Backup and Restore of ISC Repository and Standby System

This chapter explains how to back up and restore your Sybase and Oracle databases and how to set up a standby system:

Backup and Restore of ISC Repository

The CCO location of scripts for these procedures is:

http://www.cisco.com/cgi-bin/tablebuild.pl/isc

The subsections are:

Data Items Included in Backup and Recovery

Most of the ISC-related data items are stored in a repository held on a relational database and the rest are stored in an operating system level file system. For ISC to function flawlessly on restart, following a crash, it is necessary that the proposed backup and recovery feature shall include various ISC-related data items as a whole. The underlying tasks involved in backup and recovery procedures would differ depending on the nature of persistence of these data items. However, these procedures shall work commonly for all the data items in a seamless and transparent manner.

The following data elements are included in ISC's backup and recovery plan:

1. **Main repository**: This repository consists of data items such as, Customers, VPNs, Policies, Devices, and Interfaces. This data is held on a RDBMS, either the embedded Sybase Adaptive Server Anywhere (ASA) database or the customer's Oracle database.
2. **SLA repository:** This repository consists of data items pertaining to Service Level Agreements (SLA) and Probes. This repository is held on a Sybase ASA database. This is the default repository for devices that do not have a Collection Server. There will be SLA repositories in each of the

collection server machines, if available. If your SLA repository is on one or more Collection Servers separate from the Main Server, you must run the backup on each Collection Server for the SLA repository.

3. Others: There are a few data items that are stored in the OS level file system under various ISC install directories, which would be part of the proposed backup and recovery plan.

Guidelines

For the backup and recovery plan to function efficiently, customers are requested to follow these guidelines:

1. Support exists for the following types of supported backups:
   a. **Full backup** is a complete backup of the ISC repository, ISC repository transaction logs, and other ISC data files held in the file system. It is recommended to have a full backup on a default weekly basis, which could be reconfigured as desired by the customer.
   b. Incremental backup is a backup of all the data from the time of the last full or incremental backup until this incremental backup. It is recommended that the full backup be interspersed with several incremental backups, by default, daily.
   c. Archive backup is a complete backup of all ISC data in respective archive files, typically on a tape drive. **Use this backup if you are backing up directly to a tape.**
   d. Live backup creates redundant copies of transaction logs to restore the ISC repositories held on a Relational Database Management System (RDBMS) and creates redundant copies of other ISC data held on the file system on the Main server machine. These redundant copies are typically set up on a secondary machine to restart ISC if the primary server machine becomes unusable.
2. The plan default schedule requires **Weekly FULL ONLINE** (while system is running) backups interspersed with **DAILY ONLINE** incremental backups of all ISC data items. An **ARCHIVE full** backup, preferably on a tape, is recommended on a MONTHLY basis. This archive tape backup should be stored in different premises to prevent any loss of backups in case of acts of physical disasters at the main server location.
3. It is important to keep more than one full backup to prevent accidental loss of backup copies.
4. Create archive backup copies on a tape device.
5. External factors such as available hardware, the size of database files, recovery medium, disk space, and unexpected errors can affect customers' recovery time. When implementing the plan, the customer shall allow additional recovery time for miscellaneous tasks that must be performed, such as entering recovery commands or retrieving, loading, and organizing tapes.

Sybase Database Backup and Restore

It is important to protect all ISC-related data by a well-defined backup and recovery plan. Data loss could occur due to the following reasons. The objective of ISC's backup and recovery plan is to greatly minimize the risk of data loss due to any of these reasons:

- Media failure
  - The disk drive holding database files and other data files becomes unusable.
  - The database files and other data files become corrupted due to hardware/software problems.

- System failure
  - A computer or operating system goes down while there are partially completed transactions.

The Sybase Backup and Restore tool provides a suite of scripts with several options to backup and restore your embedded Sybase database.

The backup script automatically detects whether a full backup is needed for this current backup week. If a full backup already exists for this current backup week, this script automatically takes an incremental backup. However, the user can force a full backup overriding this default behavior by changing the configuration setting.

Installing

Step 1 From the location http://www.cisco.com/cgi-bin/tablebuild.pl/isc, download the tar file iscBRToolASA.tar.gz and untar this file as follows:

mkdir -p $ISC_HOME/backup/Sybase

gzip -d < iscBRToolORA.tar.gz | tar xf -

Step 2 **chmod +x install**

Run install from where the tar file is unpacked. The install script takes command line arguments. Because "install" is also a system command, to differentiate between the system command and this installation script, run the script as follows:

```
./install -t <BACKUP_INSTALL_DIR>
```

For help in the install script, use **-h(elp)** as a command line argument

Sample Install Prompts and User Responses

The following is a sample install session:

```
#./install -t /users/yourname/iscBRToolInstall
```

When the install script is invoked as above, if the specified target install directory already exists, the user is prompted as follows:

```
Looks like the installation already exists
Do you want to continue installation - it might remove the existing contents [y,n,?]
removing the previous installation
Enter the Sybase User Name: DBA (user input)
Enter the Sybase User Password: SQL (user input)
Enter the Primary ISC Host Name: yourname-u10 (user input, the host name of the machine
running ISC)
Enter Primary ISC user/owner name: yourname (user input, the user/owner name of ISC on the
above host)
```

Post Install Status

The installation creates an env.sh script under <BACKUP_INSTALL_DIR>/BackupRestore/config directory.

Editing the env.sh script is NOT RECOMMENDED. This env.sh script sets the necessary environment variables needed to run ISC backup and restore scripts.

Functionality of Backup and Restore Tool

Step 1 One time configuration is needed before the first backup is carried out. Invoke the asa_configs.sh script to configure the backup and restore process. Execute this script from the directory **BACKUP_INSTALL_DIR/BackupRestore/scripts** as follows:

```
# ./asa_configs.sh
```

A sample configuration session is as follows, with the configuration prompt on the LHS and sample user response on the RHS of the prompt.

```
Starting backup Configuration for Main ISC database
DB server Name...yourname_yourname-u10

ISC Backup script invoked with the following parameters:
------------------------------------------------------------------------
Backup directory: /users/yourname/iscBRToolInstall/BackupRestore/Backups
Number of weeks to keep: 2
Backups archived to tape (0=no, 1=yes): 0
Tape device: /dev/rmt/0
Fail backup if there is not enough space for a full backup (0=no, 1=yes): 1
Delete old backups if not archived to tape (0=no, 1=yes): 0
Run validation routines on backup files (0=no, 1=yes): 0
Force full backup (0=no, 1=yes): 0
------------------------------------------------------------------------
The ISC backup configuration file is nonexistent ... creating new file
Modifying ISC backup configuration settings ...
Enter new ISC backup directory path (a subdirectory ISC will be added
automatically) [/users/yourname/iscBRToolInstall/BackupRestore/Backups] [?]
/users/yourname/iscBackup
Backup directory for ISC specified is "/users/yourname/iscBackup/ISCMain".
Is this correct? [y]  [y,n,?] y
Enter the number of weeks to keep [2] [?] 3
Number of weeks specified is "3".
Is this correct? [y]  [y,n,?] y
Old backups archived to tape (0=no, 1=yes) [0] [?]
Archive to tape option specified is "0".
Is this correct? [y]  [y,n,?] y
Enter tape device [/dev/rmt/0] [?]
Tape device specified is "/dev/rmt/0".
Is this correct? [y]  [y,n,?] y
Fail backup if there is not enough space for a full backup (0=no,1=yes) [1] [?]
Fail backup if not enough space specified is "1".
Is this correct? [y]  [y,n,?] y
Delete old backups if not archived to tape (0=no, 1=yes) [0] [?]
Delete old backups specified is "0".
Is this correct? [y]  [y,n,?] y
Run validation routines on backup files (0=no, 1=yes) [0] [?] 1
Run validation routines specified is "1".
Is this correct? [y]  [y,n,?]
Force full backup (0=no, 1=yes) [0] [?] 0
Force full backup specified is "0".
Is this correct? [y]  [y,n,?] y
ISC Backup configuration settings have been modified ...
If you wish to verify the values or modify them again then re-run the script
asa_configs.sh again
The ISC backup engine is now exiting without backing up the database.You must run the
asa_backup.sh script for the backup to take place.
ISC Backup Configuration Successfully completed
ISC Backup Configuration script ending.
```

```
Starting backup Configuration for SLA database
DB server Name...rpokalor_rpokalor-u10
SLA Backup script invoked with the following parameters:
-------------------------------------------------------------------------------
Backup directory: /users/yourname/iscBRToolInstall/BackupRestore/Backups
Number of weeks to keep: 2
Backups archived to tape (0=no, 1=yes): 0
Tape device: /dev/rmt/0
Fail backup if there is not enough space for a full backup (0=no, 1=yes): 1
Delete old backups if not archived to tape (0=no, 1=yes): 0
Run validation routines on backup files (0=no, 1=yes): 0
Force full backup (0=no, 1=yes): 0
-------------------------------------------------------------------------------
The SLA backup configuration file is nonexistent ... creating new file
Modifying SLA backup configuration settings ...
Enter new SLA backup directory path (a subdirectory SLA will be added
automatically) [/users/yourname/iscBRToolInstall/BackupRestore/Backups] [?]
/users/yourname/iscBackup
Backup directory for SLA specified is "/users/yourname/iscBackup/SLA".
Is this correct? [y]  [y,n,?] y
Enter the number of weeks to keep [2] [?] 3
Number of weeks specified is "3".
s this correct? [y]  [y,n,?] y
Old backups archived to tape (0=no, 1=yes) [0] [?]
Archive to tape option specified is "0".
Is this correct? [y]  [y,n,?] y
Enter tape device [/dev/rmt/0] [?]
Tape device specified is "/dev/rmt/0".
s this correct? [y]  [y,n,?] y
Fail backup if there is not enough space for a full backup (0=no,1=yes) [1] [?]
Fail backup if not enough space specified is "1".
Is this correct? [y]  [y,n,?] y
Delete old backups if not archived to tape (0=no, 1=yes) [0] [?]
Delete old backups specified is "0".
Is this correct? [y]  [y,n,?] y
Run validation routines on backup files (0=no, 1=yes) [0] [?]
Run validation routines specified is "0".
Is this correct? [y]  [y,n,?]
Force full backup (0=no, 1=yes) [0] [?]
Force full backup specified is "0".
Is this correct? [y]  [y,n,?]
LA Backup configuration settings have been modified ...
If you wish to verify the values or modify them again then re-run the script
asa_configs.sh again
The SLA backup engine is now exiting without backing up the database. You must run the
asa_backup.sh script for the backup to take place.
SLA Backup Configuration Successfully completed
SLA Backup Configuration script ending.
```

Post Configuration status

```
-------------------------------------------------------------------------
The configuration creates backupISC.config and backupSLA.config files under
BACKUP_INSTALL_DIR/BackupRestore/config directory.

To modify the initial configuration settings, users can either re-run the asa_configs.sh
script or simply modify the contents of these .config files. For example, if the user
wants to suppress the validation of the database after each backup, the config file
setting validateDB property to 0 instead of 1. Similarly, if the user wants to force full
backup, set the property fullBackup=1.
```

How to Use the Backup Script

Step 1 Run the **BACUP_INSTALL_DIR/BackupRestore/script/asa_backup.sh** script to initiate the backup task.

- The backup should be made while the ISC database server is running. There is no need to stop ISC to backup the database.
- The backup directory path specified during the configuration process should ideally be on an NFS device.

  It is important to keep the backup copies on an external storage device to protect the backup copies if the main ISC system crashes.
- Install the Backup and Restore tool and implement the periodic backup tasks from the primary ISC host machine. However, the backup task can be carried out from a secondary system, provided the following conditions are met:
  - The main ISC and SLA repository files should be placed on an NFS device accessible from the primary ISC host system and the secondary ISC host system.
  - The hardware and software configuration of the secondary system should be the same as the ISC primary host system.
  - The same version of ISC should be installed on both the primary the secondary systems.
  - The Backup and Restore tool should be installed on the secondary ISC system.

Step 2 Re-run the config script to make changes to the initial configuration settings, if needed.

Behavior of the Backup Process

1. The backup scripts follow a weekly backup scheme; the backup week begins on Sunday.
2. A full backup (both .db and .log files) are taken the first time the backup script is run during the backup week. Only incremental (only .log file) backup will be taken for the reminder of the current backup week.
3. You can force a full backup instead of an automatic incremental backup by setting the fullBackup property to 1 in the backupISC.config and backupSLA.config file, before running the asa_backup.sh script.
4. A new subdirectory (under the user specified backup directory) is created for each backup week. This directory is named as MM-DD-YYYY, where MM is the month and DD is the date of the Sunday of this backup week and YYYY is the year.
5. A subdirectory is created for each full backup and all the associated incremental backups under the above weekly directory. Each time a forced full backup is made for the current backup week, there is a new subdirectory created to contain this full backup and its associated incremental backups.The full backup directory for the current backup week is named full_0*n*.dir, where *n* is 1,2...9.

How to Restore the Database from the Backup

The asa_restore.sh script supports following types of database restore:

1. A restore of a previous Full or incremental backup.
2. A recovery from a media failure on the database file.

Note Note The main ISC repository consists of repository.db and repository.log files and the SLA consists of sla.db and sla.log files. ISC does not support placing the .db and .log files in different locations. Thus, if there is a media failure on the .db file, then the associated .log file also becomes unusable as well and thus this option may not be useful.

3. Run BACUP_INSTALL_DIR/BackupRestore/script/asa_restore.sh script to initiate the restore task after being sure to follow these pre-conditions:
   a. The database server of ISC should not be running. Failing to stop the database server will result in an inconsistent database after the restore.
   b. Follow the instructions and prompts carefully while running the scripts.
   c. Do not copy over or move or delete the repository files under $REPOSITORY_HOME.

Oracle Database Backup and Restore

From the location http://www.cisco.com/cgi-bin/tablebuild.pl/isc, download the tar file iscBRToolORA.tar.gz and untar this file as follows:

mkdir -p $ISC_HOME/backup/Oracle

gzip -d < iscBRToolORA.tar.gz | tar xf -

Oracle databases have a backup and restore Recovery Manager (RMAN) tool. To use this tool for online backup, the Oracle database must be in ARCHIVELOG mode, as explained in the "Turn On ARCHIVELOG Mode" section on page C-8. RMAN maintains the book keeping intelligence of backup and recovery files and backs up at the block level. Therefore, RMAN can significantly speed up backups and reduce the server load by using incremental backups.

RMAN for Oracle 8i is explained in the user guide, which is available as follows:

http://download-west.oracle.com/docs/cd/A87862_01/NT817CLI/server.817/a76990/toc.htm

Note RMAN is convenient to use. However, it only provides a command line interface. And it still demands database analyst knowledge when recovery is needed.

Be sure that the backup data and RMAN catalog are located on a different disk from where the Oracle database (data files, redo logs, and control files) are located. Both may reside on the same ISC database server.

RMAN configuration is explained in the following areas that should be implemented sequentially:

1. Turn On ARCHIVELOG Mode, page C-8
2. Create RMAN Catalog Database, page C-8
3. Create RMAN User, page C-8
4. Create RMAN Catalog, page C-8
5. Register the ISC Database with the RMAN Catalog, page C-9
6. Modify ISC Database Initial Parameter File, page C-9
7. Backup Database, page C-9
8. Recover Database, page C-10

Turn On ARCHIVELOG Mode

Oracle allows manual backup when turning on ARCHIVELOG mode. This makes the database log all transactions into the redo logs. When one log is full, a task is started to copy the redo log to an archive log directory and at the same time the system starts logging to a different redo log. This requires the user to manage and purge archive logs that are no longer needed.

Step 1 First, turn on the archive log mode:

- **startup mount;**
- **alter database archivelog;**
- **archive log start;**

Check archive log using ‘archive log list’.

Step 2 Copy the data files regularly:

- turn the tablespace into ‘backup’ mode
- show data files, as follows:

  SVRMGR> **select file_name from dba_data_files;**

Step 3 To recover, enter the following:

SQL> **recover datafile** <*file_number_or name*>**;**

where <*file_number_or name*> is the file number, however a file name can be placed here. Recovery will be from ‘/var/tmp/oracle/backup’ and the specified data file, where the recover command determined that the redo is needed for the recovery in the archive log.

Create RMAN Catalog Database

The catalog database holds the recovery catalogs. This database typically is set up on a different server from any database being registered in it. It also works if this database is set up on the same database server as the ISC database.

Use the Oracle utility **dbassist** to create a catalog database. (This is the same as ISC database creation, except you should name the RMAN global name ‘rman’, and you should name the SID ‘rman’.)

Create RMAN User

Creating an RMAN user is the same as creating an ISC user on an ‘rman’ database. Name the RMAN user ID ‘rman’ and name the password ‘rman’. Make sure ‘rman’ has proper privileges. For example:

SQL> **grant connect, resource, recovery_catalog_owner to rman;**

Create RMAN Catalog

Create a catalog from the RMAN command prompt:

RMAN> **connect catalog rman/rman@rman**

RMAN> **create catalog;**

Register the ISC Database with the RMAN Catalog

Set the ORACLE_SID environment variable = isc.

%**rman**

RMAN > **connect catalog rman/rman@rman**

RMAN > **connect target sys/change_on_install**

RMAN > **register database;**

The default password for an Oracle sys account after Oracle installation is 'change_on_install'. Replace this sys account password with the correct sys account password for the ISC database.

Modify ISC Database Initial Parameter File

To modify the ISC database initial parameter file, do the following:

Step 1 Enter the following:

log_archive_start = true;

log_archive_dest_1 = "location=/var/tmp/oradata/arch"

log_archive_format = arch_%t_%s.arc

Step 2 Restart the ISC database server with the ARCHIVELOG mode turned on, as follows:

startup mount

alter database archivelog;

alter database open

Step 3 Check the archive log mode as follows:

SQL> **archive log list;**

Backup Database

To backup the database, do the following:

Step 1 Before you run the backup scripts, make sure you update the $ISC_HOME/backup/Oracle/backupenv.properties file.

Use a text editor to open this file and read the directions on how to update each property.

Step 2 To perform a full database backup, execute the following:

$ISC_HOME/backup/Oracle/oracle_backup.sh -f

Step 3 You can perform incremental backups after a minimum of one full backup. To perform an incremental backup, execute the following:

$ISC_HOME/backup/Oracle/oracle_backup.sh -i

Note These backup scripts can be run as cron jobs or scheduled by the ISC task manager.

Recover Database

To recover a database, do the following:

Step 1 Stop the ISC watchdog before recovering a database, as follows:

stopall

Step 2 To recover a database, you can execute the following from the location $ISC_HOME/backup/Oracle/oracle_recover.sh

%**oracle_recover.sh** [**"**<*date_time*>**"**]

The **"**<*date_time*>**"** is optional. The format is **"**mmm dd yyyy hh**:**mm**:**ss**"**, where the first mmm is the month and must be alphabetic characters with an initial capitalization, for example:

"Oct 09 2003 15:25:00"

If you do not specify <*date_time*>, the script does a full database recovery.

Note Note: Do not stop the Oracle Listener during restore.

Standby System for ISC (Secondary System)

This section explains how to set up Sybase and Oracle standby systems for ISC.

The subsections are:

Sybase Standby System Set Up

The explanation of setting up a Sybase standby system is explained as follows:

Running Live Backup of ISC Databases

1. Run BACKUP_INSTALL_DIR/BackupRestore/scripts/asa_liveBackup.sh to start the live backup after being sure to follow these pre-conditions:
   a. First set up a standby ISC system.

b. The standby system should be similar to the primary ISC host system in hardware and software configurations.

c. The ISC primary and standby systems should be on the same LAN.

d. ISC software should be installed on the secondary system and the version of ISC on the primary and standby systems should be the same.

e. The backup and restore tool should be installed on the primary and the secondary systems.

f. The live backup should be started from the secondary system only, you should not run the live backup from ISC primary system.

g. The storage device where the regular backup copies are placed should be accessible from the standby system.

h. You *must* run BACKUP_INSTALL_DIR/BackupRestore/scripts/asa_liveBackupConfig.sh to configure the live backup on the standby system before starting the live backup for the first time.

i. The ISC database server must be running on the primary ISC host before starting the live backup on the standby system.

j. The live backup will stop when the ISC database server is stopped and should be restarted after restarting ISC.

k. At least one full backup must be taken before starting the live backup.

l. Regular periodic full/incremental backups should be taken even if the live backup is running on the secondary system.

m. There should not be more than one live backup running simultaneously.

How to Restore the Database from the Live Backup

When the primary ISC host fails, the standby system restores the database from the latest available full backup, the latest incremental backup, and the live backup.

1. Run the BACUP_INSTALL_DIR/BackupRestore/script/asa_restoreFromLiveBackup.sh script on the standby system to restore the database after being sure to follow these pre-conditions:
   a. At least one full backup copy should be available to restore the database.
   b. If more than one backup copy is available, use only the latest full backup and the latest associated incremental backup.
   c. Run the restore from the standby machine

Oracle Standby System Set Up

Oracle standby is explained in the following sequential three subsections:

1. Oracle 8i Setup, page C-12
2. Activate Oracle Standby Database, page C-14
3. Restart ISC, page C-14

Oracle 8i Setup

Oracle 8i supports the standby database. The standby database is identical to the primary database, including the database name. When archive log files are generated on the primary database, they are transferred and applied to the standby database. If the primary database has a failure that cannot be resolved quickly, the standby database can be activated as a failover solution.

The ISC secondary server should copy the entire ISC directory on the ISC primary server machine. Follow the steps below to set up the standby Oracle database.

Note The standby database machine *must* have the same machine architecture and operating system as the primary database machine.

Refer to the Oracle Standby document for details:

http://download-east.oracle.com/docs/cd/A87860_01/doc/server.817/a76995/toc.htm

There are several Oracle standby setups. In this document, only one setup is described (that is, one standby instance on a remote machine).

Step 1 Shutdown ISC, as follows:

stopall

Step 2 Generate a standby control file.

a. On the primary machine, start sqlplus as a database administrator (DBA).

SQL> **alter database create standby controlfile as '/var/opt/oracle/8.1.7/dbs/standby.ctl' reuse;**

This creates a control file to be used by the standby system. Copy this file to the same location on the standby system.

Step 3 Copy files from the primary database machine.

a. Bring down the primary Oracle server, as follows:

sqlplus > **shutdown immediate**

b. Copy the entire Oracle directory from the primary system to the same location on the standby system. If the data files reside on a directory other than a subdirectory of $ORACLE_HOME, also copy these files.

Step 4 Enable the primary database to write to the secondary database.

a. On the primary machine, add the following to the parameter file:

log_archive_dest_2 = "service=standby mandatory reopen=60"

log_archive_dest_state_2 = enable

b. Also on the primary machine, add the following to $(ORACLE_HOME)/network/admin/tnsname.ora.:

```
STANDBY =
 (DESCRIPTION =
   (ADDRESS_LIST=(PROTOCOL = TCP)(HOST=<standbyhost>)(PORT=1521))
   )
 (CONNECT_DATA =
   (SERVICE_NAME = standby)
```

)
)

where: <*standbyhost*> is the standby host name.

c. Restart listener after the update.

Step 5 Configure the standby parameter file, as follows:

a. Edit the Oracle parameter file (init*SID*.ora, where **isc** is the *SID* for ISC).

service_names = standby

control_files = ("*/var/opt/oracle/8.1.7/dbs/standby.ctl*")

standby_archive_dest = */var/tmp/oradata/isc/arch*

log_archive_dest_1 = "location=*/var/tmp/oradata/isc/arch*"

log_archive_start = false

Step 6 Update the listener.ora on the standby server, as follows:

a. Add the following to the listener.or:

STANDBY_LISTENER = (ADDRESS_LIST =

(ADDRESS=(PROTOCOL=TCP)(PORT=1521)(HOST=<*standbyhost*>)))

SID_LIST_STANDBY_LISTENER =

(SID_LIST =

(SID_DESC =

(ORACLE_HOME=*/var/opt/oracle/8.1.7*)

(SID_NAME = isc)

)

)

where <*standbyhost*> is the standby host name.

b. Restart listener after the update.

Step 7 Start the standby database, as follows:

On the standby machine, enter:

% **sqlplus**

SQL> **connect internal;**

SQL> **startup nomount;**

SQL> **alter database mount standby database;**

SQL> **recover managed standby database;**

Note The last command never returns because it consistently applies the archived logs.

Step 8 Start up the primary database server, as follows:

a. Enter the following:

% **sqlplus**

SQL> **connect internal;**

SQL > **startup;**

b. Check that the standby database is receiving archived logs in the standby_archive_dest directory specified in the parameter file.

Activate Oracle Standby Database

To activate an Oracle standby database, do the following:

Step 1 In a test environment, if you want to get the very latest data flushed out to the standby system, you need to execute the following on the primary system:

SQL> **alter system archive log current;**

Step 2 Open another window on the standby machine and enter:

% **sqlplus**

SQL> **connect internal;**

SQL> **recover managed standby database cancel;**

SQL> **alter database activate standby database;**

SQL> **shutdown immediate;**

SQL> **startup;**

After you activate the standby database, it ceases to be a standby database and becomes a fully functional primary database. Because standby database activation is a unidirectional operation, you cannot return the new primary database to the standby mode. In other words, you cannot perform a failover and then undo it.

Step 3 After solving the problem at the original primary site that necessitated the failover operation, you have the option of re-creating the primary database on the original primary site. Perform the following steps, assuming the original primary site was on node A and the activated standby site is on node B.

a. Make a consistent backup of the activated standby database on node B.

b. Restore the backup created on node B to node A.

c. Shut down the activated standby database on node B.

d. Open the restored database on A. It is now the primary database.

e. Make a backup of the database on node A.

f. Use the backup of A to re-create the standby database on node B.

Restart ISC

When the standby database is activated, use the following commands to point ISC to the new database server:

stopall -y

update $ISC_HOME/etc/install.cfg and replace <old_db_server> with <new_db_server>.

execute applycfg.sh

initdb.sh

startwd