



Backup and Restore of Prime Provisioning Repository

This chapter explains how to back up and restore your Sybase and Oracle databases. When upgrading from an earlier release of Prime Provisioning, the existing backup script will no longer work. From Prime Provisioning 6.8, you can install the new backup script that is available along with the downloaded version of the product. You can download the latest version of the product from the Software Download page.

This chapter contains the following sections:

- [Data Items Included in Backup and Recovery, page 4-1](#)
- [Guidelines, page 4-2](#)
- [Sybase Backup and Restore Process Overview, page 4-2](#)
- [Sybase Database Backup and Restore, page 4-15](#)
- [Oracle Database Backup and Restore, page 4-23](#)

Data Items Included in Backup and Recovery

Most of the Prime Provisioning-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 Prime Provisioning to function flawlessly on restart, following a crash, it is necessary that the proposed backup and recovery feature include various Prime Provisioning-related data items as a whole. The underlying tasks involved in backup and recovery procedures 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 Prime Provisioning's backup and recovery plan:

1. **Main repository:** This repository consists of data items such as Customers/Organizations, VPNs, Policies, Devices, and Interfaces. This data is held on an RDBMS, either the embedded Sybase 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.
3. **Others:** There are a few data items that are stored in the OS level file system under various Prime Provisioning install directories, which would be part of the proposed backup and recovery plan.

Guidelines

This section explains how to use the supported backup methods in Prime Provisioning.

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

-
- Step 1** Support exists for the following types of supported backups:
- a. **Full backup** is a complete backup of the Prime Provisioning repository, Prime Provisioning repository transaction logs, and other Prime Provisioning 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 Prime Provisioning 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 Prime Provisioning repositories held on a Relational Database Management System (RDBMS) and creates redundant copies of other Prime Provisioning data held on the file system on the Main server machine. These redundant copies are typically set up on a secondary machine to restart Prime Provisioning if the primary server machine becomes unusable.
- Step 2** The plan default schedule requires **Weekly FULL ONLINE** (while system is running) backups interspersed with **DAILY ONLINE** incremental backups of all Prime Provisioning 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.
- Step 3** It is important to keep more than one full backup to prevent accidental loss of backup copies.
- Step 4** Create archive backup copies on a tape device.
- Step 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 Backup and Restore Process Overview

This section describes how to back up and restore Sybase ASA for an Prime Provisioning installation. This section contains the following sections:


- [Overview of the Backup and Restore Process, page 4-3](#)
- [Planning your Backup and Restore Process, page 4-3](#)
- [Configuring Sybase Credentials, page 4-4](#)
- [Installing the Backup and Restore Tool, page 4-4](#)
- [Configuring the Backup and Restore Process, page 4-6](#)
- [Understanding the Backup Process Flow, page 4-7](#)

- [Understanding the Restore Process Flow, page 4-10](#)

Overview of the Backup and Restore Process

Figure 4-1 shows an overview of the Sybase ASA backup and restore process.

Figure 4-1 Overview - Sybase ASA Backup and Restore



Planning your Backup and Restore Process

Before backing up and restoring your Sybase installation, you must first prepare a plan. To prepare your plan, follow these steps:

-
- Step 1** Determine the frequency for full backups.
 - Step 2** Determine the frequency for incremental backups.
 - Step 3** Determine the location for storing the backups.



Note The file system must be accessible by the primary Prime Provisioning production machine and the secondary system (if you want to run the restore process from the secondary system or you want to perform a live backup).

- Step 4** Document the information for [Step 1](#) to [Step 3](#).
- Step 5** Set up the proper bookkeeping for your backup and restore procedure.

Configuring Sybase Credentials

In Prime Provisioning, `changeDBPassword.sh` script has been introduced to update Sybase database password. This script is available in `<PRIMEF_HOME>/bin`. It allows only the installation owner to execute it with the initial credentials. It also validates if the database is running properly.

Adhere to the rules below, while updating the Sybase password:

- You should not enter a simple password or any of your old passwords.
- Password must contain at least: one uppercase letter, one lowercase letter, and one number.
- Password length should be in the range of 8 to 30 characters.
- Special characters and ASCII character set are not allowed.

Once the password is updated, you need to validate the backup and restore process, and upgrade Paths. Use the following steps to execute the script:

-
- Step 1** Navigate to the directory where the latest version of Prime Provisioning software is installed.
- Step 2** In `<PRIMEF_HOME>/bin`, execute the following script:
- `./changeDBPassword.sh`**
- Step 3** Enter the valid username and password.
- Step 4** After the validation, you will be prompted with the following message “Enter database new password:”.
- Step 5** Enter a valid new password that adhere to the guidelines mentioned above.

Prime Provisioning gets restarted automatically once the new password is updated in the database and the properties file.

As a part of this feature, few scripts related to Backup and Restore tool installation has been modified.




Note You will be no longer prompted with the Sybase credentials during the installation of the Backup and restore tool.

Installing the Backup and Restore Tool

Figure 4-2 shows the process flow for installing the backup and restore tool.

Figure 4-2 Installing the Backup and Restore Tool

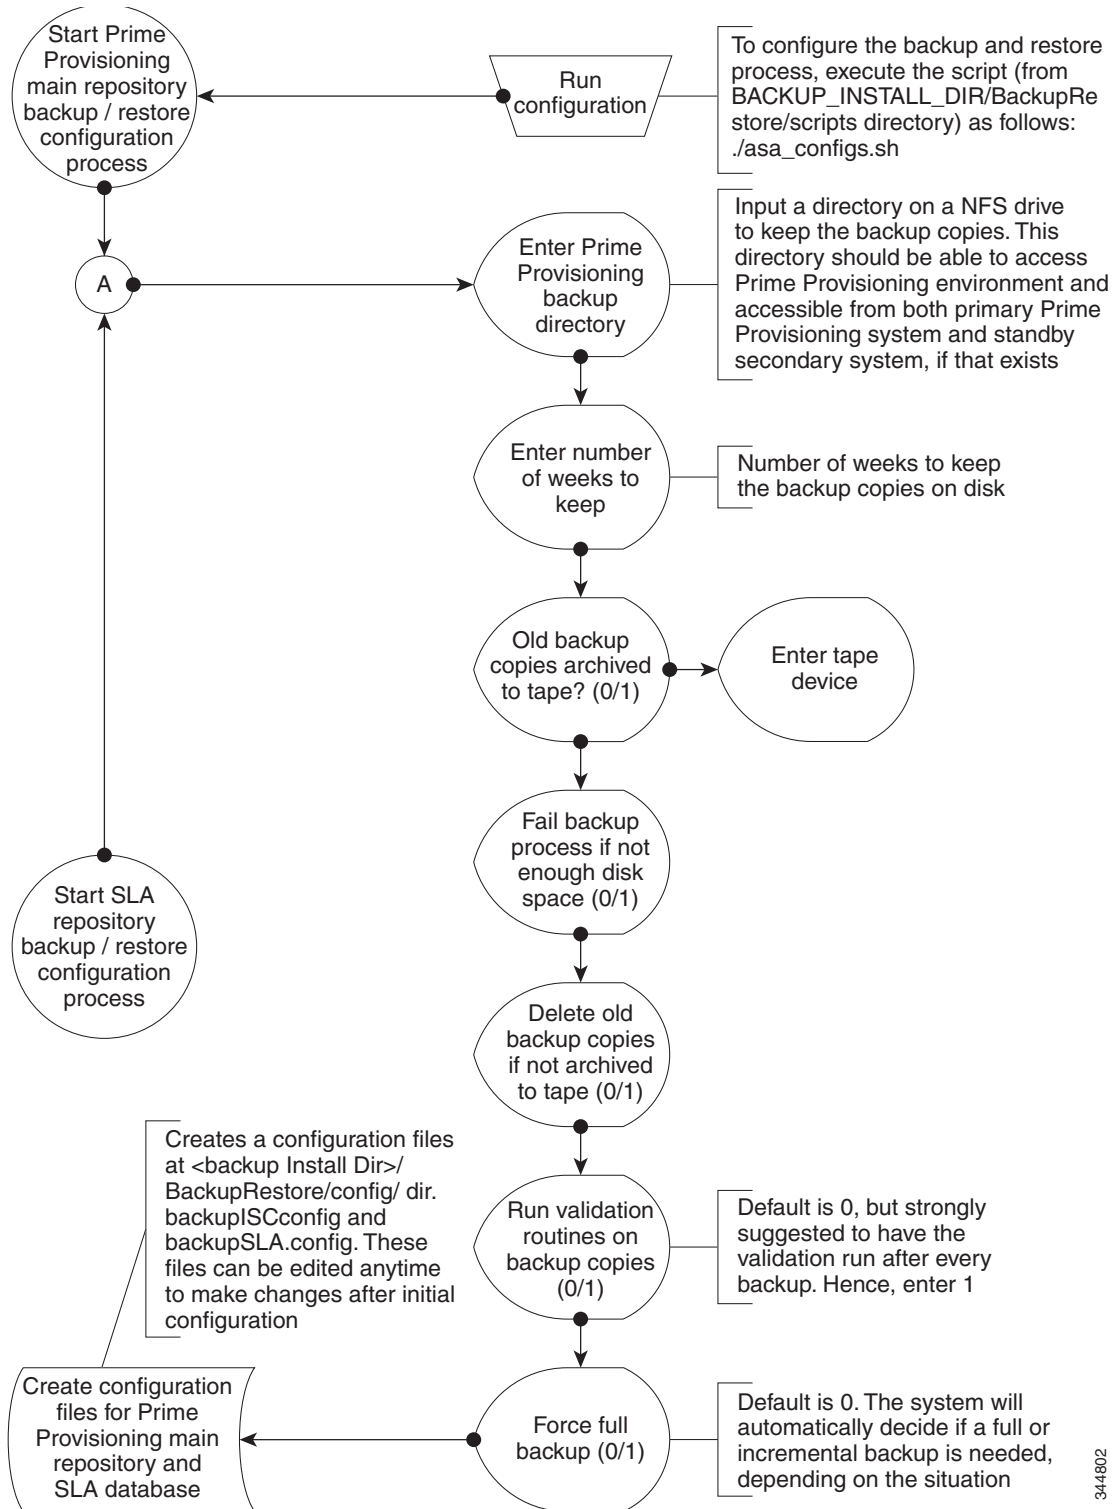


344801

Configuring the Backup and Restore Process

Figure 4-3 shows the one-time configuration process for the backup and restore.

Figure 4-3 One-Time Configuration Process Flow



344802

Understanding the Backup Process Flow

This section contains the following sections:

- [Preconditions, page 4-7](#)
- [Functions, page 4-7](#)
- [Full Backup Scheme, page 4-8](#)
- [Incremental Backup Scheme, page 4-8](#)
- [Typical Backup Directory Structure, page 4-9](#)

Preconditions

Before backing up your Sybase installation, you must observe the following preconditions:

1. The backup task must be carried out while the Prime Provisioning database server is running.
2. The backup directory path that you specify during the configuration must be on a Network File System (NFS) drive.
3. The backup and restore tool must be installed and accessible by both the primary and secondary systems.
4. The backup and restore tasks must be carried out from the Prime Provisioning primary machine. However, the live backup and restore is done from the secondary system.
5. You must not modify, rename, or move the backup directory structure after you configure it.

Functions

1. The backup follows a weekly scheme.
2. The backup week begins every Sunday.
3. A full backup occurs automatically the first time a backup is run for the backup week.
4. After the full backup, only incremental backups occur for the remainder of the week.
5. You can force a full backup during the week by changing the configuration setting to `fullBackup=1` before running the backup script.
6. A new subdirectory is created for every backup week under the backup directory specified during the configuration. The name has the format `mm-dd-yyyy`, where the date is Sunday of the current backup week.
7. A new subdirectory is created for each full backup created during the backup week. All the associated incremental backup copies are also kept under this directory. If a full backup is forced during the same backup week, a new subdirectory is created for the full backup and after associated incremental backups.



Note Do not modify, rename, delete, or move the directory structure created by the backup tool.


8. Both the database and the transaction log are backed up in a full backup.
9. Only the transaction log is backed up in an incremental backup.

10. The transaction log is truncated after each backup, either full or incremental. In other words, the transaction log is started fresh after each backup.
11. The name of the log file after backup will be of the form `yymmddnn.log`, where `yy` is the year, `mm` is the month, and `dd` is the day on which the backup is taken and `nn` is the serial number of this backup on a given day.

Full Backup Scheme

Figure 4-4 shows a full backup scheme.

Figure 4-4 Full Backup Scheme




116522

Incremental Backup Scheme

Figure 4-5 shows an incremental backup scheme.

Figure 4-5 Incremental Backup Scheme



Typical Backup Directory Structure

To create a backup directory structure on an NFS drive, you can use the following procedure.

Assume the Backup Week is 03/14/2010 through 03/20/2010 and the Backup Dir as specified during configuration is /auto/PrimeProvisioningBackups (NFS drive). The system creates two subdirectories under user specified backup dir, ISCMMain and SLA.

1. First backup run on 03/15/2010 Monday, default full backup. Creates a sub dir /03-14-2010/full_01.dir under ISCMMain and SLA directories.
2. Second backup run on the same date 03/15/2010, default incremental backup.
3. Third backup run on 03/17/2010, default incremental backup.
4. Fourth backup, Forced FULL backup (after changing configuration file setting, fullBackup to 1) on 03/18/2010. Creates a new sub dir /03-14-2010/full_02.dir under ISCMMain and SLA directories.



Note Configuration setting, full backup reset to 0.

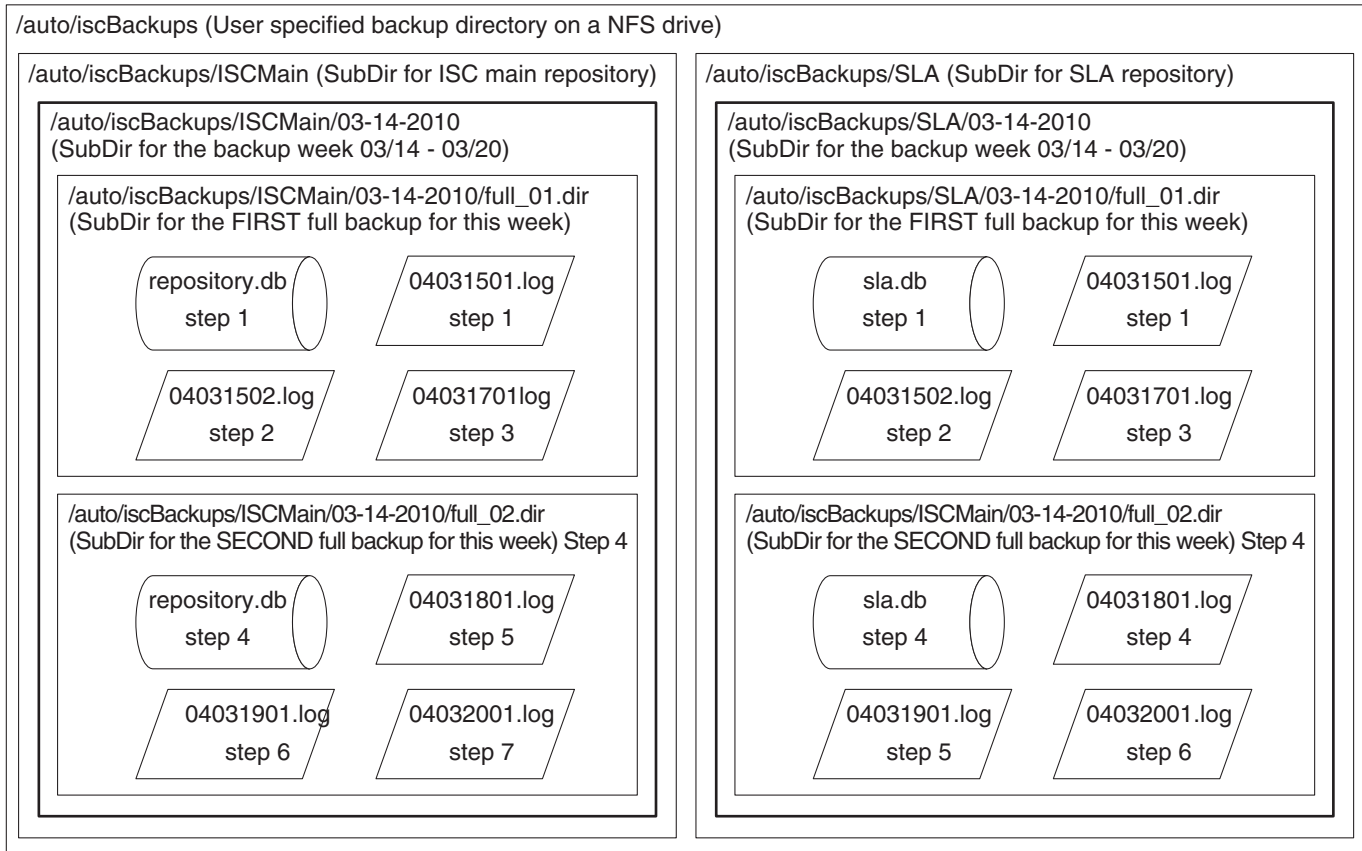
5. Fifth backup, run on 03/19/2010, default incremental backup.
6. Sixth backup, run on 03/20/2010, default incremental backup.



Note Backup Week ended on 03/20/2010.

Figure 4-6 shows a typical backup directory structure on an NFS drive.

Figure 4-6 Typical Backup Directory Structure



Understanding the Restore Process Flow

This section contains the following sections:

- [Preconditions, page 4-10](#)
- [Functions, page 4-11](#)
- [Restore from Media Failure, page 4-11](#)
- [Restore to a Desired Point-in-Time, page 4-13](#)

Preconditions

Before restoring your Sybase installation, you must observe the following preconditions:

1. The Prime Provisioning database server should be stopped while running the Restore task.

2. The backup directory path that you specify during the configuration must be on a Network File System (NFS) drive.
3. The backup and restore tool must be installed and accessible by both the primary and secondary systems.
4. The backup and restore tasks must be carried out from the Prime Provisioning primary machine. However, the live backup and restore is done from the secondary system.
5. The user running the restore script needs write permissions on the \$REPOSITORY_HOME directory.
6. The repository files shall have write permission for the user running the restore.
7. Do not modify, rename, or move the backup directory structure after configured.
8. Do not rename, move, or delete the backup copies of the repository files.
9. Do not move, rename, or delete the production repository files under \$REPOSITORY_HOME.

Functions

1. Restores the repository from existing full and incremental backup copies.
2. At least one full backup copy should be available to restore the repository.
3. The repository can be restored to a desired point in time using the available backup copies.
4. The restore process can recover the repository if there is a media failure on the database file, repository.db and/or sla.db.
5. The restore process cannot recover the repository if there is a media failure on the transaction log file. In this case, one of the following should be done to recover the database until the most recent checkpoint (partial recovery only):
 - a. Using the available backup copies, the repository can be restored to a desired point in time. Use the Prime Provisioning restore script to do this.
 - b. Make an extra backup copy of the database file immediately. When the transaction log is gone, the only record of the changes between the last backup and the most recent checkpoint is in the database file. Delete or rename the transaction log file. Restart the database with the -f switch. For example, \$SYBASE_HOME/bin/dbsrv8 \$REPOSITORY_HOME/repository.db -f



Note Please see Sybase ASA documentation for more information.




Note This option should be done by an authorized database administrator only.

Restore from Media Failure

Figure 4-7 shows the process flow for how to restore from a media failure on the database file (.db).

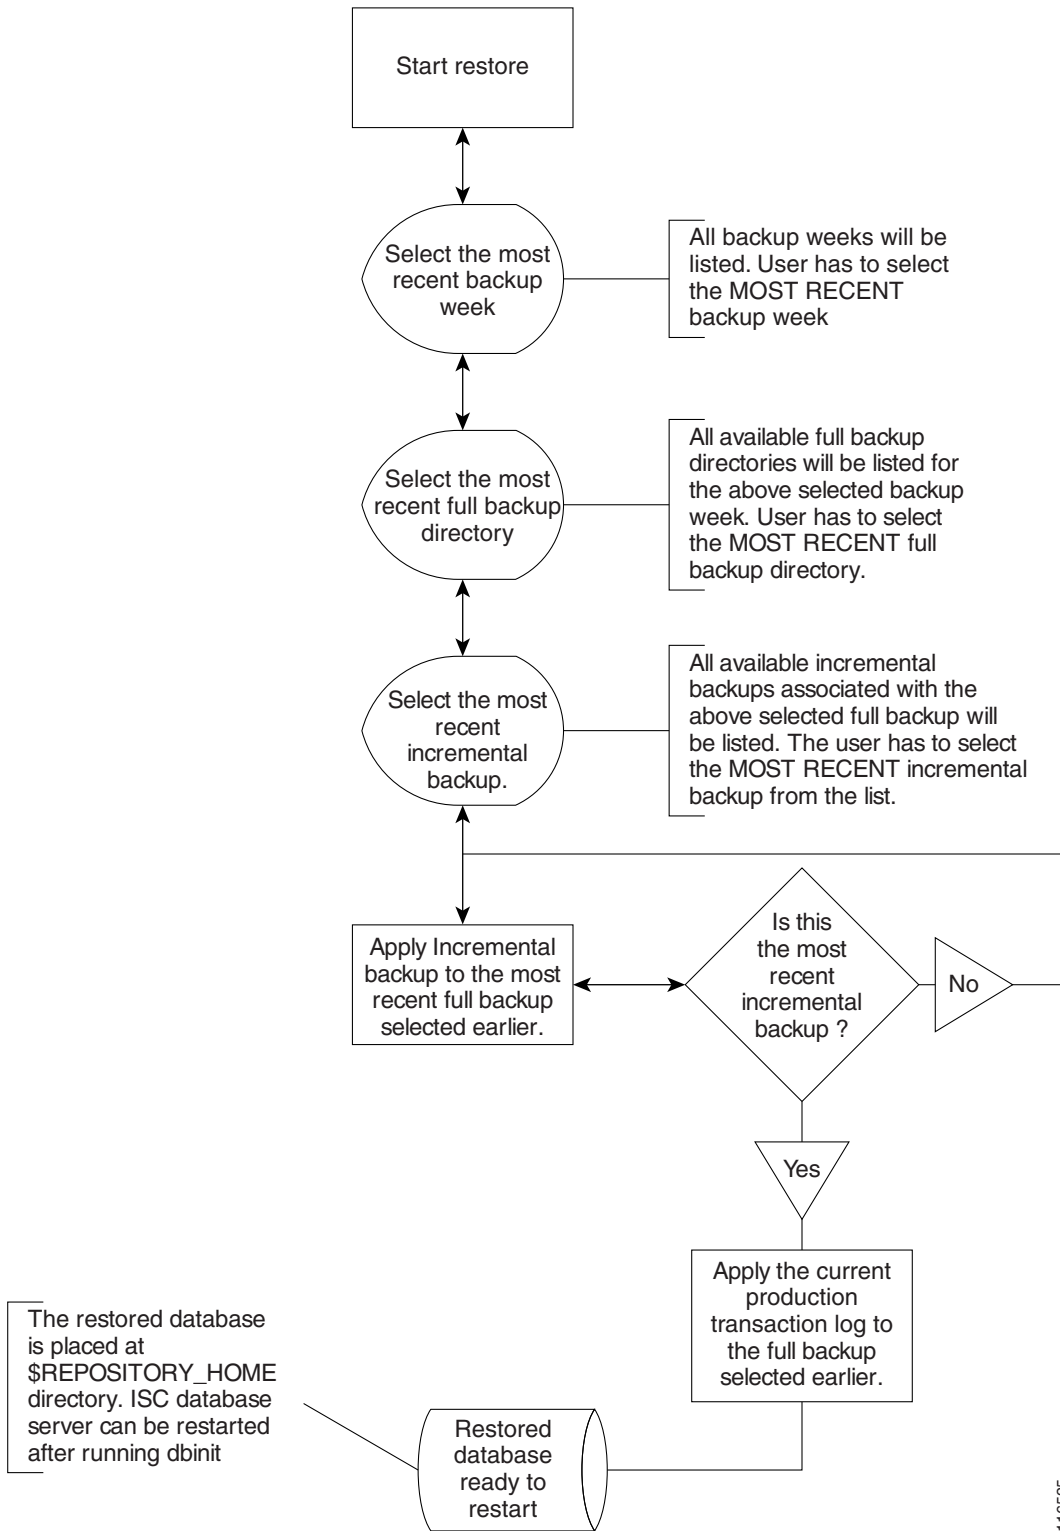
Figure 4-7 Restore from Media Failure on the Database File (.db)



Restore to a Desired Point-in-Time

Figure 4-8 shows the process flow for how to restore from a desired point-in-time.

Figure 4-8 Restore the Database to a Desired Point-in-Time



Sybase Database Backup and Restore

It is important to protect all Prime Provisioning-related data by a well-defined backup and recovery plan. Data loss could occur due to the following reasons. The objective of Prime Provisioning'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 or 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 back up 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 the Sybase Backup and Restore Tool

Step 1 In Prime Provisioning 6.8, Sybase backup and restore tool for Linux is available as `iscBRToolASA.tar` file in the new Prime Provisioning installation directory.



Note

In Prime Provisioning 6.8.1, Sybase backup and restore tool (`iscBRToolASA.tar`) location has changed due to Sybase 16 upgrade and it needs to be copied from Prime Provisioning 6.8.1 image location. Earlier the location was Prime Provisioning installation directory.

Prime Provisioning 6.8.1 image location is available at:

<https://software.cisco.com/download/release.html?i=!y&mdfid=286308515&softwareid=284525453&release=6.8.1&os=>

Step 2 Untar this file as follows:

```
mkdir -p $PRIMEP_HOME/backup/Sybase
```

```
gzip -d < iscBRToolASA.tar | tar xf -
```

Step 3 `./prime.sh` shell

Run `install` from where the tar file is unzipped. 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 "path"
```

where:

"path" is the location the installation files will be available after successful installation.

`<BACKUP_INSTALL_DIR>` must be NFS accessible by both the primary and secondary systems.

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 /opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup

/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup not exists do you want to
create it and continue? [y/n] [y]
Please Enter "y" for yes or "n" for no: y

install directory is: /opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup
Give the Database Details to Generate and source the env file
Host name: pff-lnx-vm022
Enter Primary Cisco Prime Provisioning user/owner name[root]:
Enter Database Server Name [root_pff-lnx-vm022]:
Cisco Prime Provisioning Primary Main Database Location:
/opt/prime-provisioning-6_7/Repository
Primary SLA Database Location:/opt/prime-provisioning-6_7/Repository
Cisco Prime Provisioning Primary Main Database Transaction
Location:/opt/prime-provisioning-6_7/Repository
Primary SLA Database Transaction Location:/opt/prime-provisioning-6_7/Repository
Enter Database Restore Directory Name:/opt/prime-provisioning-6_7/Repository
Note- By default the "Database Restore Directory Name" is <PRIMEF_HOME>/Repository.

```

Post Install Status

The installation creates an env.sh script under the <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 Prime Provisioning backup and restore scripts.

Adding PATH Statement

After installing the Prime Provisioning Backup and Restore tool and before configuring it, the PATH statement:

```
PATH=$PATH:/BackupRestore/scripts:/BackupRestore/config:/BackupRestore/bin export PATH
```

should be added to the login .profile file of the user iscadm.

Without this permanent addition, later runs of the backup and restore may fail.

Configuring the Sybase Backup and Restore Tool

A one-time configuration is needed before the first backup is carried out.

- Step 1** 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 Prime Provisioning database
DB server Name...yourname_yourname-u10
```


Cisco Prime Provisioning Backup script invoked with the following parameters:

```
-----
Backup directory:
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/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 Cisco Prime Provisioning backup configuration file is nonexistent... creating new file
Modifying Prime Provisioning backup configuration settings ...
Enter new Prime Provisioning backup directory path (a subdirectory Cisco
Prime Provisioning will be added automatically.)
[/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups] [?]
Enter the number of weeks to keep [2] [?] 2
Old backups archived to tape (0=no, 1=yes) [0] [?]
Enter 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] [?]
Delete old backups specified is "0".
Run validation routines on backup files (0=no, 1=yes) [0] [?] 1
Force full backup (0=no, 1=yes) [0] [?] 0
Confirm these values
[1]Cisco Prime Provisioning backup directory path =
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/ISCMai
n
[2]number of weeks to keep = 2
[3]Old backups archived to tape = 0
[4]tapeDevice = /dev/rmt/0
[5]Fail backup if there is not enough space for a full backup= 1
[6]Delete old backups if not archived to tape = 0
[7]Run validation routines on backup files = 0
[8]Force full backup = 0
    Do you want to modify any values ? [n] [y,n,?]
Cisco Prime Provisioning 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 Prime Provisioning backup engine is now exiting without backing up the database.You
must run the asa_backup.sh script for the backup to take place.
Prime Provisioning Backup Configuration Successfully completed
Prime Provisioning Backup Configuration script ending.
Starting backup Configuration for SLA database
DB server Name...yourname_yourname-u10
SLA Backup script invoked with the following parameters:
-----
Backup directory:
/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/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)
[/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups] [?]
```

```

Enter the number of weeks to keep [2] [?] 3
Old backups archived to tape (0=no, 1=yes) [0] [?]
Archive to tape option specified is "0".
Enter 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] [?]
Confirm these values
[1]SLA backup directory path =
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/SLA
[2]number of weeks to keep = 2
[3]Old backups archived to tape = 0
[4]tapeDevice = /dev/rmt/0
[5]Fail backup if there is not enough space for a full backup= 1
[6>Delete old backups if not archived to tape = 0
[7]Run validation routines on backup files = 0
[8]Force full backup = 0
Do you want to modify any values ? [n] [y,n,?]
SLA 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

The backup script is used as follows:

-
- Step 1** Run the <BACKUP_INSTALL_DIR>/BackupRestore/script/asa_backup.sh script to initiate the backup task.
- a. The backup should be made while the Prime Provisioning database server is running. There is no need to stop Prime Provisioning to back up the database.
 - b. The backup directory path specified during the configuration process *must* 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 Prime Provisioning system crashes.
 - c. Install the Backup and Restore tool and implement the periodic backup tasks from the primary Prime Provisioning host machine. However, the backup task can be carried out from a secondary system, provided the following conditions are met:

- The main Prime Provisioning and SLA repository files should be placed on an NFS device accessible from the primary Prime Provisioning host system and the secondary Prime Provisioning host system.
- The hardware and software configuration of the secondary system should be the same as the Prime Provisioning primary host system.
- The same version of Prime Provisioning should be installed on both the primary and secondary systems.
- The Backup and Restore tool should be installed on the secondary Prime Provisioning system.

Step 2 Rerun the config script to make changes to the initial configuration settings, if needed.

Behavior of the Backup Process

- Step 1** The backup scripts follow a weekly backup scheme; the backup week begins on Sunday.
- Step 2** A full backup (both .db and .log files) is taken the first time the backup script is run during the backup week. Only incremental (only .log file) backups are taken for the remainder of the current backup week.
- Step 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.
- Step 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.
- Step 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_0n.dir, where *n* is 1,2...9.
-

How to Restore the Database from the Backup

The **asa_restore.sh** script supports the 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

The main Prime Provisioning repository consists of repository.db and repository.log files and the SLA consists of sla.db and sla.log files. Prime Provisioning 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 and thus this option might not be useful.

- Step 1** Run **<BACKUP_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 Prime Provisioning should not be running. Failing to stop the database server results in an inconsistent database after the restore.

- b. Follow the instructions and prompts carefully while running the scripts.
- c. Do not copy, move, or delete the repository files under **\$REPOSITORY_HOME**.

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 restore of Main Cisco Prime Provisioning database
DB Server Name...root_pff-lnx-vm022
Restore Script
This script will restore the Sybase database from a previous backup
that is saved on the system. The database backup consists of two
files; the primary database file and the transaction log file. The
transaction log is a separate file from the database file. It
stores all changes to the database. Inserts, updates, deletes,
commits, rollbacks, and database schema changes are all logged. The
transaction log is also called the forward log or the redo log.
For a incremental backup the transaction log file is the only file saved.
This script can restore a database from a previous backup as desired or if you
had a media failure (database crash) on the database file located in
/opt/prime-provisioning-6_7/Repository/repository.db.

This restore utility is not able to recover the database if you had
a media failure on the transaction log file. You can perform a partial
recovery by renaming the current transaction log file and starting the
database with the -f switch.
(eg: /opt/prime-provisioning-6_7/thirdparty/sybase/bin/dbsrv11
/opt/prime-provisioning-6_7/Repository/repository.db -f)

If you are trying to recover from a media failure(database crash) on the database file
then you need to select the *latest full backup* and the *latest incremental backup*.
Then this script will first restore the database up to the point of latest incremental
backup
and then will attempt to recover the transactions between the last (latest) incremental
and the point of failure. Please answer "y" if are recovering from a media
failure (Database crash) or "n" if you are restoring from a previous backup.
Are you recovering from a Media failure (Default: n)?

asa_restoreEngine.sh[497]: =: not found [No such file or directory]
Restoring backup from
/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups/ISCMa
in
Checking to see if the database is running . . .
The database is not running.
-----
/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups/ISCMa
in
1) 06-15-2014
Please select the backup week you wish to restore? [1,2,...]?
1
Is this 06-15-2014 the week you wish to restore[yes or no]?
yes
The backup week of 06-15-2014/ has 1 full backups
1) Date: Jun 17 06:56 Backup DB File: full_01.dir/repository.db
Please select the full backup you wish to restore? [1,2,...]?
1
Is this 1 the backup you wish to restore[yes or no]?
yes
In addition to the full backup there is one or more incremental backups that can be
restored
1) Date: Jun 17 06:56 Incremental DB File: 14061701.log
Please select the incremental backup you wish to restore? [1,2,...]?
1
Is this 1 the backup you wish to restore [yes or no]?
```

```

yes
Restoring to incremental backup
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/ISCMai
n//full_01.dir/14061701.log
Restoring database . . .
Copying
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/ISCMai
n/06-15-2014//full_01.dir/repository.db to
/opt/prime-provisioning-6_7/Repository/repository.db
Restoring Incremental
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/ISCMai
n/06-15-2014//full_01.dir/14061701.log . . .
Copying
/opt/Backup/PP67Database_Backup/BackupRestore/Backups/ISCMMain/06-15-2014//full_01.dir/1406
1701.log /opt/prime-provisioning-6_7/Repository
/opt/prime-provisioning-6_7/thirdparty/sybase/bin/dbsrv11
/opt/prime-provisioning-6_7/Repository/repository.db -a
/opt/prime-provisioning-6_7/Repository/14061701.log -o /tmp/restoredbISC.065845
-----
SQL Anywhere Network Server Version 11.0.1.2044
OEM Authenticated Edition, licensed only for use with authenticated OEM applications.

Copyright (c) 2001-2009, iAnywhere Solutions, Inc.
Portions copyright (c) 1988-2009, Sybase, Inc. All rights reserved.
Use of this software is governed by the Sybase License Agreement. Refer to
http://www.sybase.com/softwarelicenses

2 logical processor(s) on 1 physical processor(s) detected.
Networked Seat (per-seat) model. Access to the server is limited to 1 seat(s).
This server is licensed to:
    Eoan Kerr
    Cisco Systems Ltd
Running Linux 2.6.18-194.8.1.el5 #1 SMP Wed Jun 23 10:52:51 EDT 2010 on X86_64
Server built for X86_64 processor architecture
39384K of memory used for caching
Minimum cache size: 8192K, maximum cache size: 7354120K
Using a maximum page size of 8192 bytes
Starting database "repository" (/opt/prime-provisioning-6_7/Repository/repository.db) at
Tue Jun 17 2014 06:59
Database recovery in progress
    Last checkpoint at Tue Jun 17 2014 06:56
    Checkpoint log...
    Transaction log: /opt/prime-provisioning-6_7/Repository/14061701.log...
    Checkpointing...
Starting checkpoint of "repository" (repository.db) at Tue Jun 17 2014 06:59
Finished checkpoint of "repository" (repository.db) at Tue Jun 17 2014 06:59
Recovery complete
Database server shutdown automatically after log applied
Database server stopped at Tue Jun 17 2014 06:59
-----
Do you want this script to bring up the database and Cisco Prime Provisioning servers?
no
Successfully completed
Restore script ending.
Starting restore of SLA database
DB Server Name...iscadm_pff-rhel5-vm40
Restore Script
This script will restore the Sybase database from a previous backup
that is saved on the system. The database backup consists of two
files; the primary database file and the transaction log file. The
transaction log is a separate file from the database file. It
stores all changes to the database. Inserts, updates, deletes,
commits, rollbacks, and database schema changes are all logged. The
transaction log is also called the forward log or the redo log.

```

For an incremental backup the transaction log file is the only file saved. This script can restore a database from a previous backup as desired or if you had a media failure (database crash) on the database file located in /opt/prime-provisioning-6_7/Repository/sla.db.

This restore utility is not able to recover the database if you had a media failure on the transaction log file. You can perform a partial recovery by renaming the current transaction log file and starting the database with the -f switch.

```
(eg: /opt/prime-provisioning-6_7/thirdparty/sybase/bin/dbsrv11
/opt/prime-provisioning-6_7/Repository/sla.db -f )
```

If you are trying to recover from a media failure (database crash) on the database file then you need to select the *latest full backup* and the *latest incremental backup*. Then this script will first restore the database up to the point of latest incremental backup

and then will attempt to recover the transactions between the last (latest) incremental and the point of failure. Please answer "y" if are recovering from a media failure (Database crash) or "n" if you are restoring from a previous backup. Are you recovering from a Media failure (Default: n)?

```
asa_restoreEngine.sh[497]: =: not found [No such file or directory]
Restoring backup from
/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups/SLA
Checking to see if the database is running . . .
The database is not running.
-----
/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups/SLA
1) 06-15-2014
Please select the backup week you wish to restore? [1,2,...]?
1
Is this 06-15-2014 the week you wish to restore[yes or no]?
yes
The backup week of 06-15-2014/ has 1 full backups
1) Date: Jun 17 06:56 Backup DB File: full_01.dir/sla.db
Please select the full backup you wish to restore? [1,2,...]?
1
Is this 1 the backup you wish to restore[yes or no]?
yes
In addition to the full backup there is one or more incremental backups that can be
restored
1) Date: Jun 17 06:56 Incremental DB File: 14061701.log
Please select the incremental backup you wish to restore? [1,2,...]?
1
Is this 1 the backup you wish to restore [yes or no]?
yes
Restoring to incremental backup
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/SLA/06
-15-2014//full_01.dir/14061701.log
Restoring database . . .
Copying
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/SLA/06
-15-2014//full_01.dir/sla.db to
/opt/prime-provisioning-6_7/backup/Sybase/Repository/sla.db
Restoring Incremental
/opt/prime-provisioning-6_7/backup/Sybase/PP661Database_Backup/BackupRestore/Backups/SLA/0
6-15-2014//full_01.dir/14061701.log . . .
Copying
/opt/prime-provisioning-6_7/backup/Sybase/PP67Database_Backup/BackupRestore/Backups/SLA/06
-15-2014//full_01.dir/14061701.log /opt/prime-provisioning-6_7/Repository
/opt/prime-provisioning-6_7/thirdparty/sybase/bin/dbsrv11
/opt/prime-provisioning-6_7/Repository/sla.db -a
/opt/prime-provisioning-6_7/Repository/14061701.log -o /tmp/restoredbSLA.065925
-----
```

```

SQL Anywhere Network Server Version 11.0.1.2044
OEM Authenticated Edition, licensed only for use with authenticated OEM applications.

Copyright (c) 2001-2009, iAnywhere Solutions, Inc.
Portions copyright (c) 1988-2009, Sybase, Inc. All rights reserved.
Use of this software is governed by the Sybase License Agreement. Refer to
http://www.sybase.com/softwarelicenses


2 logical processor(s) on 1 physical processor(s) detected.
Networked Seat (per-seat) model. Access to the server is limited to 1 seat(s).
This server is licensed to:
    Eoan Kerr
    Cisco Systems Ltd
Running Linux 2.6.18-194.8.1.el5 #1 SMP Wed Jun 23 10:52:51 EDT 2010 on X86_64
Server built for X86_64 processor architecture
8192K of memory used for caching
Minimum cache size: 8192K, maximum cache size: 7354116K
Using a maximum page size of 4096 bytes
Starting database "sla" (/opt/prime-provisioning-6_7/Repository/sla.db) at Tue Jun 17 2014
06:59
Database recovery in progress
    Last checkpoint at Tue Jun 17 2014 06:56
    Checkpoint log...
    Transaction log: /opt/prime-provisioning-6_7/Repository/14061701.log...
    Checkpointing...
Starting checkpoint of "sla" (sla.db) at Tue Jun 17 2014 06:59
Finished checkpoint of "sla" (sla.db) at Tue Jun 17 2014 06:59
Recovery complete
Database server shutdown automatically after log applied
Database server stopped at Tue Jun 17 2014 06:59
-----
Do you want this script to bring up the database and Cisco Prime Provisioning servers?
no
Successfully completed
Restore script ending.

```

Oracle Database Backup and Restore

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 [“Create RMAN Catalog Database” section on page 4-25](#). RMAN maintains the bookkeeping 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.

[Figure 4-9](#) shows an Oracle Database Backup Diagram.

Figure 4-9 Oracle Database Backup

RMAN for Oracle 10g is explained in the quick start guide and reference manual, which are available from Oracle's website.

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 can reside on the same Prime Provisioning database server.

Oracle Enterprise manager (GUI) can be used to set up RMAN.

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

-
- Step 1** [Create RMAN Catalog Database, page 4-25.](#)
 - Step 2** [Create RMAN User, page 4-25.](#)
 - Step 3** [Create RMAN Catalog, page 4-25.](#)
 - Step 4** [Register the Prime Provisioning Database with the RMAN Catalog, page 4-25.](#)
 - Step 5** [Add PATH Statement, page 4-25](#)
 - Step 6** [Modify Prime Provisioning Database Initial Parameter File, page 4-25.](#)
 - Step 7** [Back up Database, page 4-26.](#)
 - Step 8** [Recover Database, page 4-27.](#)
-

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 Prime Provisioning database.

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

Create RMAN User

Creating an RMAN user is the same as creating an Prime Provisioning user on an **rcat** database. Name the RMAN user ID **rmanuser** and name the password **rmanpassword**. Make sure **rmanuser** has proper privileges. For example:

```
SQL> grant connect, resource, recovery_catalog_owner to rmanuser;
```

Create RMAN Catalog

Create a catalog from the RMAN command prompt:

```
RMAN> connect catalog rmanuser/rmanpassword@rcat
RMAN> create catalog;
```

Register the Prime Provisioning Database with the RMAN Catalog

Set the ORACLE_SID environment variable = prime.

```
% rman
```

```
RMAN > connect catalog rmanuser/rmanpassword@rcat
```

```
RMAN > connect target sys/change_on_install
```

```
RMAN > register database
```

```
RMAN> configure controlfile autobackup on;
```

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 Prime Provisioning database.

Add PATH Statement

After installing the Prime Provisioning Backup and Restore tool and before configuring it, the PATH statement:

```
PATH=$PATH:/BackupRestore/scripts:/BackupRestore/config:/BackupRestore/bin export PATH
```

should be added to the login .profile file of the user iscadm.

Without this permanent addition, later runs of the backup and restore may fail.


Modify Prime Provisioning Database Initial Parameter File

To modify the Prime Provisioning database initial parameter file, do the following:

-
- Step 1** To ensure the database is in archive log mode, enter the following:
- ```
SQL> alter system set log_archive_dest_1 = 'location= </var/tmp/oradata/arch>' SCOPE=BOTH;
SQL> alter system archive log start;
```
- where `</var/tmp/oradata/arch>` is the location of the archive destination.
- Step 2** Restart the Prime Provisioning 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;
```
- 

## Back up Database

To back up the database, do the following:

- 
- Step 1** Download the software for backup and restore from:
- <http://software.cisco.com/download/navigator.html?mdfid=284127465&flowid=37682>
- Step 2** Before you run the backup scripts, make sure you update the file `$PRIMEP_HOME/backup/Oracle/backupenv.properties`
- Use a text editor to open this file and read the directions on how to update each property.
- 
- Note**  The file `$PRIMEP_HOME/backup/Oracle/backupenv.properties` contains `BACKUP_DEST`, which must point to a directory that is writable by the owner of the Oracle database. To do this, specify `chmod atw <file_defined_by_BACKUP_DEST>`
- 
- Step 3** To perform a full database backup, execute the following:
- ```
$PRIMEP_HOME/backup/Oracle/oracle_backup.sh -f
```
- Step 4** You can perform incremental backups after a minimum of one full backup. To perform an incremental backup, execute the following:
- ```
$PRIMEP_HOME/backup/Oracle/oracle_backup.sh -i
```



- 
- Note** These backup scripts can be run as cron jobs and cannot be scheduled by the Prime Provisioning task manager.
- 

## Back up Non-database Files

On the Prime Provisioning server machine, to backup non-database related files, such as task logs or Prime Provisioning system properties, execute the script: `non_db_backup.sh`.

## Recover Database

To recover a database, do the following:

**Step 1** Stop the Prime Provisioning watchdog before recovering a database, as follows:

```
./prime.sh stop
```

**Step 2** To recover a database, you can execute the following from the location  
\$PRIMEP\_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:

```
“Mar 09 2010 15:25:00”
```

If you do not specify <date\_time>, the script does a full database recovery.



**Note**

Do not stop the Oracle Listener during restore.

## Standby System for Prime Provisioning (Secondary System)

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

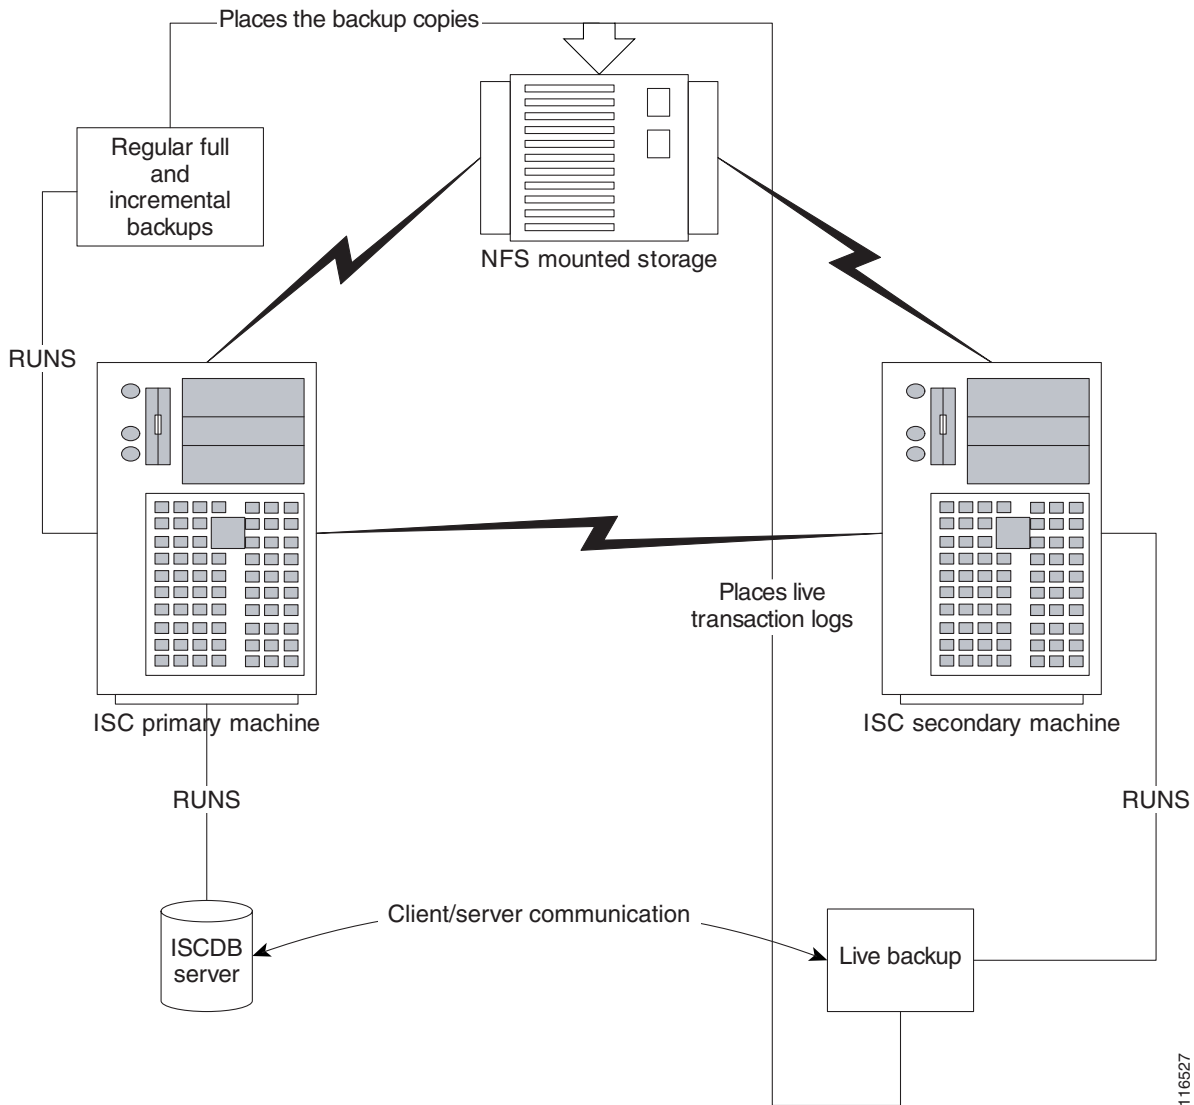
The subsections are:

- [Sybase Standby System Process Overview, page 4-28](#)
- [Sybase Standby System Set Up, page 4-30](#)
- [Oracle Standby System Set Up, page 4-31](#)

## Sybase Standby System Process Overview

Figure 4-10 shows a live backup scheme.

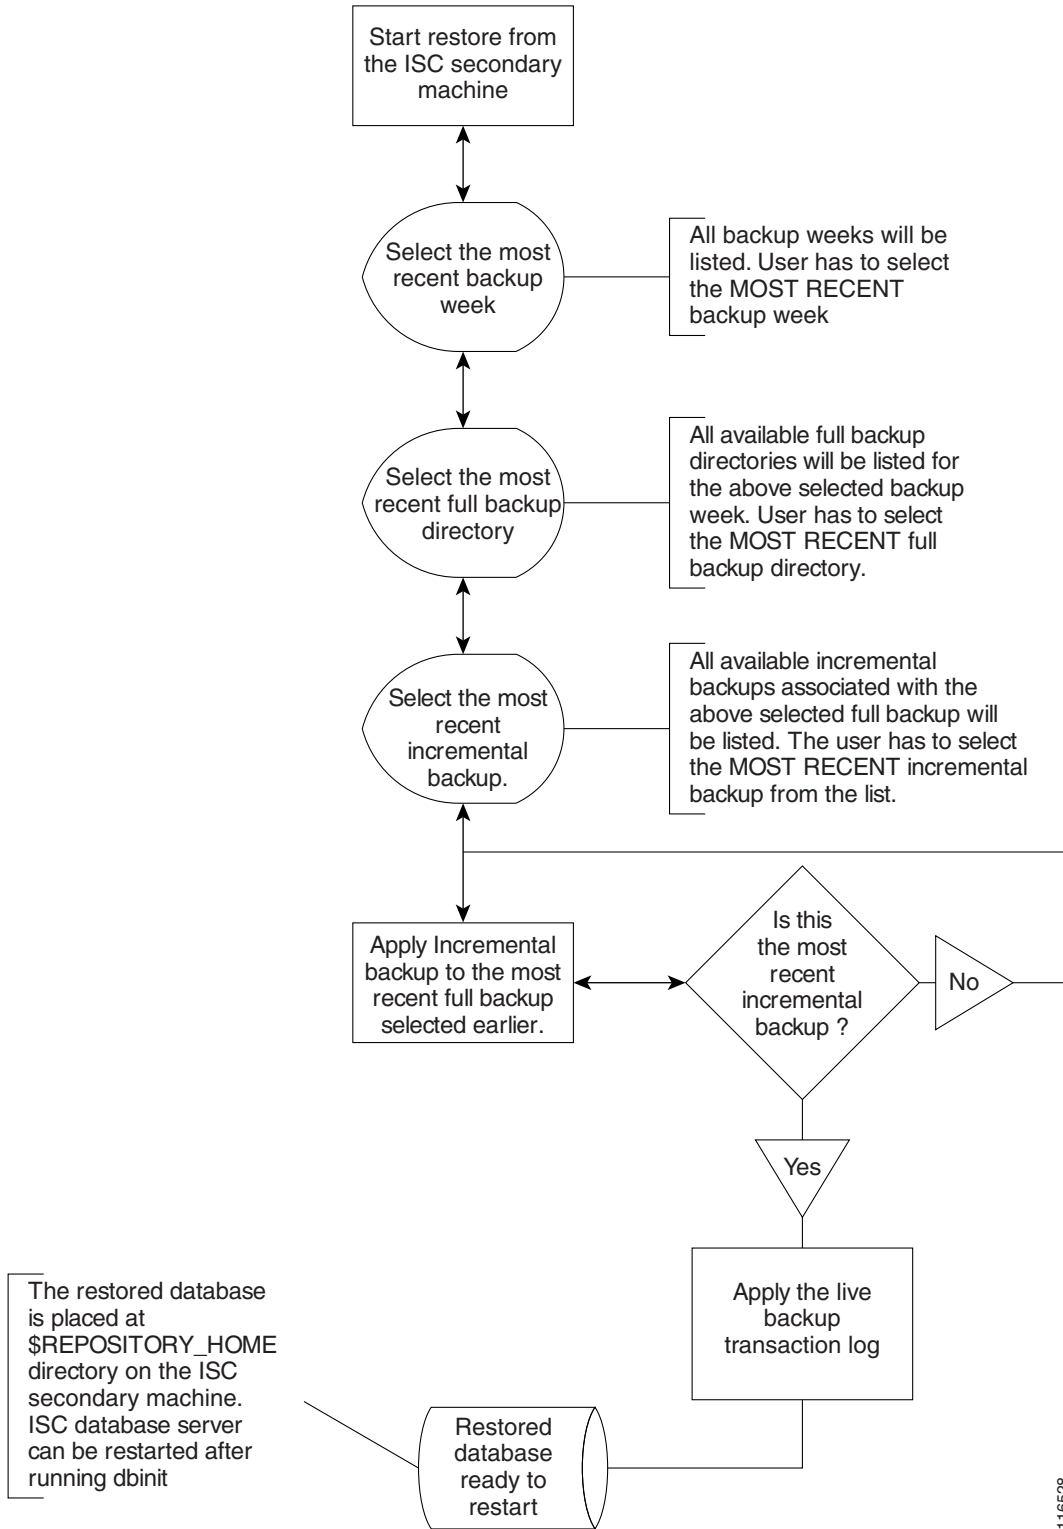
**Figure 4-10 Live Backup Scheme**



## Restore from Live Backup

Figure 4-11 shows the process flow for how to restore from a live backup.

Figure 4-11 Restore from Live Backup



## Sybase Standby System Set Up

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

- [Running Live Backup of Prime Provisioning Databases, page 4-30](#)
- [How to Restore the Database from the Live Backup, page 4-30](#)

### Running Live Backup of Prime Provisioning Databases

Run `<BACKUP_INSTALL_DIR>/BackupRestore/scripts/asa_liveBackup.sh` from the Prime Provisioning secondary system to start the live backup after being sure to follow these preconditions:

- 
- Step 1** Set up a standby Prime Provisioning system.
  - Step 2** The standby system should be similar to the primary Prime Provisioning host system in hardware and software configurations.
  - Step 3** The Prime Provisioning primary and standby systems should be on the same LAN.
  - Step 4** Prime Provisioning software should be installed on the secondary system and the version of Prime Provisioning on the primary and standby systems should be the same.
  - Step 5** The backup and restore tool should be installed on the primary and the secondary systems.
  - Step 6** The live backup should be started from the secondary system only, you should not run the live backup from Prime Provisioning primary system.
  - Step 7** The storage device where the regular backup copies are placed should be accessible from the standby system.
  - Step 8** 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.
  - Step 9** The Prime Provisioning database server must be running on the primary Prime Provisioning host before starting the live backup on the standby system.
  - Step 10** The live backup stops when the Prime Provisioning database server is stopped and should be restarted after restarting Prime Provisioning.
  - Step 11** At least one full backup must be taken before starting the live backup.
  - Step 12** Regular periodic full/incremental backups should be taken even if the live backup is running on the secondary system.
  - Step 13** There should not be more than one live backup running simultaneously.



#### Note

When backup is taken using standby server in a different subnet, the host parameter (**host=\$ISCHOST**) need to be added after `tcPIP` in `dbping` command which exists in the backup script. For example:

```
RES=`$SYBASE_HOME/bin/dbping -c "uid=$DBUSER;pwd=$DBPASSWD;dbn=$dbName;eng=$dbEng;links=tcPIP(host=$ISCHOST)"`
```

### How to Restore the Database from the Live Backup

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

Run the `<BACKUP_INSTALL_DIR>/BackupRestore/script/asa_restoreFromLiveBackup.sh` script on the standby system to restore the database after being sure to follow these preconditions:

- 
- Step 1** At least one full backup copy should be available to restore the database.
  - Step 2** If more than one backup copy is available, use only the latest full backup and the latest associated incremental backup.
  - Step 3** Run the restore from the standby machine.
- 

## Oracle Standby System Set Up

Prime Provisioning 6.0 supports both physical standby and logical standby in Oracle 10g Data Guard. For information about the Oracle 10g standby concept and configuration, see the *Oracle Data Guard Concept and Administration 10g Release 1 (10.1)* Part No. B10823-01. The document can be found at Oracles' website.

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

```
./prime.sh stop -y
```

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

```
execute applycfg.sh
```

```
./prime.sh initdb.sh
```

```
./prime.sh startwd
```

where:

`<old_db_server>` is the name of the old database server

`<new_db_server>` is the name of the new database server.

