archive_dest_1` parameter is added to the `init<SID>.ora` file on the auxiliary instance.
5. Restart the Oracle Services, if using Windows clients.
6. Add the destination instance name in the `Listener.ora` and `Tnsnames.ora` files. If using a different host, add the duplicate database instance name in the `Listener.ora` file on the destination host and `Tnsnames.ora` files on the destination and source hosts. Also, add the original database name in the `Tnsnames.ora` file on the destination host.
7. Restart the Listener.

Windows Clients:

```
DB_FILE_NAME_CONVERT=
('source_of_df_path/', 'dup_of_df_path/',
'source_of_temp_path/', 'dup_of_temp_path/', ...)
LOG_FILE_NAME_CONVERT=('source_of_log_path/
redo', 'dup_of_log_path/redo')
```

Unix Clients:

```
DB_FILE_NAME_CONVERT=
(source_of_df_path/, dup_of_df_path/
, source_of_temp_path/, dup_of_temp_path/, ...)
LOG_FILE_NAME_CONVERT=
(source_of_log_path/redo, dup_of_log_path/redo)
```

```
DUPDB = (DESCRIPTION =
(ADDRESS = (PROTOCOL = TCP) (HOST = powerpc02) (PORT =
1521))
(CONNECT_DATA = (SERVER = DEDICATED)
(SERVICE_NAME = dupdb) (UR=A) ) )
$!snrctl reload
```

8. Ensure that the auxiliary instance is in NOMOUNT mode.

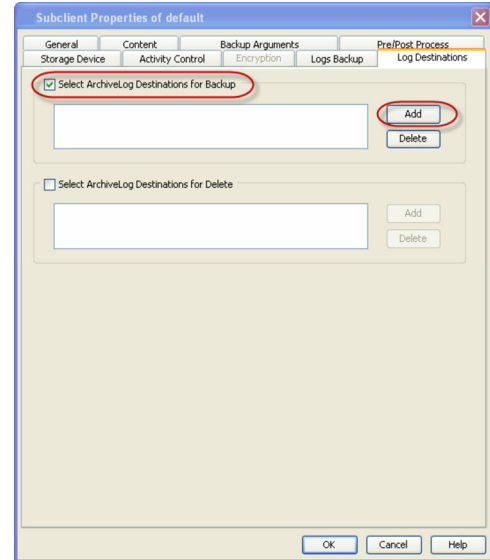
```
sql> startup nomount;
```

CONFIGURING ARCHIVE LOG DESTINATIONS

When you backup archive logs, you can specify the locations from where the log backups should be performed. This capability enables you to schedule backup operations from different log destinations on the same subclient. If necessary, you can also delete the logs after the backup.

SETTING UP THE LOG DESTINATION FOR DELETION

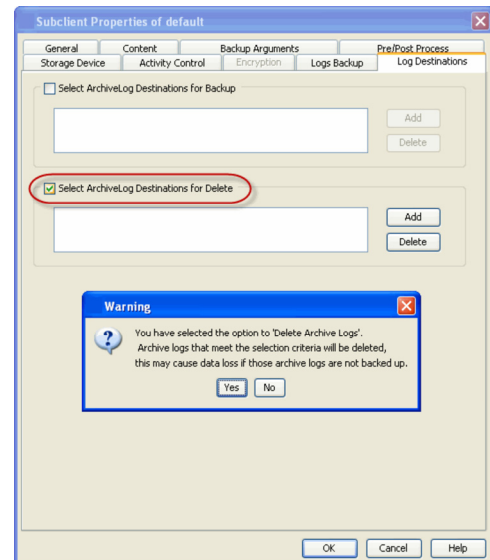
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Log Destinations** tab.
4. Select the **Select ArchiveLog Destinations for Backup** check box.
5. Click **Add**, type or select the Archive log file destination, and then click **OK**.
6. Click **OK**.



DELETING THE LOGS AFTER A BACKUP

Ensure that the logs in the specified location are backed up. This may cause data loss if those archive logs are not backed up.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Log Destinations** tab.
4. Select the **Select ArchiveLog Destinations for Delete** check box.
5. Click **Yes** to close the warning message.
6. In the Confirmation box, type confirm.
7. Click **OK**.
8. Click **Add**.
9. Type or select the Archive log file destination.
10. Click **OK**.
11. Click **OK**.

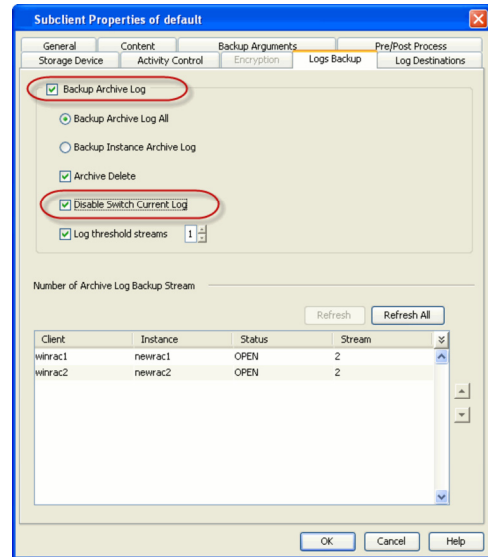


DISABLING LOG SWITCH

When performing archive log backups, the current redo log file is closed (even if it is not filled up completely) and the next redo log file is used for writing using a log switch. The closed redo log file is then archived during the log phase. Use the following steps to disable the log switch:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.

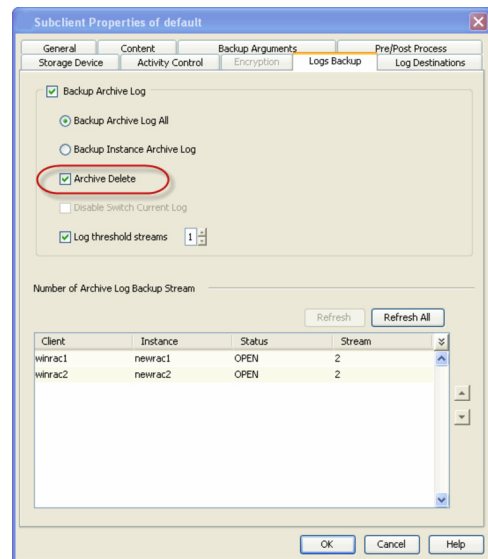
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Content** tab.
4. Clear the **Data** checkbox.
5. Click the **Logs Backup** tab.
6. Select the **Backup Archive Log** check box.
7. Select the **Disable Switch Current Log** check box.
8. Click **OK**.



ENABLING LOG DELETION AFTER BACKUP

When backing up archive logs, by default the logs are not deleted after the backup operation. Use the following steps to enable deletion of logs soon after the backup.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Content** tab.
4. Select the **Archive Delete** check box.
5. Click **OK**.

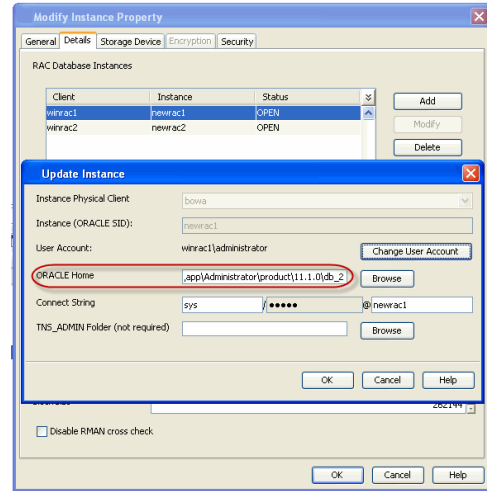


See [Deleting Archive Logs After a Specific Backup](#) for more information.

CHANGING THE ORACLE HOME DIRECTORY

Oracle HOME directory refers to the location where the Oracle application resides. When instances are discovered, the home path is automatically assigned to the instance. However, you can modify the home path location for an instance at any point of time. Use the following steps to configure the Oracle HOME directory:

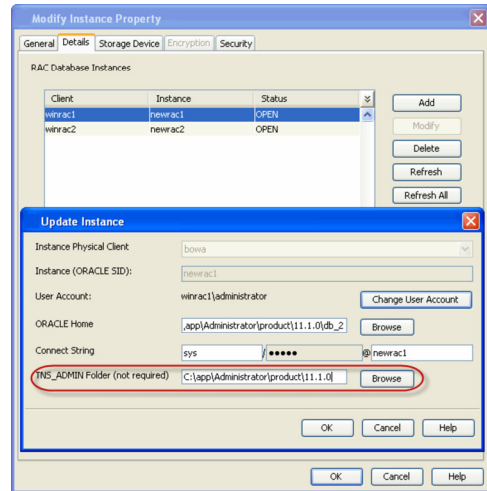
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, and then click **Properties**.
3. Click the **Details** tab, and select the desired RAC database instance, and then click **Modify**.
4. In the **ORACLE HOME** box, type the path to Oracle Home directory. Alternatively, you can use **Browse** to locate the path.
5. Click **OK**.



CHANGING THE TNS ADMIN DIRECTORY

When an instance is configured, by default, the TNS_Admin directory is created in the Oracle HOME\network\admin path. Use the following steps to change the path for the TNS_ADMIN directory:

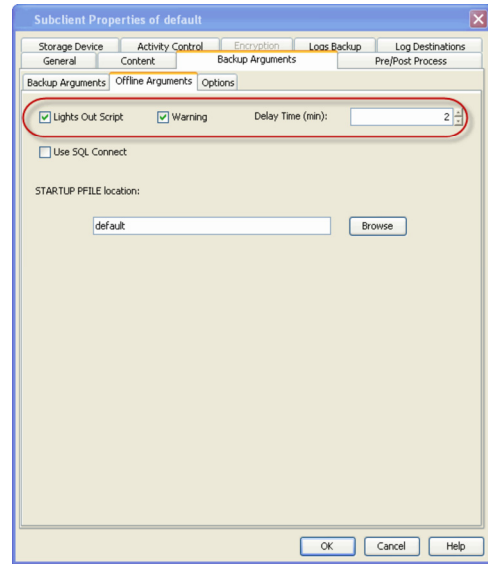
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, and then click **Properties**.
3. Click the **Details** tab.
4. Select the desired database instance, and then click **Modify**.
5. In the **TNS_ADMIN** folder box, type the location for the TNS Admin directory. Alternatively, you can use **Browse** to locate the path.
6. Click **OK**.



CONFIGURING LIGHTS OUT SCRIPT FOR OFFLINE BACKUPS

In order to perform offline backups of the Oracle database, all users must be logged out of the database and it must be completely shut down. When the database is online, you can enable a graceful shutdown of the database using lights out script that will notify the users about the shutdown and set a specific wait time for the users to logout of the database. Use the following steps to configure lights out script for offline backups:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Content** tab, and then click **Offline Database**.
4. Click the **Backup Arguments** tab, and then click the **Offline Arguments** tab.
5. Select the **Lights Out Script** check box.
6. Select the **Warning** checkbox to display a warning message on the physical node.
7. In the **Delay Time** box, select the delay time interval in minutes.
8. Select the **Use SQL Connect** checkbox to connect to the oracle database using the SQL command `connect <connect string> as sysdba`.
9. Click **OK**.

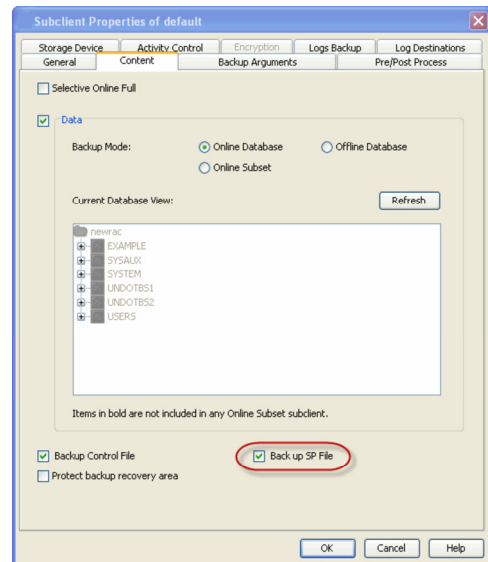


INCLUDING SERVER PARAMETER (SP) FILE DURING BACKUPS

The Server Parameter file (SPFile) contains the database startup information. By default, the backups do not include the SPFile.

Use the following steps to include SPFile during backups:

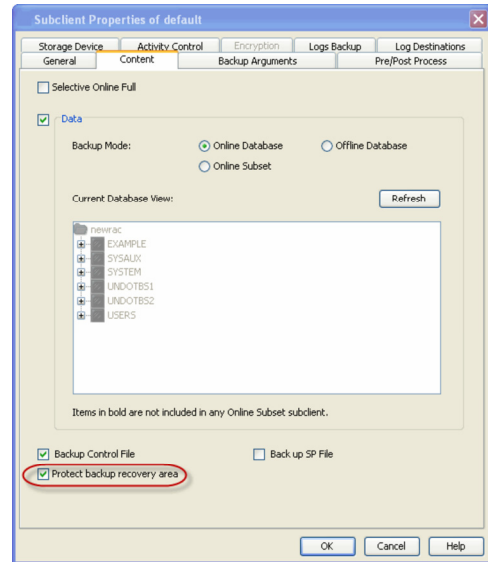
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Content** tab.
4. Select the **Back up SP FILE** check box.
5. Click **OK**.



ENABLING BACKUPS OF FLASH RECOVERY AREA

Flash Recovery Area is a specific area in the disk storage that exclusively holds a copy of all backup-related components, such as image copies, redo logs, and control file auto backups. This facilitates faster restores of backup data and minimizes restores from tapes during restore operations. Use the following steps to enable backup of flash recovery area:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Content** tab.
4. Select the **Protect backup recover area** check box.
5. Click **OK**.



CONFIGURING USER ACCOUNTS FOR BACKUPS

In order to perform backup and restore operations, you need to create and configure the following user accounts on the Oracle client:

- User Account with administrative privileges to access the Oracle application.

On Unix clients, the user should be part of the user group assigned during the iDataAgent install. You can also use the operating system user account to verify the rights to perform all backup and restore operations for the associated Oracle instance.

On Windows clients, the user should be part of local administrator group and also part of the ora_dba group with read/write permissions on Calypso folder. You can use the Impersonate user account to verify the rights to perform all backup and restore operations for the associated Oracle instance.

- User account to access the Oracle database. You can use separate accounts to access the standard database and the Recovery Catalog database. The account information is provided as a connect string with the following information:
 - Database user ID.
 - Password for the user ID.
 - Oracle Service name.

The user account for standard database should have the following privileges:

- Administration privileges (default) or Oracle database administration privileges.
- SYSDBA and ALTER SYSTEM system privileges.

Alternatively, instead of using administrator user account with SELECT ANY TABLE privilege, you can also create user accounts with lesser privileges for accessing specific objects:

For example:

- SELECT ON "SYS"."V_\$DATABASE"
- SELECT ON "SYS"."V_\$DATAFILE"
- SELECT ON "SYS"."DBA_TABLESPACES"
- GRANT SELECT ON "SYS"."V_\$ARCHIVE_DEST" TO "USER_NAME"

The Recovery Catalog database user account must have recovery catalog owner privileges.

Additional accounts (except Impersonate User) should be established by Oracle database administrator.

CONFIGURING USER ACCOUNT TO ACCESS THE ORACLE APPLICATION

Use the following steps to configure the user account to access the Oracle application:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, and then click **Properties**.
3. Click the Details tab.
4. Select the desired Instance, and then click **Modify**.

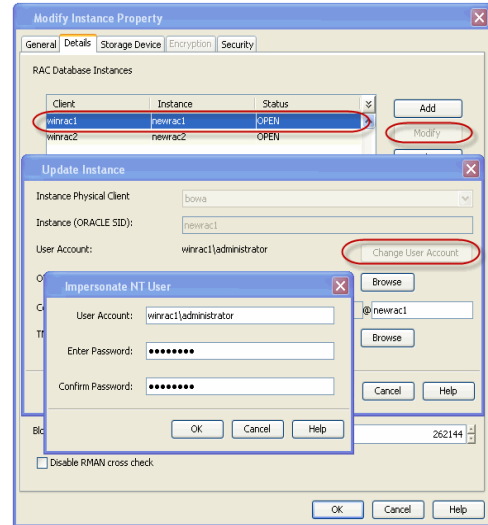
On Windows clients:

- o Click **Change User Account**.
- o In the **User Account** box, type the user account name.
- o In the **Enter Password** box, type the password for the user account.
- o In the **Confirm Password** box, retype the password, and then click **OK**.

On Unix Clients:

- o Click **Change User Account**.
- o In the **User Name** box, type the user name, and then click **OK**.

5. Click **OK**.
6. Click **OK**.



CONFIGURING USER ACCOUNT TO ACCESS THE ORACLE DATABASE

Use the following steps to configure the user account privileges to access the Oracle database:

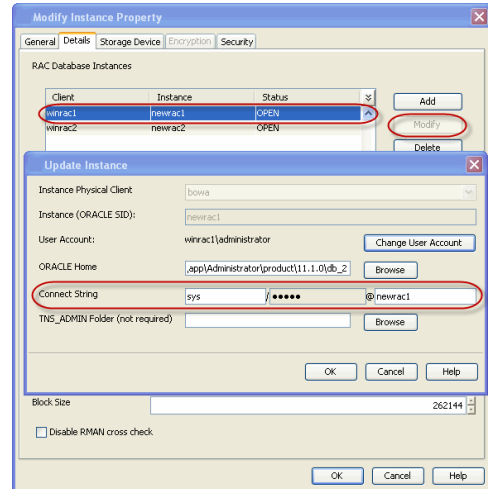
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, and then click **Properties**.
3. Click the **Details** tab.
4. Select the desired Instance, and then click **Modify**.
5. In the **Connect String** box, type the connect string to connect to the Oracle database as following:
 - o Type the Database user ID.
 - o Click the Password box, in the **Enter Password** box, type the password for the user ID.
 - o In the **Confirm Password** box, retype the password, and then click **OK**.
 - o Type the Oracle service name.

For example:

```
sys/<password>@winrac1
```

where, *sys* is the Database User ID, *<password>* is the password of the Database User ID, and *winrac1* is the Oracle service name.

6. Click **OK**.



USING RECOVERY CATALOG FOR BACKUPS

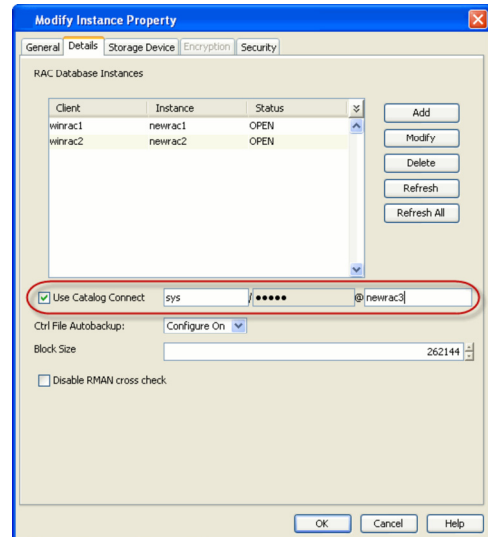
By default, recovery catalog is not used for backup and restore. As the recovery catalog contains metadata about RMAN operations for each registered database, it is recommended to include recovery catalog in backups. The metadata information is useful when you want to restore and recover the database after a crash. Use the following steps to use the Recovery Catalog for backups:

1. Verify that you have a recovery catalog database created using the following command:
2. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
3. Right-click the **<Instance>**, and then click **Properties**.
4. Click the **Details** tab.
5. Select the **Use Catalog Connect** check box.

In the Catalog Connect string box, type the credentials to connect to the desired Recovery Catalog. For example, *sys2/pwd34@orc12*

where *sys2* is database user ID, *pwd34* is password for the user ID and *orc12* is Oracle service name.

6. Click **OK**.

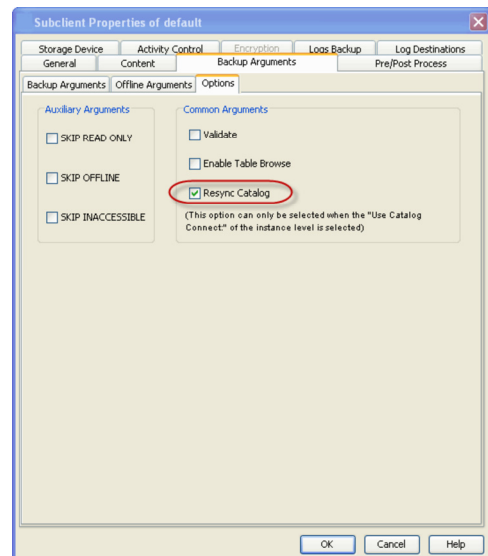


SYNCHRONIZING THE CONTROL FILE WITH THE RECOVERY CATALOG

It is a good practice to synchronize the recovery catalog with the control file, as the control file contains the latest backup information, which is need to perform restores.

Use the following steps to synchronize the control file with the recovery catalog:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab, and then click the **Options** tab.
4. Select the **Resync Catalog** check box.
5. Click **OK**.



CONFIGURING STREAMS FOR BACKUPS

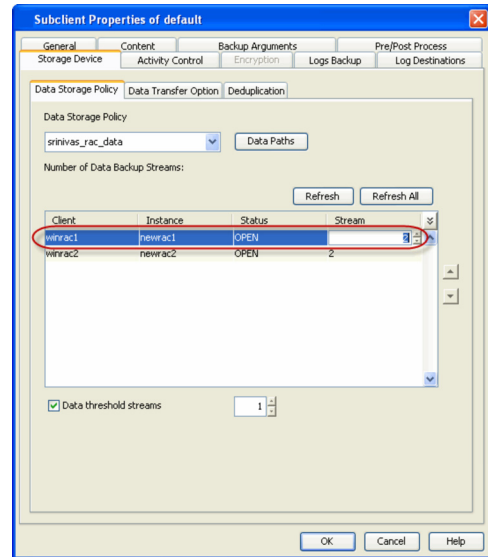
The Oracle RAC iDataAgent allows you to perform parallel backups across multiple RAC nodes as a matter of design. Unlike other agents, this agent requires just one subclient and one storage policy per type used (i.e., data, log) along with the requisite number of allocated streams to perform parallel backups. As such, you can load-balance your backups across all or selected nodes of your RAC database and thereby take advantage of the strengths of each node.

Keep in mind that the same number of streams used for performing a backup will be needed to restore the data.

DATA BACKUPS

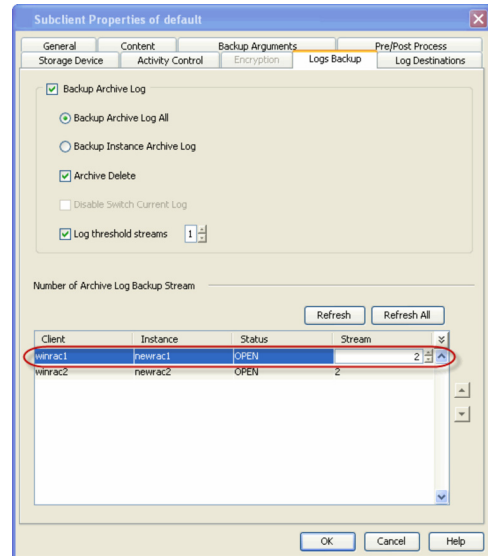
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Storage Device** tab.
4. Double click the Stream box of the desired instance, and then type or select the number of streams.

5. Click **OK**.



LOG BACKUPS

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, and then click **Properties**.
3. Click the **Logs Backup** tab.
4. Double click the Stream box of the desired instance, and then type or select the number of streams.
5. Click **OK**.



ENHANCING BACKUP PERFORMANCE

Several options are available for enhancing backup performance and reduce the network bandwidth used for performing backups. These options include:

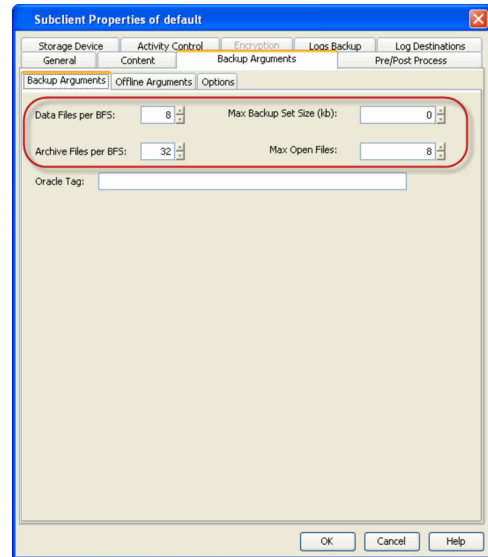
- Specifying the number of datafiles and archive files to be included in each RMAN backup set. The default value is 32.
- Limiting the maximum size for each RMAN backup set. By default no limits are defined.
- Limiting the number of open datafiles that can be read by RMAN during the backup. The default value is 8.
- Limiting the maximum size of data blocks used during backups. The default value is 262144 Kb.

In addition to the above configurations, you can also enable distribution of data across disks during backup operations. See Enhancing Backup Performance for more details

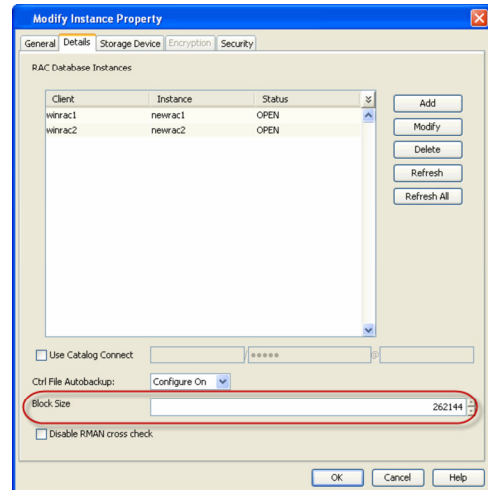
Use the following steps to enhance the backup performance:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab.
4. In the **Data Files per BFS** box, type or select the number of datafiles to be bundled in each RMAN backup set.

5. In the **MAX Backup Set Size (kb)** box, type or select the size of backup set allowed for RMAN backup set.
6. In the **Archive Files per BFS** box, type or select the number of archive files in each RMAN backup set.
7. In the **Max Open Files** box, type or select the number of datafiles that RMAN can read from simultaneously during a backup operation.



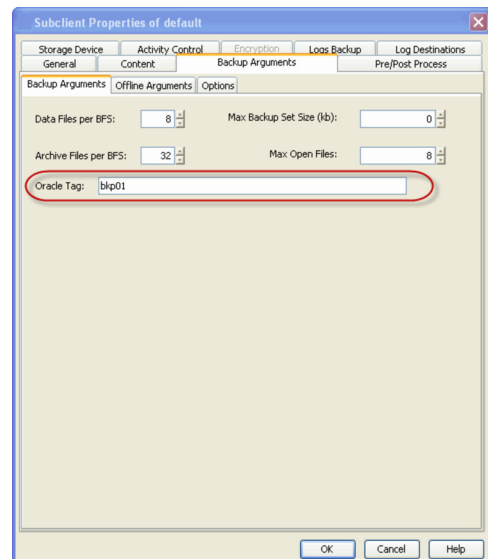
8. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
9. Right-click the **<instance>**, and then click **Properties**.
10. Click the **Details** tab.
11. In the **Block Size** box, type or select the size of the block for backup and restore operations.
12. Click **OK**.



ASSIGNING UNIQUE IDENTIFICATION TAGS FOR BACKUPS

You can assign unique identification tags for all backup operations from a specific subclient. These tags can be used during a restore operation to easily identify a particular backup. Use the following steps to assign unique identification tag for backups.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab.
4. In the **Oracle Tag** box, type the tag name.
5. Click **OK**.



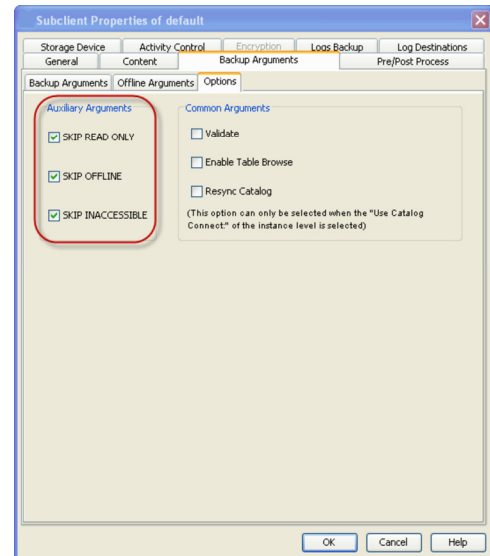
EXCLUDING DATA DURING BACKUPS

You can configure a subclient to skip the following data during backup operations:

- Read-only tablespaces
- Offline tablespaces
- Inaccessible Datafiles and Archived redo log files

Use the following steps to exclude data during backup operations:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab, and then click the **Options** tab.
4. Select the **SKIP READ ONLY** check box to exclude the read only tablespaces.
5. Select the **SKIP OFFLINE** check box to exclude the offline tablespaces.
6. Select the **SKIP INACCESSIBLE** check box to exclude inaccessible data and log files.
7. Click **OK**.



SIMULATING A BACKUP

Prior to running backup operations, you can check datafiles for physical and logical corruption.

Prior to running backup operations, you can validate a backup job, which will cause RMAN to simulate the backup for the purpose of determining whether the backup can be successfully restored. Once a validate job is completed, you can view the log files of the job to identify and correct any validation issues.

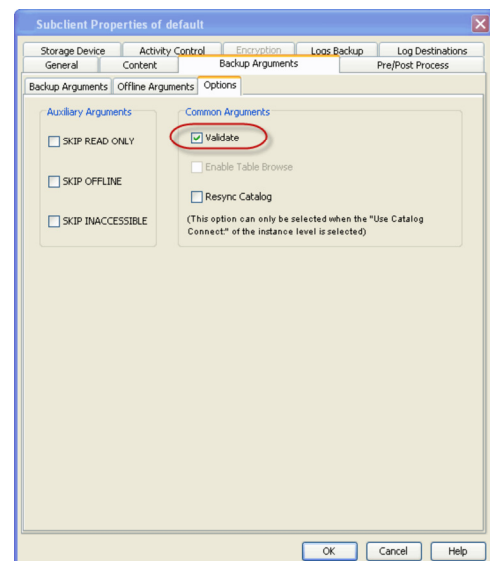
During validation, the backup jobs are simulated without the media, you can view the log files of the job to identify and correct any validation issues.

Use the following steps to enable validation of backup jobs:

Prior to running a backup, you can check the following:

- Datafiles for physical and logical block corruption
- Database files exist and are in the correct location.

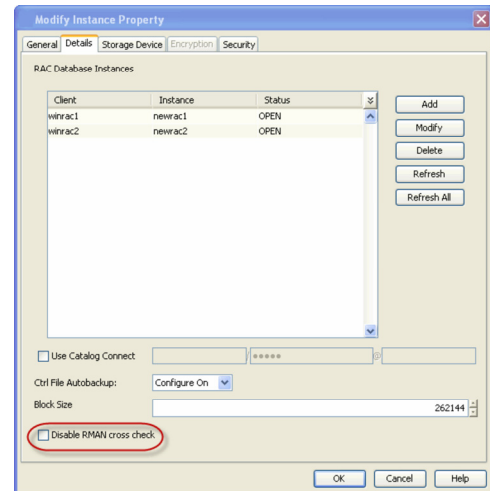
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab, and then click the **Options** tab.
4. Select the **Validate** check box.
5. Click **OK**.



DISABLING THE RMAN CROSSCHECK

By default, during a data aging operation, an Oracle CROSSCHECK is performed by the system to synchronize the entries in the CommServe database with the RMAN catalog. Use the following steps to disable the cross check operation:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**.
2. Right-click the **<Instance>**, and then click **Properties**.
3. Click the **Details** tab.
4. Select the **Disable RMAN cross check** check box.
5. Click **OK** to close the warning message.
6. Click **OK**.



ENABLING MULTIPLE BACKUP COPIES

You can take multiple copies of the data or log backups using RMAN command line. During restores, even if one of the copies is missing or corrupted, the restore operation will automatically failover to the other copy and restore the data.

1. To utilize the PARALLELISM option, you need to set the initial parameter in pfile or spfile.
2. Configure device type.
3. Specify two copies of data.
4. Specify two copies of log backups.
5. Add the environmental variables for the client and instance on which the iDataAgent is installed.
6. Once you set the parameters, restart the Oracle database and perform backup operations from the CommCell Console.

Example:

```
BACKUP_TAPE_IO_SLAVES=TRUE
```

Example:

```
RMAN> CONFIGURE DEVICE TYPE DISK PARALLELISM 2
BACKUP TYPE TO BACKUPSET;
```

Example:

```
RMAN> CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE
TYPE 'SBT_TAPE' TO 2;
```

Example:

```
RMAN> CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE
TYPE 'SBT_TAPE' TO 2;
```

Example:

```
allocate channel ch1 type 'sbt_tape'
PARMS="<software_install_path>/Base/libobk.so,
ENV=(CvClientName=<client_name>,
CvInstanceName=<instance_name>)"
```

Example: To restart the database,

1. Connect to the database.

```
SQL> connect admin/admin@orcl as sysdba
```

2. Shutdown the database.

```
SQL> shutdown
```

3. Mount the database.

```
SQL> startup mount;
```

4. Change to archive log mode.

```
SQL> alter database archivelog;
```

5. Open the database.

```
SQL> alter database open;
```

COMMAND LINE OPERATIONS

You can add, modify, or delete several configurable properties for the Oracle RAC iDataAgent from the command line interface.

Command line configuration enables you to:

- configure the same properties across multiple clients simultaneously.
- reuse the same configurations for additional entities.

The following sections describe the available command line configurations:

AVAILABLE PARAMETERS

PARAMETER	DESCRIPTION OF PARAMETER VALUES
appName	Name of the application. In this case it would be 'Oracle RAC'
clientName	Name of the RAC pseudo client as displayed in the CommCell Browser (eg., RACPseudoClient)
instancePhysicalClient/clientName	Name of RAC client node (eg., racnode1)
instanceName	Name of the Oracle instance (eg., dbname)
userAccount/userName	Name of the user that has local administrator rights to administer the Oracle application. This is the account used by the user to log in to and use the application to run jobs. (eg., oracle)
oracleRACInstance/connectString/userName	Name of the Oracle recovery catalog database (eg., cataloguser)
oracleRACInstance/connectString/password	Password for the catalog connect user (eg., catalogpasswd)
oracleRACInstance/connectString/serviceName	Name of the Oracle recovery catalog database (eg., catalogdb)
racDBInstance/connectString/userName	Name of the user with dba privileges (eg., sys)
racDBInstance/connectString/password	Password for the user with dba privileges
instanceOracleSID	Name of the RAC node database. (eg., dbname1 for racnode1 and dbname2 for racnode2)
oracleHome	Path for the Oracle application software (eg., /oracle/oraclehome/)
tnsAdminFolder	Path to the TNS Admin directory (eg., /oracle/oraclehome/network/admin/)
dataBackupStoragePolicy/storagePolicyName	Name of the data storage policy (eg., data_sp)
logBackupStoragePolicy/storagePolicyName	Name of the log storage policy (eg., log_sp)
blockSize	Block size value for backup and restore operations on the selected instance (eg., 163840)
enableDeduplication	Option to enable or disable deduplication. Valid values are True/False.
networkAgents	Number of Network Agents (eg., 2)
useCatalogConnect	Option to choose whether RMAN will use recovery catalog database at the time of backup. Valid values are True/False. If recovery catalog database need to be used choose true else false.
softwareCompression	Option to enable compression on the Client or MediaAgent computer. Valid values are: <ul style="list-style-type: none"> • ON_CLIENT, to enable software compression on the client. • ON_MEDIAAGENT, to enable software compression on the MediaAgent. • USE_STORAGE_POLICY_SETTINGS, to use the software compression options defined on the storage policy. • OFF, to disable software compression.

LOG ON TO THE COMMSERVE

To run command line operations you must first login to the CommServe as follows:

- From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
qlogin -cs <commserve name> -u <user name>
```

- For example, to log on to CommServe 'server1' with username 'user1':

```
qlogin -cs server1 -u user1
```

CREATING A RAC PSEUDO CLIENT

1. Download the CreateRACClient.xml file and save it on the computer from where the command will be executed.
2. From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
qoperation execute -af <xml_path>\CreateRACClient.xml -clientInfo/clientType 'RAC' -entity/clientName '<RACpseudoclientname>'
```

ADDING FIRST RAC NODE

1. Download the CreateRACInstance.xml file and save it on the computer from where the command will be executed.
2. From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
qoperation execute -af <xml_path>\CreateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName '<RACpseudoclientname>' -instance/instanceName '<RACinstancename>' -useCatalogConnect true -
```

```
oracleRACInstance/connectString/userName '<catalog_db_user_name>' -oracleRACInstance/connectString/password
'<catalog_db_password>' -oracleRACInstance/connectString/domainName '<catalog_db_name>' -ctrlFileAutoBackupType
NOT_CONFIGURED -blockSize 163840 -disableRMAN true -networkAgents 2 -softwareCompression OFF -
dataBackupStoragePolicy/storagePolicyName '<data_storage_policy>' -logBackupStoragePolicy/storagePolicyName
'<log_storage_policy>' -enableDeduplication true -generateSignature ON_CLIENT -racDBOperationType ADD -
instancePhysicalClient/clientName '<RACnode1name>' -instanceOracleSID <node1instancename> -userAccount/userName
'<oracleusername>' -oracleHome '<oracle_home_path>' -racDBInstance/connectString/userName '<racdb_username>' -
racDBInstance/connectString/password '<racdb_password>' -racDBInstance/connectString/serviceName '<racdbname>' -
tnsAdminFolder '<tnsadmin_path>'
```

ADDING SECOND RAC NODE

1. Download the UpdateRACInstance.xml file and save it on the computer from where the command will be executed.
2. From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
operation execute -af <xml_path>\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName
'<RACpseudoclientname>' -instance/instanceName '<RACinstancename>' -racDBOperationType ADD -
instancePhysicalClient/clientName '<RACnode2name>' -instanceOracleSID <node2instancename> -userAccount/userName
'<oracleusername>' -oracleHome '<oracle_home_path>' -racDBInstance/connectString/userName '<racdb_username>' -
racDBInstance/connectString/password '<racdb_password>' -racDBInstance/connectString/ServiceName '<racdbname>' -
tnsAdminFolder '<tnsadmin_path>'
```

GETTING RAC INSTANCE PROPERTIES

1. Download the GetRACInstance.xml file and save it on the computer from where the command will be executed.
2. From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
operation execute -af <xml_path>/GetRACInstance.xml -appName 'Oracle RAC' -clientName '<RACpseudoclientname>' -
instanceName '<RACinstancename>'
```

CLICK HERE TO SEE A SAMPLE OUTPUT

MODIFYING RAC INSTANCE PROPERTIES

1. Download the UpdateRACInstance.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
operation execute -af <xml_path>\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName
'<RACpseudoclientname>' -instance/instanceName '<RACinstancename>' -racDBOperationType OVERWRITE -
instancePhysicalClient/clientName '<RACnodename>' -instanceOracleSID <instancename_for_node> -userAccount/userName
<username> -oracleHome '<homepath>' -racDBInstance/connectString/userName '<connect_string_username>' -
racDBInstance/connectString/password '<connect_string_password>' -racDBInstance/connectString/ServiceName
'<connect_string_name>' -tnsAdminFolder '<tnsadmin_path>'
```

DELETING RAC INSTANCE

1. Download the DeleteRACInstance.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
operation execute -af <xml_path>DeleteRACInstance.xml -appName 'Oracle RAC' -clientName '<RACpseudoclientname>' -
instanceName '<RACinstancename>'
```

EXAMPLES

Creating a RAC Client	operation execute -af C:\XML\CreateRACClient.xml -clientInfo/clientType 'RAC' -entity/clientName 'RacPseudoClient'
Adding First RAC Node	operation execute -af C:\XML\CreateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -useCatalogConnect true -oracleRACInstance/connectString/userName 'cataloguser' -oracleRACInstance/connectString/password 'catalogpasswd' -oracleRACInstance/connectString/domainName 'catalogdb' -ctrlFileAutoBackupType NOT_CONFIGURED -blockSize 163840 -disableCATALOGUSER true -networkAgents 2 -softwareCompression OFF -dataBackupStoragePolicy/storagePolicyName 'data_sp' -logBackupStoragePolicy/storagePolicyName 'log_sp' -enableDeduplication true -generateSignature ON_CLIENT -racDBOperationType ADD -instancePhysicalClient/clientName 'racnode1' -instanceOracleSID dbname1 -userAccount/userName 'oracle' -oracleHome '/oracle/oraclehome/' -racDBInstance/connectString/userName 'sys' -racDBInstance/connectString/password 'syspasswd' -racDBInstance/connectString/serviceName 'dbname1' -tnsAdminFolder '/oracle/oraclehome/network/admin/'

Adding First RAC Node without Catalog Connect	<pre>qoperation execute -af C:\XML\CreateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -useCatalogConnect false -ctrlFileAutoBackupType NOT_CONFIGURED -blockSize 163840 -disableCATALOGUSER true -networkAgents 2 -softwareCompression OFF -dataBackupStoragePolicy/storagePolicyName 'data_sp' -logBackupStoragePolicy/storagePolicyName 'log_sp' -enableDeduplication true -generateSignature ON CLIENT -racDBOperationType ADD -instancePhysicalClient/clientName 'racnode1' -instanceOracleSID dbname1 -userAccount/userName 'oracle' -oracleHome '/oracle/oraclehome/' -racDBInstance/connectString/userName 'sys' -racDBInstance/connectString/password 'oracle' -racDBInstance/connectString/serviceName 'dbname1' -tnsAdminFolder '/u01/app/oracle/product/11.2.0.3/db_1/network/admin/'</pre>
Adding Second RAC Node	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -racDBOperationType ADD -instancePhysicalClient/clientName 'racnode2' -instanceOracleSID dbname2 -userAccount/userName oracle -oracleHome '/oracle/oraclehome/' -racDBInstance/connectString/userName 'sys' -racDBInstance/connectString/password 'passwd' -racDBInstance/connectString/ServiceName 'dbname2' -tnsAdminFolder '/oracle/oraclehome/network/admin/'</pre>
Configure Control File Auto Backup	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -ctrlFileAutoBackupType CONFIGURE ON</pre>
Set the Block Size	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -blockSize 163840</pre>
Modify Log Storage Policy and Disable Catalog User	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -logBackupStoragePolicy/storagePolicyName 'log_sp' -disableCATALOGUSER true</pre>
Change Network Agents	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -networkAgents 2</pre>
Change Software Compression and Network Agents	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -softwareCompression OFF -networkAgents 1</pre>
Enable Deduplication	<pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -enableDeduplication true -generateSignature ON_MEDIAAGENT</pre>
Delete RAC Database Instance	Use the following command to delete the instance dbname1: <pre>qoperation execute -af C:\XML\UpdateRACInstance.xml -instance/appName 'Oracle RAC' -instance/clientName 'RacPseudoClient' -instance/instanceName 'dbname' -racDBOperationType DELETE -instancePhysicalClient/clientName 'racnode1' -instanceOracleSID 'dbname1'</pre>
Delete RAC Instance	<pre>qoperation execute -af <xml_path>DeleteRACInstance.xml -appName 'Oracle RAC' -clientName 'RacPseudoClient' -instanceName 'dbname'</pre>

MODIFYING AN AGENT, INSTANCE, OR SUBCLIENT

There are several configurable properties available for your agent that can be modified from the agent, instance, or subclient level as per need.

It is recommended that you do not modify the properties of a subclient when a job is in progress for that specific subclient. If a job is in progress, either wait for the job to complete or kill the job from the Job Controller.

The following table describes the properties that can be configured from the agent, instance, and subclient levels.

OPTION	DESCRIPTION	RELATED TOPICS
Change Storage Policies	<p>You can modify the storage policies in any of the following situations:</p> <ul style="list-style-type: none"> To include a different media for the backup operation. To use a storage policy with a different retention criteria. <p>You can change the storage policies from the subclient level.</p> <ol style="list-style-type: none"> From the CommCell Browser, right-click the subclient. Click Properties. Click Storage Device. Select the Storage policy from the drop-down menu. Click OK. 	Refer to Storage Policies.
Rename a Subclient	<p>You can rename a subclient:</p> <ol style="list-style-type: none"> From the CommCell Browser, right-click the subclient. Click Properties. Type the new name in the Subclient name field. 	

	<p>4. Click OK.</p>	
<p>Data Transfer Options</p>	<p>You can efficiently configure the available resources for transferring data secured by data protection operations from the subclient level. This includes the following:</p> <ul style="list-style-type: none"> • Enable or disable Data Compression either on the client or the MediaAgent. • Configure the transfer of data in the network using the options for Network Bandwidth Throttling and Network Agents. <p>You can configure the data transfer options.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the subclient. 2. Click Properties. 3. Click Storage Device. 4. Click Data Transfer Option tab. 5. Choose the appropriate software compression option for this subclient. 6. Select Throttle Network Bandwidth and set the required bandwidth. 7. Click OK. 	<p>Refer to Data Compression and Network Bandwidth Throttling.</p>
<p>View Data Paths</p>	<p>You can view the data paths associated with the primary storage policy copy of the selected storage policy or incremental storage policy. You can also modify the data paths including their priority from the subclient level.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Storage Device. 4. Select Storage Policy from the drop-down menu. 5. Click Data Paths. 	
<p>Configure a Subclient for Pre/Post Processing of Data Protection</p>	<p>You can add, modify or view Pre/Post processes for the subclient. These are batch files or shell scripts that you can run before or after certain job phases.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Pre/Post Process. 4. Click one of the following phases and type the full path of the process that you want to execute during that phase. Alternatively, click Browse to locate the process (applicable only for paths that do not contain any spaces). <ul style="list-style-type: none"> ○ PreBackup Process ○ PostBackup Process ○ PreSnap Process ○ PostSnap Process 5. Click OK. 6. Select Run Post Backup Process for all attempts to run a post backup process for all attempts. 7. For subclients on Windows platforms, Run As displays Not Selected. If you want to change the account that has permission to run these commands, click Change. <ol style="list-style-type: none"> a. In the User Account dialog box, select Use Local System Account, or select Impersonate User and enter the user name and password. click OK. b. If you selected Local System Account, click OK to the message advising you that commands using this account have rights to access all data on the client computer. 	<p>Refer to Pre/Post Processes.</p>
<p>Configure Activity Control</p>	<p>You can enable backup and restore operations from the agent and subclient level. However, you can enable restore operations only from the agent level.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Activity Control, select or clear option(s) as desired. 4. Click OK. 	<p>Refer to Activity Control.</p>
<p>Configure User Security</p>	<p>You can configure user security from the agent or subclient level.</p> <p>You can perform the following functions:</p> <ul style="list-style-type: none"> • Identify the user groups to which this CommCell object is associated. • Associate this object with a user group. • Disassociate this object from a user group. <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 	<p>Refer to User Administration and Security.</p>

	<ol style="list-style-type: none"> Click Security. Select the appropriate user groups to which you want to associate to the CommCell object from the Available Groups pane, and then move the user group to the Associated Groups pane. Click OK. 	
Enable/Disable Data Encryption	<p>You can enable data encryption from the subclient level. Encryption must be enabled at the client level prior to configuring any instances residing on that client.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. Click Properties. Click Encryption. Select the desired encryption. Click OK. 	Refer to Data Encryption.
View Software Version and Installed Updates	<p>The Version tab, at the Agent level displays the software version of the component.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the agent. Click Properties. Click Version. Click OK. 	
CommCell Configuration Report	<p>The CommCell Configuration Report provides the properties of the CommServe, MediaAgents, clients, agents, SRM agents, subclients, and storage policies within the CommCell based on the selected filter criteria.</p> <ol style="list-style-type: none"> From the CommCell browser, click Reports icon. Select CommCell Configuration. Click Run. 	Refer to CommCell Configuration.

DELETING AN AGENT, INSTANCE, OR SUBCLIENT

The following sections describe the steps involved in deleting an agent, instance, or subclient.

When you delete an instance or backupset, the associated data is logically deleted and you can no longer access the corresponding data from CommCell Console for recovery purposes.

Refer to the troubleshooting article on Recovering Data Associated with Deleted Clients and Storage Policies for information on how to recover data if you accidentally delete an entity.

DELETING AN AGENT

You need to uninstall or DeConfigure the agent software from the client computer before deleting from CommCell Browser. After you delete the client software, you can either leave the corresponding data intact for appropriate action or you can remove the data immediately. If you choose to remove the data immediately, you must delete the agent from the CommCell Browser. If you delete the agent, all of the agent's data is irretrievably lost.

- You cannot delete an agent while operations for that agent are running.
- From the CommCell Browser, navigate to **Client Computers | <Client>**.
 - Right-click the **<Agent>**, and then click **Delete**.
 - A confirmation message is displayed with the following message:

This operation will permanently delete the data backed up from this level and it cannot be restored.

- Click **OK** to continue with the deletion operation., or click **No** to abort the deletion.

DELETING AN INSTANCE

Data that is backed up for any Oracle instance configured on the Oracle RAC node is lost if the configured Oracle instance is hard-deleted.

Consider the following before deleting an instance:

- When you delete a specific instance all job schedules and job histories that pertain to any of the levels within the deleted instance are deleted.
- You cannot delete an instance if it is being backed up. Attempts to delete an instance under such conditions cause the deletion to fail. If a backup is in progress, either wait for the backup to complete or kill the backup job using the Job Manager. Once the backup is no longer in progress, you can delete the instance level.
- You cannot delete an instance if there is only one instance present for an agent. To delete the final instance, you must remove the agent software from the client computer.

1. From the CommCell Browser, right-click the instance that you want to delete, click **All Tasks** and then click **Delete**.
2. click **Yes** to confirm the deletion. (clicking **No** cancels the deletion and retains the node.)
3. Type the requested phrase in the **Enter Confirmation Text** dialog box and click **OK**. This should delete the instance.

DELETING A SUBCLIENT

Consider the following before deleting a subclient:

- You cannot delete a default subclient.
 - Schedules associated with the subclient are also automatically deleted.
1. From the CommCell Browser, navigate to **Client Computers** | **<RAC Client>** | **<Instance>**.
 2. Right-click the **<Subclient>** that you want to delete, and then click **Delete**.
 3. A confirmation message is displayed, asking if you want to delete the subclient.
Click **No** to cancel the deletion and retain the subclient, or click **Yes** to continue the deletion.

Advanced Backup - Oracle RAC iDataAgent

TABLE OF CONTENTS

Full Backups

Incremental Backups

Cumulative Incremental Backups

Archive Log Backups

All Archive Logs

By Number of Days

Within a Specific Time Range

By Log Sequence Number

By System Change Number

By a Specific String Pattern

From a Specific Location

Logs Not Backed Up a Specified Number of Times

Deleting Archive Logs for a Specific Backup

Control File Backups

Command Line Backups

Running Backups using QCommands

Running RMAN Scripts from Third Party Command Line

Multi Stream Backups from Third Party Command Line

Previewing RMAN Scripts from CommCell Console

Customizing RMAN Scripts from CommCell Console

Scheduling Backups

Marking Backups with a Unique Identification Tag

For a Specific Backup

For All Backups

Enhancing Backup Performance

Validating Backups for Restore

Managing Jobs

Restarting Jobs

Controlling Jobs

Additional Options

FULL BACKUPS

Full backups provide the most comprehensive protection of data.


Backups for any client start with a full backup. If the database is in NOARCHIVELOG mode, you should perform offline backup only.

You can perform a full backup of an online or offline database.

Use the following steps to run a full backup:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** and click **Backup**.
3. Select **Full** as the backup type and click **Immediate**.
4. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup job has completed, the **Job Controller** displays **Completed**.



Data that is backed up for any Oracle instance configured on the Oracle RAC node is lost if the configured Oracle instance is hard-deleted.

INCREMENTAL BACKUPS

Incremental backups can be performed when the database is online.

The incremental backups will fail if the database is offline.

An incremental backup contains only data that is new or has changed since the last

RELATED TOPICS

Scheduling

Provides comprehensive information on scheduling jobs.

Job Management

Provides comprehensive information on managing jobs.


Command Line Interface

Provides comprehensive information on running jobs from the command line.

backup, regardless of the type. On average, incremental backups consume less media and use less resources than full backups.


The illustration on the right clarifies the nature of full and incremental backups.

Follow steps given below to perform an incremental backup:



1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** and click **Backup**.
3. Select **Incremental** as the backup type and click **Immediate**.
4. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup job has completed, the **Job Controller** displays **Completed**.



CUMULATIVE INCREMENTAL BACKUPS

In a cumulative level n backup, all the data changes since the most recent backup at level $n-1$ or lower are backed up.

For example, in a cumulative level 2 backup, data changes since the most recent level 1 backup are backed up. If no level 1 backup is available, data changes since the base level 0 backup are backed up.


Cumulative incremental backups reduce restore times because you need one incremental backup from any particular level. However, cumulative backups require more space and time because they duplicate the work done by previous backups at the same level.

By default, the incremental level for cumulative backups is 1.

Follow the steps below to perform a cumulative incremental backup:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** and click **Backup**.
3. Select **Incremental** as the backup type, and then type or select the incremental level in the **Oracle Incremental Level** box.
4. Select the **Cumulative** check box and click **Immediate**.
5. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup job has completed, the **Job Controller** displays **Completed**.



ARCHIVE LOG BACKUPS

An archive log backup captures the archive redo logs generated during database transactions.

Archive log backups are useful when you want to recover database transactions that have been lost due to an operating system or disk failure. You can apply these archive logs to an online backup in order to recover a database.

By default full backups include both data and archive logs. However, you can also perform separate archive log backups.

In order to perform a backup of the archive logs:

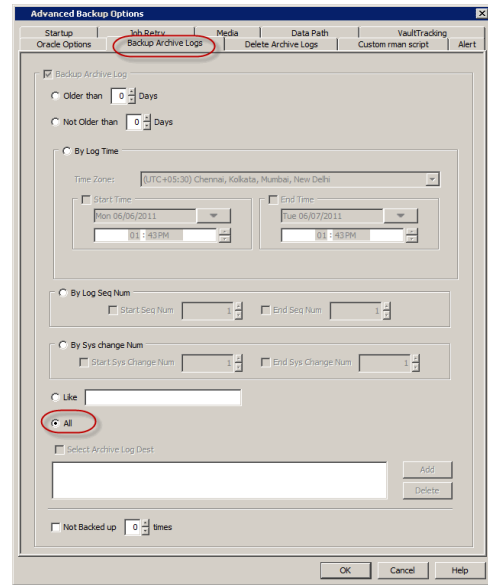
- The database has to be in ARCHIVELOG mode.
- Subclient should have been configured for archive log backups. See *Creating A Subclient for Log Backups* for step-by-step instructions to create a separate subclient for archive logs.

ALL ARCHIVE LOGS

Use the following steps to backup all archive logs. Note that this is the default option:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.

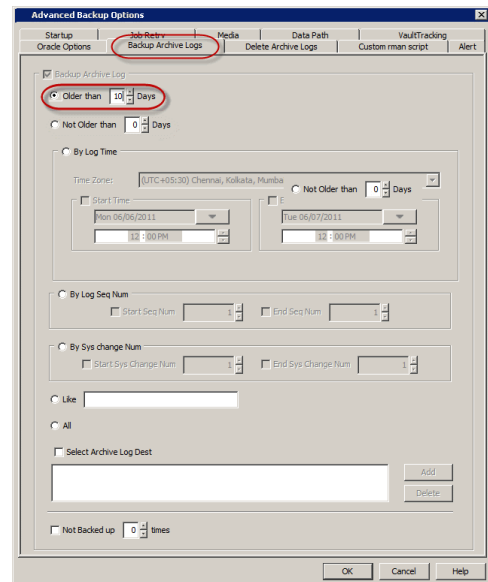
4. Click the **Backup Archive Logs** tab.
5. Click **All**.
6. Click **OK**.



BY NUMBER OF DAYS

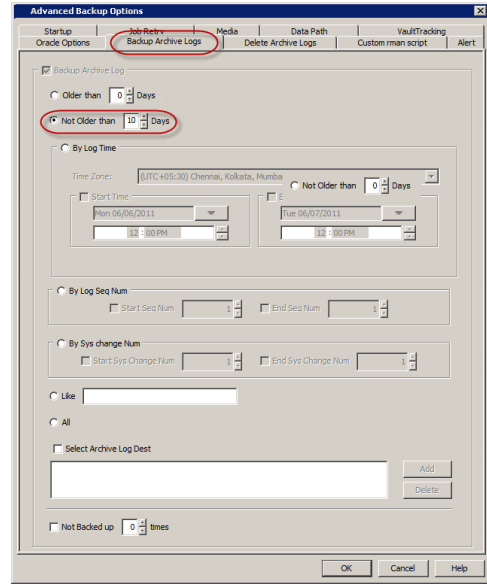
Use the following steps to backup archive logs older than a specified number of days.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Backup Archive Logs** tab.
5. Click **Older than n Days**, and type or select the number of days older than which the archive logs are to be backed up.
For example, to backup logs older than 10 days, type 10.
6. Click **OK**.



Use the following steps to backup archive logs not older than a specified number of days.

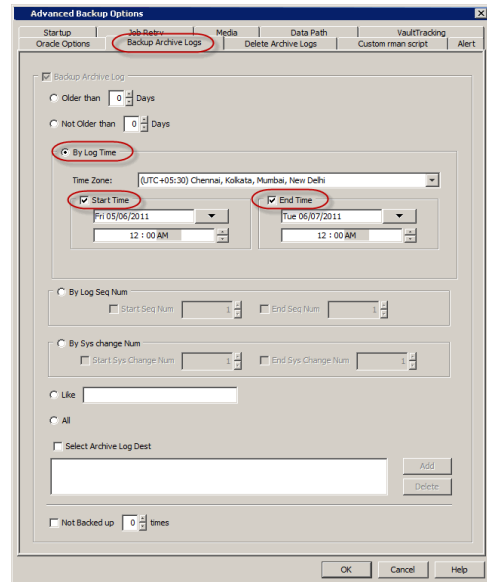
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Backup Archive Logs** tab.
6. Click **Not Older than n Days**, and type or select the number of days not older than which the archive logs are to be backed up.
For example, to backup logs not older than 10 days, type 10.
7. Click **OK**.



WITHIN A SPECIFIC TIME RANGE

Use the following steps to backup archive logs between a specified time:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Backup Archive Logs** tab.
5. Click **By Log Time**.
6. In the **Time Zone** box, select the time zone.
7. Select the **Start Time** check box and enter the start time after which the logs were generated.
8. Select the **End Time** check box and enter the time before which the logs were generated.
9. Click **OK**.



BY LOG SEQUENCE NUMBER


Log Sequence Number uniquely identifies an archive log. For example, if you create a database with two online log files, then the first file is assigned log sequence number 1. When the first file fills Oracle switches to the second file and assigns a log sequence number of 2; when it switches back to the first file, it assigns log sequence number 3, and so forth.

Use the following steps to backup archive logs within a specific range of log sequence numbers:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Backup Archive Logs** tab.
5. Click **By Log Seq Num**.
6. Select the **Start Seq Num** check box and enter the start sequence number of logs to

be backed up.

7. Select the **End Seq Num** check box and enter the end sequence number of logs to be backed up.
8. Click **OK**.




BY SYSTEM CHANGE NUMBER

System Change Number (SCN) is a stamp that defines a committed version of a database at a point in time. Oracle assigns every committed transaction a unique SCN. For example, SCNs of two successive transactions committed could be 576601 and 576799.

Use the following steps to backup archive logs within a specific range of system change numbers:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Backup Archive Logs** tab.
5. Click **By Sys change Num**, and then specify the **Start Sys Change Num** and **End Sys Change Num**.
6. Click **OK**.



BY A SPECIFIC STRING PATTERN

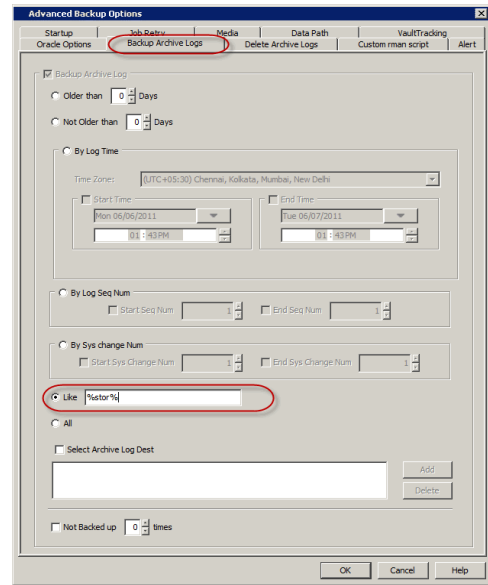
Use the following steps to backup archive log files whose name match a specific naming pattern. Note that if you do not specify any pattern, all the logs from the specified destination will be backed up.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Backup Archive Logs** tab.
5. Click **Like**, and then enter the desired string pattern in the text box.

For example, to backup all archive logs whose names start with 'arch' type

'arch%.log'.

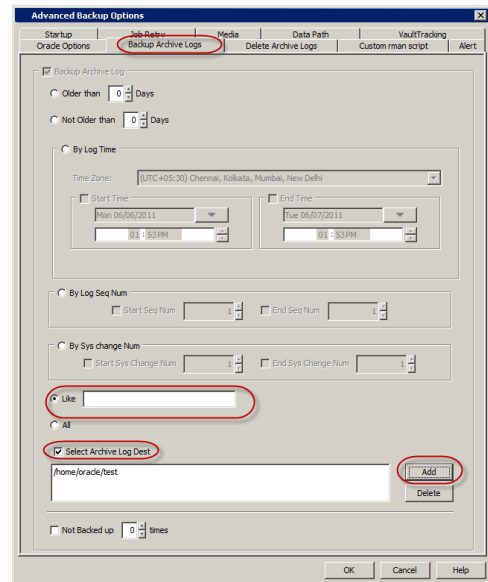
- Click **OK**.



FROM A SPECIFIC LOCATION

Use the following steps to backup archive logs from a specific path or location. Note that the path or location specified at the backup level will override the archive log location defined at the subclient level.


- From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
- Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
- Click **Advanced**.
- Click the **Backup Archive Logs** tab.
- Click **Like**.
- Select **Select Archive Log Dest** check box.
- Click **Add** to select the archive log location.
- Click **OK**.



LOGS NOT BACKED UP A SPECIFIED NUMBER OF TIMES

Use the following steps to backup archive logs that have failed to backup a specified number of times earlier:

- From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
- Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
- Click **Advanced**.
- Click the **Backup Archive Logs** tab.
- Select the **Not Backed Up n times** check box, and type or select the number so that the logs that meet the criteria 'Not Backed Up n times' are backed up.
- Click **OK**.



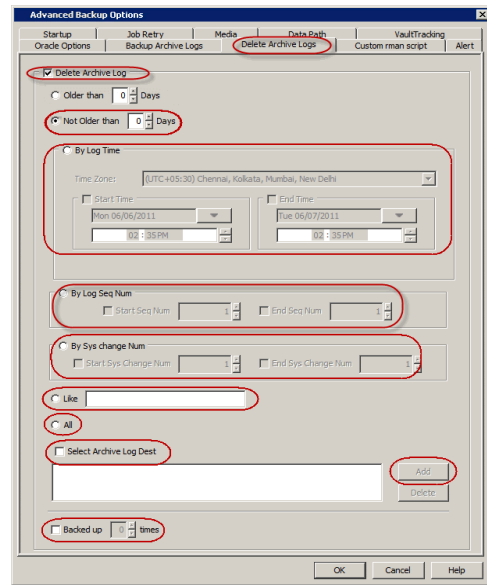
DELETING ARCHIVE LOGS FOR A SPECIFIC BACKUP

Once the archive logs are backed up you can choose to delete them from the destination location. This can be done by configuring the subclient to delete archive logs soon after the backup. See [Enabling Log Deletion After Backup](#) for step-by-step instructions on enabling deletion of logs.

Always ensure that the archive logs are backed up before they are deleted to prevent data loss.

However, you can also choose to delete the archive logs for a specific backup job. Moreover, you can also specify additional criteria to delete the archive logs.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Delete Archive Logs** tab.
5. Select the **Delete Archive Log** check box.
6. Click **Yes** on the warning dialog.
7. Type **confirm** in the **Enter Confirmation text** Dialog, and then click **OK**.
8. Specify the desired criteria for archive log deletion.
9. Click **OK**.



CONTROL FILE BACKUPS

The control file contains metadata about the physical structure of the database including the location of all files, the current database state etc. Each control file is associated with only one database. The control file backups are used for creating standby databases from the CommCell Console. Standby databases are used during database recovery.

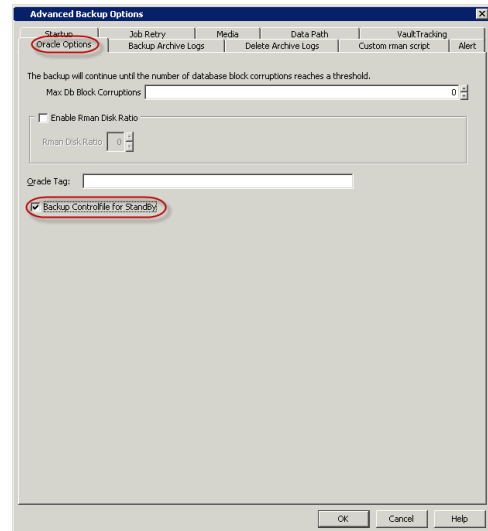
The subclients need to be configured prior to running control file backups; see [Enable/ Disable Control File Backups for a Specific Subclient](#) for step-by-step instructions.

Use the following steps to backup the control file:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> |**

<Instance>.

2. Right-click the <Subclient> configured for control file backup, and click **Backup**.
3. Click **Advanced**.
4. Click the **Oracle Options** tab.
5. Select the **Backup Controlfile for Standby** check box.
6. Click **OK**.



COMMAND LINE BACKUPS

In addition to using CommCell Console, you can also perform backups from the Command Line Interface.

RUNNING BACKUPS USING QCOMMANDS

You can submit RMAN scripts from the Command Line Interface using QCommands. The RMAN scripts are submitted through argument files.

When you submit RMAN scripts using QCommands:

- One job ID is used in the CommServe. The same Job ID is also used across different streams and attempts.
- The job can be resumed from the point of failure from the CommCell Console or Command Line.
- The job history can be viewed for these jobs.
- A list of media can be obtained for the job in primary or secondary copy.
- Job-based storage policies can be used.
- Multiple streams can be allocated before the job starts.

1. Create an argument file on the RAC client.

Example: Argument file for full backup *argfile.txt*

```
[client]
machine1_cn
[dataagent]
Q_ORACLE
[instance]
orcl
[subclient]
default
[backuptype]
Q_FULL
```

2. From the command prompt, login to the CommServe using the `qlogin` command.

Example: To log on to CommServe *server1* with user name *user1*:

```
qlogin -cs server1 -u user1
Password:
```

3. Run the backup operation using `qoperation backup`.

Example: To run a full backup on RAC client *client1* using the parameter file *argfile1.txt*:

```
D:\>qoperation backup -af D:\argfile.txt
```

RUNNING RMAN SCRIPTS FROM THIRD PARTY COMMAND LINE

Backup operations can also be performed from the third-party command line using the RMAN utility. The RMAN executable is located in the ORACLE_HOME/bin directory.

SBT PARAMETERS

PARAMETER NAME	DESCRIPTION	EXAMPLE
[CvClientName]	Name of the client defined in the CommCell Console. This parameter is optional. It is primarily used in a clustered environment.	[CvClientName] client_name
[CvInstanceName]	Name of the Calypso instance installed. This parameter is optional. In cases of multiple instances of the software, the first installed instance would be 'Instance001'.	[CvInstanceName] Instance001
[CvOraSID]	Name of the Oracle System ID (SID). This parameter is used during multi stream backups and also when the Oracle instance name is different from Oracle SID.	[CvOraSID] DB1
[-hn]	RAC control host name for the backup/restore job.	[-hn] dbserve10.mycompany.com
[CvOraRacDBName]	Oracle RAC database name for the backup/restore job.	[CvOraRacDBName] asmdb
[CvOraRacDBClientName]	Oracle RAC database client name for cross machine and/or duplicate database restore job. If you are restoring the database to a computer that is not part of the RAC pseudo client (cross machine restore), you need to define the environmental variable CVOraRacDBClientName to specify a client that is part of the RAC pseudo client.	[CvOraRacDBClientName] dbserve99

Use the following steps to run backups from the third-party command line:

1. Create an RMAN script file on the RAC client.
2. From the RMAN command prompt on the RAC client, add the environmental variables for the RAC client and instance on which the iDataAgent is installed.

On RAC clients for Unix, add the SBT_LIBRARY path.

Example: RMAN script file *backup.txt*

Example:

```
allocate channel ch1 type 'sbt_tape'
PARMS="ENV=(CvClientName=<client_name>,
CvOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>, -hn=<hostname>)"
```

Example:

```
allocate channel ch1 type 'sbt_tape'
connect <sql_connect_string>
PARMS="SBT_LIBRARY=<software_install_path>/Base/libobk.so,
ENV=(CvClientName=<client_name>,
CvInstanceName=<instance_name>,
CvOraRacDBName=<rac_database_name>)"
```

The SBT_LIBRARY path for the various platforms are listed below:

- AIX with 64 bit Oracle - <Client Agent Install Path>/Base64/libobk.a(shr.o)
- HP UX PA RISC 64 bit Oracle - <Client Agent Install Path>/Base64/libobk.sl
- Solaris with 64 bit Oracle - <Client Agent Install Path>/Base64/libobk.so
- Linux on System Z with 64 bit Oracle - <Client Agent Install Path>/Base64/libobk.so
- All Other Unix platforms - <Client Agent Install Path>/Base/libobk.so

3. Add the RMAN script for backup to the file *backup.txt* on the RAC client.

Example: RMAN script file *backup.txt*

```
run {
allocate channel ch1 type 'sbt_tape'
PARMS="BLKSIZE=262144,
SBT_LIBRARY=/opt/calypso/Base/libobk.so,
ENV=(CvClientName=<client_name>,
CvOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>);
backup database;
release channel ch1;
}

rman target sys/sys@<databasename>
```

4. Connect to the target database.

- Execute the RMAN script.

```
@backup.txt
```

MULTI STREAM BACKUPS FROM THIRD PARTY COMMAND LINE

The Oracle RAC iDataAgent allows you to perform parallel backups across multiple RAC nodes as a matter of design. Unlike other agents, this agent requires just one subclient and one storage policy per type used (i.e., data, log) along with the requisite number of allocated streams to perform parallel backups. As such, you can load-balance your backups across all or selected nodes of your RAC database and thereby take advantage of the strengths of each node.

Oracle third party command line operations running on multiple streams will share the same Job ID in the Job Manager. If all the streams return failure, then the job is marked as failed. However, if one of the streams fail, it is submitted to the other stream for completion.

- When you multiplex Oracle third party operations with multiple streams, each stream uses different drives by default
- When performing third party command line operations, a new `C1OraControlAgent.log` file is generated along with `ORASBT.log` to record the command line jobs
- When you use different storage policies for archive log backups and command line backups and if control file autobackup is configured in RMAN, when you submit a command line backup, the archive logs are backed up using log backup storage policy and control file autobackups use the command line storage policy.

Use the following steps to run multi stream backups from the third party command line:

- From the RMAN command prompt on the RAC client, set the number of automatic channels for a specific device type.

Note that to utilize the PARALLELISM option, set the initial parameter in pfile or spfile. Eg.,
`BACKUP_TAPE_IO_SLAVES=TRUE`

- If you are using the OEM application with multiple channels, include the RMAN settings in the Oracle Enterprise Manager.

In the below example, RMAN allocates two channels for the device type when using automatic channels.

```
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 2 BACKUP TYPE TO BACKUPSET;
```

Example:

Unix:

```
SBT_LIBRARY=<software_install_path>
/Base64/libobk.so,
BLKSIZE=262144,ENV=(
CvClientName=<rac_client_name>,
CVOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>)
```

Windows:

```
ENV=(CvClientName=<
CvClientName=<rac_client_name>,
CVOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>,
BLKSIZE=262144)
```

- Create RMAN script file to run the backup operation with a single job ID, and save it in the desired <location_path>/<file_name>. For example, `D:\backup1.txt`.

When creating the RMAN script, the `CvClientName` and `CvOraSID` parameters can be used optionally for backup jobs.

If you use both the RMAN PARALLELISM configure parameter and set multiple streams from RMAN script, the backup job will utilize double the number of streams. For example, if PARALLELISM is set to 2 and 2 streams are set from RMAN script, the backup job will utilize 4 streams.

RMAN script to run backup using two channels for node 1 and two channels for node 2.

```
run
{allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
```

Example:

```
rman target sys/sys@racdb5 << EOF
run {
allocate channel ch1 type 'sbt_tape'PARMS=
"SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac1,CVOraRacDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.devemc.Bull Calypso.com)";
allocate channel ch2 type
'sbt_tape'PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac2,CVOraRacDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.devemc.Bull Calypso.com)";
allocate channel ch3 type 'sbt_tape' connect
sys/sys@racdb6PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
hn=rac1.devemc.Bull Calypso.com)";
allocate channel ch4 type 'sbt_tape' connect
sys/sys@racdb6PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac2,CVOraRacDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.devemc.Bull Calypso.com)";
```

```

ENV=(CvClientName=<rac_client_name1>,
CvOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>");

allocate channel ch2 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name1>,
CvOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>");

allocate channel ch3 type 'sbt_tape'
connect <connect_string>
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name2>,
CvOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>");

allocate channel ch4 type 'sbt_tape'
connect <connect_string>
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name2>,
CvOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>");

setlimit channel ch1 maxopenfiles 8;
setlimit channel ch2 maxopenfiles 8;
backup
incremental level = 0
filesperset = 8
database
include current controlfile;
}
EOF

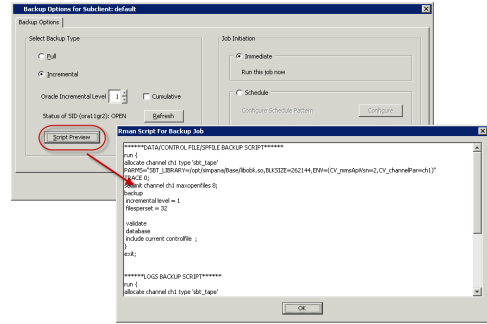
```

4. Connect to the target database. `rman target sys/sys@<databasename>`
5. Navigate to the saved location and execute the RMAN `@<file_path>backup1.txt` script.

PREVIEWING RMAN SCRIPTS FROM COMMCELL CONSOLE

Prior to running a backup operation from the CommCell Console, you can preview the corresponding RMAN script for the backup job. This is useful to determine whether the selected backup options will yield the desired result in the script. You can also manually copy and save the generated RMAN script to your computer and later execute the script from the command line.

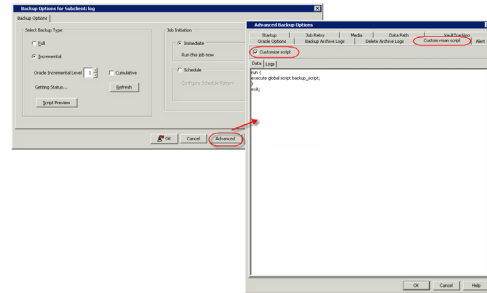
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** and click **Backup**.
3. Click **Script Preview**.
4. Click **OK**.



CUSTOMIZING RMAN SCRIPTS FROM COMMCCELL CONSOLE

In addition to previewing the RMAN script, you can also modify the script from the CommCell Console. This is useful when you want to include the RMAN commands that are not supported by the software.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** configured for archive log backup, and then click **Backup**.
3. Click **Advanced**.
4. Click the **Custom rman script** tab.
5. Select the **Customize script** check box.
6. Edit the RMAN script as required, and then click **OK**.





SCHEDULING BACKUPS

Follow the steps given below to schedule a backup:

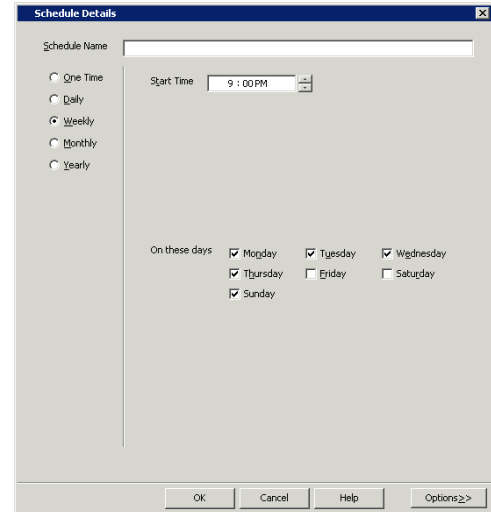
Follow the steps given below to schedule a backup:

- 1..
 - From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
 - Right-click the **<Subclient>** and click **Backup**.
2.
 - Select the **Backup type**.
 - Click **Schedule** to schedule the backup for a specific time.
 - Click **Configure** to set the schedule for the backup job. The Schedule Details dialog displays.



3. Select the appropriate scheduling options. For example:
 - Click **Weekly**.
 - Check the days you want the run the backup job.
 - Change the Start Time to 9:00 PM.
 - Click **OK** to close the Schedule Details dialog.
 - Click **OK** to close the Backup Options dialog.

The backup job will execute as per the schedule.

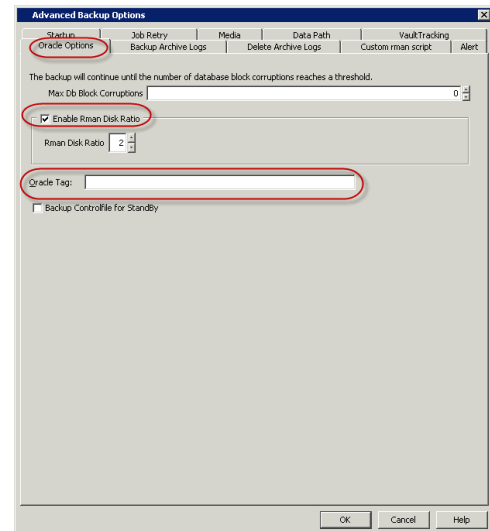


MARKING BACKUPS WITH A UNIQUE IDENTIFICATION TAG

You can configure tags for backups to uniquely identify a particular backup copy. These tags can be used later during restore operation to restore from a specific backup. Note that the tag defined at the backup job level will override the tag defined at the subclient level.

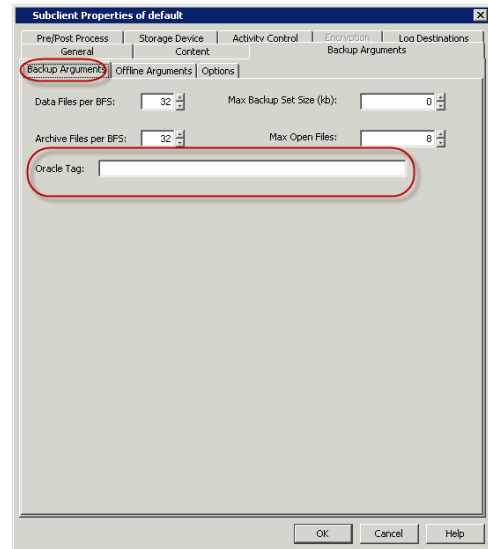
FOR A SPECIFIC BACKUP

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** and click **Backup**.
3. Click **Advanced**.
4. Click the **Oracle Options** tab.
5. In the **Oracle Tag** box, type the tag name.
For Example, weekly_backup.
6. Click **OK**.



FOR ALL BACKUPS

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab.
4. In the **Oracle Tag** box, type the tag name.
5. Click **OK**.



ENHANCING BACKUP PERFORMANCE

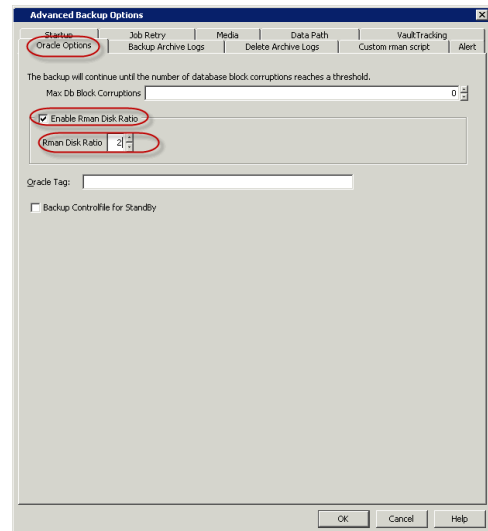
Backup performance can be increased by setting the RMAN Disk Ratio value for distributing the backup load across disks. By default, the RMAN disk ratio is not set.

For example, consider data files distributed across 10 disks that supply data at 10 bytes/second and a tape drive that needs 40 bytes/second to keep streaming. In this case, you can set the disk ratio value to 4, which will direct RMAN to include data files from 4 disks in each backup set.

Backup performance can also be enhanced by setting additional configurations at the subclient level. See Enhancing Backup Performance for more details.

Use the following steps to set the RMAN disk ratio:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>** and click **Backup**.
3. Click **Advanced**.
4. Click the **Oracle Options** tab.
5. Select the **Enable Rman Disk Ratio** check box.
6. Type the number of disks in the **Rman Disk Ratio** box.
7. Click **OK**.

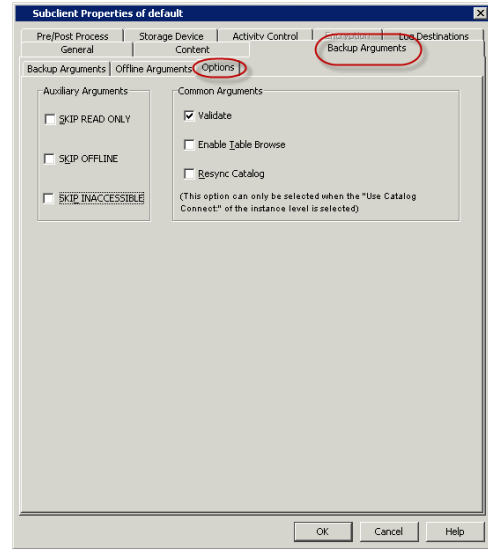


VALIDATING BACKUPS FOR RESTORE

You can validate previously run backups to ensure the integrity of the data for successful restores.

Use the following steps to validate the backup jobs of a subclient:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the **<Subclient>**, and then click **Properties**.
3. Click the **Backup Arguments** tab, and then click the **Options** tab.
4. Select the **Validate** check box.
5. Click **OK**.



MANAGING JOBS

Jobs can be managed in a number of ways. The following sections provide information on the different job management options available:

RESTARTING JOBS

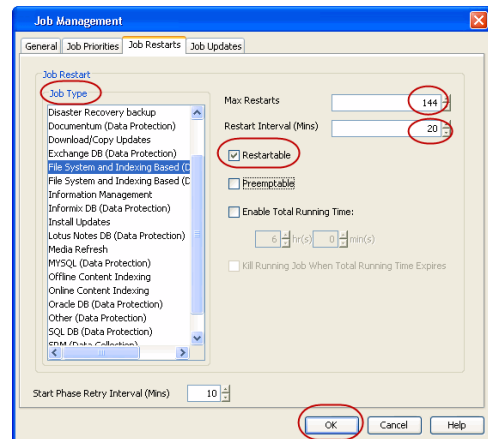
Jobs that fail to complete successfully are automatically restarted based on the job restartability configuration set in the Control Panel. Keep in mind that changes made to this configuration will affect all jobs in the entire CommCell.

To Configure the job restartability for a specific job, you can modify the retry settings for the job. This will override the setting in the Control Panel. It is also possible to override the default CommServe configuration for individual jobs by configuring retry settings when initiating the job. This configuration, however, will apply only to the specific job.

Backup jobs for this Agent are resumed from the point-of-failure.

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

1. From the CommCell Browser, click **Control Panel** icon.
2. Select **Job Management**.
3. Click **Job Restarts** tab and select a **Job Type**.
 - o Select **Restartable** to make the job restartable.
 - o Change the value for **Max Restarts** to change the maximum number of times the Job Manager will try to restart a job.
 - o Change the value for **Restart Interval (Mins)** to change the time interval between attempts for the Job Manager to restart the job.
4. Click **OK**.




CONFIGURE JOB RESTARTABILITY FOR AN INDIVIDUAL JOB

1. From the CommCell Console, navigate to **Client Computers | <RAC Client> | <Instance>**.
2. Right-click the Subclient and select **Backup**.
3. Click **Advanced**.
4. In the **Advanced Backup Options** dialog box, click the **Job Retry** tab.
5. Select **Enable Total Running Time** and specify the maximum elapsed time before a

job can be restarted or killed.

Select **Kill Running Jobs When Total Running Time Expires** to kill the job after reaching the maximum elapsed time.

6. Select **Enable Number Of Retries** and specify the number of retries.
7. Click **OK**.



CONTROLLING JOBS

The following controls are available for running jobs in the Job Controller window:

SUSPEND	Temporarily stops a job. A suspended job is not terminated; it can be restarted at a later time.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running. The status depends on the availability of resources, the state of the Operation Windows, or the Activity Control setting.
KILL	Terminates a job.

SUSPENDING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Suspend**.
2. The job status may change to **Suspend Pending** for a few moments while the operation completes. The job status then changes to **Suspended**.

RESUMING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Resume**.
2. As the Job Manager attempts to restart the job, the job status changes to **Waiting, Pending, or Running**.

KILLING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Kill**.
2. Click **Yes** when the confirmation prompt appears if you are sure you want to kill the job. The job status may change to **Kill Pending** for a few moments while the operation completes. Once completed, the job status will change to **Killed** and it will be removed from the Job Controller window after five minutes.

ADDITIONAL OPTIONS

The following table describes the available additional options to further refine your backup operations:

OPTION	DESCRIPTION	RELATED TOPICS
Startup Options	<p>The Startup Options are used by the Job Manager to set priority for resource allocation. This is useful to give higher priority to certain jobs. You can set the priority as follows:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and click Startup tab. 4. Select the Change Priority checkbox. 5. Enter a priority number - 0 is the highest priority and 999 is the lowest priority. 6. Select the Start up in suspended State check box to start the job in a suspended state. 7. Click OK. 	Refer to Job Priority and Priority Precedence.
Alerts	<p>This option enables users or user groups to get automatic notification on the status of the data protection job. Follow the steps given below to set up the criteria to raise notifications/alerts:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Alert tab. 4. Click Add Alert. 5. From the Add Alert Wizard dialog box, select the required threshold and notification criteria and click Next. 	Refer to Alerts.

	<ol style="list-style-type: none"> 6. Select the required notification types and click Next. 7. Select the recipients and click Next. 8. Click Finish. 9. Click OK. 	
Vault Tracker	<p>This feature provides the facility to manage media that is removed from a library and stored in offsite locations. Depending on your VaultTracker setup, select the required options. Use the following steps to access and select the VaultTracker options.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the VaultTracking tab. 4. Select the required options. 5. Click OK. 	Refer to VaultTracker or VaultTracker Enterprise.
Extended Data Retention	<p>This option allows you to extend the expiration date of a specific job. This will override the default retention set at the corresponding storage policy copy. Follow the steps given below to extend the expiration date:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Media tab. 4. Select one of the following options: <ul style="list-style-type: none"> o Infinite - Select this option to extend the expiration date by infinite number of days o Number of day - Select this option to specify the number of days to extend the expiration date and then enter the number of days. 5. Click OK. 	Refer to Extended Retention Rules.
Allow Other Schedules to Use Media Set	<p>The Allow Other Schedules to use Media Set option allows jobs that are part of the schedule or schedule policy and using the specific storage policy to start a new media. It also prevents other jobs from writing to the same set of media.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Media tab. 4. Select the Allow Other Schedules To Use Media Set check box. 5. Click OK. 	Refer to Creating an Exportable Media Set.
Mark Media Full	<p>This option marks the media as full, two minutes after the successful completion of the data protection job. This option prevents another job from writing to this media. Follow the steps given below:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select Media tab. 4. Select the Mark Media Full on Success check box. 5. Click OK. 	Refer to Start New Media.
Start New Media	<p>The Start New Media option enables you to start the data protection operation on a new media. This feature provides control over where the data physically resides. Use the following steps to start the data protection operation on a new media:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Media tab. 4. Select the Start New Media check box. 5. Click OK. 	Refer to Start New Media.
Data Path Options	<p>Data Protection operations use a default Library, MediaAgent, Drive Pool, and Drive as the Data Path. You can use this option to change the data path if the default data path is not available. Follow the steps given below to change the default data path:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> <Instance>. 	Refer Change Data Path.

	<ol style="list-style-type: none"> 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Data Path tab. 4. Select the MediaAgent and Library. 5. Select the Drive Pool and Drive for optical and tape libraries. 6. Click OK. 	
CommCell Readiness Report	<p>The CommCell Readiness Report provides you with vital information, such as connectivity and readiness of the Client, MediaAgent and CommServe. It is useful to run this report before performing the data protection or recovery job. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports CommServe CommCell Readiness. 3. Click the Client tab and click the Modify button. 4. In the Select Computers dialog box, clear the Include All Client Computers and All Client Groups check box. 5. Select the client from the Exclude list. 6. Click the Include > button. 7. Click OK. 8. Click the MediaAgent tab. 9. Clear the Include All MediaAgents checkbox. 10. Select the MediaAgent from the Exclude list. 11. Click Include >. 12. Click Run. <p>The generated report is displayed.</p>	Refer to CommCell Readiness Report.
Backup Job Summary Report	<p>The Backup Job Summary Report provides you with information about all the backup jobs that are run in last 24 hrs for a specific subclient. You can get information such as status, time, data size etc. for each backup job. It is useful to run this report after performing the backup. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports Jobs Job Summary. 3. Click Data Management on the General tab in the right pane. 4. Select the Computers tab. 5. Click Subclient and select the Edit tab. 6. Navigate to Client Computers <RAC Client> File System <BackupSet> <Subclient>. 7. Click Run. 	Refer to Backup Job Summary Report.

Advanced Restore - Oracle RAC iDataAgent

TABLE OF CONTENTS

Restoring and Recovering an Entire Database

- To the Same Host
- To a New Host with the Same Directory Structure
- To a New Host with a Different Directory Structure
- To a Previous Point-in-Time
- Using a Tag

Restoring Individual Datafiles/Tablespaces

- Redirecting to a New Location
- Using a Tag

Restoring Archive Logs

- All Logs
- To a Specific Log Time
- To a Specific Log Serial Number
- Using a Tag

Restoring Control File/SP File

- From the Latest Backup
- From a Specific Backup
- From a Copy of the Control File
- To a Point-in-Time
- To a New Location

Recovering a Database

- To the Current Time
- To a Point-in-Time
- To a Specific System Change Number

Creating a Duplicate Database

- On a Different Host with the Instance Configured
- On the Same Host with the Instance Configured
- Excluding Read-only Tablespaces During Restore
- Opening the Database in Restricted Mode after Restore
- Excluding Tablespaces from a Restore
- Creating Online Redo Log Files after a Restore
- Setting up the Redo Log Files
- Disabling File Name Validation during a Restore

Creating a Standby Database

- On a Different Host with the Instance Configured

Restoring Database Tables

- To the Source Database
- To a Different Database on the Same Host
- To a Different Host
- Using a User-defined Auxiliary Instance
- Disabling Clean-up of Auxiliary Instance after Restore
- Exporting Table Objects
- Selecting/De-Selecting Dependent/Referenced Tables
- Including all Dependencies to the Dependent/Referenced Tables
- Deleting Existing Tables during a Restore

Command Line Restores

- Using the Same Block Size Value
- Viewing RMAN Scripts from the CommCell Console
- Customizing RMAN Scripts from the CommCell Console
- Running RMAN Scripts from the RMAN Interface
- Restoring Multi Streamed Backups
- Running RMAN Scripts for Multiple Streams
- Restoring from a Secondary Copy using RMAN Interface

Setting the Restore Order for Streams

Automatically Switching the Database Mode before a Restore

Opening the Database after a Restore

Setting the Log State after a Restore

Setting the Database Incarnation

Resetting the Database after a Restore

RELATED TOPICS

Scheduling

Provides comprehensive information on scheduling jobs.

Job Management

Provides comprehensive information on managing jobs.

List Media

Provides comprehensive information on media management and prediction.

Disabling Failovers during Restores**Validating Restores****Setting the Database Identifier (DBID)****Setting Up Online Redo Log Files****Enhancing Restore Performance****Scheduling a Restore****Managing Restore Jobs**

- Restarting Jobs
- Controlling Jobs
- Resubmitting Jobs

Additional Restore Options**RESTORING AND RECOVERING AN ENTIRE DATABASE**


When the RAC database is corrupted or lost, you can restore and recover it from the latest offline or online full backup (depending on how the subclient was configured for backups). Make sure to bring all the RAC instances offline except on the client configured in CommCell when you restore an entire database.

TO THE SAME HOST

By default, the database is restored to the same location from where it was backed up. Once the database is restored, it is recovered to the current time.

Use the following steps to restore and recover a database to the same host:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Verify that the **Restore Data** and **Recover** options are selected.
4. Select the **Restore Control File** check box.
5. Click **OK**.

**TO A NEW HOST WITH THE SAME DIRECTORY STRUCTURE**

If the shared computer on which you hosted a RAC database is damaged or destroyed, you can restore and recover the lost database with the same directory structure on to a new host.

By default, the database is restored in the ARCHIVELOG mode. You can also choose to restore the db in NOARCHIVELOG mode.

Use the following steps to restore and recover a database to a new host with the same directory structure:

PREREQUISITES

1. Verify the following in both the source and destination computers:
 - The connection specifications (host, service name, port) in the `tnsnames.ora` file on both the source and destination computers should be different.
 - The `<username>` you use for the destination computer is different than the username for the source computer.
 - Sufficient disk space is available on the destination computer to accommodate the restored database.
 - Both the source and destination computers should have the following similar features:
 - Operating systems
 - Oracle version
 - ORACLE_SID
 - `init <SID>.ora` file
 - Processor (32-bit or 64-bit)
 - Datafile directory structures

SETTING UP THE SOURCE AND DESTINATION HOSTS

2. Create a new user account with recovery catalog owner permission within the Recovery Catalog for the destination computer. Use a different `<username>`

Example:

```
SQL>create user <username> identified by <password>
2>temporary tablespace <temp_tablespace_name>
3>default tablespace <default_tablespace_name>
4>quota unlimited on <default_tablespace_name>;
```

Statement processed.

3. Manually transfer the Oracle password file **orapw**<Oracle SID name> from the source computer to the destination computer. Usually, this file resides in ORACLE_HOME/dbs.
4. Export the recovery catalog data for the catalog user.
For example, if the user ID for the recovery catalog owner is **user1**, you need to export the database backup information for **user1**.
5. Import the recovery catalog data to the new user account for the destination computer.
6. Copy the recovery catalog's connect string entry in the `tnsnames.ora` file from the source host to the destination host.
7. Make sure that the ORACLE_SID and ORACLE_HOME are appropriately configured on the destination computer.
8. Create the Oracle RAC client and configure it as a client in the same CommServe in which the source machine resides.
9. Create and configure a new Oracle instance in the RAC Client on the destination computer, similar to the one existing in the source computer. Ensure that this instance is in NOMOUNT mode.

RESTORING THE DATABASE

10. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
11. Right-click the **<Instance>** point to **All Tasks** and then click **Restore**.
12. Select the name of the client computer from the **Destination Client** list.
13. Verify that the **Restore Data** and **Recover** options are selected.
14. Select **Restore Control File** check box.
15. Click **Advanced**.
16. Click the **Options** tab.
17. If the database is in NOARCHIVELOG mode, then select **No Redo Logs**.
18. Click **OK**.

```
SQL>grant connect, resource, recovery_catalog_owner to <username>;
```

Statement processed.

Example using IMPORT CATALOG Command::

```
RMAN>IMPORT CATALOG user1/user1@src;
```

```
<service_name> =
(DESCRIPTION =
(AADDRESS = (PROTOCOL = <protocol>)(HOST = <host>) (PORT = <##>))
(CONNECT_DATA = (SID = <Recovery Catalog database>)))
```

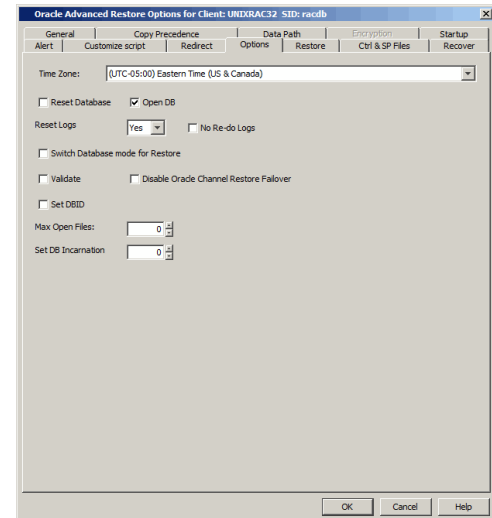
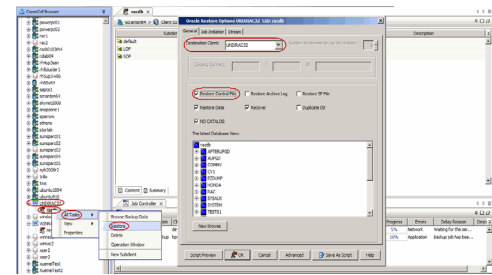
Example::

For Unix:

```
#export ORACLE_SID= <target database SID>
#export ORACLE_HOME= <Oracle home directory>
```

For Windows:

```
C:\set ORACLE_SID= <target database SID>
C:\set ORACLE_HOME= <Oracle home directory>
```



TO A NEW HOST WITH A DIFFERENT DIRECTORY STRUCTURE

If the shared computer on which you hosted a RAC database is damaged or destroyed, you can restore and recover the lost database on a new standalone host computer with a different directory structure. You can restore a database either in ARCHIVELOG or NOARCHIVELOG mode on a new host.

By default, the database is restored in the ARCHIVELOG mode, You can also choose to restore the db in NOARCHIVELOG mode.

Use the following steps to restore and recover a database to a new host with a different directory structure:

PREREQUISITES

1. Verify the following in both the source and destination computers:
 - The connection specifications (host, service name, port) in the `tnsnames.ora` file on both the source and destination computers should be different.
 - The `<username>` you use for the destination computer is different than the username for the source computer.
 - Sufficient disk space is available on the destination computer to accommodate the restored database.
 - Both the source and destination computers should have the following similar features:
 - Operating systems
 - Oracle version
 - ORACLE_SID
 - `init <SID>.ora` file
 - Processor (32-bit or 64-bit)
 - Datafile directory structures

CONFIGURING THE INIT <SID>.ORA FILE

2. Copy the `init<SID>.ora` from the old host to the new host.
3. Edit the `init<SID>.ora` file on the new host to reflect all the directory structure changes (i.e., change the path for control files, archive log destination and *dump destinations).
4. Create the directory structures as defined in `init<SID>.ora` file for all paths.

SETTING UP THE SOURCE AND DESTINATION HOSTS

5. Create a new user account with recovery catalog owner permission within the Recovery Catalog for the destination computer. Use a different `<username>`
6. Manually transfer the Oracle password file `orapw<Oracle SID name>` from the source computer to the destination computer. Usually, this file resides in `ORACLE_HOME/dbs`.
7. Export the recovery catalog data for the catalog user.
For example, if the user ID for the recovery catalog owner is `user1`, you need to export the database backup information for `user1`.
8. Import the recovery catalog data to the new user account for the destination computer.
9. Copy the recovery catalog's connect string entry in the `tnsnames.ora` file from the source host to the destination host.
10. Make sure that the `ORACLE_SID` and `ORACLE_HOME` are appropriately configured on the destination computer.
11. Create the Oracle RAC client and configure it as a client in the same CommServe in which the source machine resides.

Example:

```
SQL>create user <username> identified by <password>
2>temporary tablespace <temp_tablespace_name>
3>default tablespace <default_tablespace_name>
4>quota unlimited on <default_tablespace_name>;
```

Statement processed.

```
SQL>grant connect, resource, recovery_catalog_owner to
<username>;
```

Statement processed.

Example using IMPORT CATALOG Command:

```
RMAN>IMPORT CATALOG user1/user1@src;
```

```
<service_name> =
(DESCRIPTION =
(AADDRESS = (PROTOCOL = <protocol>)(HOST = <host>)(PORT
= <##>))
(CONNECT_DATA = (SID = <Recovery Catalog database>)))
```

Example:

For Unix:

```
#export ORACLE_SID= <target database SID>
#export ORACLE_HOME= <Oracle home directory>
```

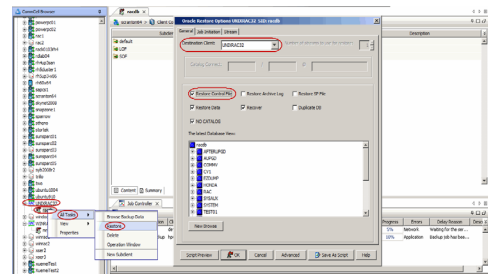
For Windows:

```
C:\set ORACLE_SID= <target database SID>
C:\set ORACLE_HOME= <Oracle home directory>
```

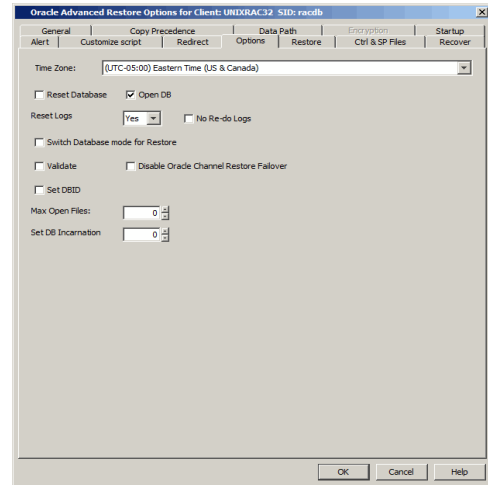
12. Create and configure a new Oracle instance in the RAC Client on the destination computer, similar to the one existing in the source computer. Ensure that this instance is in NOMOUNT mode.

RESTORING THE DATABASE

13. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
14. Right-click the **<Instance>** point to **All Tasks** and then click **Restore**.
15. Select the name of the client computer from the **Destination Client** list.
16. Select **Restore Control File** check box.
17. Click **Advanced**.



18. Click the **Options** tab.
19. If the database is in NOARCHIVELOG mode, then select **No Redo Logs**.
20. Click **OK**.



TO A PREVIOUS POINT-IN-TIME

The point-in-time restore is useful in the following scenarios:


- If any undesired transaction occurs in the database, you can revert the database to a state just before the transaction.
- If a database fails, you can restore to the state just before the point of failure.

When you restore and recover an entire database to a previous point-in-time from an online backup or offline backup (depending on how the subclient was configured for backups) to the original host, it is recommended to use the control files.

When you perform a point-in-time restore for a database, the next scheduled backup for that database will automatically convert to a full backup.

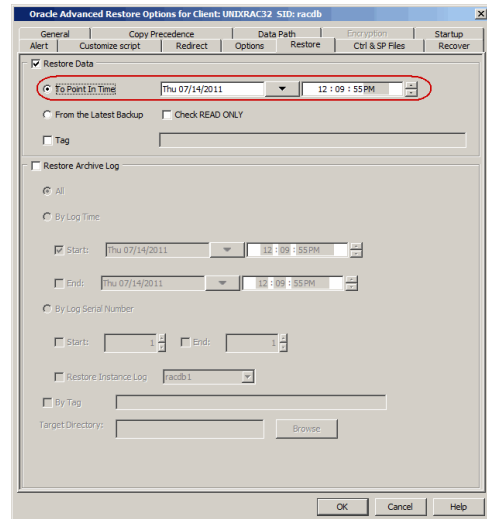
Use the following steps to restore and recover a database to a previous point-in-time:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select **Restore Control File** check box, if you want to restore the control file(s).
4. Click **Advanced**.

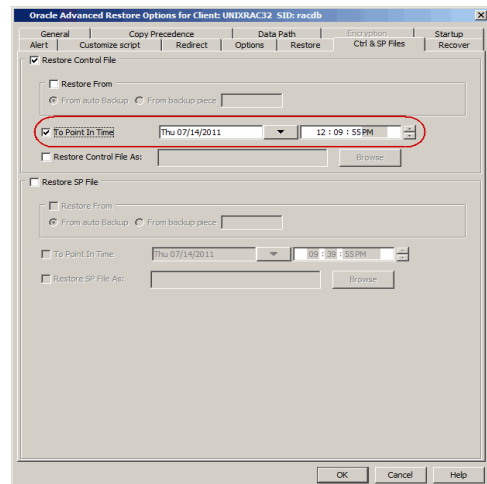
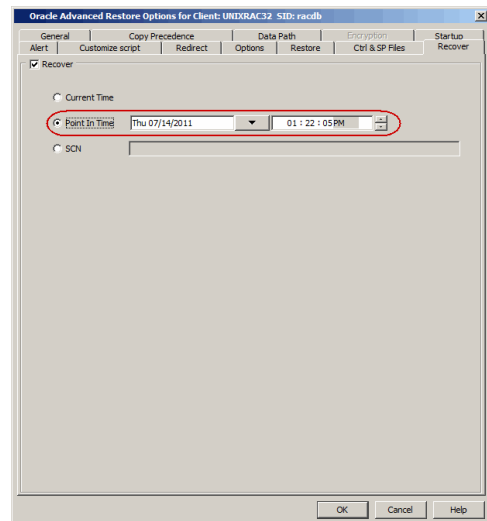


5. Click the **Restore** tab.
6. Click **To Point-In-Time** and select the data and time.

7. Click the **Recover** tab.
8. Click the **Point-In-Time** and select the data and time.



9. Click **Ctrl&SP Files** Tab, if you have selected to restore the control file(s).
10. Click the **Point-In-Time** and select the data and time.
You must restore the control files to a point-in-time later than or equal to the point-in-time set in the **Restore** tab.
11. Click **OK**.

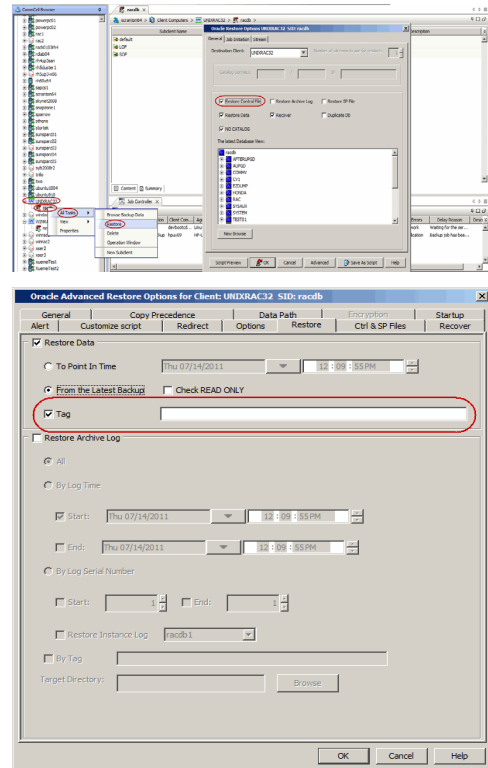


USING A TAG

If you have assigned unique identification tags for the data, you can restore and recover an entire database from a specific backup using the tags. Use the following steps to restore the database with specific identification tags:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Oracle**.

- Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
- Select **Restore Control File** check box.
- Click **Advanced**.
- Click the **Restore** tab.
- Click **By Tag** and type the Tag name that is assigned for a specific data backup that you want to restore.
- Click **OK**.

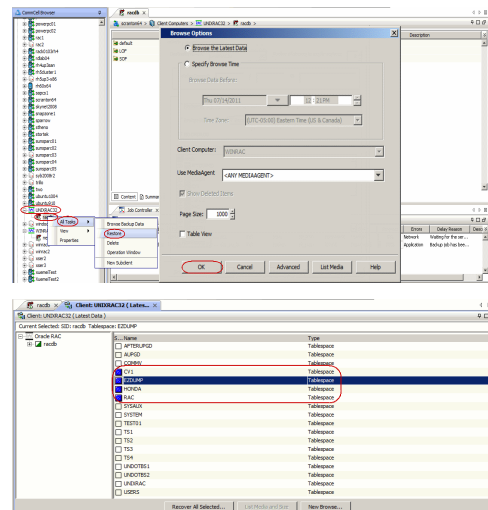


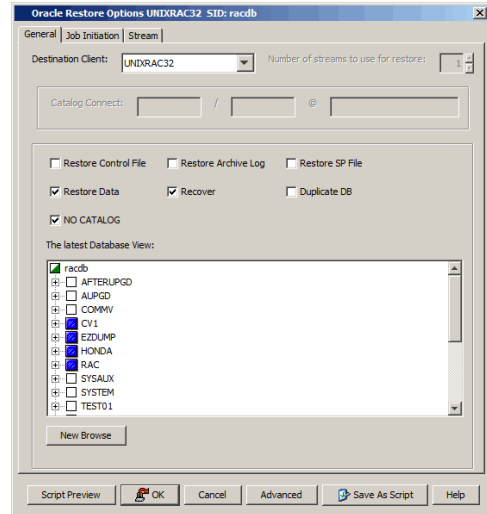
RESTORING INDIVIDUAL DATAFILES/TABLESPACES

In addition to restoring a RAC database, you can also restore specific tablespaces or datafiles that were lost due to an error or corruption. By default, the selected tablespaces/datafiles are restored to the original location from the latest online backup.

Use the following steps to restore the datafile(s) or tablespace(s):

- From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
- Right-click the **<Instance>**, point to **All Tasks** and then click **Browse Backup Data**.
- Click **OK**.
- In the right pane of the Browse window, select the datafiles or tablespaces you want to restore and click **Recover All Selected**.
- Click **OK**.

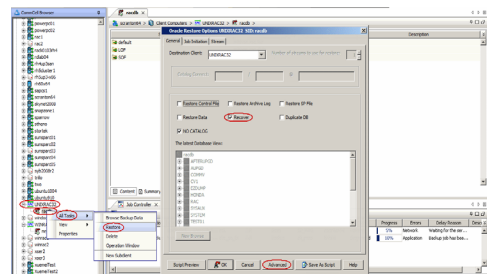
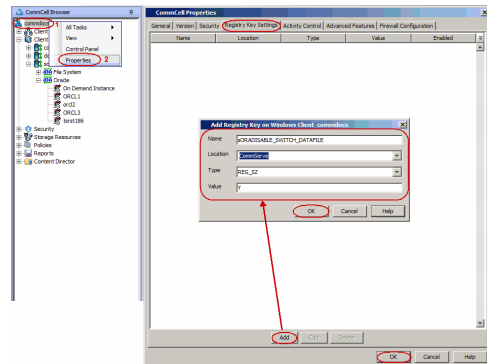


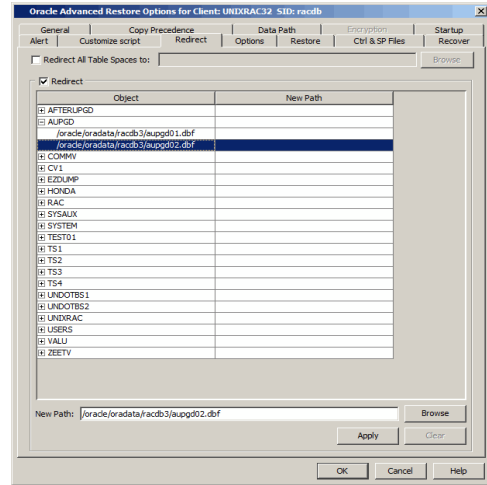


REDIRECTING TO A NEW LOCATION

By default, the control file will be automatically updated with the new location when you redirect tablespaces/datafiles to a new location. Use the following steps to redirect the tablespaces/datafiles to a new location without updating the control file:

1. From the CommCell Browser, right-click the <CommServe> and then click **Properties**.
2. Click the **Registry Key Settings** tab.
3. Click **Add**.
4. In the **Name** field, type sORADISABLE_SWITCH_DATAFILE.
5. In the **Location** list, select CommServe from the list.
6. In the **Type** list, select **String**.
7. In the **Value** field, type **x**.
8. Click **OK**.
9. From the CommCell Browser, navigate to **Client Computers | <Client> | Oracle**.
10. Right-click the <Instance>, point to **All Tasks** and then click **Restore**.
11. Clear the **Recover** check box.
12. Click **Advanced**.
13. Click the **Redirect** tab.
14. Select the **Redirect** check box.
15. Select each object and type the new path in **New Path** box or click **Browse** to locate the new path.
16. Click **Apply** to confirm the new path.
17. Select the **Redirect All Table Spaces** to redirect all the objects displayed in the object column to a new location.
18. Click **OK**.

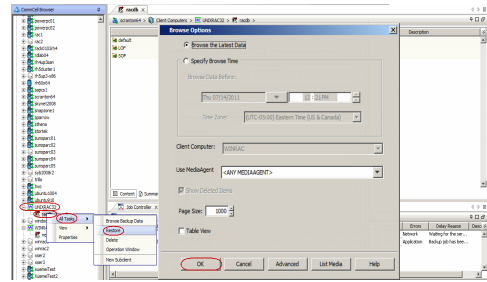





USING A TAG

If you have assigned unique identification tags during backups, you can restore from a specific backup using the tag. Use the following steps to restore the datafile(s) or tablespace(s) using a specific tag:

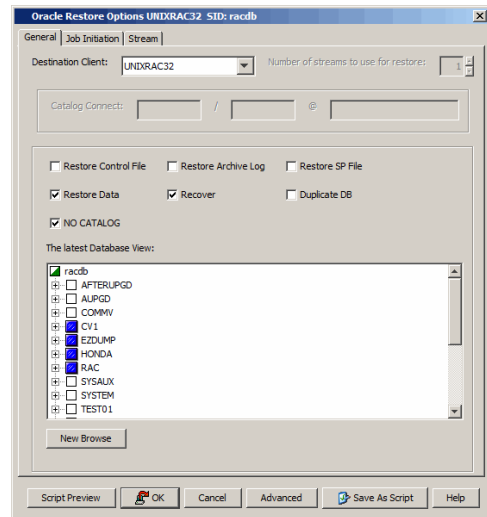
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Browse Backup Data**.
3. Click **OK**.



4. In the right pane of the Browse window, select the datafiles or tablespaces you want to restore and click **Recover All Selected**.



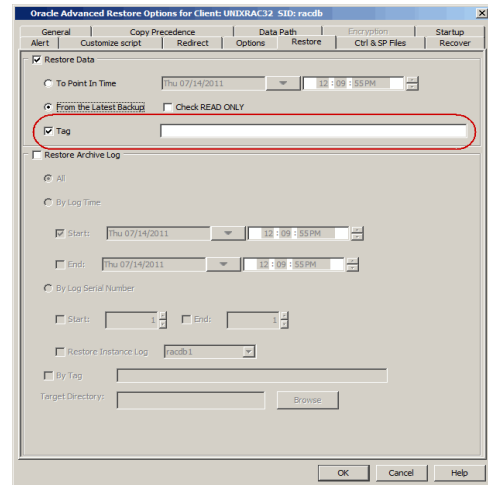
5. Click **Advanced**.



6. Click the **Restore** tab.
7. Select the **Tag** check box and type the Oracle Tag name that is assigned for a specific

backup that you want to restore.

8. Click **OK**.



RESTORING ARCHIVE LOGS

Archive logs can be restored separately or along with the database. Archive Log restores are useful in the following scenarios:

- If there is a database failure and you need to recover the database to the recent state, you will restore all the logs along with the database.
- If the logs from a specific time range were lost due to a hard disk corruption, you can restore them by performing a point-in-time restore of the logs.
- In certain cases, you might need to restore only specific logs that are missing in the database. Such logs can be identified and then restored using a serial number or identification tag.

While restoring archive logs, ensure that the archive logs are shareable (i.e., accessible from each node). Perform either of the following to enable sharing:

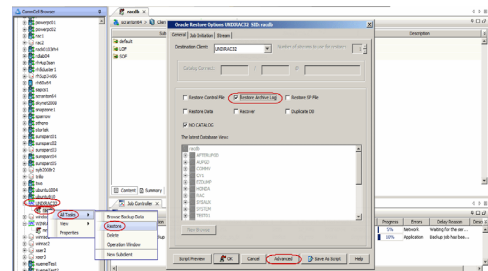
- Mount an NFS drive on all the nodes
- Copy the archive files to all the nodes

If the archive logs are located only on the local disk, the archive logs will be accessible only from the corresponding node. The log file restores initiated by any other node will fail.

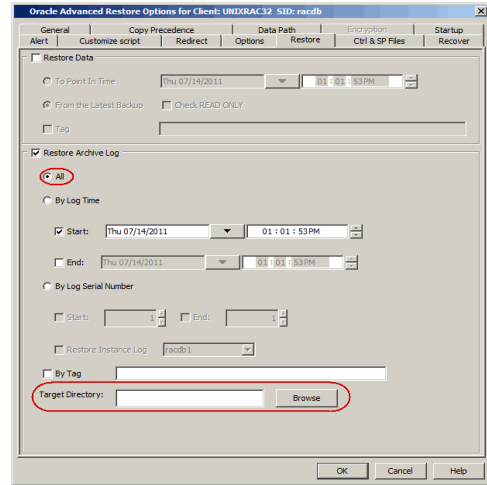
ALL LOGS

Use the following steps to restore all the archived logs (note that this is the default option):

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the **Restore Archive Log** check box.
4. Click **Advanced**.



5. Click the **Restore** tab.
6. In the **Target Directory** box, type the path or click **Browse** to specify the path to restore all the logs.
7. Click **OK**.

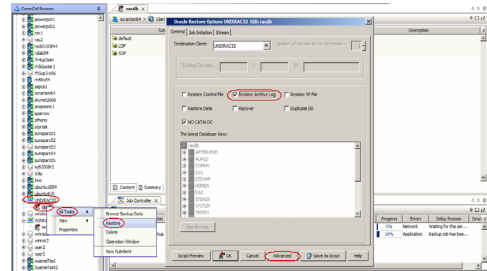


TO A SPECIFIC LOG TIME

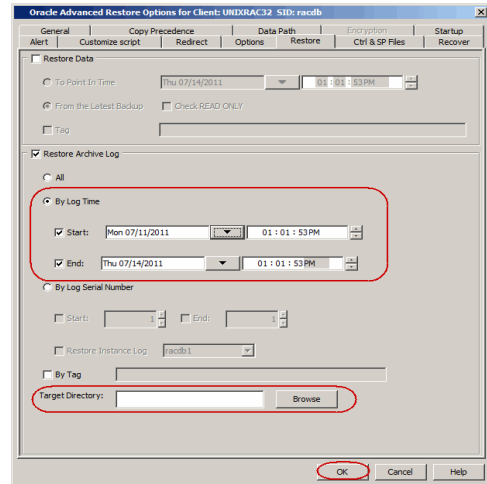
When you browse between a specific point of time range, the logs pertaining to all the cycles within the specified time range will be listed. When restoring the archive logs based on the log time, if the data is also included in the restore, ensure that the point-in-time range for the restore is the same for both the data and logs.

If you are including the database in the restore, see point-in-time restore to restore the database to a specific point-in-time. Use the following steps to restore the logs to a specific log time:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the **Restore Archive Log** check box.
4. Click **Advanced**.



5. Click the **Restore** tab.
6. Click **By Log Time** and specify the point-in-time (date and time) restore of archived log files.
7. In the **Target Directory** box, type the path or click **Browse** to specify the path to restore all the logs.
8. Click **OK**.



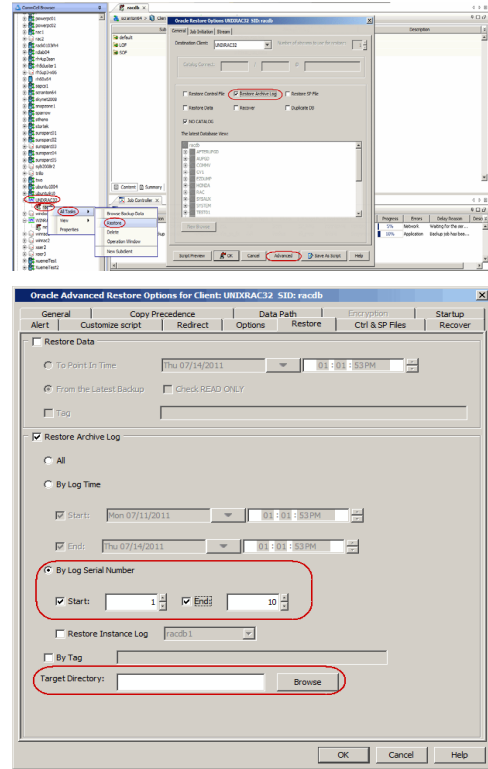
TO A SPECIFIC LOG SERIAL NUMBER

Note the following when you restore the logs based on the log serial number:

- The end serial number must be greater than or equal to the start serial number. For example, S.No 1 to 15.
- The start and the end serial number cannot be 0.
- The range of serial numbers to be restored could be from one or more backup jobs or from a part of a job and from one or more cycles.
- If only the start sequence number is provided then the logs starting from that sequence number will be restored.

Use the following steps to restore the log identified by a serial number:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select **Restore Archive Log** check box.
4. Click **Advanced**.
5. Click the **Restore** tab.
6. Click **By Log Serial Number** and type or select the **Start** and **End** log serial numbers.
7. In the **Target Directory** box, type the path or click **Browse** to specify the path to restore all the logs.
8. Click **OK**.

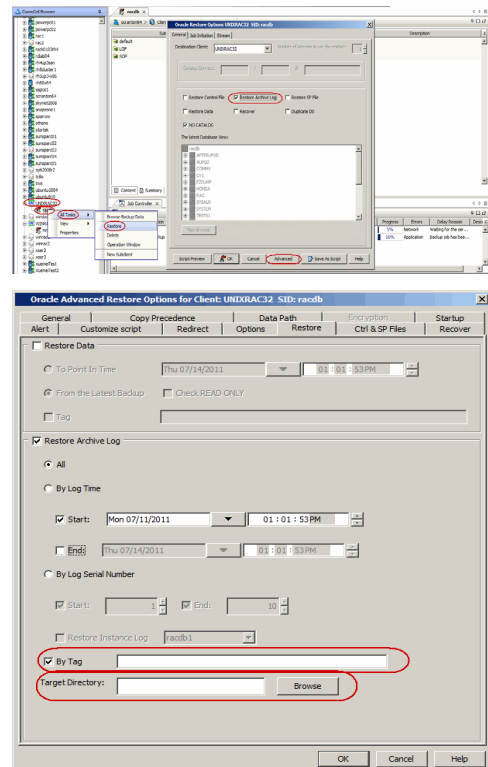


USING A TAG

If you have assigned unique identification tags during log backups, you can restore the logs using the tags.

Use the following steps to restore the logs with specific identification tags:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select **Restore Archive Log** check box.
4. Click **Advanced**.
5. Click the **Restore** tab.
6. Click **By Tag** and type the Tag name that is assigned for a specific log backup that you want to restore.
7. In the **Target Directory** box, type the path or click **Browse** to specify the path to restore all the logs.
8. Click **OK**.



RESTORING CONTROL FILE/SP FILE

Control and SP files are required to recover a database to the current state. Restoring a control/sp file is useful in the following scenarios:

- If you want to restore the backup repository contained in the control file when the Control file is lost.
- If the recovery catalog is lost.
- If the recovery catalog was never used.
- If the catalog connect string is not specified for the instance during the backup.

Ensure that the database is in NOMOUNT mode when you restore the control/sp files. Ensure that you have previously configured auto backup of control files to restore the control file from auto backup. Restoring a control file will destroy all the previous backups. Hence, you need to perform a full backup after you restore a control file.

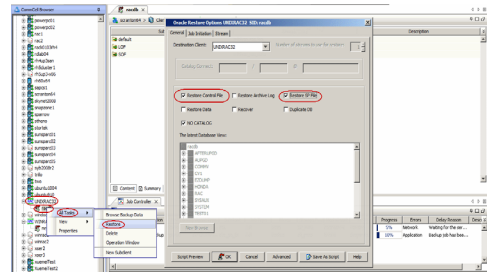
The database will be in MOUNT mode after you restore the control/SP file.

When you restore control files, be sure to select only one node for the restore. During control file restores, Oracle puts the database in a state where other processes may not be able to access the database. Therefore, if you select more than one node, the job may fail if the database is inaccessible by all the nodes. Make sure to bring all the RAC instances offline except on the client configured in CommCell when you restore a Control/SP file.

FROM THE LATEST BACKUP

By default, the Control/SP files are restored from the latest backup. Use the following steps to restore a Control/SP file from a latest backup:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select **Restore Control File** and **Restore SP File** check boxes.
4. Click **OK**.

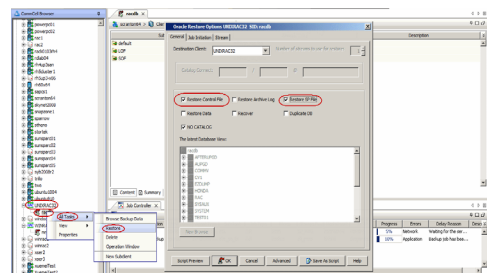


FROM A SPECIFIC BACKUP

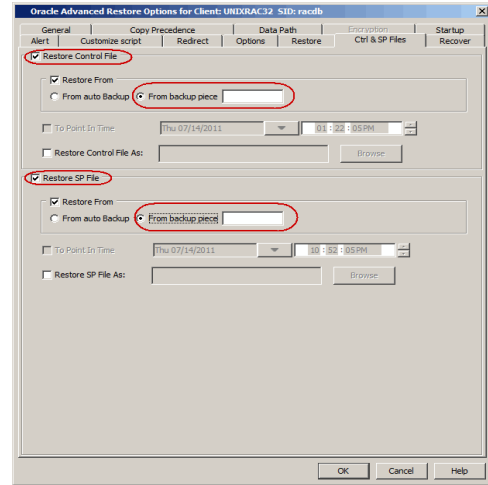
If the control file and recovery catalog are lost, you can restore the control file from a specific backup piece using a backup piece number. (Backup piece references the backup of one or several database files)

You can obtain the backup piece value from the RMAN logs of the backup job. Use the following steps to restore a control/sp file from a specific backup:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the **Restore Control File** and **Restore SP File** check boxes.
4. Click **Advanced**.



5. Click the **Ctrl&SP Files** tab.
6. Under the Restore Control File, select **Restore From** check box.
7. Click **From Backup Piece** box and type the Backup Piece value.
8. Under the Restore SP File, select **Restore From** check box.
9. Click **From Backup Piece** box and type the Backup Piece value.
10. Click **OK**.



FROM A COPY OF THE CONTROL FILE

If your database consists of multiple copies of control files, you can replace a corrupted control file by manually copying from an existing control file and restoring it.

Follow the steps given below to restore a control file from an existing control files:

1. Shutdown the database.

```
SQL> shutdown
```

2. Manually copy one of the existing control files to the missing control file location.
3. Assign the same owner, group and file permissions to the new control file as that of the original missing control file.
4. Startup the database.

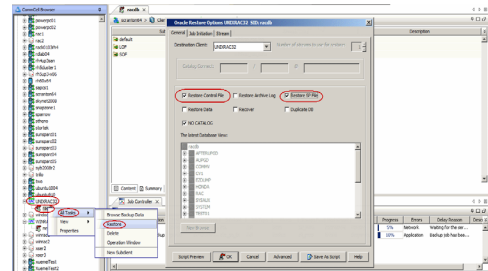
```
SQL> startup nomount;
```

TO A POINT-IN-TIME

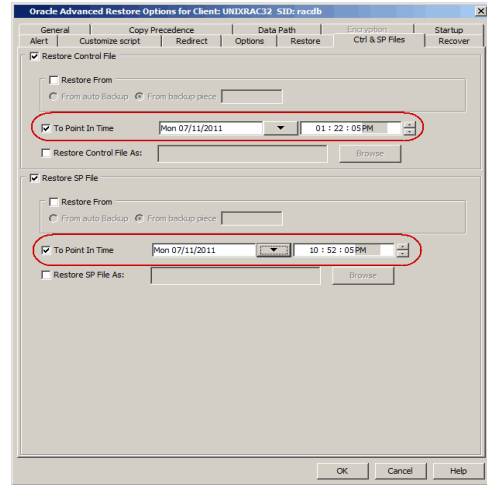
If you are certain that the control file was corrupted, lost at or after a specific time, you can restore the control file to such a point-in-time.

Follow the steps given below to restore a control/sp file to a point-in-time:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the **Restore Control File** and **Restore SP File** check boxes.
4. Click **Advanced**.



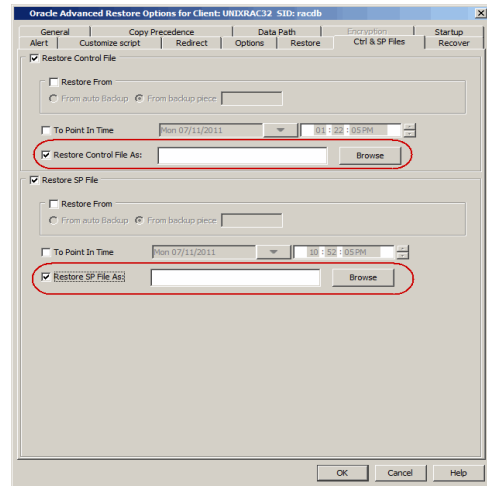
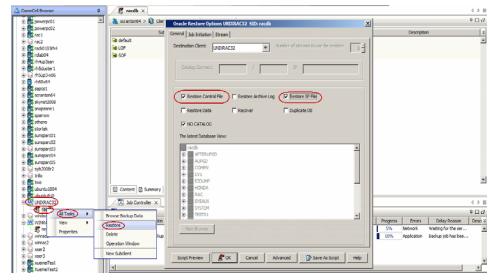
5. Click **Ctrl&SP Files** tab.
6. Under the Restore Control File, select the **To Point-In-Time** check box and select the start and end time.
7. Under the Restore SP File, select the **To Point-In-Time** check box and select the start and end time.
8. Click **OK**.



TO A NEW LOCATION

By default, the Control file/SP File is restored to the original location. Use the following steps to restore the control files to a new location:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the **Restore Control File** and **Restore SP File** check boxes.
4. Click **Advanced**.
5. Click the **Ctrl&SP Files** tab.
6. Select the **Restore Control File As** check box.
7. Type the path or click **Browse** to specify the new location to restore the control file.
8. Select the **Restore SP File As** check box.
9. Type the path or click **Browse** to specify the new location to restore the SP file.
10. Click **OK**.



RECOVERING A DATABASE

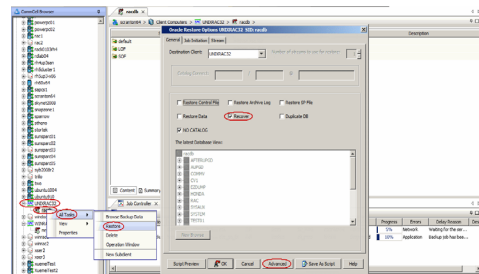
By default, the database is recovered along with the restore. However, you can also restore the data and then recover the database at a later point-in-time.

RECOVERING A DATABASE TO THE CURRENT TIME

You can recover a database to the current time either to the original host or to a different host. Use the following steps to recover a database to the current time:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Clear the **Restore Data** check box.

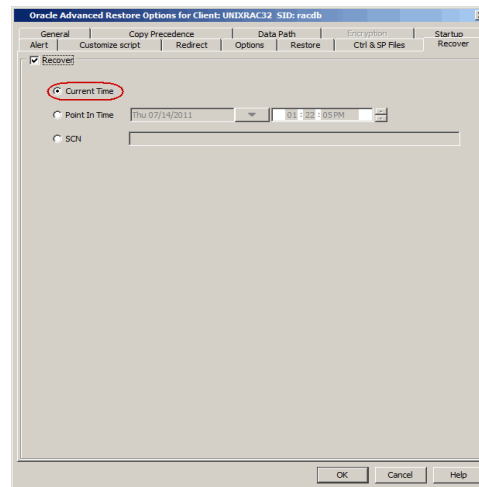
4. Click **Advanced**.



5. Click the **Options** tab.

6. In the **Reset Logs** box, select **Yes**.

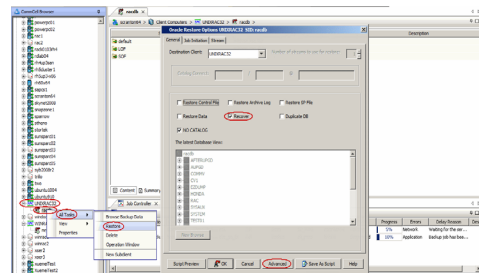
7. Click **OK**.



RECOVERING A DATABASE TO A POINT-IN-TIME

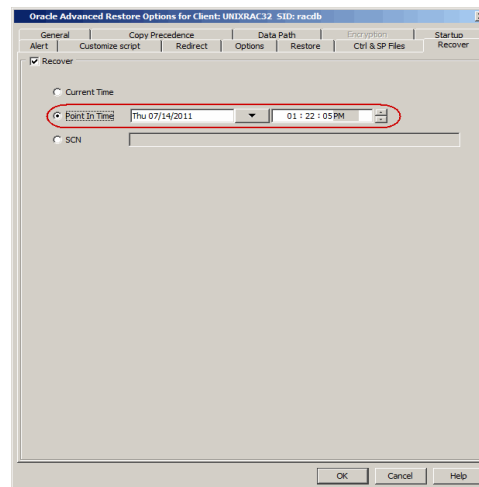
You can apply archived logs and recover a database to a previous point-in-time where it is consistent and stable. Use the following steps to recover a database to a point-in-time:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Clear the **Restore Data** check box.
4. Click **Advanced**.



5. Click **Recover** tab.

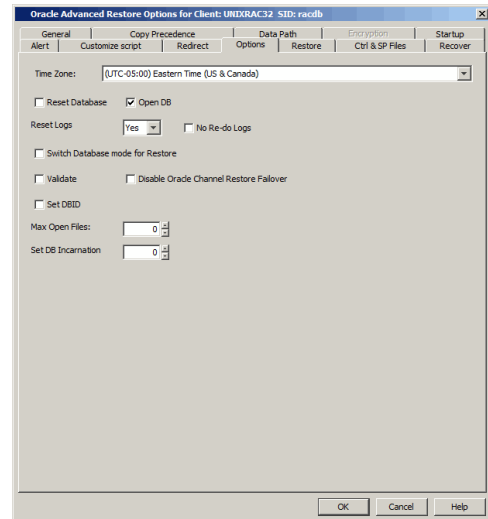
6. Click **Point-In-Time** and select the date and time.



7. Click the **Options** tab.

8. In the **Reset Logs** box, select **Yes**.

- Click **OK**.

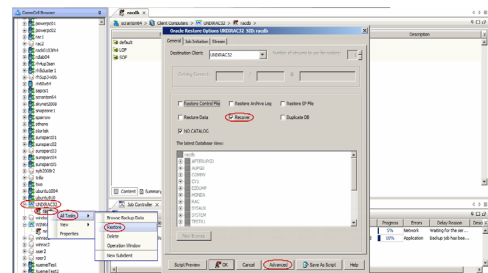


RECOVERING A DATABASE USING THE SYSTEM CHANGE NUMBER

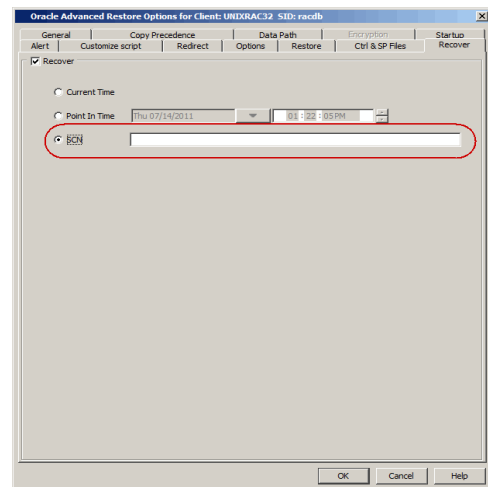
The System Change Number (SCN) keeps track of the timing of transactions in the oracle database. The SCN's are stored in the control files and the datafile headers. You can recover the database to the last existing SCN number in the control file. (The last SCN number denotes the last consistent state of the database.)

Use the following steps to recover a database using SCN:


- From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
- Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
- Clear the **Restore Data** check box.
- Click **Advanced**.



- Click **Recover** tab.
- Click **SCN** and type the System Change Number.



- Click the **Options** tab.
- In the **Reset Logs** box, select **Yes**.
- Click **OK**.



CREATING A DUPLICATE DATABASE

A duplicate (auxiliary) database is a copy or subset of the target database and has a unique DBID. It is independent of the primary database and can be registered in the same recovery catalog as the primary database. The duplicate database will be useful for testing and demo purposes.

If the duplicate (auxiliary) database already exists in the destination computer, it will be overwritten. Duplicate database is created from the full backup of the database with the logs. If you want the latest data in the duplicate database, make sure to perform a full backup with the log files before creating the duplicate database.

ON A DIFFERENT HOST WITH THE INSTANCE CONFIGURED

Use the following steps to create a duplicate database on a different host with the configured instance. Make sure that the instance is configured on a different host in the CommCell Console.

1. Perform a full backup along with the log files on the original database.
2. On the destination host, make sure to remove the temp.dbf file from the existing database instance.
3. Manually, copy the init<SID>.ora file from the source computer to the destination computer.
4. Update the database name, dump files, archive logs and the control file locations in the init<SID>.ora file for the duplicate database instance.
5. Add the DB_FILE_NAME_CONVERT and LOG_FILE_NAME_CONVERT parameters in the init<SID>.ora file. These parameters will redirect the datafiles, temp files, and log files to the auxiliary instance.

Make sure that all the other parameters in the init<SID>.ora file are same as that in the original database.

On Unix:

```
$ORACLE_HOME
```

On Windows:

```
%ORACLE_HOME%
```

On Unix:

```
DB_FILE_NAME_CONVERT=
(source_of_df_path/,dup_of_df_path/,
source_of_temp_path/,dup_of_temp_path/,...)
LOG_FILE_NAME_CONVERT=
(source_of_log_path/redo,dup_of_log_path/redo)
```

On Windows:

```
DB_FILE_NAME_CONVERT=
('source_of_df_path/','dup_of_df_path/',
'source_of_temp_path/','dup_of_temp_path/',...)
LOG_FILE_NAME_CONVERT=
('source_of_log_path/redo','dup_of_log_path/redo')
```

(When using these parameters on a Windows computer, the file paths should be entered in uppercase.)

6. On Windows clients, restart Oracle services.

Skip this step, if you are using an Unix client.

7. Add the duplicate database instance name in the Listener.ora file on the destination host and add TNS entry on Tnsnames.ora files on the source and destination hosts.
8. Add the source database name in the Tnsnames.ora file on the destination host.
9. Restart the Listener.
10. Provide a valid connect string for the auxiliary channel.
11. Startup the duplicate database instance in NOMOUNT mode.

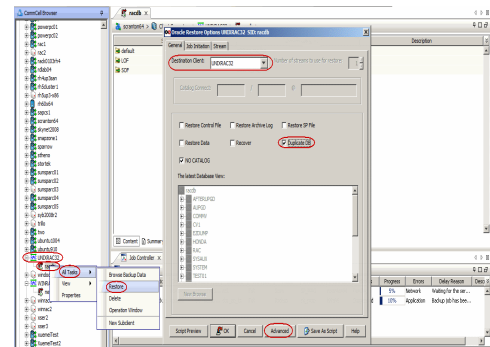
```
DUPDB =
DESCRIPTION =
ADDRESS = (PROTOCOL = TCP) (HOST = powerpc02) (PORT =
1521)
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dupdb)
)
)
```

\$ lsnrctl reload

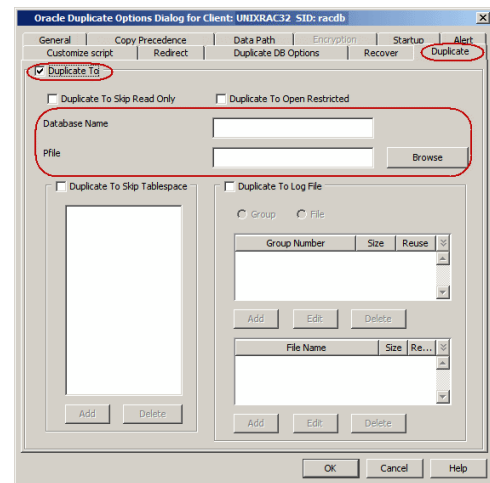
Example:

```
sys/sys@<SID name>
```

12. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
13. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
14. Select **Duplicate DB** check box.
15. Click **Advanced**.



16. Click **Duplicate** tab.
17. Click **Duplicate To**.
18. Type the name of duplicate database in **Database Name** box.
19. Click **OK**.
20. Select the name of the client computer from the **Destination Client** list.
21. Click **Advanced**.
22. Type the name of startup Parameter file in **Pfile** box or click **Browse** to locate it.
23. Click **OK**.
24. Under **The latest Database View** box, click **Refresh** to verify that the status of the duplicate database is **STARTED**.
25. Click **OK**.



ON THE SAME HOST WITH THE INSTANCE CONFIGURED

Use the following steps to create a duplicate database on the same host with the configured instance:

1. Perform a full backup along with the log files on the original database.
2. Create a duplicate database instance on the destination host. If the database already exists on the destination host, make sure to remove the temp.dbf file before performing a restore operation.
3. Manually, copy the init<SID>.ora file from the source computer to the destination computer.

On Unix:

```
$ORACLE_HOME
```

On Windows:

```
%ORACLE_HOME%
```

4. Update the database name, dump files, archive logs and the control file locations in the `init<SID>.ora` file for the duplicate database instance.
5. Add the `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT` parameters in the `init<SID>.ora` file. These parameters will redirect the datafiles, temp files, and log files to the auxiliary instance.

Make sure that all the other parameters in the `init<SID>.ora` file are same as that in the original database.

On Unix:

```
DB_FILE_NAME_CONVERT=
(source_of_df_path/,dup_of_df_path/,
source_of_temp_path/,dup_of_temp_path/,...)
LOG_FILE_NAME_CONVERT=
(source_of_log_path/redo,dup_of_log_path/redo)
```

On Windows:

```
DB_FILE_NAME_CONVERT=
('source_of_df_path/', 'dup_of_df_path/',
'source_of_temp_path/', 'dup_of_temp_path/',...)
LOG_FILE_NAME_CONVERT=
('source_of_log_path/redo', 'dup_of_log_path/redo')
```

(When using these parameters on a Windows computer, the file paths should be entered in uppercase.)

6. On Windows clients, restart Oracle services.
Skip this step, if you are using an Unix client.
7. Add the duplicate database instance name in the `Listener.ora` file and `Tnsnames.ora` files.
8. Add the source database name in the `Tnsnames.ora` file on the destination host.
Make sure to use the password `change_on_install` when you create the password file for auxiliary database.

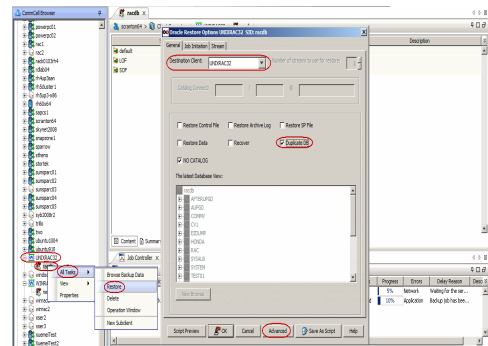
```
DUPDB =
DESCRIPTION =
ADDRESS = (PROTOCOL = TCP) (HOST = powerpc02) (PORT = 1521)
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dupdb)
)
)
```

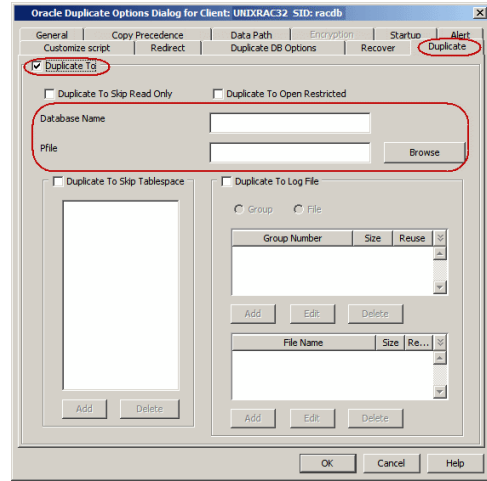
9. Restart the Listener.
10. Provide a valid connect string for the auxiliary channel.
11. Startup the duplicate database instance in NOMOUNT mode.
12. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
13. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
14. Select **Duplicate DB** check box.
15. Under **The latest Database View** box, click **Refresh** to verify that the status of the duplicate database is **STARTED**.
16. Click **Advanced**.

Example:

`sys/sys@<SID name>`

17. Click **Duplicate** tab.
18. Click **Duplicate To**.
19. Type the name of duplicate database in **Database Name** box.
20. Type the name of startup Parameter file in **Pfile** box or click **Browse** to locate it.
21. Click **OK**.



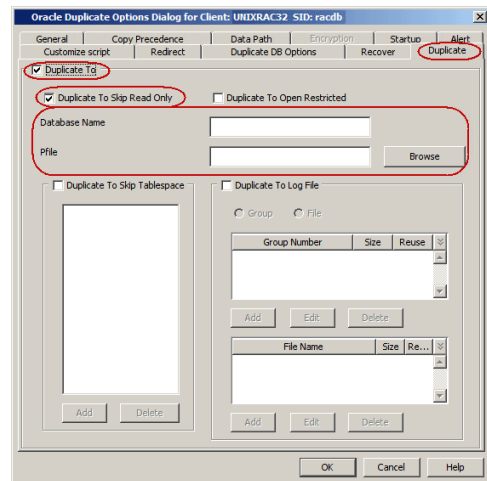
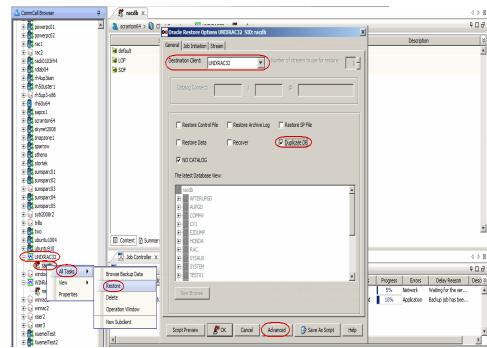


EXCLUDING READ-ONLY TABLESPACES DURING RESTORE

By default, the read only tablespaces are not verified for consistency and are restored from the backup. You can skip the consistent tablespaces and restore the tablespaces that are not consistent or missing. This will save the time taken for the restore.

Use the following steps to exclude read only table spaces during restore:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the name of the client computer from the **Destination Client** list.
4. Select **Duplicate DB** check box.
5. Click **Advanced**.
6. Click **Duplicate** tab.
7. Click **Duplicate To**.
8. Type the name of duplicate database in **Database Name** box.
9. Type the name of startup Parameter file in **Pfile** box or click Browse to locate it.
10. Select **Duplicate To Skip Read Only** check box.
11. Click **OK**.




OPENING THE DATABASE IN RESTRICTED MODE AFTER A RESTORE

By default, a duplicated database is opened without any restricted access. If necessary, you can open the db in restricted mode for administrative tasks. This will restrict access to other users.

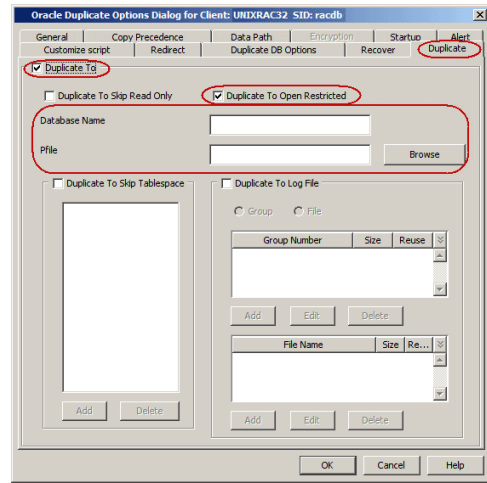
Use the following steps to open the duplicate database in restricted mode:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**

2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the name of the client computer from the **Destination Client** list.
4. Select **Duplicate DB** check box.
5. Click **Advanced**.




6. Click **Duplicate** tab.
7. Click **Duplicate To**.
8. Type the name of duplicate database in **Database Name** box.
9. Type the name of startup Parameter file in **Pfile** box or click Browse to locate it.
10. Select **Duplicate To Open Restricted** check box.
11. Click **OK**.



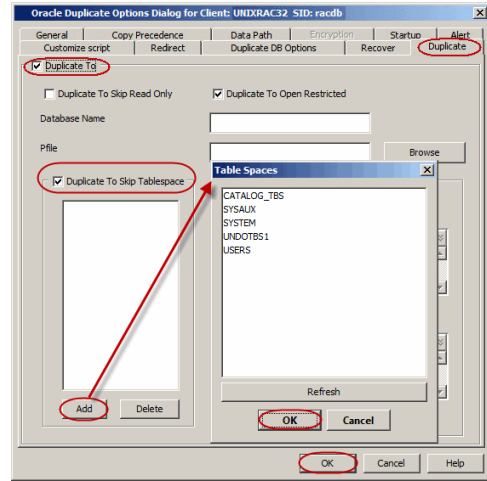
EXCLUDING TABLESPACES FROM A RESTORE

While creating a duplicate database, you can exclude some tablespaces from the duplicate database. Use the following steps to exclude the tablespaces from the duplicate database:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the name of the client computer from the **Destination Client** list.
4. Select **Duplicate DB** check box.
5. Click **Advanced**.



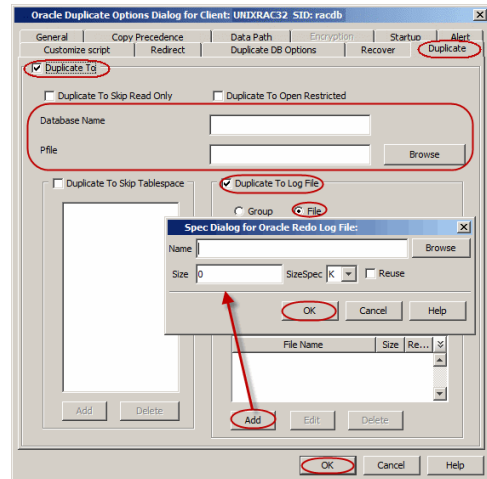
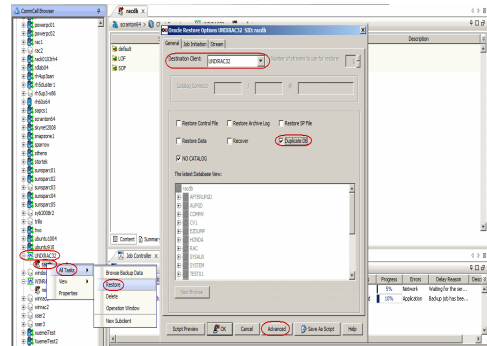
6. Click **Duplicate** tab.
7. Click **Duplicate To**.
8. Type the name of duplicate database in **Database Name** box.
9. Type the name of startup Parameter file in **Pfile** box or click Browse to locate it.
10. Select the **Duplicate To Skip TableSpaces** check box.
11. Click **Add**
12. Select the tablespaces that appear in the **TableSpaces** box to exclude from the duplicate database.
13. Click **OK**.



SETTING UP THE REDO LOG FILES

You can create online redo logs for duplicate database and apply them to restore the database in case of corruption. Use the following steps to create an online redo log file:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the name of the client computer from the **Destination Client** list.
4. Select **Duplicate DB** check box.
5. Click **Advanced**.
6. Click **Duplicate** tab.
7. Click **Duplicate To**.
8. Type the name of duplicate database in **Database Name** box.
9. Type the name of startup Parameter file in **Pfile** box or click **Browse** to locate it.
10. Select **Duplicate To Log File** check box.
11. Click **File** to select a file containing the online redo log.
12. Click **Add** to include the specifications for an online redo log file.
13. In the **Spec Dialog for Oracle Redo Log File** box, type the name or click **Browse** to select the redo log file.
14. Type the **Size** of the online redo log file.
15. Select the Size Specifications of the file from **SizeSpec** list.
16. Select **Reuse** check box to allow the database to reuse an existing file.
17. Click **OK**.

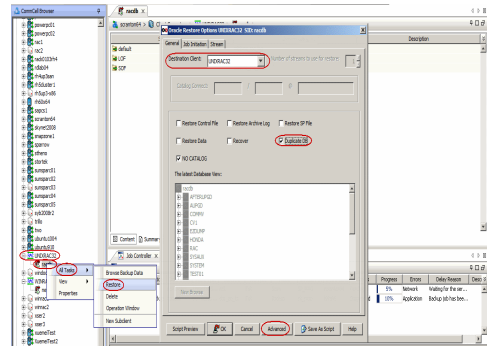


TO ADD A REDO LOG FILE FOR A GROUP

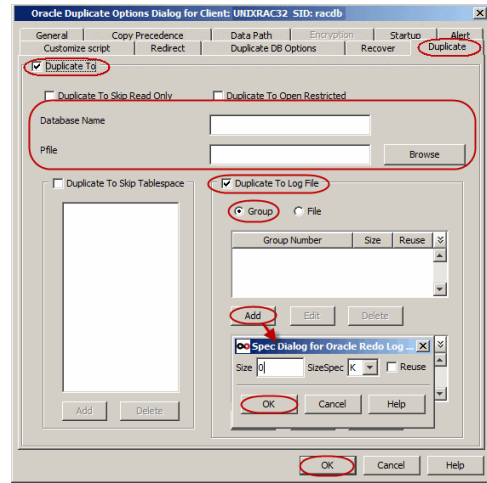
By default, groups are created to include specific online redo log members. Use the following steps to add specifications for each of these online redo log groups:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the name of the client computer from the **Destination Client** list.
4. Select **Duplicate DB** check box.

5. Click **Advanced**.




6. Click **Duplicate** tab.
7. Click **Duplicate To**.
8. Type the name of duplicate database in **Database Name** box.
9. Type the name of startup Parameter file in **Pfile** box or click Browse to locate it.
10. Select **Duplicate To Log File** check box.
11. Click **Group** to select a group containing the online redo log members.
12. Click **Add**.
13. In the **Spec Dialog for Oracle Redo Log Group** box, add the **Size** of the Group.
14. Select the Size Specifications from **SizeSpec** list.
15. Select **Reuse** check box to allow the database to reuse an existing file.
16. Click **OK**.



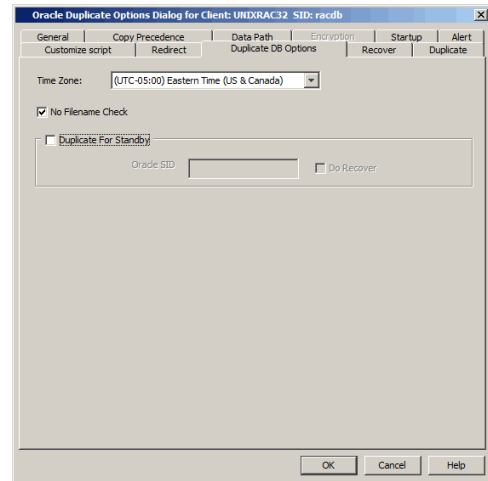
DISABLING FILE NAME VALIDATION DURING A RESTORE

By default, when you create a duplicate database, RMAN will verify the target datafiles for duplicate files (files sharing the same names). This verification job may consume more time. Hence, use the following steps to prevent RMAN from performing this verification:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Type the connect string name in the **Catalog String** box.
4. Select **Duplicate DB** check box.
5. Click **Advanced**.



6. Click **Duplicate DB Options** tab.
7. Select **No FileName Check** check box.
8. Click **OK**.



CREATING A STANDBY DATABASE

Standby databases are useful when a primary database experiences a disaster such as hardware related failure or data corruption and it is not configured for a cluster failover. A standby database is a replicated copy of the primary database. It is updated by applying archived redo logs from the primary database. A standby database will not have a unique DBID.

ON A DIFFERENT HOST WITH THE INSTANCE CONFIGURED

Use the following steps to create a standby database on a different host with the configured instance:

Setting up a Standby Database:

1. Perform an online full backup with current control file.
2. Set the following initialization parameters in the primary initialization parameter file `init<standbydb>.ora` (i.e., Startup PFile) on source (primary) host.
3. Create an initialization parameter file, `init<standbydb>.ora` for the standby database and set the following initialization parameters on Standby host:

```
LOG_ARCHIVE_DEST_1='LOCATION=C:\Oracle\Oradata\TSH1
\Archive MANDATORY REOPEN=30'
```

```
LOG_ARCHIVE_DEST_2='SERVICE=stby1 LGWR SYNC AFFIRM'
```

```
*.log_archive_dest_1='LOCATION=/
LOCATION_OF_DESTINATION_DB_ARCHIVELOG /arch'
```

```
*.STANDBY_ARCHIVE_DEST='/
LOCATION_OF_STANDBY_DB_ARCHIVELOG /stdbyarch'
```

```
*.STANDBY_FILE_MANAGEMENT=auto
```

```
DB_FILE_NAME_CONVERT=
('/DATA_FILE_PATH_ON_SOURCE_MACHINE/stdby',
'/ORACLE_DATA_FILE_PATH_ON_STANDBY_MACHINE/stdby/')
```

```
LOG_FILE_NAME_CONVERT=('/
DATA_FILE_PATH_ON_SOURCE_MACHINE /stdby',
'/ORACLE_DATA_FILE_PATH_ON_STANDBY_MACHINE /stdby/')
```

On Source computer:

```
<Standbydb1_use_diff_connection_string> =
(DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP)(HOST = destination_host_name)
(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = <standbydb_use_same_db_name_as_primary>)
(UR = A)
)
)
<Standbydb_primary> =
(DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP)(HOST = primary_host_name) (PORT
= 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
```

4. Perform any Oracle Net setup and configuration required to connect to the standby database.

For example, set up the listener and configure the destination computer's `$TNS_ADMIN/tnsnames.ora` file for the standby database.

```
(SERVICE_NAME = <standbydb_same_as_destination>)<br>)<br>)</pre>

```

On standby computer:

```
<standbydb1_destination_connection_string> =<br>(DESCRIPTION =<br>(ADDRESS = (PROTOCOL = TCP) (HOST =<br>destination_standby_host_name) (PORT = 1521))<br><br>(CONNECT_DATA =<br><br>(SERVER = DEDICATED)<br><br>(SERVICE_NAME =<br><standbydb_destination_db_name_same_as_primary>)<br><br>(UR = A)<br>)<br>)<br><br><standbydb_primary_connection> =<br>(DESCRIPTION =<br>(ADDRESS = (PROTOCOL = TCP) (HOST = primary_host_name) (PORT<br>= 1521))<br><br>(CONNECT_DATA =<br><br>(SERVER = DEDICATED)<br><br>(SERVICE_NAME = <standbydb_primary_db>)<br>)<br>)</pre>

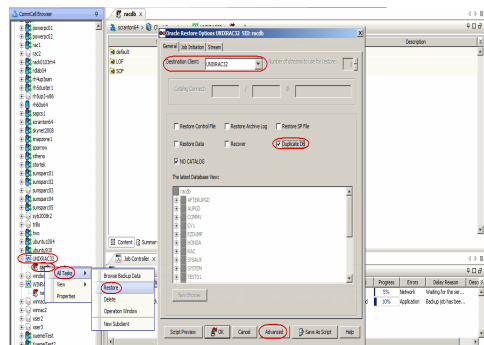
```

5. Use a connect string to connect to the auxiliary database when the instance is not configured. Make sure to use the password `change_on_install` when you create the password file for auxiliary channel while creating password file for standby database on destination host.
6. Startup the standby database instance in NOMOUNT mode. Configure the Oracle instance for the destination host in the CommCell Browser.
7. Ensure to run the `Ora_install.sh` on the auxiliary client.
8. Create the Oracle RAC client and configure it as a client on the Destination host.
9. Ensure that you set the Standby Role Initialization parameter, `DB_FILE_NAME_CONVERT`, to add all the temp datafiles from the primary database location to the standby database location.

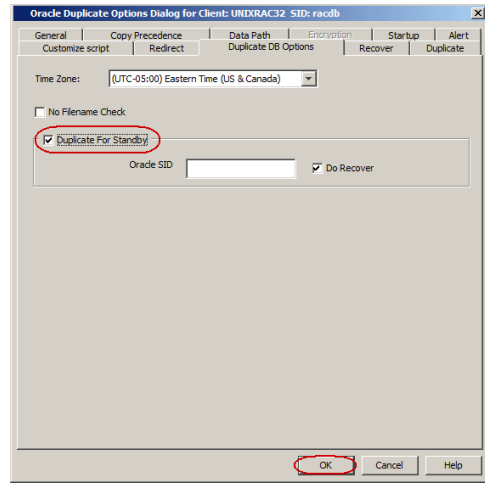
`DB_FILE_NAME_CONVERT='<primary_database_temp_datafile_old_location>','<standby_database_temp_datafile_new_location>'`

Creating a Standby Database:

10. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
11. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
12. Type the connect string name in the **Catalog String** box.
13. Select **Duplicate DB** check box.
14. Under **The latest Database View** box, click **Refresh** to verify that the status of the standby database is **STARTED**.
15. Click **Advanced**.
18. Click **Duplicate DB Options** tab.
19. Click **Duplicate For Standby**.
20. Type the Oracle <SID> name in **Oracle SID** box.
21. Select **Do Recover** check box to recover the standby database.



22. Click **OK**.



RESTORING DATABASE TABLES

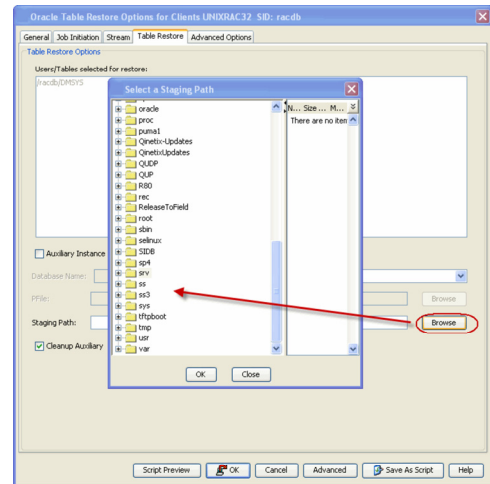
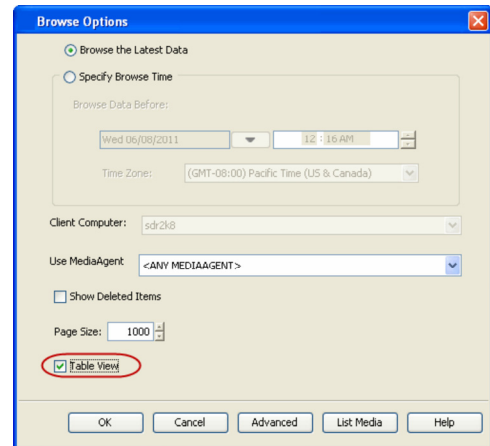
By default, the database tables can be restored from an online full backup, provided the table browse was enabled in the associated subclient before performing the backup. See Enabling Table Browse for Restores for information on configuring the subclient for table browse.

When restoring database tables, by default an auxiliary instance is automatically created. Hence, make sure that there is enough disk space on the client for the auxiliary instance.

RESTORING TABLES TO THE SOURCE DATABASE

If some of the tables in the database are lost or corrupted, you can restore those tables back to the same database using the following steps:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
 2. Right-click the **<Instance>**, point to **All Tasks** and select **Browse Backup Data**.
 3. Select the **Table View** check box and click **OK**.
 4. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
-
5. Click the **Table Restore** tab.
 6. In the **Staging Path** box, click **Browse** and select the location where the auxiliary instance will be created.
 7. Click the **Advanced Options** tab.

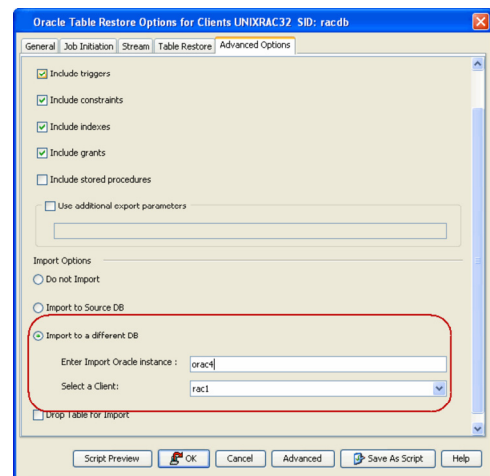


8. Select **Import to Source DB**.
9. Click **OK**.

RESTORING TABLES TO A DIFFERENT DATABASE ON THE SAME HOST

Use the following steps to restore tables to a different database on the same host:

1. Add the destination instance name in the `Listener.ora` and `Tnsnames.ora` files.
2. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
3. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
4. Select the **Table View** check box and click **OK**.
5. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
6. Click the **Table Restore** tab.
7. In the **Staging Path** box, type the location where the tables will be restored.
8. Click the **Advanced Options** tab.
9. Select **Import to a Different DB**.
10. In the **Enter Import Oracle Instance:** box, type the destination instance name.
11. In the **Select a Client** box, select the source client.
12. Click **OK**.



RESTORING TABLES TO A DIFFERENT HOST

When restoring the tables to a different host, ensure the following:

- Both the source and the destination host should have the same database schema.
- Add the duplicate database instance name in the `Listener.ora` file on the destination host and `Tnsnames.ora` files on the destination and source hosts. ,
- Add the source database name in the `Tnsnames.ora` file on the destination host.
- Ensure that both the source and destination clients use a different connection name in the `tnsnames.ora` file.

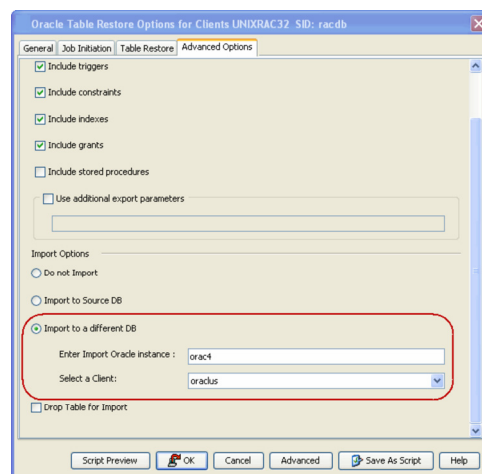
For example:

```
// db1_table = <--- Name of the source database
db1_table =
(DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP) (HOST = M1) (PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = db1)
)
)
//db1 = <--- Name of the destination database
db1 =
```

```
(DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP)(HOST = M2)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = db1)
)
)
```

Use the following steps to restore the tables to a different host:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.
4. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
5. In the **Destination Client** box, select the destination client name.
6. Click the **Table Restore** tab.
7. In the **Staging Path** box, type the location where the auxiliary instance will be created.
8. Click the **Advanced Options** tab.
9. Select **Import to a Different DB**.
10. In the **Enter Import Oracle Instance:** box, type the destination instance name.
11. In the **Select a Client** box, select the destination client.
12. Click **OK**.



RESTORING TABLES USING A USER-DEFINED AUXILIARY INSTANCE

By default, when you restore database tables to a target instance, the system automatically duplicates the source database to an auxiliary instance in the specified temporary staging location. Once the database is duplicated, you can import the tables to the target instance.

However, if required, you can also use an user-defined auxiliary instance for the restore operation. This is used when you want to restore a table to a specific failure point.

When restoring tables to a different host, if an user-defined auxiliary instance option is selected for the restore, you need to recover the database to a specified point-in-time or SCN number. You cannot recover the database to the current time using an user-defined auxiliary instance.

SETTING UP THE AUXILIARY INSTANCE

1. Copy the `init<SID>.ora` file from the source database to the auxiliary database instance.
2. Update the database name and the database file locations in the `init<SID>.ora` file for the auxiliary database instance.
3. Add the `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT` parameters in the `init<SID>.ora` file. These parameters will redirect the datafiles, temp files, and log files to the auxiliary instance.
4. Add the `log_archive_dest_1` parameter is added to the `init<SID>.ora` file on the auxiliary instance.
5. Restart the Oracle Services, if using Windows clients.
6. Add the destination instance name in the `Listener.ora` and `Tnsnames.ora` files. If using a different host, add the duplicate database instance name in the `Listener.ora` file on the destination host and `Tnsnames.ora` files on the destination and source hosts. Also, add the original database name in the

Windows Clients:

```
DB_FILE_NAME_CONVERT=
('source_of_df_path','dup_of_df_path','source_of_temp_path','dup_of_temp_path',...)
LOG_FILE_NAME_CONVERT=('source_of_log_path/redo','dup_of_log_path/redo')
```

Unix Clients:

```
DB_FILE_NAME_CONVERT=
(source_of_df_path,dup_of_df_path,source_of_temp_path,dup_of_temp_path,...)
LOG_FILE_NAME_CONVERT=(source_of_log_path/redo,dup_of_log_path/redo)
```

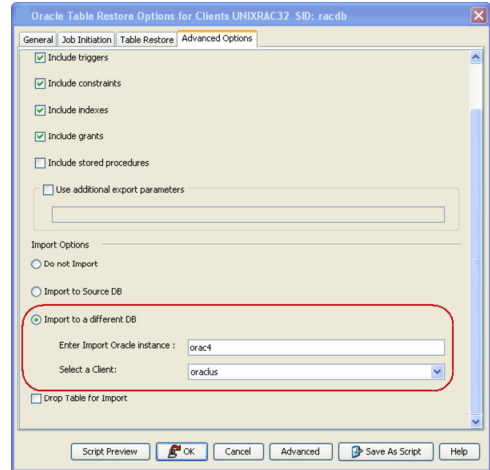
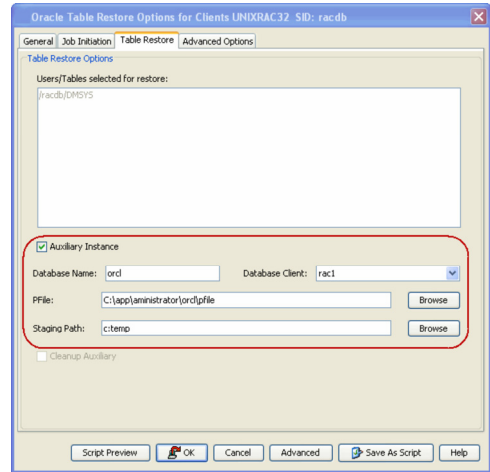
```
DUPDB = (DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP)(HOST = powerpc02)(PORT = 1521))
(CONNECT_DATA = (SERVER = DEDICATED)
(SERVICE_NAME = dupdb) (UR=A) ) )
```

Tnsnames.ora file on the destination host.

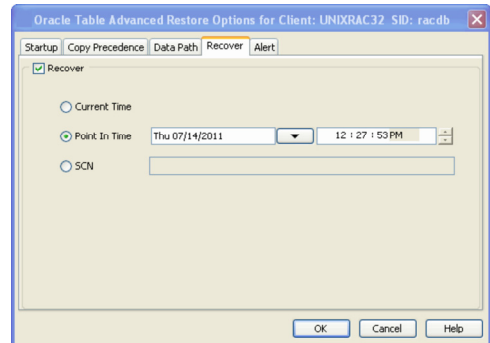
7. Restart the Listener. `$lsnrctl reload`
8. Ensure that the auxiliary instance is in NOMOUNT mode. `sql> startup nomount;`

RESTORING THE TABLES USING THE AUXILIARY INSTANCE

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.
4. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
5. Click the **Table Restore** tab.
6. Select the **Auxiliary Instance** checkbox.
7. In the **Database Instance** box, type the auxiliary instance name.
8. In the **Database Client** box, select the destination client for the auxiliary instance.
9. In the **PFile** box, type the path to the PFile of the auxiliary instance. Alternatively, click **Browse** to select the path.
10. In the **Staging Path** box, type the location where the auxiliary instance will be created. Alternatively, click **Browse** to select the path.
11. Click the **Advanced Options** tab.
12. Select **Import to a Different DB**.
13. In the **Enter Import Oracle Instance:** box, type the destination instance name.
14. In the **Select a Client** box, select the destination client.
15. Click **Advanced**.



16. Select the **Recover** tab.
17. Select **Point-In-Time** checkbox and specify the time range to which the the database need to be recovered.
18. Click **OK**.




DISABLING CLEAN-UP OF AUXILIARY INSTANCE AFTER RESTORE

By default, the system generated auxiliary instance is deleted automatically once the tables are imported to the destination instance.

Use the following steps to disable the clean-up of auxiliary instance after the restore:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.



4. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
5. Click the **Table Restore** tab.
6. In the **Staging Path** box, type the location where the tables will be restored.
7. Clear the **Cleanup Auxiliary** checkbox.
8. Click **OK**.

EXPORTING TABLE OBJECTS

During table restores, the tables are exported from the auxiliary instance to the destination client and later imported to the target database. By default, the following data objects are exported along with the tables:

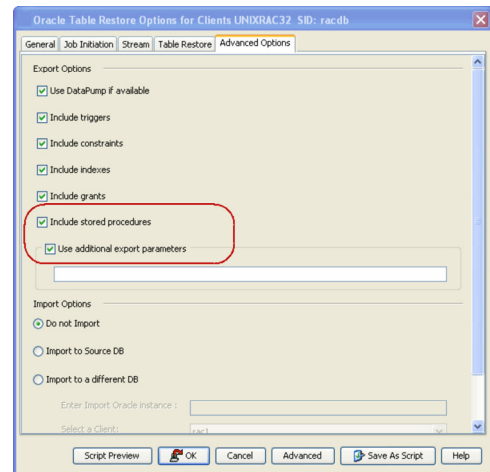
- Triggers
- Constraints
- Indexes
- Grants

However, the stored procedures associated with the selected tables are not exported by default. Use the following steps to export the stored procedures and additional export parameters, such as (COMPRESS or PARALLEL):

Stored procedures are restored from the Schema level. Schema is the collection of data objects created by the user to contain or reference their data. Hence, if one of the table within the schema is selected for restore, all the stored procedures for that schema will also get restored.

When exporting the tables, the datapump export utility is used if it is supported by the Oracle application. The datapump utility facilitates the export of stored procedures. In oracle versions that do not support datapump export utility, you will not be able to include stored procedures during export.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.
4. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
5. Click the **Table Restore** tab.
6. In the **Staging Path** box, type the location where the auxiliary instance will be restored.
7. Click the **Advanced Options** tab.
8. Select the **Include Stored Procedures** checkbox.
9. Select **Use additional export parameters** checkbox and type the parameters to be exported.
10. Click **OK**.




SELECTING/DE-SELECTING DEPENDENT/REFERENCED TABLES

When you browse using the table view, you can also view the dependent and referenced tables associated with the tables selected for the restore.

Dependent tables are the parent tables (containing the primary key) that the selected table (containing the foreign key) depends upon. Similarly, Referenced tables are the child tables (containing the foreign key) that references the selected table (containing the primary key).

By default, all the dependent and referenced tables will be included in the restore operation. Use the following steps to exclude the dependent/referenced tables:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.
4. From the **Browse** window, navigate to the table to be restored.
5. Right-click the **<table>** and click **Select/Deselect All Dependent Tables** to



exclude all the dependent tables.

Similarly, click **Deselect All Referenced Tables** to exclude all the referenced tables.

6. Click **Restore All Selected**.
7. Click the **Table Restore** tab.
8. In the **Staging Path** box, type the location where the auxiliary instance will be restored.
9. Click **OK**.

INCLUDING ALL DEPENDENCIES TO THE DEPENDENT/REFERENCED TABLES


When restoring tables, you can include recursive dependency relationship of all the dependent/referenced tables.

Use the following steps to include all the dependent/referenced tables recursively:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.
4. From the **Browse** window, navigate to the table to be restored.
5. Right-click the **<table>** and click **Select All Dependent Tables Recursively** to include recursive dependency of dependent tables.

Similarly, click **Deselect All Referenced Tables Recursively** to include recursive dependency of referenced tables.

6. Click **Restore All Selected**.
7. Click the **Table Restore** tab.
8. In the **Staging Path** box, type the location where the auxiliary instance will be restored.
9. Click **OK**.



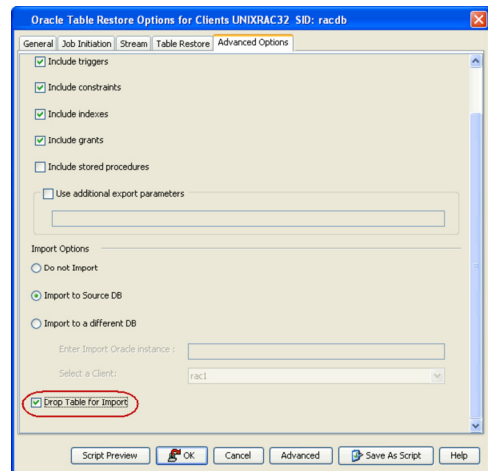
DELETING EXISTING TABLES DURING A RESTORE

By default, the restore operation will overwrite the existing tables in the destination database during the restore. You can also configure the restore operation to delete the existing tables before performing the restore.

Manually drop/delete the existing tables at the destination instance and then import the tables.

Use the following steps to automatically delete existing tables on the destination instance during restore. Note that you can also manually drop/delete the existing tables at the destination instance and perform the restore without enabling this option.

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to All Tasks and select **Browse Backup Data**.
3. Select the **Table View** check box and click **OK**.
4. From the **Browse** window, navigate and select the tables to be restored and click **Recover All Selected**.
5. Click the **Table Restore** tab.
6. In the **Staging Path** box, type the location where the tables will be restored.
7. Click the **Advanced Options** tab.
8. Select **Import to Source DB**.
9. Click **Drop Table for Import** checkbox.
10. Click **OK**.



PERFORMING THIRD PARTY COMMAND LINE OPERATIONS

In addition to restores from CommCell Console, you can also perform restores from the Command Line Interface. Note that the backups taken from CommCell Console can be restored using Command Line and vice versa. However, backups taken from a previous version of CommCell Console can be restored only from the Command Line.

USING THE SAME BLOCK SIZE VALUE

When you run a third party command line restore for Oracle RAC iDataAgents, make sure that you use the same block size value that was used for the corresponding backup job.

For example, if the Oracle backup job used a non-default block size value (say, 32768), then during the third party command line restore operation, you need to specify the same block size value.

If you perform a backup job using the default block size value from the CommCell Console, you need to specify the same block size value for the corresponding third party command line restore operation. This is because, the default block size value used by the Oracle iDataAgent in the CommCell Console is 262144, whereas the default block size value used by Oracle RMAN is 65536. Hence, if you run a third party command line restore for Oracle without specifying the block size, the restore fails. In such cases, specify the block size as follows:

```
allocate channel ch1 type 'sbt_tape'

PARMS=" SBT_LIBRARY=<software_install_path>/Base/libobk.so, ENV=
(CvClientName=clientname,CvInstanceName=instancename),BLKSIZE=262144";
```

However, in the case of Oracle10g and higher, if the backup job (performed from the CommCell Console) used the default block size value, and you perform a third party command line restore operation without specifying the block size, the restore will succeed since the default block size value for Oracle iDataAgent and Oracle RMAN is the same (i.e., 262144).

USING SYNTAX ON A UNIX CLUSTER

In order to perform Oracle Command Line backups on a Unix cluster, you need to use the following syntax for the allocate command within the RMAN script:

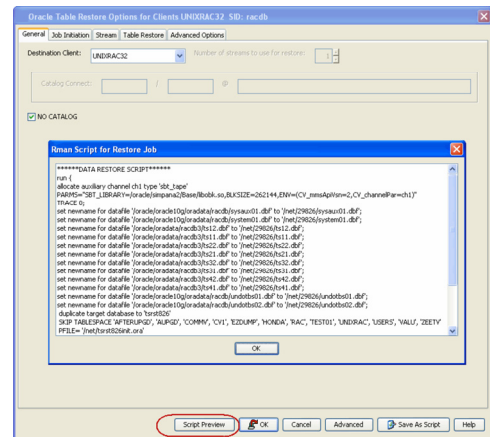
```
allocate channel ChannelName type 'sbt_tape';
```

where *ChannelName* can be chN (N is a stream number: 1, 2, ...)

VIEWING RMAN SCRIPTS FROM THE COMMCELL CONSOLE

Prior to running a restore operation from the CommCell Console, you can preview the corresponding RMAN script for the restore job. This is useful to determine whether the selected restore options will yield the desired result in the script. You can also manually copy and save the generated RMAN script to your computer and later execute the script from the command line.

1. From the CommCell Browser, navigate to **Client Computers** | **<RAC Client>**
2. Right-click the **<Instance>**, and select **Restore**.
3. Click **Script Preview**.
4. Click **OK**.




CUSTOMIZING RMAN SCRIPTS FROM THE COMMCELL CONSOLE

In addition to previewing the RMAN script, you can also modify the script from the CommCell Console. This is useful when you want to include the RMAN commands that is not supported by the software.

1. From the CommCell Browser, navigate to **Client Computers** | **<RAC Client>**
2. Right-click the **<Instance>**, and select **Restore**.
3. Click **Advanced**.
4. Click the **Customize Script** tab.
5. Select the **Customize Script** checkbox.

The script for the control file restore will be generated.



6. Click the **Database** tab to view the script for the database restore.
7. Click **OK**.

RUNNING RMAN SCRIPTS FROM THE RMAN INTERFACE

SBT PARAMETERS

PARAMETER NAME	DESCRIPTION	EXAMPLE
[CvClientName]	Name of the client defined in the CommCell Console. This parameter is optional. It is primarily used in a clustered environment.	[CvClientName] client_name
[CvInstanceName]	Name of the Calypso instance installed on the client from where the RMAN script runs. This parameter is optional. In cases of multiple instances of the software, the first installed instance would be 'Instance001'.	[CvInstanceName] Instance001
[CvOraSID]	Name of the Oracle System ID (SID). This parameter is used during multi stream backups and also when the Oracle instance name is different from Oracle SID.	[CvOraSID] DB1
[-hn]	RAC control host name for the backup/restore job.	[-hn] server_host_name.mycompany.com
[CVOraRacDBName]	Oracle RAC database name for the backup/restore job.	[CVOraRacDBName] asmdb
[CvOraRacDBClientName]	The Oracle RAC database client name for cross machine and/or duplicate database restore job. If you are restoring the database to a computer that is not part of the RAC pseudo client (cross machine restore), you need to define the environmental variable CVOraRacDBClientName to specify a client that is part of the RAC pseudo client.	[CvOraRacDBClientName] dbserve99

Prior to running the RMAN scripts from the RMAN command line, do the following:

1. Add the environmental variables for the client and instance on which the iDataAgent is installed.

```
allocate channel ch1 type 'sbt_tape'
PARMS="<software install path>/Base/libobk.so,ENV=
(CvClientName=rac1,CVOraRacDBName=racdb,
CvInstanceName=instance001 -
hn=rac1.devemc.commvault.com) "
```

2. On Unix clients, add the SBT LIBRARY path.

```
For example,
allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software install
path>/Base/libobk.so,ENV=
(CvClientName=rac1,CVOraRacDBName=racdb,
CvInstanceName=instance001, -
hn=rac1.devemc.commvault.com) "
```

The SBT_LIBRARY path for the various platforms are listed below:

- **AIX with 64 bit Oracle** - <Client Agent Install Path>/Base64/libobk.a (shr.o)
- **HP UX PA RISC 64 bit Oracle** - <Client Agent Install Path>/Base64/libobk.sl
- **Solaris with 64 bit Oracle** -<Client Agent Install Path>/Base64/libobk.so
- **Linux on System Z with 64 bit Oracle** - <Client Agent Install Path>/Base64/libobk.so
- **All Other Unix platforms** -<Client Agent Install Path>/Base/libobk.so

3. Add the same block size value that was used for the corresponding backup job. You can skip this step if the default block size was used for the backup.

```
allocate channel ch1 type 'sbt_tape'
PARMS=" SBT_LIBRARY=<software install
path>/Base/libobk.so, ENV=(
(CvClientName=rac1,CVOraRacDBName=racdb,
CvInstanceName=instance001 -
hn=rac1.devemc.commvault.com,BLKSIZE=32768) "
```

4. From the RMAN command prompt, connect to the target database.
5. Execute the RMAN script.

```
rman target sys/sys@<databasename>
@restore.txt
```

The restore and recover processes are run as separate RMAN run blocks and hence when resumed, the job is restarted from the last failed RMAN run block.

RESTORING MULTI STREAMED BACKUPS

In order to restore using multiple streams from RMAN interface, set the following parameters in the RMAN script.

1. Set the number of automatic channels for a specific device type using the `CONFIGURE DEVICE TYPE ... PARALLELISM` command.

```
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 2
BACKUP TYPE TO BACKUPSET;
```

In the above example, RMAN allocates two channels for the device type when using automatic channels.

2. Run the restore operation.

Make sure to set the `CVOracDBName` and `-hn` parameters to run the restore operation with a single job ID.

`-hn` parameter specifies the node from where the restore operation is initiated.

```
rman target sys/sys@racdb5 << EOF
run {
  allocate channel ch1 type 'sbt_tape'
  PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac1,CVOracDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.company.com)";
  allocate channel ch2 type 'sbt_tape'
  PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac1,CVOracDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.company.com)";
  allocate channel ch3 type 'sbt_tape' connect sys/sys@racdb6
  PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac2,CVOracDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.company.com)";
  allocate channel ch4 type 'sbt_tape' connect sys/sys@racdb6
  PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac2,CVOracDBName=racdb,CvInstanceName=Instance001,-
hn=rac1.company.com)";
  setlimit channel ch1 maxopenfiles 8;
  setlimit channel ch2 maxopenfiles 8;
  setlimit channel ch3 maxopenfiles 8;
  setlimit channel ch4 maxopenfiles 8;
  restore database;
  recover database;
}
EOF
```

3. If you are using the OEM application, include the RMAN settings in the Oracle Enterprise Manager.

For Unix:

```
SBT_LIBRARY=<software_install_path>
/Base|Base64/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=client_name,
CvInstanceName=Instance_name,CVOracDBName=database_name)
```

Windows:

```
ENV=(CvClientName=client_name,
CvInstanceName=Instance_name,CVOracDBName=database_name),BLKSIZE=262144
```

RUNNING RMAN SCRIPTS FOR MULTIPLE STREAMS

Oracle third party command line operations running on multiple streams will share the same Job ID in the Job Manager. If all the streams return failure, then the job is marked as failed. However, if one of the streams fail, it is submitted to the other stream for completion.

When you multiplex Oracle third party operations with multiple streams, each stream uses different drives by default.

When performing third party command line operations, a new `ClOraControlAgent.log` file is generated along with `ORASBT.log` to record the command line jobs.

Oracle third party command line restore operations with single job ID require the `CVOracDBName` parameter, or else multiple jobs IDs will be generated.

For example:

```
run
{allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base64/libobk.so,ENV=(CvClientName=rclab04,
CVOracDBName=DB1,CvInstanceName=Instance001)";
allocate channel ch2 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
```

```

/Base64/libobk.so,ENV=(CvClientName=rdlab04,
CVOraRacDBName=DB1,CvInstanceName=Instance001)";
allocate channel ch3 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base64/libobk.so,ENV=(CvClientName=rdlab04,
CVOraRacDBName=DB1,CvInstanceName=Instance001)";
restore database ;
recover database ;
sql "alter database open"; }

```

If you are restoring the database to a computer that is not part of the RAC pseudo client (cross machine restore), you need to define the environmental variable `CVOraRacDBClientName` to specify a client that is part of the RAC pseudo client

Example:

```

run
{allocate auxiliary channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=clientname,
CVOraRacDBClientName=RACDBsrcClientName,
CvInstanceName=instancename,
CVOraRacDBName=databasename)" TRACE 2;
duplicate target database to <databasename>;
}

```

where, `RACDBsrcClientName` is the physical client name on which the RAC database runs.

For example:

```
rman target sys/sys@racdb5 << EOF
```

```

run {
allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac1,CVOraRacDBName=racdb,CvInstanceName=Instance001,-hn=rac1.company.com)";
allocate channel ch2 type 'sbt_tape'
PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac1,CVOraRacDBName=racdb,CvInstanceName=Instance001,-hn=rac1.company.com)";
allocate channel ch3 type 'sbt_tape' connect sys/sys@racdb6
PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac2,CVOraRacDBName=racdb,CvInstanceName=Instance001,-hn=rac1.company.com)";
allocate channel ch4 type 'sbt_tape' connect sys/sys@racdb6
PARMS="SBT_LIBRARY=/oracle/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CvClientName=rac2,CVOraRacDBName=racdb,CvInstanceName=Instance001,-hn=rac1.company.com)";
setlimit channel ch1 maxopenfiles 8;
setlimit channel ch2 maxopenfiles 8;
setlimit channel ch3 maxopenfiles 8;
setlimit channel ch4 maxopenfiles 8;
restore database;
recover database;
}
EOF

```

RESTORING ON THE SAME HOST (INPLACE RESTORE):

Run the following commands to perform a restore on the same computer using multiple streams:

If you use both the RMAN PARALLELISM configure parameter and set multiple streams from RMAN script, the restore job will utilize double the number of streams. For example, if PARALLELISM is set to 2 and 2 streams are set from RMAN script, the restore job will utilize 4 streams.

RMAN script to perform restore using two channels for node 1 and two channels for node 2.

```

run
{allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name1>,
CVOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>);
allocate channel ch2 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name1>,
CVOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>);
allocate channel ch3 type 'sbt_tape'
connect <connect_string>
PARMS="SBT_LIBRARY=<software_install_path>

```

```

/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name2>,
CVOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>);

allocate channel ch4 type 'sbt_tape'
connect <connect_string>
PARMS="SBT_LIBRARY=<software_install_path>
/Base/libobk.so,BLKSIZE=262144,
ENV=(CvClientName=<rac_client_name2>,
CVOraRacDBName=<rac_database_name>,
CvInstanceName=<instance_name>,
-hn=<host_name>);

```

SETTING TARGET DATABASE TO A DUPLICATE DATABASE

When you perform a command line restore or create duplicate database of an Oracle RAC database to a single instance on the same or different computer, make sure that the SBT_LIBRARY path and the environment variables for CvClientName, CvInstanceName, and CVOraRacDBName are defined in the RMAN script.

Example:

```

rman target sys/sys@asmdb1 auxiliary sys/sys@dupdb << EOF
RUN
{
allocate auxiliary channel ch1 type 'sbt_tape' PARMS="SBT_LIBRARY=/oracle/
calypso/Base/libobk.so,
BLKSIZE=262144,
ENV=(CvClientName=dbserve5,
CVOraRACDBClientName=dbserve10,
CvOraSID=dupdb,CVOraRacDBName=asmdb) "
TRACE 2;
duplicate target database to dupdb
PFIL= /oracle/oracle10g/dbs/initdupdb.ora
SKIP TABLESPACE 'DBTEAM'
nofilenamecheck;
}
EOF

```

RAC TO RAC CROSS COMPUTER RESTORE

Example:

```

rman target sys/sys@racdb5 catalog unixrac/unixrac@rcat10g2 << EOF
run { allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=/opt/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CVOraRACDBClientName=racnode1,CvInstanceName=Instance001,CvClientName=rac1,CVOraRacDBName=
allocate channel ch2 type 'sbt_tape'
PARMS="SBT_LIBRARY=/opt/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CVOraRACDBClientName=racnode1,CvInstanceName=Instance001,CvClientName=rac1,CVOraRacDBName=
allocate channel ch3 type 'sbt_tape' connect sys/sys@racdb6;
PARMS="SBT_LIBRARY=/opt/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CVOraRACDBClientName=racnode1,CvInstanceName=Instance001,CvClientName=rac2,CVOraRacDBName=
allocate channel ch4 type 'sbt_tape' connect sys/sys@racdb6
PARMS="SBT_LIBRARY=/opt/calypso/Base/libobk.so,BLKSIZE=262144,ENV=
(CVOraRACDBClientName=racnode1,CvClientName=rac2,CvInstanceName=Instance001,CVOraRacDBName=
restore database;
recover database;
sql "alter database open resetlogs";
release channel ch1;
release channel ch2;
release channel ch3;
release channel ch4;
}
exit;

```

RAC TO SINGLE NODE

Example:

```

run {
set newname for datafile '/oradata/oeldd/oeldd/example01.dbf' to
'/database/oracle/10GR/oradata/oeldd/example01.dbf';
set newname for datafile '/oradata/oeldd/oeldd/sysaux01.dbf' to
'/database/oracle/10GR/oradata/oeldd/sysaux01.dbf';
set newname for datafile '/oradata/oeldd/oeldd/system01.dbf' to
'/database/oracle/10GR/oradata/oeldd/system01.dbf';
set newname for datafile '/oradata/oeldd/oeldd/undotbs01.dbf' to
'/database/oracle/10GR/oradata/oeldd/undotbs01.dbf';
set newname for datafile '/oradata/oeldd/oeldd/undotbs02.dbf' to
'/database/oracle/10GR/oradata/oeldd/undotbs02.dbf';

```

```

set newname for datafile '/oradata/oeldd/oeldd/users01.dbf' to
'/database/oracle/10GR/oradata/oeldd/users01.dbf';
set newname for datafile '/oradata/oeldd/oeldd/lamba02.dbf' to
'/database/oracle/10GR/oradata/oeldd/lamba02.dbf';
set newname for datafile '/oradata/oeldd/oeldd/lamba01.dbf' to
'/database/oracle/10GR/oradata/oeldd/lamba01.dbf';
allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=/opt/calypso/Base/libobk.so,
BLKSIZE=262144,ENV=(CVOracDBClientName=rac1,CvClientName=ernie,
CVOrasID=oeldd,CvInstanceName=Instance001,CVOracDBName=oeldd)"
TRACE 2;
restore database until time = "TO_DATE('01/08/2010 14:46:24','MM/DD/YYYY HH24:MI:SS)";
switch datafile all;
recover database until scn 6024452;
sql "alter database open resetlogs";
release channel ch1;
release channel ch2;
release channel ch3;
release channel ch4;
}

```

RESTORING FROM A SECONDARY COPY USING RMAN INTERFACE

During restores, if the primary copy of the database is not available, you can restore from a secondary copy using the RMAN command line.

1. Include the environmental variable for copy precedence in the RMAN script.


```

allocate channel ch1 type 'sbt_tape'
PARMS=" SBT_LIBRARY=<software install path>/Base/libobk.so,
ENV=
(CV_restCopyPrec=2,CvClientName=rdlab04,CVOracDBName=DB1,CvInstanceName=Instance001) ,BLKSIZE=32768";

```
2. From the RMAN command prompt, connect to the target database.


```

rman target sys/sys@<databasename>

```
3. Execute the RMAN script.


```

@restore.txt

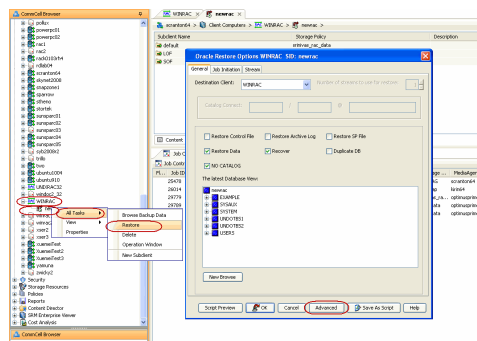
```

SETTING THE RESTORE ORDER FOR STREAMS

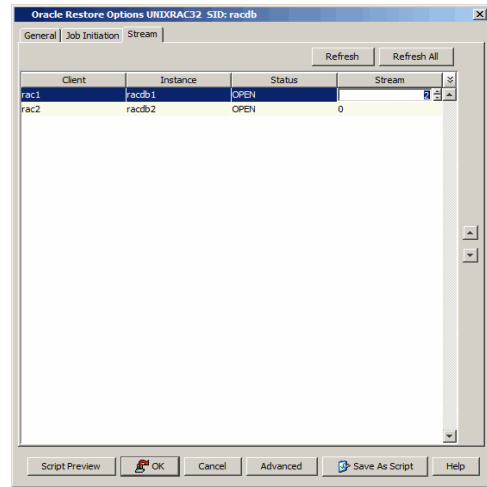
The Oracle RAC iDataAgent allows you to perform parallel restores across multiple RAC nodes. You can set a requisite number of allocated streams for each node and set the order to perform parallel restores. This will allow you to load-balance your restores across all or selected nodes of your RAC database.

Use the following steps to allocate the streams and set the restore order for each client (node):

1. From the CommCell Browser, navigate to **Client Computers | < RAC Client >**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.



3. Click the **Streams** tab.
4. Select the RAC Client (node) for which you want to allocate streams and set the restore order.
5. Double-click the **Streams** box and type the number of streams.
6. Click **OK**.

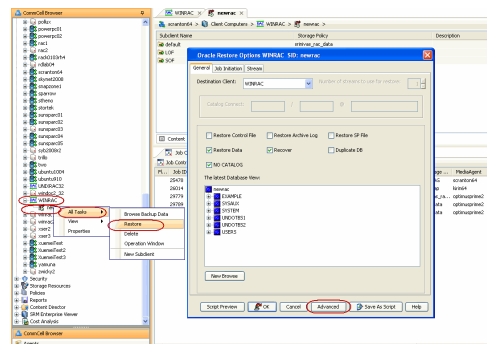


AUTOMATICALLY SWITCHING THE DATABASE MODE BEFORE A RESTORE

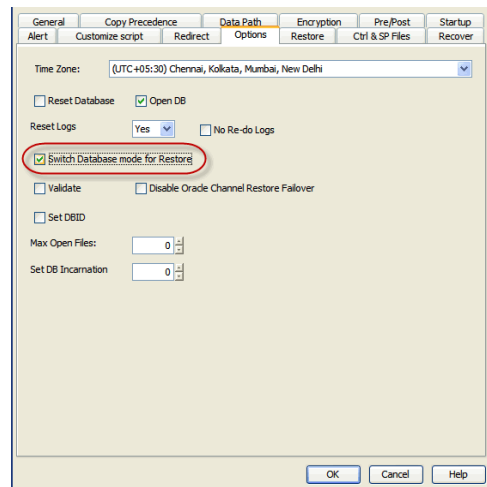
In order to perform a restore operation, the database should be in the MOUNT mode. If the database is not in mounted state, you are prompted to switch the database to the mounted state and then perform the restore.

Use the following steps to automatically switch the database to mount mode prior to restore:

1. From the CommCell Browser, navigate to **Client Computers | < RAC Client >**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



4. Click the **Options** tab.
5. Select **Switch Database mode for Restore**.
6. Click **OK**.



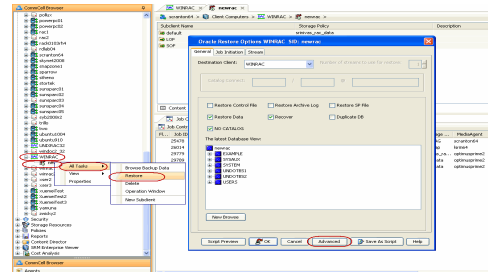
OPENING THE DATABASE AFTER A RESTORE

After a restore operation, you need to open the database for recording further transactions.

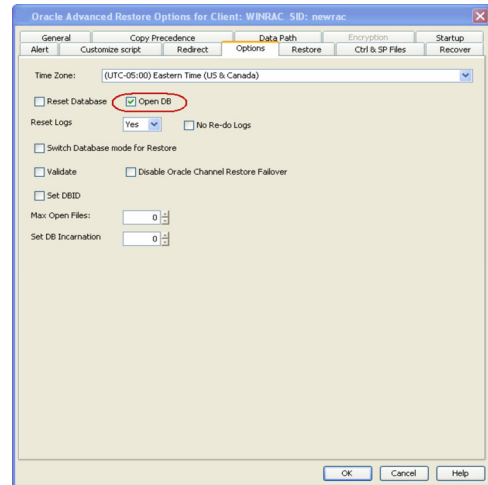
Use the following steps to open the database:

1. From the CommCell Browser, navigate to **Client Computers | < RAC Client >**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.

3. Click **Advanced**.



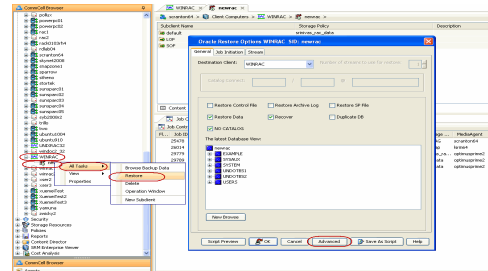
4. Click **Options** tab.
5. Select **Open DB** check Box.
6. Click **OK**.



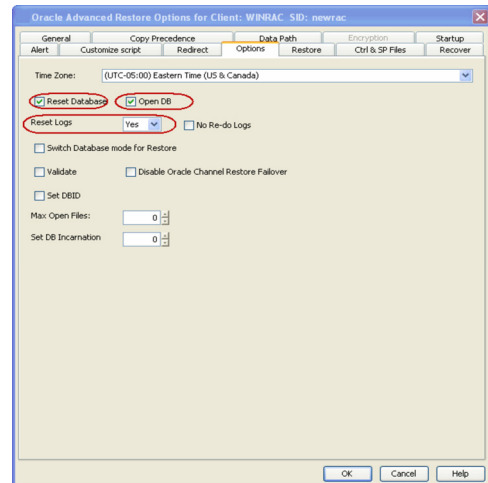
SETTING THE LOG STATE AFTER A RESTORE

By default, the database is automatically set to open and the logs are reset. Use the following steps to prevent resetting the logs:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



4. Click **Options** tab.
5. Select **Open DB** check Box.
6. Select the following options from **Reset Logs** list.
 - o **None** - Open the database without RESETLOGS option.
 - o **No** - Open the database with NORESETLOGS option.
7. Click **OK**.

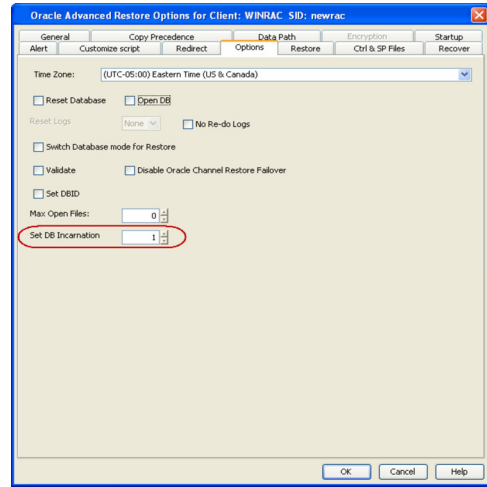
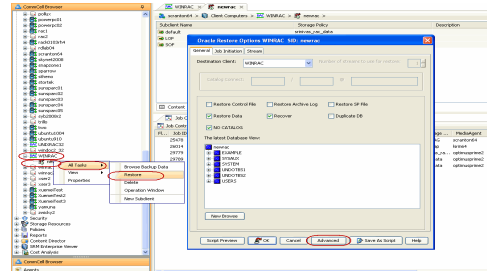


SETTING THE DATABASE INCARNATION

When you perform a point-in-time recovery of an Oracle database with RESETLOGS, a new incarnation of the database is created. All archive log files generated

after resetting the logs will be associated to the new incarnation. However, in order to perform a point-in-time recovery from a backup of a previous incarnation, you need to reset the current incarnation to the previous incarnation value. Use the following steps to set the incarnation value:

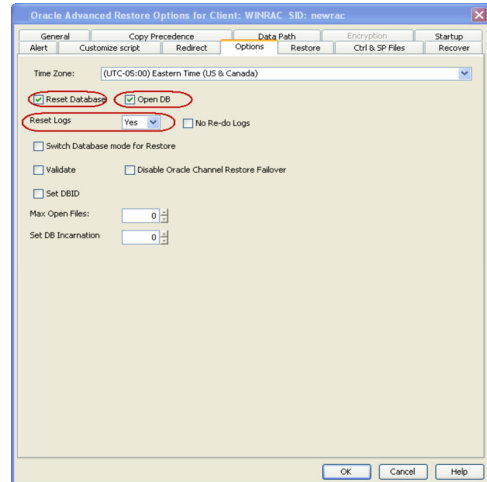
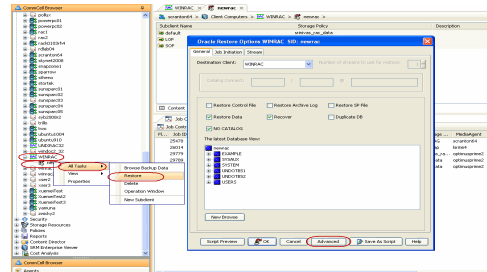
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.
4. Click **Options** tab.
5. Select the database incarnation value from **Set DB Incarnation** list.
6. Click **OK**.



RESETTING THE DATABASE AFTER A RESTORE

By default, the database is not reset. After resetting the logs to open state, you can reset the database. Use the following steps to reset the database after a restore:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.
4. Click **Options** tab.
5. Select **Open DB** check Box.
6. Select **Yes** from the Reset Logs list.
7. Select **Reset Database** check box.
8. Click **OK**.



DISABLING FAILOVERS DURING RESTORES

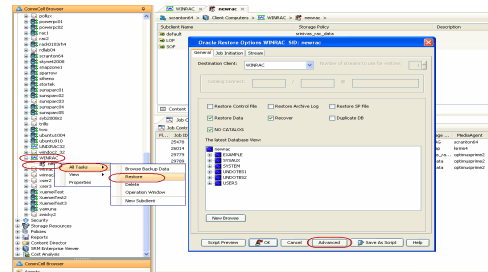
During restore operations, RMAN automatically looks for another copy of the file under the following circumstances:

- a backup piece is corrupted or deleted
- a media agent is offline
- a block in the backup is corrupted within the latest full backup

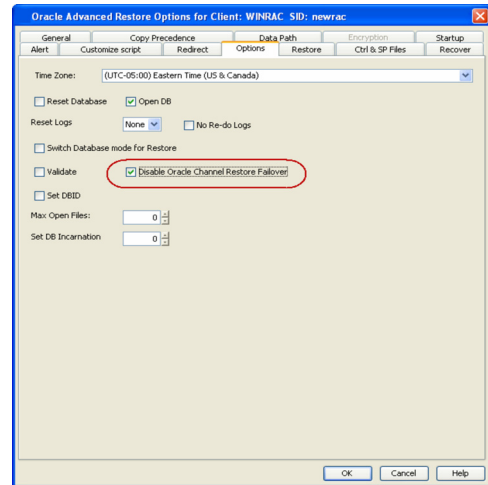
If it is not available in the other copy, RMAN will use older versions of the file, if available. When multiple channels are available for the same device type, RMAN will automatically retry on another channel. RMAN continuously searches all prior backups until it has exhausted all possibilities. This process will delay the restore jobs.

Use the following steps to disable the failovers during restore and prevent job delays:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



4. Click **Options** tab.
5. Select **Disable Oracle Channel Restore Failover** check box.
6. Click **OK**.

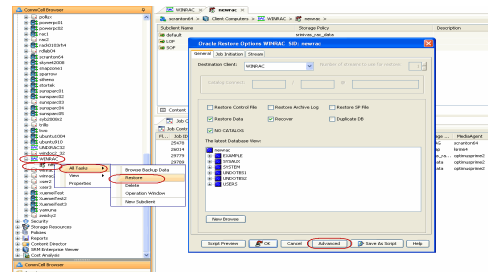


VALIDATING RESTORES

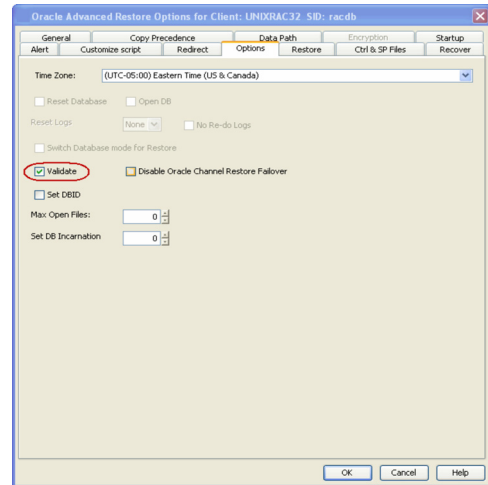
If you perform a validating restore job, the RMAN will stimulate a restore job and verifies whether the backup copies of data and logs required for the restore are intact and usable.

Use the following steps to validate a restore job:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



4. Click **Options** tab.
5. Select **Validate** check Box.
6. Click **OK**.

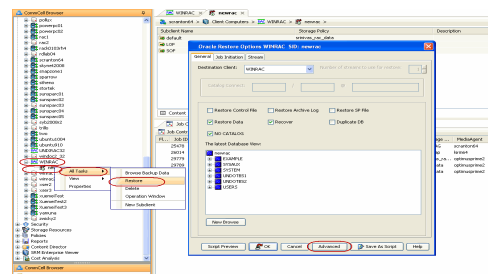


SETTING THE DATABASE IDENTIFIER (DBID)

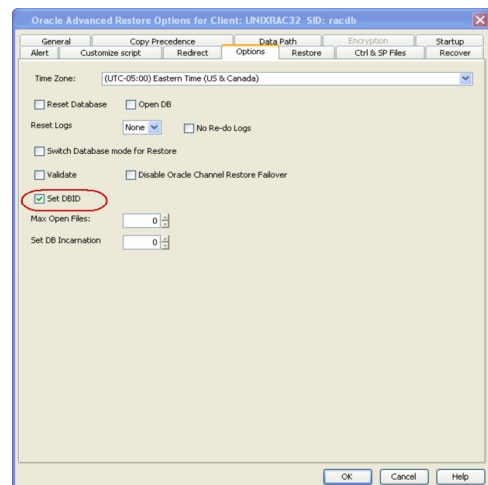
The Database Identifier (DBID) is an internal, uniquely generated number that will distinguish the target database from the rest of the databases that have the same name in the recovery catalog. Oracle creates this number automatically when you create the database. The DBID is set while restoring the control file to differentiate the database for which the control file is restored.

Use the following steps to set the DBID and differentiate the database:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



4. Click **Options** tab.
5. Select the **Set DBID** check box.
6. Click **OK**.



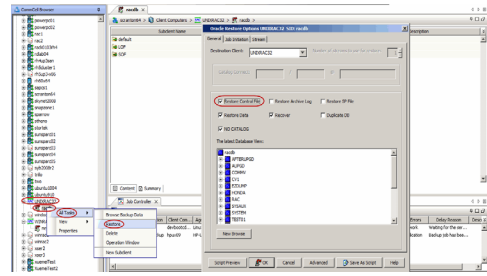
SETTING UP ONLINE REDO LOG FILES

The redo log records all changes made in datafiles. When you run your database in NOARCHIVELOG mode, you disable the archiving of the redo log.

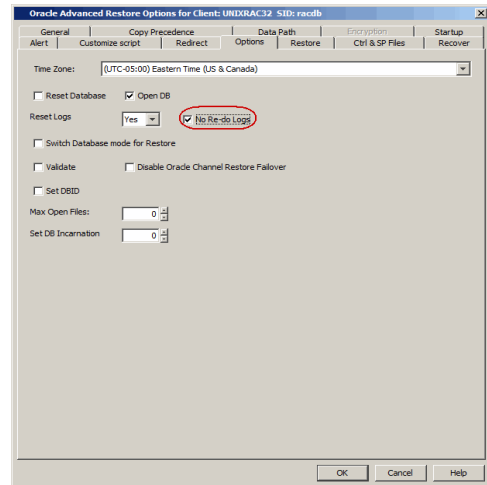
By default, the No Redo Log is disabled. Hence, While restoring the database, RMAN will search for archived re-do logs after applying incremental backup data. Setting No Re-do Logs will enable the RMAN to suppress the archived re-do logs so that only data from incremental backups is restored. Enable No Re-do logs when you perform a point-in-time restore of a database that was backed up in NOARCHIVELOG mode.

Use the following steps to enable No Re-do Logs and perform a restore operation:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



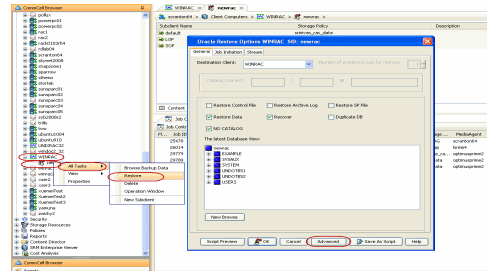
4. Click **Options** tab.
5. Select the **No Re-do Logs** check box.
6. Click **OK**.



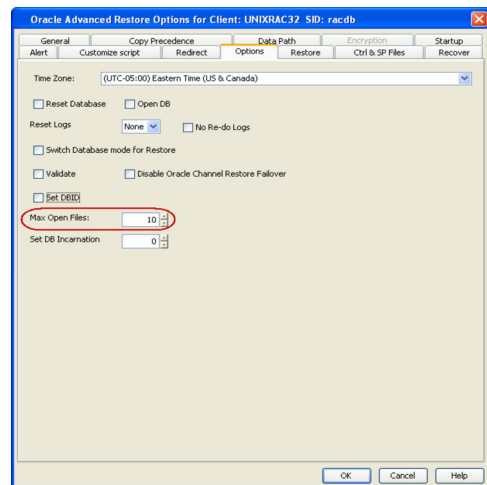
ENHANCING RESTORE PERFORMANCE

You can perform a restore operation faster when you set a maximum number of concurrent open datafiles for RMAN to read simultaneously. Use the following steps to enhance your restore operation:

1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Click **Advanced**.



4. Click **Options** tab.
5. Select the number of open files from **Max Open Files** list.
6. Click **OK**.
7. Click **OK** to start the restore.

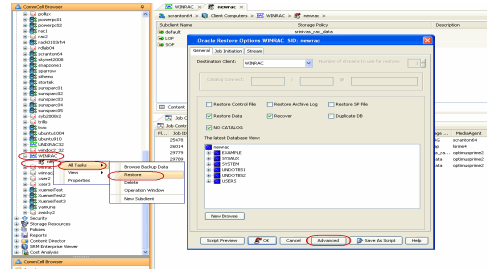


REVERTING A RESTORE

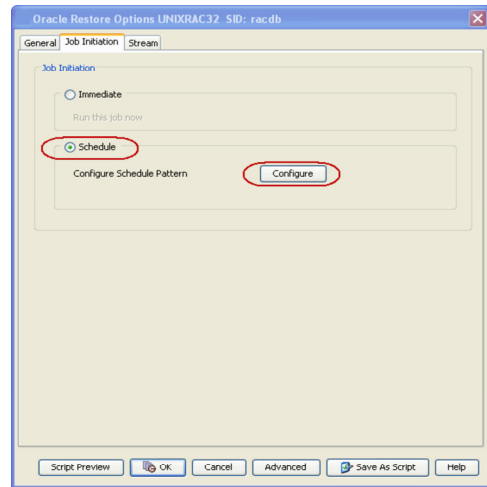
SCHEDULING A RESTORE

Follow the steps given below to schedule a restore job:

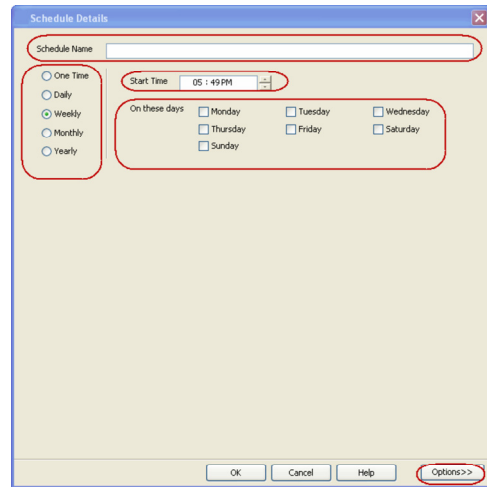
1. From the CommCell Browser, navigate to **Client Computers | <RAC Client>**
2. Right-click the **<Instance>**, point to **All Tasks** and then click **Restore**.
3. Select the required restore options.




4. Click **Job Initiation** tab and then click **Schedule**.
5. Click **Configure**.



6. Select the appropriate scheduling options. For example:
 - Click **Weekly**.
 - Check the days you want the run the restore job.
 - Change the Start Time to 9:00 PM



7. Click Options for the following advanced schedule options:
 - Range of recurrence
 - Repeat times
 - Time Zone



8. Click **OK**.

The restore job will execute as per the schedule.

MANAGING RESTORE JOBS

Once you initiate the restore operation, a restore job is generated in the Job Controller. Jobs can be managed in a number of ways. The following sections

provide information on the different job management options available:

RESTARTING JOBS


Jobs that fail to complete successfully are automatically restarted based on the job restartability configuration set in the Control Panel. Keep in mind that changes made to this configuration will affect all jobs in the entire CommCell.

The restore and recover processes are run as separate RMAN run blocks, and when the job is resumed, it resumes from the last failed RMAN block. If the recover encounters problems after media recovery is completed, the jobs gets completed with errors. In other cases, the job goes to pending and resumes from the last failed RMAN block.

To Configure the job restartability for a specific job, you can modify the retry settings for the job. This will override the setting in the Control Panel. It is also possible to override the default CommServe configuration for individual jobs by configuring retry settings when initiating the job. This configuration, however, will apply only to the specific job.

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

1. From the CommCell Browser, click **Control Panel** icon.
2. Select **Job Management**.
3. Click **Job Restarts** tab and select a **Job Type**.
 - o Select **Restartable** to make the job restartable.
 - o Change the value for **Max Restarts** to change the maximum number of times the Job Manager will try to restart a job.
 - o Change the value for **Restart Interval (Mins)** to change the time interval between attempts for the Job Manager to restart the job.
4. Click **OK**.



CONTROLLING JOBS

The following controls are available for running jobs in the Job Controller window:

SUSPEND	Temporarily stops a job. A suspended job is not terminated; it can be restarted at a later time.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running. The status depends on the availability of resources, the state of the Operation Windows, or the Activity Control setting
KILL	Terminates a job.

SUSPENDING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Suspend**.
2. The job status may change to **Suspend Pending** for a few moments while the operation completes. The job status then changes to **Suspended**.

RESUMING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Resume**.
2. As the Job Manager attempts to restart the job, the job status changes to **Waiting, Pending, or Running**.

KILLING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Kill**.
2. Click **Yes** when the confirmation prompt appears if you are sure you want to kill the job. The job status may change to **Kill Pending** for a few moments while the operation completes. Once completed, the job status will change to **Killed** and it will be removed from the Job Controller window after five minutes.

RESUBMITTING JOBS

If a restore job fails to complete successfully, you can resubmit the job without the need to reconfigure the original job's restore options using the **Resubmit Job** feature. When a job is resubmitted, all the original options, restore destinations, and other settings configured for the job remain in tact.

RESUBMIT A RESTORE JOB

1. From the CommCell Browser, right-click a client computer whose data recovery history you want to view, click **View**, then click to view a job history.

2. From the **Job History Filter** dialog box, select **Restore**.
 - o If you want to view more advanced options for restores, from the Job History Filter, select **Restore**, then click **Advanced**.
 - o From the **Data Recovery History Advanced Filter** select the destination client computer of the restores you would like to view, then click **OK**.
3. The system displays the results of the options you selected in the **Data Recovery Job History** window.
4. Right-click on any job, and select **Resubmit**.
5. Select the job options.
6. Click **OK**.

ADDITIONAL RESTORE OPTIONS

Several additional options are available to further refine your restore operations. The following table describes these options, as well as the steps to implement them.

Be sure to read the overview material referenced for each feature prior to using them.

OPTION	DESCRIPTION	RELATED TOPICS
Use hardware revert capability if available	<p>This option allow you to revert the data to the time when the snapshot was created. Selecting this option brings back the entire LUN to the point when the snapshot was created, overwriting all modifications to the data since the snapshot creation. This option is only available if the storage array used for SnapProtect Backup supports the revert operation.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> 2. Right-click the <Instance> point to All Tasks, and then click Browse Backup Data. 3. Click OK. 4. In the right pane of the Browse window, select the data you want to restore and click Recover All Selected. 5. Click Advanced. 6. Select Use hardware revert capability if available check box. 7. Click OK. 	
Startup Options	<p>The Startup Options are used by the Job Manager to set priority for resource allocation. This is useful to give higher priority to certain jobs. You can set the priority as follows:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> 2. Right-click the <Instance> point to All Tasks, and then click Browse Backup Data. 3. Click OK. 4. In the right pane of the Browse window, select the data you want to restore and click Recover All Selected. 5. Click Advanced. 6. Click Startup tab. 7. Select Change Priority. 8. Select a priority number - 0 is the highest priority and 999 is the lowest priority. 9. Select the Start up in suspended State check box to start the job in a suspended state. 10. Click OK. 	Refer to Job Priority and Priority Precedence.
Copy Precedence	<p>By default, the system retrieves data from the storage policy copy with the lowest copy precedence. If the data was pruned from the primary copy, the system automatically retrieves data from the other copies of the storage policy in the lowest copy precedence to highest copy precedence order. Once the data is found, it is retrieved, and no further copies are checked.</p> <p>You can retrieve data from a specific storage policy copy (Synchronous Copy or Selective Copy). If data does not exist in the specified copy, the data retrieve operation fails even if the data exists in another copy of the same storage policy. Follow the steps given below to retrieve the data from a specific storage policy copy:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> 2. Right-click the <Instance> point to All Tasks, and then click Browse Backup Data. 3. Click OK. 4. In the right pane of the Browse window, select the data you want to restore and click Recover All Selected. 5. Click Advanced. 6. Click Copy Precedence tab. 7. Select the Restore from copy precedence check box. 	Refer to Recovering Data From Copies.

	<ol style="list-style-type: none"> 8. Select the number from Copy Precedence list. 9. Click OK. 	
<p>Data Path Options</p>	<p>The data recovery operations use a default Library, MediaAgent, Drive Pool, and Drive as the Data Path. You can use this option to change the data path if the default data path is not available. Follow the steps given below to change the default data path:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> 2. Right-click the <Instance> point to All Tasks, and then click Browse Backup Data. 3. Click OK. 4. In the right pane of the Browse window, select the data you want to restore and click Recover All Selected. 5. Click Advanced. 6. Click Data Path tab. 7. Select the MediaAgent from Use MediaAgent list. 8. Select the Library from Use Library list. 9. Select the drive pool and drive from Use Drive Pool and Use Drive lists for optical and tape libraries. 10. Select the name of the Proxy server from Use Proxy list, if you wish to restore using a proxy server. 11. Click OK. 	<p>Refer to Change Data Path.</p>
<p>Encryption</p>	<p>If the client's data is encrypted with a pass phrase, you must enter the pass-phrase to start the data recovery operation. Follow the steps given below to enter the pass-phrase:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> 2. Right-click the <Instance> point to All Tasks, and then click Browse Backup Data. 3. Click OK. 4. In the right pane of the Browse window, select the data you want to restore and click Recover All Selected. 5. Click Advanced. 6. Click Encryption tab. 7. Type the pass phrase in Pass Phrase box. 8. Re-type the pass phrase in Re-enter Pass Phrase box. 9. Click OK. 	<p>Refer to Data Encryption.</p>
<p>Alerts</p>	<p>This option enables users or user groups to get automatic notification on the status of the data recovery job. Follow the steps given below to set up the criteria to raise notifications/alerts:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <RAC Client> 2. Right-click the <Instance> point to All Tasks, and then click Browse Backup Data. 3. Click OK. 4. In the right pane of the Browse window, select the data you want to restore and click Recover All Selected. 5. Click Advanced. 6. Click Alerts tab. 7. Click Add Alert. 8. From the Add Alert Wizard box, select the required threshold and notification criteria and click Next. 9. Select the required notification types and click Next. 10. Select the recipients and click Next. 11. Click Finish. 12. Click OK. 	<p>Refer to Alert.</p>
<p>CommCell Readiness Report</p>	<p>The CommCell Readiness Report provides you with vital information, such as connectivity and readiness of the Client, MediaAgent and CommServe. It is useful to run this report before performing the data protection or recovery job. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports CommServe CommCell Readiness. 3. Click the Client tab and click the Modify button. 	<p>Refer to CommCell Readiness Report.</p>

	<ol style="list-style-type: none"> 4. In the Select Computers dialog box, clear the Include All Client Computers and All Client Groups check box. 5. Select the client from the Exclude list. 6. Click the Include > button. 7. Click OK. 8. Click the MediaAgent tab. 9. Clear the Include All MediaAgents checkbox. 10. Select the MediaAgent from the Exclude list. 11. Click Include >. 12. Click Run. <p>The generated report is displayed.</p>	
Restore Job Summary Report	<p>The Restore Job Summary Report provides you with information about all the data recovery jobs that are run in last 24 hours for a specific client and agent. You can get information such as failure reason, failed objects, job options etc. It is useful to run this report after performing the restore. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports Jobs Job Summary. 3. Click Data Recovery on the General tab in the right pane. 4. On the Computers tab, select the client and the agent for which you want to run the report. 5. Click Run. 	Refer to Restore Job Summary Report.

Data Aging - Oracle RAC iDataAgent

TABLE OF CONTENTS

Getting Started

Extended Retention Rules

Data Aging for Transaction, Archive, and Logical Log Backups

Data Aging of the Oracle Recovery Catalog Database

Timeout for Oracle Crosscheck Per Instance During Data Aging

Data Aging Rules for Oracle Archive Index

Disable Oracle RMAN Crosschecks During Data Aging

Data Aging Rules for Selective Online Full Backups

Data Aging Rules for Command Line Backups

Data Aging Rules for On Demand and Customized RMAN Script Backups

Oracle RMAN Retention Policy

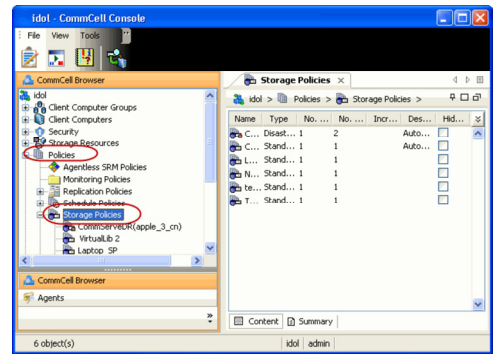
Data Aging Rules for Jobs Completed with Errors

GETTING STARTED

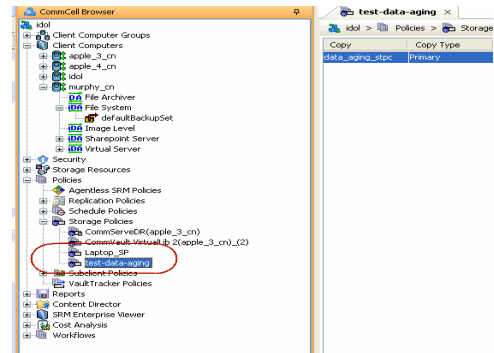
Data Aging is the process of removing old data from secondary storage to allow the associated media to be reused for future backups.

By default, all backup data is retained infinitely. However, you should change the retention of your data based on your needs. Note that if you continue to have infinite retention, you will also need infinite storage capacity.

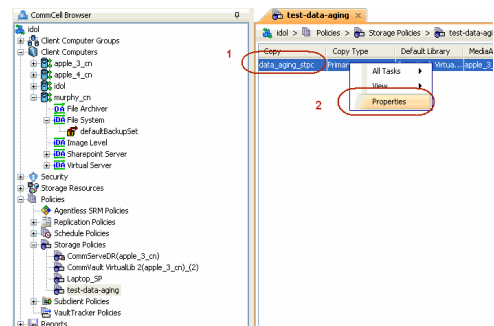
1. From the CommCell Browser, navigate to **Policies | Storage Policies**.



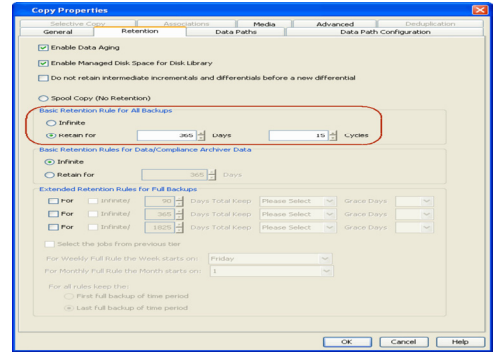
2. Highlight the **Storage Policy**.



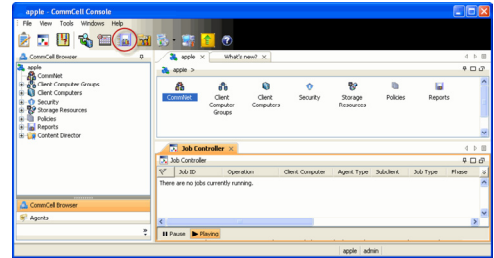
3. From the right pane, right-click the **Storage Policy Copy** and click the **Properties**.



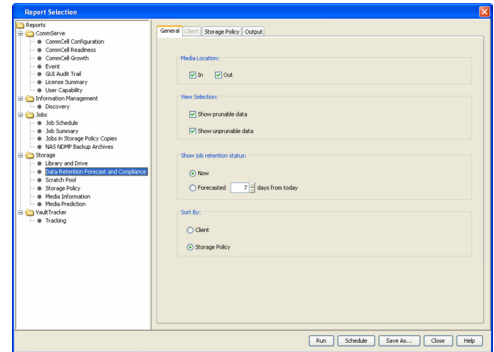
- Click the **Retention** tab.
 - Click the **Retain For** in the **Basic Retention Rules for All Backups** area.
 - Enter number of days to retain the data.
 - Enter number of cycles to retain the data.
 - Click **OK**.



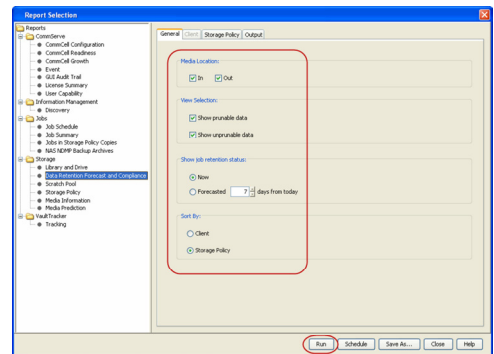
- From the CommCell Browser, click the **Reports** icon.



- Expand Reports and select **Data Retention Forecast and Compliance**.



- Click **Run**.

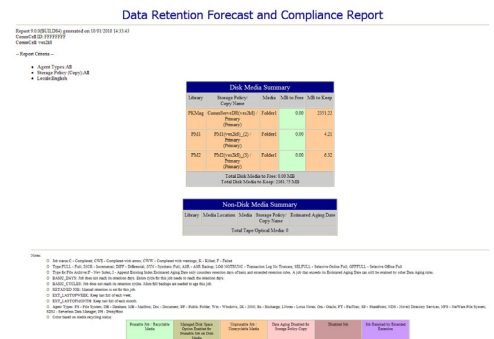


- The report will display the data to be pruned when a data aging job is run.

To ensure only data intended for aging is actually aged, it is important to identify the data that will be aged based on the retention rules you have configured. Hence, ensure this report includes only the data you intend to age.

If necessary, fine-tune your rules so that only the intended data is aged.

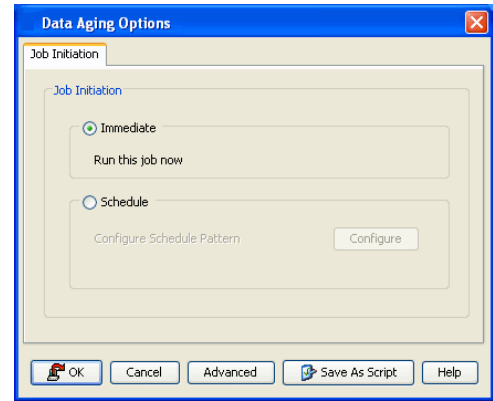
Once you run a data aging job, the data will be lost.



- From the CommCell Console, right click the CommServe icon and click **All Tasks | Data Aging**.




10. Select **Immediate** in the Job Initiation section and click **OK**.



11. You can track the progress of the job from the **Job Controller** window. When the job has completed, the Job Controller displays **Completed**.

Make sure that the job completes successfully. If the job did not complete successfully, re-run the job.



EXTENDED RETENTION RULES

Extended retention rules allow you to keep specific full (or synthetic full) backups for longer periods of time.

Extended retention rules can be used in the following circumstances:

- If you have a single drive tape library
- If you want to create a hierarchical retention scheme (grandfather-father-son tape rotation)

In all other cases, it is recommended that the Auxiliary Copy feature be used for extended storage as it actually creates another physical copy of the data, thereby reducing the risk of data loss due to media failure.

UNDERSTANDING EXTENDED RETENTION RULES

Extended retention allows you to retain a specific full (or synthetic full) backup for an additional period of time. For example, you may want to retain your monthly full backups for 90 days.

Extended retention rules allow you to define three additional "extended" retention periods for full (or synthetic full) backups. For example:

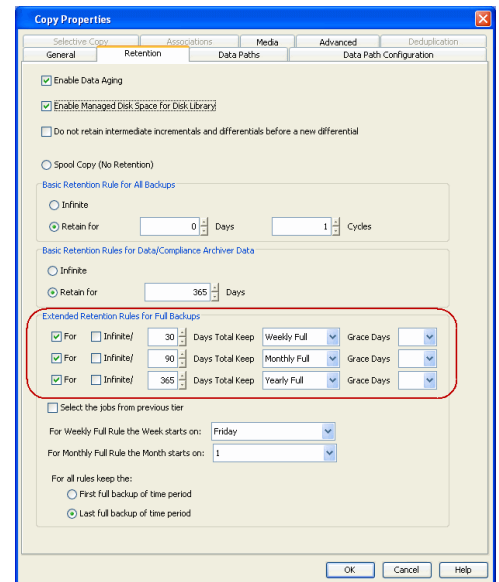
- You may want to retain your weekly full backups for 30 days.
- You may want to retain your monthly full backup for 90 days.
- You may want to retain your yearly full backup for 365 days.

A backup job will be selected for extended retention based on its start time. For example: If a backup job starts at 11:55 pm on August 31st and ends at 1 am on September 1st, then it will be selected as the last full backup for the month of August and will be picked up for extended retention.

SETTING UP EXTENDED RETENTION RULES

Use the following steps for setting up the extended retention rules:

1. Right-click the storage policy copy and click **Properties**.
2. Click the **Retention** tab.
3. Set the basic retention rules by clicking **Retain for** and entering the number of days and cycles appropriate for your organization.
4. Set the extended retention rules as follows:
 - Click the **For** button.
 - Enter the number of **Days Total** to retain the backup.
 - Click the **Keep** drop-down list, and select the desired backup criteria (e.g., Monthly Full).
 - Click the **Grace Days** drop-down list and select the number of days (e.g., 2).
5. Repeat Step 4 to configure additional extended retention.
6. Click **OK**.



DATA AGING FOR TRANSACTION, ARCHIVE, AND LOGICAL LOG BACKUPS

Log Backups (transaction, archive, or logical logs) are not considered part of the backup cycle. Therefore, storage policy cycle retention parameters do not apply to them. However, log backups may be linked to data backup operations, which can affect their retention as follows:

- Log backups are linked to a full backup if they are run at the same time. This is regardless of whether the full backup included data only or data and logs. Such backups follows the standard data aging rules.
- If a full backup job is run on data and logs, then the next log backup will not be linked to this full backup job. These are unlinked log backups and by default, this will follow the unique data aging rules for log backups as given below:
 - Logs that need to be copied to secondary copies will not be aged both on primary and non-primary source copy
 - Logs that exist only on one copy will be aged when they are older than the oldest data
 - When logs exist on multiple copies, the logs on the copy with longest retention days will be retained with the data and will be aged after the oldest data. The log jobs on the remaining copies will be aged according to copy retention days without checking if the oldest data exists or not.
 - Partial, disabled logs will be aged when they are older than the oldest data
- If a full backup job is run on data, then the next log backup job will be linked to this full backup job. These are considered as linked or chained log backups and are not aged until the linked data is aged. In addition, these log backups will also follow the unique data aging rules for log backups.

PRUNING ALL LOG BACKUPS BY DAYS RETENTION RULE

Use the following steps to enable unlinked log backups to be aged according to the defined days retention rule for the data:

1. From the CommCell Browser, select **Tools | Control Panel**.
2. Double-click **Media Management**
3. Click the **Data Aging** tab.
4. Enable the **Prune All Database Agent Logs Only By Days Retention Rule** option.
5. Click **OK**.

DATA AGING OF THE ORACLE RECOVERY CATALOG DATABASE

When a Data Aging job is run, the **BackupPieceName UNAVAILABLE** command is automatically issued to RMAN to disable specific backup pieces in the Oracle Recovery Catalog database that were pruned from the Media Manager CommServe tables. Any backup pieces that were aged from the system's database that have exceeded their retention criteria will be marked as unavailable in the Oracle Recovery Catalog database through this methodology. You can delete these specific backup pieces by creating and enabling the `OracleDeleteAgedBackupPiece` registry key.

TIMEOUT FOR ORACLE CROSSCHECK PER INSTANCE DURING DATA AGING

By default the timeout for Oracle CROSSCHECK per instance is 600 seconds during data aging operation. You can modify this value (or disable the option) by using the `OraCrossCheckTimeOut` registry key.

EFFECTS OF DOWNED SERVICES

When data aging is running, if the Oracle services go down, the data aging operation will complete successfully. However, you need to manually execute Oracle

CROSSCHECK to synchronize the Oracle Recovery Catalog database with that of the CommServe database.

EFFECTS ON ORACLE ARCHIVE LOGS

Oracle archive logs get deleted for those clients/instances where the Oracle CROSSCHECK has been completed successfully. However, if the timeout for the Oracle CROSSCHECK is small (between 1 - 300) and if there are many archive logs, then the crosscheck will fail with a timeout error (or any other error). In such cases, the archive logs will get deleted from the CommServe database during the next data aging operation.

EFFECTS OF UNINSTALLING THE SOFTWARE

When uninstalling the iDataAgent software, CROSSCHECK will no longer be performed by the system to synchronize entries in the CommServe Database with the RMAN catalog. If either of these iDataAgents is later re-installed, then the next data aging job will synchronize the RMAN catalog with the CommServe Database unless the data on tape has been deleted (such as the case where the tape/volume was used for other backups and has been pruned).

DATA AGING RULES FOR ORACLE ARCHIVE INDEX

Oracle archive index is deleted when the associated backup data is deleted. This applies to SnapProtect Backup and Table Level Backup.

DISABLE ORACLE RMAN CROSSCHECKS DURING DATA AGING

By default, during a data aging operation, an Oracle CROSSCHECK is performed by the system to synchronize the entries in the CommServe database with the RMAN catalog. If required, you can disable this CROSSCHECK operation using the **Disable RMAN Cross Check** option in the Instance Properties (Details) for the specific Oracle RAC instance. For step-by-step instructions, see [Disable RMAN Cross Check](#).

DATA AGING RULES FOR SELECTIVE ONLINE FULL BACKUPS

A selective online full operation that consists of archive logs and oracle data can also be linked to the logs of a separate job, which was initiated within the time frame of the selective online full operation. These logs and the selective online full are then considered as one entity within the software, regardless of whether or not separate jobs have the same job ID. Therefore, they are copied to synchronous and selective copies together during auxiliary copy operations and are aged together. If any part of the selective online full is missing from a copy, the full will not be considered as a valid full and will not be counted as a cycle during data aging. Consider the following:

- Data from selective online full backups are considered the same as regular full and offline full backups for each subclient in terms of basic retention rules of cycles and days. However, if any logs on a primary copy have not been fully copied to a secondary copy, the selective online full cannot be aged.
- Data from selective online full backups are considered the same as offline full backups for each subclient in terms of extended retention rules of days. Selective online full backups and all logs linked with it must be retained together on the same storage policy copy.
- Those Logs that are linked with a selective online full (and the logs of the selective online full) can be aged only if they are older than the oldest data that can be aged and the corresponding data of the selective online full that can be or have been aged.
- Selective online full backup jobs that are completed with errors will not be retained by extended retention rules during data aging operation.

DATA AGING RULES FOR COMMAND LINE BACKUPS

- Oracle third party command line log backups can be linked to third party data backups as well as any other kind of backup data as per regular data link rule.
- Data from third-party command line (i.e., RMAN) backups for Oracle RAC is aged differently than data from backups initiated through the CommCell Console. Retention cycles are not used for copies involved in operations from the third-party command line. For such operations, data is aged according to the associated retention time. However, you can manually set the retention time for each third party command line job from the storage policy copy. The command line log backups will be aged according to the retention time set for its associated command line data backup job.
- The `qoperation agedata` command can age data and logs simultaneously based on the Job ID, and it is especially useful for aging each of these items separately.

DATA AGING RULES FOR CUSTOMIZED RMAN SCRIPT BACKUPS

Data Aging for Customized RMAN Script backup jobs uses days/time, and ignores cycles, as the determining factor for pruning the data. Therefore, once the retention time criteria has been met, all data (for both data and logs) is pruned that was backed up using the storage policy specified in the RMAN script that was run through the Command Line Interface.

An effective storage policy strategy for Customized RMAN Script backups is as follows:

- The same storage policy should not be used for regular Oracle backups or Customized RMAN Script backups.
- The storage policy copy containing logs of Customized RMAN Script backups should have a much longer retention time than other storage policies used by regular Oracle backups for the same instance. This is to prevent the logs of Customized RMAN Script backups from being pruned before the data of regular Oracle backups, and allow the database to be fully restored and recovered using the data of old regular Oracle backups and logs afterwards.

ORACLE RMAN RETENTION POLICY

An Oracle RMAN retention policy can be configured for each database. When RMAN retention rules are in effect, RMAN considers the backup jobs comprising data files and control files as obsolete, that is, no longer needed for recovery, according to criteria that you specify in the `CONFIGURE RETENTION POLICY` command. When you run `DELETE OBSOLETE` or `CROSS CHECK` operations, RMAN ages data by freeing disk and tape space used by backups that are no longer

needed.

Do not configure RMAN retention policy if you want to retain data using the data aging feature provided in the CommCell console. To disable the RMAN retention policy, use the following command: `CONFIGURE RETENTION POLICY TO NONE`. This ensures that data will only be aged according to the retention rules specified in the associated storage policy copy.

DATA AGING RULES FOR JOBS COMPLETED WITH ERRORS

Jobs that are completed with errors are not treated as a valid full backup job and hence are pruned based on basic retention rules. However, in case if you require to apply extended retention rules to these jobs, you can exclude jobs that completed with errors during extended retention calculations. Note that this option is applicable only for Selective Online full backup jobs.

1. From the CommCell Browser, select **Tools | Control Panel**.
2. Double-click **Media Management**
3. Click the **Data Aging** tab.
4. Change the value for the **Ignore Completed With Errors job option for Extended Retention calculations** option from 1 to 0.
5. Click **OK**.

ADVANCED TOPICS

Data Aging - Advanced

Provides comprehensive information on additional Data Aging capabilities.

Disaster Recovery - Oracle RAC iDataAgent

TABLE OF CONTENTS

Planning for a Disaster Recovery

Rebuilding the Operating System

Restoring the Oracle Database

PLANNING FOR A DISASTER RECOVERY

Before you begin a disaster recovery, make sure to perform the following:

- Perform frequent Oracle backups with control files and the latest file system.
- Make sure that the recovery catalog is available on a separate computer.


REBUILDING THE OPERATING SYSTEM

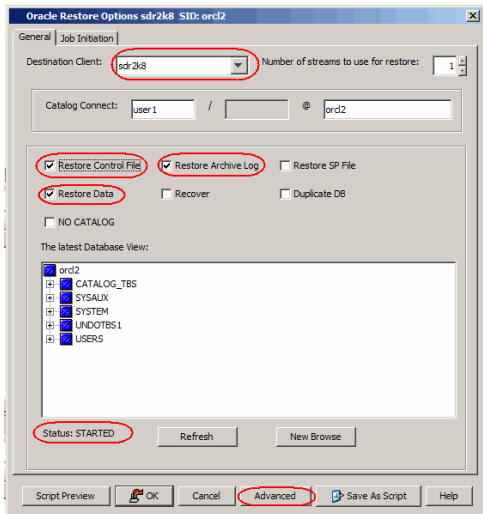
In the case of disaster recovery, where a full system restore is required, you must first rebuild the system to exactly the state as it existed before the problem. Use the following steps to rebuild the operating system:

1. Rebuild the hardware if needed.
2. Install the same version of the operating system with the same patches that were previously installed.
3. Edit the host file of the client to include the CommServe name.
4. Install the File System iDataAgent on the client.
5. Perform a full system restore of the File System iDataAgent. This will restore the Oracle application available under the <oracle install> directory. If there are any problems with the Oracle iDataAgent, we recommend that you un-install and re-install the software in the same drive as it was previously installed. Use the same installation parameters, passwords, and backup type settings as were previously used. In addition to restoring your operating system, this operation will restore the Oracle application files and the Oracle iDataAgent.
6. Make sure to exclude the Oracle Database files while performing the file system backup. You can achieve this by setting a backup filter for **.dbf** files in the appropriate file system subclients.
7. Reboot the computer.
8. Manually copy the init<SID>.ora file on to the computer you are rebuilding.
9. Perform the following steps if the recovery catalog is on the same computer that you are rebuilding:
 - o Create the recovery catalog database as it existed before the crash.
 - o Create the user who was the owner of the recovery catalog with the same user privileges that existed for the user.
 - o Import the user which was exported to a flat file and was backed up as a part of the file system using the Oracle **Import** command.
10. Start the database in the **NOMOUNT** mode.

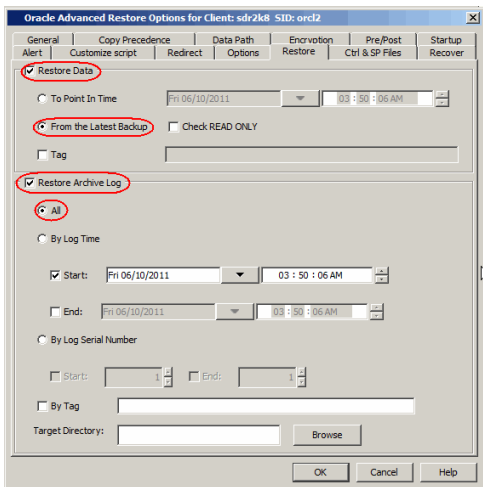
RESTORING THE ORACLE DATABASE

1.
 - From the CommCell Browser, navigate to **Client Computers | <Client> | Oracle**.
 - Right-click the <Instance> point to **All Tasks** and then click **Restore**.
2.
 - Select the name of the client computer from the **Destination Client** list.
 - Clear **Recover** check box.
 - Select **Restore Archive Log** check box.
 - Select **Restore Control File** check box.
 - Verify that the **Status** of the database is displayed as **STARTED**
 - Click **Advanced**.

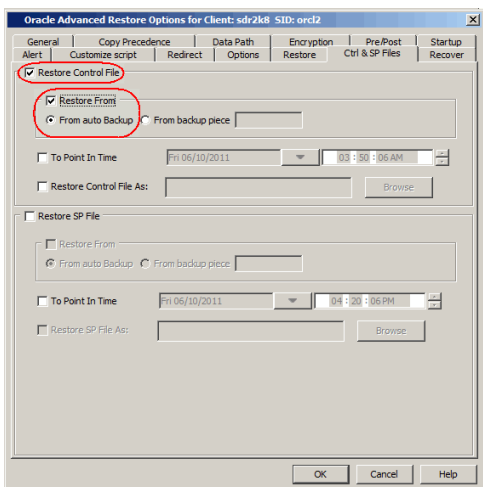




3.
 - Click the **Restore** tab.
 - Ensure that the **Restore Data with From the Latest Backup and Restore Archive Log** with **All** options are selected.

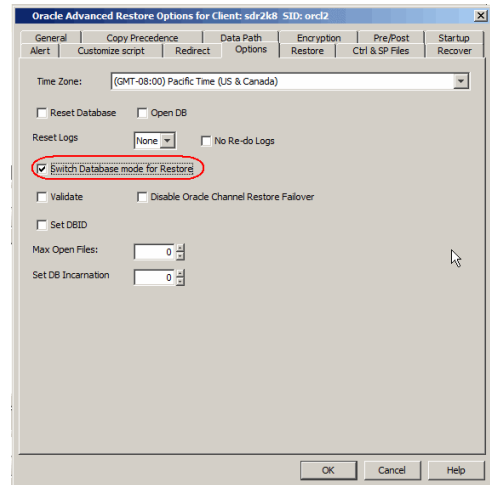


4.
 - Click the **Control&SP files** tab.
 - Ensure that the **Restore Control File with From Auto Backup** options are selected.

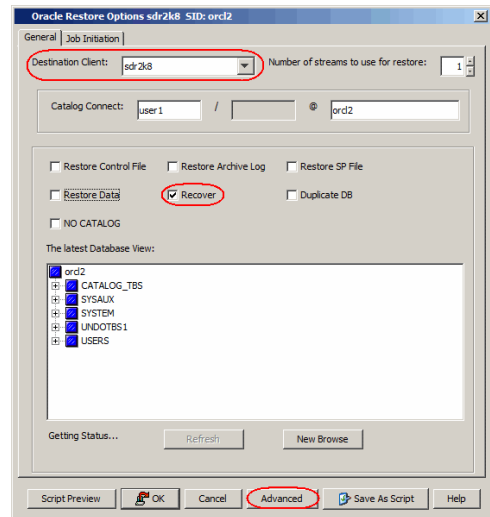



5.
 - Click the **Options** tab.
 - Select the **Switch Database mode for Restore** check box.
 - Click **OK**.

6.
 - From the CommCell Browser, navigate to **Client Computers | <Client> | Oracle**.
 - Right-click the **<Instance>** point to **All Tasks** and then click **Restore**.



7.
 - Select the name of the client computer from the **Destination Client** list.
 - Clear the **Restore Data** check box.
 - Click **Advanced**.



8.
 - Click **Recover** tab.
 - Ensure that the **Recover** with **Current Time** options are selected.
 - Click **OK**.

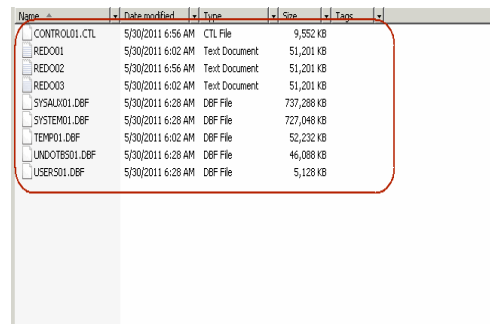
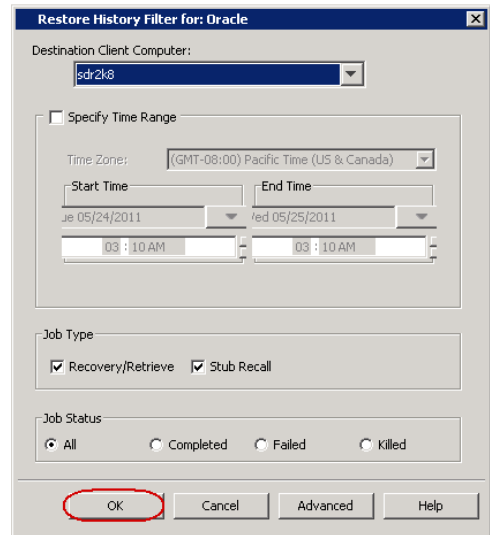
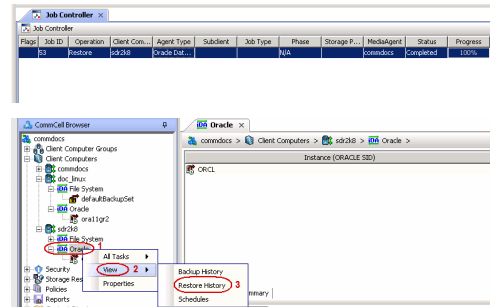
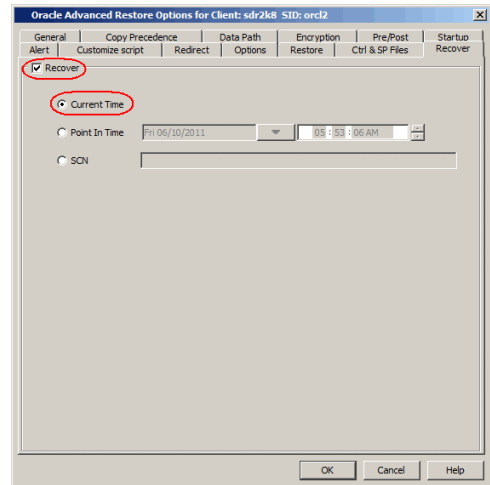
9. You can monitor the progress of the restore job in the **Job Controller**.

10. Once the restore job has completed, right-click the agent and click **View | Restore History**.

11. Click **OK**.

12. Once the database is restored, verify that the restored database and log files are available in the original location.

After the Database is restored, perform a full backup of the Oracle Database.



Additional Operations - Oracle RAC iDataAgent

TABLE OF CONTENTS

Audit Trail

Auxiliary Copy

License Administration

Online Help Links

Operating System and Application Upgrades

Operation Window

Schedule Policy

Storage Policy

Uninstalling Components

AUDIT TRAIL

The Audit Trail feature allows you to track the operations of users who have access to the CommCell. This capability is useful if a detrimental operation was performed in the CommCell and the source of that operation needs to be determined.

Audit Trail tracks operations according to four severity levels:

- **Critical:** This level records operations that will result in imminent loss of data.
- **High:** This level records operations that may result in loss of data.
- **Medium:** This level records changes to the general configuration of one or more entities. Such changes may produce unintended results when operations are performed.
- **Low:** This level records changes to status, addition of entities, and other operations that have minimal impact on existing CommCell functions.

To set Audit Trail retention periods:

1. From the **Tools** menu in the CommCell Console, click **Control Panel**, and then select **Audit Trail**.
2. From the **Audit Trail** dialog box, select the desired retention time (in days) for each severity level.
3. Click **OK**.

AUXILIARY COPY

An auxiliary copy operation allows you to create secondary copies of data associated with data protection operations, independent of the original copy.

1. Expand the **Policies** node, right-click storage policy for which you want to perform an auxiliary copy, click **All Tasks**, and then click **Run Auxiliary Copy**.
2. If you are starting the auxiliary copy operation from the CommServe level, select the storage policy for which you wish to perform the auxiliary copy.

If you are starting the auxiliary copy operation from the Storage Policy level, the Storage Policy field is already populated with the name of the Storage Policy you selected.

3. If the source copy is configured with a shared library, select the **Source MediaAgent** for the auxiliary copy.
4. Click **OK** to start the auxiliary copy operation. A progress bar displays the progress of the operation.

LICENSE ADMINISTRATION

REQUIRED LICENSES

The Oracle RAC iDataAgent requires the **iDataAgent for Oracle RAC** license.

CONVERTING EVALUATION LICENSES TO PERMANENT LICENSES

If you installed the software using an Evaluation License, you can upgrade to a Permanent License as follows:

1. From the CommCell Browser, right-click the CommServe icon, click **Control Panel**, and then click **License Administration**.
2. Select the **Update License** tab and then click **Convert**.
3. Check the box that corresponds to the evaluation license you would like to upgrade and then click **Convert**.

The license information is updated in the **License Administration** window.

USING CAPACITY-BASED LICENSING

License Usage by Capacity is a licensing mechanism that allows you to obtain licenses based on the amount of data you back up. It provides the following features:

- Flexibility of usage on all agents, rather than being tied to number of server, agents, etc.
- Allows you to purchase licenses based on your data protection needs

Both **Core** and **Enterprise** license types are available. Refer to License Usage by Capacity for comprehensive information on utilizing this method.

RELEASING A LICENSE

If you no longer require a license on a computer, such as cases where the computer is being retired, you can release the license and use it later for another computer. Backup data from the retired computer can still be restored after the license is released provided the data is not aged.

1. From the CommCell Browser, right-click the name of the client from which you want to release a license, click **All Tasks**, and then click **Release License for Client**.
2. Click **OK** to continue releasing the license
3. Click **Yes** to confirm you want to release the license or **No** to abort.

ONLINE HELP LINKS

Use the following links to view the online help for the corresponding tabs in the CommCell Console:

OPERATIONS	ENTITY	ONLINE HELP LINKS	SUB LINKS
CONFIGURATION	Client	Oracle RAC Properties (General) Oracle RAC Properties (Security) Oracle RAC Properties (Activity Control) Oracle RAC Properties (Groups) Oracle RAC Properties (Network Throttling) Oracle RAC Properties (Encryption) Oracle RAC Properties (Client Side Deduplication)	
	Instance	Properties of Oracle RAC : <instance name> (General) Properties of Oracle RAC : <instance name> (Details) Properties of Oracle RAC : <instance name> (Storage Device) Properties of Oracle RAC : <instance name> (Encryption) Properties of Oracle RAC : <instance name> (Security)	
	Subclient	Subclient Properties of <Subclient Name> (General) Subclient Properties of <Subclient Name> (Content) Subclient Properties of <Subclient Name> (Backup Arguments) Subclient Properties of <Subclient Name> (Pre/Post Process) Subclient Properties of <Subclient Name> (Log Destinations) Subclient Properties of <Subclient Name> (Logs Backup) Subclient Properties of <Subclient Name> (Encryption) Subclient Properties of <Subclient Name> (Activity Control) Subclient Properties of <Subclient Name> (Storage Device)	
BACKUP	Backup Options	Backup Options for Subclient <Subclient Name>	Save As Script Command Line XML Options Schedule Details
	Advanced Backup Options	Advanced Backup Options (Startup) Advanced Backup Options (Job Retry) Advanced Backup Options (Media) Advanced Backup Options (Data Path) Advanced Backup Options (VaultTracking) Advanced Backup Options (Alert) Advanced Backup Options (Custom RMAN Script)	

		Advanced Backup Options (Backup Archive Logs) Advanced Backup Options (Delete Archive Logs) Advanced Backup Options (Oracle Options)	
RESTORE	Restore Options	Oracle RAC Restore Options (General) Oracle RAC Restore Options (Job Initiation) Oracle RAC Restore Options (Stream)	Save As Script Command Line XML Options Schedule Details
	Advanced Restore Options	Advanced Restore Options (General) Advanced Restore Options (Copy Precedence) Advanced Restore Options (Data Path) Advanced Restore Options (Encryption) Advanced Restore Options (Pre/Post) Advanced Restore Options (Startup) Advanced Restore Options (Recover) Advanced Restore Options (Ctrl & SP Files) Advanced Restore Options (Restore) Advanced Restore Options (Options) Advanced Restore Options (Redirect) Advanced Restore Options (Customize Script) Advanced Restore Options (Alert) Advanced Restore Options (Duplicate) Advanced Restore Options (Duplicate DB Options)	

OPERATING SYSTEM AND APPLICATION UPGRADES

Operating system upgrades are only supported when upgrading from one version of an OS to a different version of the same OS (e.g., Win2003 to Win2008). The two methods of upgrading are:

- **Seamless Upgrade** - This involves uninstalling the Agent software, upgrading the operating system, and then re-installing the Agent software.
- **Full OS Replacement** - This involves performing a clean install of a new version of the OS, re-installing any application software, then re-installing the CommServe, MediaAgent, and/or Agent software.

For Full OS Replacement, the client computer must be configured to have the CommServe, MediaAgent, and/or Client software re-installed to the same location, the same Fully Qualified Domain Name or short domain name, the same partitions, disk drive format (FAT, NTFS, et. al.), and IP configuration as previously.

If it is necessary to remove Agent software to facilitate an operating system or application upgrade, do not delete the icon for the Agent from the CommCell Console, or all associated backed up data will be lost.

Use the following strategy to upgrade the operating system software:

- Identify the computers you want to upgrade and the CommCell components installed on each of these computers.
- Choose the type of upgrade procedure you want to use on each computer: seamless or full replacement.
- CommServe, MediaAgent, and Client computers can be upgraded in any order.

OPERATION WINDOW

By default, all operations in the CommCell can run for 24 hours. To prevent certain operations from running during certain periods of the day, you can define operation rules so that these operations are disabled during those times.

When operation rules are configured, operations that are started within the time window specified will go to a queued (as opposed to pending) state. Once the time window specified in the operation rule has elapsed, these queued or running operations will resume automatically.

1. In the CommCell Browser, right-click the appropriate entity, click **All Tasks**, and then click **Operation Window**.
2. Click **Add**.
3. From the **Operation Window** dialog box:
 - o Enter the name of the rule in the **Name** field.
 - o Select either an administration, data protection (either full or non-full), and/or a data recovery operation from the **Operations** pane.
4. Click **OK**.

SCHEDULE POLICY

A schedule policy is a defined schedule or group of schedules for specific operations to be performed on associated objects within the CommCell. When the schedules from a policy are run, the specified operations, (e.g., auxiliary copy, backup, etc.,) will be performed on the associated CommCell objects.

1. Expand the **Policies** node, right-click **Schedule Policies** and click **Add**.
2. Type the **Name** of the schedule policy.
3. Select the **Type** of schedule policy.
4. Select the **Agent Type**.
5. Type a description of the schedule policy.
6. Click **Add**.
7. Enter a Schedule Name in the **Schedule Pattern** tab.
8. Click **OK**.
9. On the **Associations** tab, select the objects to be associated with the schedule policy.
10. Click **OK**.

STORAGE POLICY

A Storage policy defines the data lifecycle management rules for protected data. Storage policies map data from its original location to a physical storage media and determine its retention period.

1. Expand the **Policies** node, right-click **Storage Policies**, and select **New Storage Policy**.
2. Click **Next**.
3. Select **Data Protection and Archiving** to create a regular storage policy or **CommServe Disaster Recovery Backup** to backup the CommServe database and click **Next**.
4. Click **Next**.
5. Enter the name of storage policy and click **Next**.
6. Enter the name of the primary copy and click **Next**.
7. From the drop down box, select the default library for the primary copy and click **Next**.
8. From the drop down box, select the MediaAgent and click **Next**.
9. Enter number of data streams and set the retention period for the policy and click **Next**.
10. Click **Next**.
11. Click **Browse**, browse to your designated deduplication store location and click **Next**.
12. Confirm your selections and click **Finish**.

UNINSTALLING COMPONENTS

You can uninstall the components using one of the following method:

Method 1: Uninstall Components Using the CommCell Console

Method 2: Uninstall Components from Add or Remove Programs

METHOD 1: UNINSTALL COMPONENTS USING THE COMMCELL CONSOLE

1. From the CommCell Browser, right-click the desired Client Computer and click **All Tasks -> Add/Remove Programs** and click **Uninstall Software**.
2. **Uninstall Software Options** dialog will appear.
3. In the **Uninstall Software** tab, select **Uninstall All** to uninstall all the software packages.
4. In the **Job Initiation** tab, select **Immediate** to run the job immediately.

You can track the progress of the job from the **Job Controller** or **Event Viewer**.

METHOD 2: UNINSTALL COMPONENTS FROM ADD OR REMOVE PROGRAMS

1. Click the **Start** button on the **Windows** task bar and then click **Control Panel**.

2. Double-click **Add/Remove Programs**.

For Windows Vista/Windows 2008, click Uninstall a Program in the **Control Panel**.

3. Click **Remove** to uninstall the components in the following sequence:

1. <Agent>
2. File System iDataAgent
3. Base Software

ADVANCED TOPICS

Provides comprehensive information about additional capabilities for the following features:

- Audit Trail
- Auxiliary Copy
- License Administration
- License Usage by Capacity
- Operation Window
- Schedule Policy
- Storage Policies
- Uninstalling Components

[Back to Top](#)

Best Practices - Oracle RAC iDataAgent

PLANNING A BACKUP

- If you are running in NOARCHIVELOG mode, you must back up the database offline.
- If you are running in ARCHIVELOG mode and database accessibility is a priority, an online backup of the database may be appropriate.
- If the database must be accessible and you have a small backup window, you can run a series of online backups in which different portions of the database are backed up at different times.

You can also combine all of these backup types in your backup strategy. Once you have determined your backup needs, you can define your backup strategy by creating one or more subclients for the database.

CREATING INSTANCES

When creating/modifying an instances, make sure that you do not enter a "/" at the end of the string for Oracle Home path. Due to limitations in the Oracle application, having a "/" at the end of the string for Oracle Home path can cause backups to fail occasionally.

CREATING SUBCLIENTS

As a best practice, it is recommended that you create separate subclients to backup data that undergo frequent changes.

For example, if the `EXAMPLE` and `USERS` tablespaces undergo frequent changes, you can create a separate subclient for each tablespace.

Example:

- User-defined subclient: Test1

Content: `EXAMPLE`

- User-defined subclient: Test2

Content: `USERS`

It is recommended that you create a separate user-defined subclient for the log files on the client.

Distributing the client data using subclients as recommended above, can help improve backup performance by organizing the workload on the client into logical groupings.

RE-CONFIGURING DEFAULT SUBCLIENT CONTENT

We recommend that you do not re-configure the content of a default subclient because this would disable its capability to serve as a catch-all entity for client data. As a result, some data will not get backed up or scanned.

ENABLING FULL SYSTEM RESTORE ON DEFAULT SUBCLIENT

The default subclient for the Oracle instance usually includes all the objects in the instance. Once the default subclient is created, synchronize the recovery catalog with the control files. This would help in successfully performing a full system restore in the event of destroyed or damaged client.

RESTORING CONTROL/SP FILES

Ensure that the database is in NOMOUNT mode when you restore the control/sp files. Ensure that you have previously configured auto backup of control files to restore the control file from auto backup. Restoring a control file will destroy all the previous backups. Hence, you need to perform a full backup after you restore a control file.

RESTORING DEPENDENT TABLES

It is recommended to include the parent tables along with all the dependent tables for a successful restore operation. This will ensure the inclusion of all the reference constraints of the dependent tables along with parent tables.

RESTORING - RECOVERY CATALOG

When restoring backups with recovery catalog, the method you use to restore the Recovery Catalog depends on the method you have used to back up the Recovery Catalog.

If you have exported the Recovery Catalog using the `exp` command, you must import the Recovery Catalog into a new database using the `imp` command. If you have backed up the Recovery Catalog from the CommCell Console, use an appropriate restore method from the CommCell Console to restore the Recovery Catalog.

RUNNING COMMAND LINE OPERATIONS

- Any RMAN command line operation can fail if the required media resource is unavailable due to a conflicting resource reservation by another backup/recovery process or when the resource is unavailable offline. If this happens, it is recommended to establish or verify the availability of the required resource and then rerun the RMAN command line operation.
- When performing backups using RMAN scripts from command line, it is recommended to use separate scripts for data and logs since only one data type can be passed through the argument file (i.e., data or log). This data type is used by the system to mark the archive files created by the backups as DATA or LOG in the CommServe database. Therefore, if you have scripts containing both data and logs, the archive files will be marked as either DATA or LOG depending on the data type mentioned in the argument file.
- Although you can combine data and log scripts in the same job, it is a best practice to run data and log scripts in separate jobs. The drive reservation and re-try mechanism will function more efficiently, if you use this method.

OBTAINING CORRECT INSTANCE STATUS

For the Oracle RAC iDataAgent, if you have custom SQL settings, you must turn off these settings to ensure that the instance browse works correctly and that the instance status is correctly secured. To do this, create and edit the `cvlogin.sql` file with the appropriate contents at `ORACLE_HOME/sqlplus/admin`. For example, run the following command:

```
edit file $ORACLE_HOME/sqlplus/admin/cvlogin.sql
```

Then include the following contents in the `cvlogin.sql` file:

```
set heading on;  
set pause off;
```

Frequently Asked Questions - Oracle RAC iDataAgent

HOW DO WE PROTECT ORACLE DATA ON A LOCAL ZONE?

In order to protect Oracle data on a local/non-global zone, you need to install the Oracle iDataAgent on the non-global zone.

CAN WE PROTECT DATA ON NON-GLOBAL ZONES IF THE FILE SYSTEM /DATAAGENT IS INSTALLED ONLY ON THE GLOBAL ZONE?

Yes. We can protect file system data on non-global zones if the File System iDataAgent is installed only on the global zone. However, in order to enable consistent backups of application specific data on the non-global zones, you will need to install the corresponding application specific iDataAgent on the non-global zone.

CAN WE PERFORM COMMAND LINE BACKUPS OF ORACLE RAC INSTANCES WITH DIFFERENT ORACLE SID AND DATABASE NAMES?

Yes. Different SID and database names for Oracle RAC instances are supported for Oracle command line backups.

CAN WE PERFORM COMMAND LINE BACKUPS IF ORACLE RAC DATABASE NAME AND INSTANCE NAME (AS DEFINED IN THE COMCELL CONSOLE) ARE DIFFERENT?

No. Whenever you run backups from the RMAN command line, ensure that the Oracle database name and the Instance Name as defined in the CommCell Console are the same, otherwise you will need to pass the name of the instance in the RMAN script as follows:

```
allocate channel ch1 type 'sbt_tape'
PARMS="SBT_LIBRARY=<software_install_path>/Base/libobk.so,ENV=
(CvClientName=<rac_client_name>,CVOraRacDBName=<rac_database_name>,CvInstanceName=<instance_name>)"
```

where <rac_client_name> and <instance_name> are the names of the RAC Client and Instance (e.g., Instance001) on which the iDataAgent is installed.

ON WINDOWS CLIENTS FOR RAC, CAN WE RUN MULTIPLE RMAN JOBS SIMULTANEOUSLY FOR THE SAME INSTANCE FROM THE COMMCELL CONSOLE?

No. You should always limit the RMAN jobs for the same instance to run one at a time in the CommCell Console and Scheduler.

CAN WE RUN THIRD-PARTY COMMAND LINE BACKUP OPERATIONS WHEN A SELECTIVE ONLINE BACKUP JOB IS RUNNING?

No. Third-party command line backup operations cannot be run when a selective online full backup job is running.

WHEN DOES A NON-FULL BACKUP AUTOMATICALLY GET CONVERTED TO A FULL BACKUP?

- When you select backup type as Full in CommCell Console and enter the backup type as Incremental on customize script, the job will run as Full.
- When you select backup type as Incremental in CommCell Console and enter the backup type as Full on customize script, the job will run as Full. However, the Job Manager will treat this job as Incremental (on job controller and job history and job report). This job will be pruned as it reaches the data retention date.
- When a storage policy is changed and if backup type is switched to full from job manager, the increment level will be changed to 0.
- If you do not enter the backup type on customize script, then the system picks up the backup type selected in CommCell Console.

WHEN DOES INCORRECT DATA TRANSFER COUNT DISPLAYED IN JOB DETAILS?

When you resume a backup job, if the subclient content includes the SP file, the backup job gets restarted from the beginning and all the data that was backed up previously gets backed up again. This causes an incorrect display of data transferred count in job details.

HOW IS THE DATA RETRIEVED DURING BROWSE AND RESTORE OPERATIONS?

When you browse or restore data, the browse function by default returns the requested data based on the latest backups available. This is usually the information that users are interested in. The browse function does this by using the current date and time as the effective date.

WHAT ARE THE TABLES THAT CAN BE RESTORED IN ORACLE 11G2?

You can restore tables under non-system tablespaces in Oracle 11g2.

HOW ARE TABLE INDEXES HANDLED DURING RESTORE?

When you include the indexes during a table restore, note that only the index definition is restored and not the indexed data.

CAN WE MANUALLY IMPORT TABLES FROM THE AUXILIARY INSTANCE?

Once the tables are restored to the auxiliary instance, you can either use the CommCell Console to export the data from the auxiliary instance and import it to a destination that you choose, or you can leave the data in the staging path and import the data manually by using Oracle import utilities.

HOW ARE STORED PROCEDURES RESTORED?

Stored procedures are restored from the Schema level. Schema is the collection of data objects created by the user to contain or reference their data. Hence, if one of the table within the schema is selected for restore, all the stored procedures for that schema will also get restored.

CAN WE RESTORE TABLES WHEN THERE IS A SCHEMA CHANGE IN THE DATABASE?

No. Table restores are not supported when there are schema changes in the database.

WHAT HAPPENS IF ONE OF THE STREAMS FAIL WHEN RUNNING THIRD-PARTY COMMAND LINE OPERATIONS WITH MULTIPLE STREAMS?

Oracle third-party command line operations running on multiple streams will share the same Job ID in the Job Manager. If all the streams return failure, then the job is marked as failed. However, if one of the streams fail, it is submitted to the other stream for completion.

HOW ARE THIRD PARTY COMMAND LINE JOBS HANDLED WHEN ORACLE SERVICES USE A DIFFERENT SERVICE ACCOUNT INSTEAD OF LOCAL SYSTEM?

When Oracle Services use a different service account other than Local System, third party command line jobs are run with multiple job IDs. In such cases, use the following steps to enable third party command line jobs to run with a single job ID:

1. From the **Start** menu, point to **Administrative Tools**, and then click **Local Security Policy**.
2. Double-click **Local Policies**, and then double-click **User Rights Assignment**.
3. In the details pane, double-click **Adjust memory quotas for a process**.
4. Click **Add User or Group**.
5. In the **Enter the object names to select** box, type the user or group name to which you want to assign the user right, and then click **OK**.
6. In the details pane, double-click **Replace a process level token**.
7. Click **Add User or Group**.
8. In the **Enter the object names to select** box, type the user or group name to which you want to assign the user right, and then click **OK**.

HOW DO WE DELETE ARCHIVE LOGS FROM A FLASH RECOVERY AREA THAT IS PROTECTED?


When the Oracle database is configured to save archive logs in the Flash recovery area, backup jobs will fail if the Oracle RAC subclients have both the Protect backup recovery area and the Archive Delete options enabled at the same time.

To avoid such backup failures you need to configure two different subclients, one for Protect backup recovery area and the other for enabling Archive Deletion option for deleting Archive logs.

HOW DO WE CONFIGURE BACKUPS ON ORACLE EXADATA DATABASE MACHINES?

In order to backup Oracle Exadata Database Machines, do the following:

1. Install the Oracle iDataAgent on all the Exadata nodes.
2. Create a RAC pseudo client that includes all the Exadata nodes.
3. Add the MediaAgent to the Infiniband network.
If using tapes for backup, make sure to include an optimal tape block size (eg.,2048K for Linux clients) in the storage policy copy. See Set the Chunk Size and Block Size for a Data Path for information on setting the block size value.
4. Add Data Interface Pair (DIPs) for Infiniband IP address between the Exadata nodes and the MediaAgent.
5. Run backups on the pseudo client.



CAN YOU RESTORE FROM A RAC ONE NODE IN A MULTIPLE NODE CONFIGURATION?

Yes. You can restore from a RAC One Node in a multiple node configuration. However, consider the following while restoring from a RAC One Node:

- In the case of emergency failover, the instance (same ORACLE_SID) will be relocated to another physical node. The emergency failover is supported during

backup. Make sure that the active ORACLE instance and client combination should be the same as configured in CommCell console during controlfile/spfile restore.

- For RAC One node configuration, it is suggested not to use SwitchDBMode during restores. It is required to keep the database in correct mode (for e.g. NOMOUNT mode for Control file restore and MOUNT mode for database restore) manually before submitting the restore.
- Select streams only for the active node where instance is running.

Troubleshooting Backup - Oracle RAC iDataAgent

Backup | Restore

BACKUP FAILURES

The following section provides information on troubleshooting backups.

Increase in sbtio.log file size	<p>Sometimes, jobs fail due to increase in the size of <code>sbtio.log</code> file in the <code>\$UDUMP</code> directory.</p> <p>To resolve this, set the size limit for the <code>sbtio.log</code> file using the <code>sMAXORASBTIOLOGFILESIZE</code> registry key. Once the specified size limit is reached, the <code>sbtio.log</code> file gets pruned automatically.</p>
Command Line Backup Failures	<ul style="list-style-type: none"> • Make sure if the required media resource is available and then run the backups once again. • For on demand backups, you can run more than one script for an instance. However, backup jobs will fail if there are more than one instance in the argument file. • For Oracle on Windows, it is recommended to avoid using a space after a comma in the argument file. A backup job will fail if you leave a space after a comma in the argument file. • RMAN command line backup fails with the following error <pre>"Unable to open lock file /opt/calypso/Base/Temp/locks/.dir_lock: Permission denied"</pre> <p>This may occur if the <code>unmask</code> parameter is set as <code>022</code> in the <code>.profile</code> file for the Oracle instance. As a workaround, change the <code>unmask</code> to <code>000</code> or <code>002</code> and try the backup again.</p>
Offline backups fails when using lights out script	<ul style="list-style-type: none"> • Offline backups using lights out script fails with the following error: <pre>RMAN error "ORA-12528 TNS listener - all appropriate instances are blocking new connections</pre> <p>As a workaround, add a reference to the database in the <code>listener.ora</code> file as shown in the example below:</p> <pre>SID_LIST_LISTENER = (SID_LIST = (SID_DESC = (SID_NAME = PLSExtProc) (ORACLE_HOME = C:\oracle\product\10.1.0\db_1) (PROGRAM = extproc)) (SID_DESC = (SID_NAME = rman10g) (ORACLE_HOME = C:\oracle\product\10.1.0\db_1) (SID = rman10g)))</pre> <p>Oracle offline backup with lights out option fails when you use the default value for retry attempts for the subclient. As a workaround, increase the retry attempts by setting the Tries number value greater than or equal to 5. See Configuring Lights Out Script for Offline Backups for more details.</p>
Time Out Failures	<p>The default time for resources to allocate streams during RMAN command line backups is 86400 seconds (i.e., 24 hours). If a backup fails due to a timeout being reached, you can configure the <code>sALLOCATESTREAMSECS</code> registry key to increase the waiting time period.</p>
Backup Failures	<ul style="list-style-type: none"> • If the following line is present in the <code>\$ORACLE_HOME/sqlplus/admin/glogin.sql</code> file, it may cause the SrvOraAgent server process on the CommServe to fail when browsing database contents or executing a backup. <pre>set linesize 80</pre> <p>To avoid such failures, comment out that line from the file and re-try the browse or backup operation.</p> • Backup fails with following error: <pre>Character conversion not supported</pre> <p>By default, the <code>NLS_LANG</code> variable on the client computer is set to <code>American_America.US7ASCII</code> character. If the Oracle instance uses <code>NLS_LANG</code> set to a non <code>American_America.US7ASCII</code> character, the Oracle backup operations will fail.</p> <p>In such cases, use the <code><oracle_SID> NLS_LANG</code> registry key to set the <code>NLS_LANG</code> environment variable to the non <code>American_America.US7ASCII</code> character on the client computer.</p> <p>In the case of Oracle RAC iDataAgent, the registry key need to be set in all the RAC instance nodes.</p>
Database block corruption	<p>Oracle backups fail with the following error:</p> <pre>LISTING 2: r_20030520213618.log RMAN-00571: ===== RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====</pre>

	<pre> RMAN-00571: ===== RMAN-03009: failure of backup command on d1 channel at 05/20/2003 21:36:26 ORA-19566: exceeded limit of 0 corrupt blocks for file /u01/app/Oracle/oradata/MRP/sales_data_01.dbf </pre> <p>Make sure that the maximum value for database block corruptions is set for the backup. It is recommended that you set this value to match the number of corrupted database blocks identified by RMAN for the database file being backed up.</p>
<p>Backups fail intermittently on Linux clients</p>	<p>On Linux clients, if the <code>libobk.so</code> library fails to load, the backups may fail.</p> <p>As a workaround, do the following steps:</p> <ol style="list-style-type: none"> 1. Log in to the Oracle client computer as <code>root</code>. 2. From the system prompt, enter the following command: <pre>ldconfig /<Base_directory_name></pre> For example: # <code>ldconfig <software installation path>/Base</code> <p>This will ensure that the <code>libobk.so</code> library is loaded so that backups for Oracle on Linux can run successfully.</p>
<p>Backup fails on Windows Clients</p>	<p>Make sure that the Oracle user is part of administrator group. If the user is not part of administrator group, assign group permissions for the user as follows:</p> <ol style="list-style-type: none"> 1. From Windows Explorer, right-click Calypso folder and then select Properties. 2. Click the Security tab. 3. Select the user and click Edit. 4. Click the Allow checkbox for Full Control permission for the user, and then click OK. 5. From the Registry Editor, navigate to HKEY_LOCAL_MACHINE SOFTWARE. 6. Right click CommVault Systems and select Permissions... 7. Select the user and click Allow checkbox for Full Control permission.
<p>Log backups failure</p>	<ul style="list-style-type: none"> • If the Oracle database is configured to save the archive logs in the Flash recovery area, and Oracle subclients having both Protect backup recovery area and Archive Delete enabled at the same time then the backup will fail. To resolve this, there should be two different subclients, one for Protect backup recovery area and the other for Archive Delete. • Log backup fails if you select the default <code>USE_DB_RECOVERY_FILE_DEST</code> entry as a log destination for the backup. To resolve this, make sure that the log destinations are included in the <code>PFile(init<SID>.ora)</code> or <code>SPFile(spfile.ora)</code> file. Also ensure that the correct log destination is selected for the backup.

COMPLETED WITH ONE OR MORE ERRORS

Backup jobs from Oracle RAC iDataAgent will be displayed as "Completed w/ one or more errors" in the Job History in the following cases:

- When RMAN Script execution for the backup job completes with warnings.
- When job is killed after backing up some data.
- During offline backups, if the database cannot be opened after a backup.

ORACLE ERRORS

If you receive an Oracle error during an Oracle backup operation, we recommend that you follow procedures published by Oracle Corporation on resolving the specific error. We also advise you to consult with your on-site Oracle database administrator, as needed.

Troubleshooting Restore - Oracle RAC iDataAgent

Backup | Restore

The following section provides information on troubleshooting restores.

BROWSE FAILURES

Point in time Table Browse Failures	<p>When you have encryption enabled for the client, point in time table browse operation fails with the following error message:</p> <p>Pass-phrase protection is on for client [80], but pass-phrase was not specified.</p> <p>Make sure that the pass phrase is exported to the MediaAgent when encryption is enabled for the client.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the client and select Properties. 2. Click the Encryption tab. 3. Click Via Pass-Phrase. 4. Click Export. 5. In the Destination Computer box, select the MediaAgent. 6. In the Pass-Phrase box, type the pass-phrase used for encryption. 7. In the Re-enter Pass-Phrase box, re-type the pass-phrase to confirm. 8. Click Export. 9. Click OK.
--	--

RESTORE FAILURES

Table Restore Failures	Make sure that the Oracle Services are running as Local System.
Database Restore Failures	After performing an Oracle restore operation from the CommCell Console where options were selected for Redirect , Rename and Recover at the same time, you must click the Refresh button on the Subclient Properties (Content) tab or run a backup after the restore operation has completed before proceeding with another restore. This is necessary to ensure that the CommCell Console recognizes the changes that were made to the Oracle database and control file, so that it reflects the current structure of the database to be restored, otherwise the restore will fail.
Unable to create a Duplicate Database	<ul style="list-style-type: none"> • When you are creating a duplicate database or using an auxiliary instance for a table restore, make sure that one of the databases use sys as the connect string. • If a duplicate database restore fails with error PLS-00553: character set name is not recognized; then make sure that the character sets are the same between the location from where you are running RMAN, and the location of the target database. As this is an Oracle related issue, please contact Oracle support for more information.
Increase in sbtio.log File Size	<p>Sometimes, jobs fail due to increase in the size of <code>sbtio.log</code> file in the \$UDUMP directory.</p> <p>To resolve this, set the size limit for the <code>sbtio.log</code> file using the <code>sMAXORASBTIOLOGFILESIZE</code> registry key. Once the specified size limit is reached, the <code>sbtio.log</code> file gets pruned automatically.</p>
Intermittent Table Level Restore Failures	<p>The table level restore operation may fail intermittently due to an error in the Oracle's DataPump utility and the following error message will be displayed:</p> <pre>UDE-00008: operation generated ORACLE error 31623 ORA-31623: a job is not attached to this session via the specified handle</pre> <p>In such cases, set the <code>sNODATAPUMPEXPORT</code> registry key to Y on the client and re-submit the job.</p>
Control File Restores Failures	Ensure that the DBID is assigned for the instance. Make sure that the DBID value for the database you are restoring is automatically displayed in Instance Properties.
Commandline Restore Failure	Verify the availability of the required resource then rerun the RMAN command line operation
Unable to create a Standby Database	<p>Standby database fails with the following error message:</p> <pre>temporary file TEMP01.DBF conflicts with file used by target database</pre> <p>Make sure that the Standby Role Initialization parameter, <code>DB_FILE_NAME_CONVERT</code>, is set to add all the temp datafiles from the primary database location to the standby database location, as follows:</p> <pre>DB_FILE_NAME_CONVERT='<primary_database_temp_datafile_old_location>', <standby_database_temp_datafile_new_location>'</pre>

OPTIMIZING MEMORY ALLOCATION FOR TABLE RESTORES

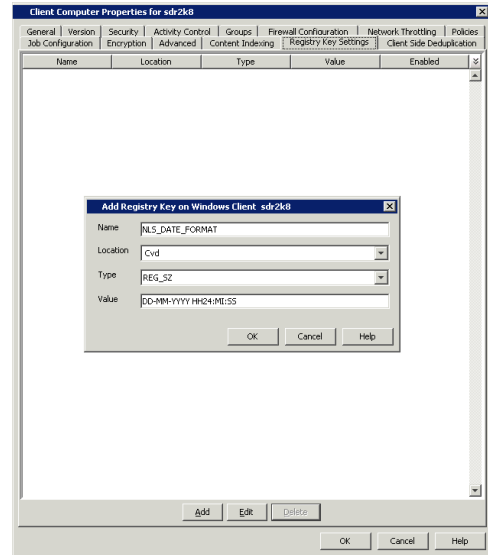
When restoring large tables, the restore operation may fail if there is insufficient memory allocation for creating the auxiliary instance.

Use the following steps to optimize the memory allocation for the auxiliary instance:

ALLOCATING MEMORY FOR AUXILIARY INSTANCE

By default, 16MB pool size is allocated for the auxiliary instance. Use the following steps to increase this size limit:


1. From the CommCell Browser, navigate to **Client Computers**.
2. Right-click the **<Client>**, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. In the **Name** box, type `sLARGEPOOLSIZE`.
6. In the **Location** box, select **iDataAgent**.
7. In the **Type** box, select **REG_SZ**.
On Unix Client, select **Value**.
8. In the **Value** box, type `<Value>`.
For example, 32M.
9. Click **OK**.



SETTING THE UNDO RETENTION PERIOD

Whenever a transaction is committed, the old undo information, is retained by default for a period of 1800 seconds. You can modify this value, using the following steps:

1. From the CommCell Browser, navigate to **Client Computers**.
2. Right-click the **<Client>**, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. In the **Name** box, type `sUNDORETENTIONSIZE`.
6. In the **Location** box, select **iDataAgent**.
7. In the **Type** box, select **REG_SZ**.
On Unix Client, select **Value**.
8. In the **Value** box, type `<Value>`.
9. Click **OK**.



VIEWING RMAN ERRORS

- If the system notifies you that there was an "RMAN error" during a backup or restore operation, we advise you to view the RMAN Output Log to identify which error was encountered. To view the RMAN Output Log, do the following:
 1. From the CommCell Browser, right-click the **<Instance>**, point to **View**, and then click either **Backup History** or **Restore History**.
 2. Click **OK**.
 3. On the Job History window, right-click the job, and then select **View Log File**. The RMAN Output Log will display for the job.

Once you have identified the specific RMAN error, consult one of the following publications from Oracle Corporation for information on resolving the specific error, in conjunction with your on-site Oracle database administrator:

- o Oracle - Backup and Recovery Guide
- o Oracle - Recovery Manager User's Guide and Reference
- o Oracle - Recovery Manager User's Guide

COMMCELL CONSOLE ERRORS

- If only command line backups have been performed, and a Browse Restore operation is subsequently attempted from the CommCell Console without first opening and closing the subclient properties, then the system will display a message indicating that no data was backed up. If this happens, ensure that you open and then close the subclient properties of the Oracle iDataAgent you are trying to restore, then try the Browse Restore operation again.
- If the line "set linesize 80" is present in the file \$ORACLE_HOME/sqlplus/admin/glogin.sql, it may cause the SrvOraAgent server process on the CommServe to fail when browsing database contents or executing a backup. To avoid such failures, comment out that line from the file and re-try the browse or backup operation.
- For Oracle 10g and higher, if the instance is configured for autobackup with flash recovery, then restoring the SPFile from the CommCell Console will not work. The following work-arounds are available in this situation:
 - Work-around 1:
Comment out the option "db_recovery_file_dest" from the PFile, then re-create the SPFile using the new PFile. Afterwards, restoring the SPFile from the CommCell Console should work.

- Work-around 2:

Restore the SPFile manually using RMAN with the script below:

```
RMAN> run {
2> ALLOCATE CHANNEL c1 DEVICE TYPE DISK;
3> restore spfile from autobackup;
4> }
```

POINT-IN-TIME RECOVERY

When you recover a database to a point in time, the RMAN command ALTER DATABASE OPEN RESETLOGS is executed which will reset the SCN (System Change Number) and time stamp on every object of the database (i.e., datafiles and control files). Also, only the archived redo logs that match the RESETLOGS SCN and timestamp value will be applied to the database, thus recovering the database to a time that is not current. This is a very useful operation if the point-in-time to which you are trying to recover is certain and known, but can be counterproductive if you are guessing at the point-in-time.

If you are not sure about the point-in-time for the recovery, it is recommended to restore the data and the control files to a point in time without recovery. This method will allow you to restore the database to a state that you can make the determination whether or not you have achieved the correct point-in-time, without invoking the "ALTER DATABASE OPEN RESETLOGS" statement that would reset SCNs and time stamps on the database objects.

After determining the correct point-in-time through this method, you can recover the database to the point in time to reset your Oracle database to the desired incarnation.

Sample scripts are provided below for your Oracle database administrator to use as reference for developing custom scripts that you can run from the RMAN command line, to perform special operations apart from the CommCell Console.

SAMPLE SCRIPT FOR RESETTNG A DATABASE AFTER RESETLOGS

The following example resets a database after performing an incomplete media recovery:

```
run {
allocate channel dev1 type disk;
set until logseq 1234 thread 1;
restore database skip tablespace readonly;
recover database;
sql "ALTER DATABASE OPEN RESETLOGS";
release channel dev1;
}

reset database;
```

SAMPLE SCRIPT FOR RESETTNG THE DATABASE TO AN OLD INCARNATION

The following command makes an old incarnation of database PROD1 current again:

```
# obtain primary key of old incarnation

list incarnation of database prod1;
List of Database Incarnations
```

DB Key	Inc Key	DB Name	DB ID	CUR	Reset SCN	Reset Time
-----	-----	-----	-----	---	-----	-----
1	2	PROD1	1224038686	NO	1	02-JUL-98
1	582	PROD1	1224038686	YES	59727	10-JUL-98

```
shutdown immediate;

# reset database to old incarnation

reset database to incarnation 2;
```

```
# recover it
run {
allocate channel dev1 type disk;
restore controlfile;
startup mount;
restore database;
recover database;
sql "ALTER DATABASE OPEN RESETLOGS";
release channel dev1;
}
```

COMPLETED WITH ONE OR MORE ERRORS

Restore jobs from Oracle iDataAgent will be displayed as "Completed w/ one or more errors" in the Job History in the following cases:

- During a table restore, if the export or import of table fails.
- RMAN recovery is completed, but an incorrect open mode is selected for restore.

RECOVERING DATA ASSOCIATED WITH DELETED CLIENTS AND STORAGE POLICIES

The following procedure describes the steps involved in recovering data associated with the following entities:

- Deleted Storage Policy
- Deleted Client, Agent, Backup Set or Instance

BEFORE YOU BEGIN

This procedure can be performed when the following are available:

- You have a Disaster Recovery Backup which contains information on the entity that you are trying to restore. For example, if you wish to recover a storage policy (and the data associated with the storage policy) that was accidentally deleted, you must have a copy of the disaster recovery backup which was performed before deleting the storage policy.
- Media containing the data you wish to recover is available and not overwritten.
- If a CommCell Migration license was available in the CommServe when the disaster recovery backup was performed, no additional licenses are required. If not, obtain the following licenses:
 - IP Address Change license
 - CommCell Migration license
See License Administration for more details.
- A standby computer, which will be used temporarily to build a CommServe.

RECOVERING DELETED DATA

1. Locate the latest Disaster Recovery Backup which contains the information on the entity (Storage Policy, Client, Agent, Backup Set or Instance) that you are trying to restore.
 - You can check the Phase 1 destination for the DR Set or use Restore by Jobs for CommServe DR Data to restore the data.
 - If the job was pruned and you know the media containing the Disaster Recovery Backup, you can move the media in the **Overwrite Protect Media** Pool. See Accessing Aged Data for more information. You can then restore the appropriate DR Set associated with the job as described in Restore by Jobs for CommServe DR Data.
 - If the job is pruned and you do not know the media containing the Disaster Recovery Backup, you can do one of the following:

If you regularly run and have copies of the Data on Media and Aging Forecast report you can check them to see if the appropriate media is available.
2. On a standby computer, install the CommServe software. For more information on installing the CommServe, see CommServe Deployment.
3. Restore the CommServe database using the CommServe Disaster Recovery Tool from the Disaster Recovery Backup described in Step 1. (See Restore a Disaster Recovery Backup for step-by-step instructions.)
4. Verify and ensure that the **Bull Calypso Client Event Manager Bull Calypso Communications Service (EvMgrS)** is running.
5. If you did not have a CommCell Migration license available in the CommServe when the disaster recovery backup was performed, apply the IP Address Change license and the CommCell Migration license on the standby CommServe. See Activate Licenses for step-by-step instructions.
6. Export the data associated with the affected clients from the standby CommServe as described in Export Data from the Source CommCell.

When you start the Command Line Interface to capture data, use the name of the standby CommServe in the `commcell` argument.

7. Import the exported data to the main CommServe as described in Import Data on the Destination CommCell.

This will bring back the entity in the CommServe database and the entity will now be visible in the CommCell Browser. (Press F5 to refresh the CommCell Browser if the entity is not displayed after a successful merge.)

8. If you have additional data that was backed up after the disaster recovery backup and before the deletion of the entity, use the procedure described in Import Metadata from a Tape or Optical Media to obtain the necessary information.
9. You can now browse and restore the data from the appropriate entity.

As a precaution, mark media (tape and optical media) associated with the source CommCell as READ ONLY before performing a data recovery operation in the destination CommCell.

RESTORE COMPLETED WITH WARNINGS

Restore jobs from Oracle for Oracle RAC iDataAgent will be displayed as "Completed with Warnings" in the Job History in the following case:

- When the database restore succeeds but it fails when recovering the database or opening the database.

ORACLE ERRORS

If you receive an Oracle error during an Oracle restore operation, we recommend that you follow procedures published by Oracle Corporation on resolving the specific error. We also advise you to consult with your on-site Oracle database administrator, as needed.

