Operations Reports Administration and Setup

This section is for administrators who are responsible for setup, maintenance, and troubleshooting of Prime Network Operations Reports. It contains the following topics:

Overview of Operations Reports Architecture

Within Prime Network, Operations Reports has a dedicated AVM (AVM 44) that processes the information received from the VNEs and sends it to the Operations Reports agent on the gateway, and from there to the Infobright database.

Operations reports fetches data as described below:

- The inventory data is handled by AVM 44, which reads and writes from Infobright database.
- The raw events that are sent by AVM 100 are processed by the VNEs and are sent to AVM 25; AVM 25 then writes the data to both Oracle and Infobright databases. AVM 25 that is present in the gateway performs a two-way communication with the fault database component in Oracle and then writes the archived faults data to Infobright database. Operations Reports uses the data that is available in Infobright database to generate reports. The archived events are read from Infobright database and the active events are read from the Oracle database.
- AVM 100 writes the raw events to both Oracle and Infobright databases. The raw events are read from Infobright database.

The following diagram illustrates how Prime Network components interact with Operations Reports and with the Infobright database.

Figure 2-1 Prime Network Operations Reports Architecture

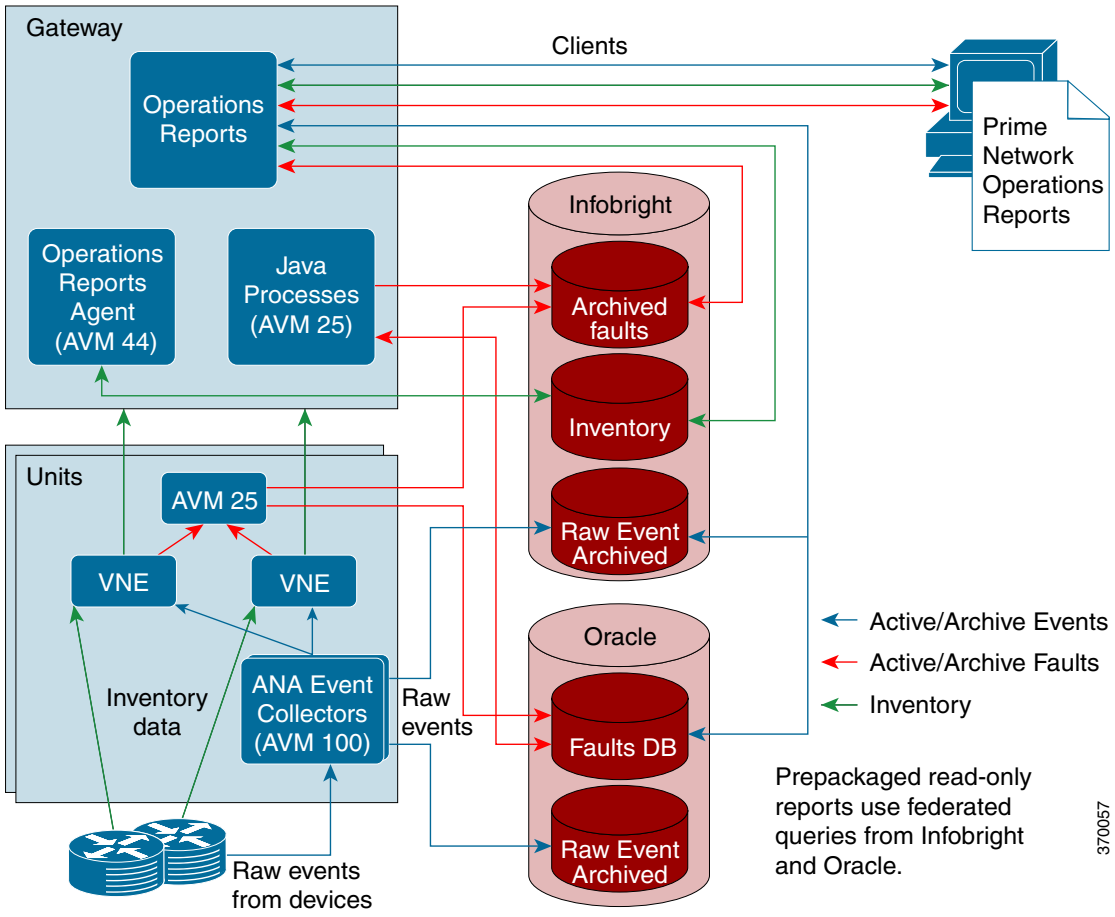

It takes up to five minutes for changes made in the Prime Network GUI clients to be reflected in Operations Reports. For example, when a VNE is added to the system, it is not shown in Operations Reports until the VNE is in one of the following states: OPERATIONAL, PARTIALLY_DISCOVERED, or CURRENTLY_UNSYNCHRONIZED. This could take up to five minutes. Other changes to devices, such as port up/down, might take less time to be reflected in the reports.

Upon adding or deleting VNEs, it takes approximately 30 minutes for the Device Selector list to be updated; to manually refresh and view the changes, see Manually Refreshing Operations Reports to Show VNE Changes, page 2-8.

Setting Up Operations Reports

This section describes the tasks that the administrator needs to perform when setting up Operations Reports. It includes the following topics:

Managing User Access Rights for Operations Reports

Operations Reports uses SSO (Single Sign On). User authentication is performed using the same method being used by the other Prime Network components (internal or LDAP).

Operations Reports uses the same RBAC (Role-Based Access Control) facilities as are used in the rest of Prime Network—device scopes and user access rights. Users can only see Operations Reports data for devices that are in their device scope. Users have permissions for certain actions depending on their user role (Administrator, Configurator, and so on). Administrators have full permissions and can grant permissions to other user role, as described in Granting User Permissions for Reports, page 2-3.

Granting User Permissions for Reports

For each report folder or individual report, the administrator can determine which actions can be performed by the various user roles (Configurator, Viewer, and so on). By default, Administrators have full permissions and all other user roles are allowed to execute reports.

An example of granting permissions on the folder level might be if you want to prevent certain users from viewing and accessing data center reports. In this case, you would remove permissions on the folder level. The folder will not be visible to those users.

On the individual report level, you might want to add privileges allowed for a certain type of user, for example, allow a Configurator user to delete reports.

Note All reports inherit the Execute permission that is set at the respective folder level. For a non-admin user, if the Execute permission at report level is unchecked, and the Execute permission for the folder is checked, the folder is visible, but the report is not visible. However, if the Execute permission is unchecked at folder level and is checked at the report level, both the folder and the report are not visible.

To assign user permissions:

Step 1 In the Browse pane, right-click on the folder or report for which you want to grant permissions and select **Properties**.

Step 2 In the Properties dialog, click the **Share** tab.

Step 3 Select the user role for which you want to grant permissions, and check the actions you want to allow for that user role for the selected report or folder.

Step 4 Click **OK**.

Enabling Reports to be Sent by E-mail

When scheduling prepackaged read-only reports, users have the option to send the report to a specified recipient by e-mail. The administrator must set up Operations Reports to communicate with the SMTP server in order to enable this functionality.

To enable the e-mail facility:

Step 1 Locate and open the email_config.xml file under the $PRIME_NETWORK_HOME/pentaho/server/biserver-ee/pentaho-solutions/system/smtp-email/email_config.xml

Step 2 Configure the SMTP server details. Figure 2-2 shows a sample email_config.xml file.

Step 3 Save the file.

Step 4 In Operations Reports, choose **Tools > Refresh > System Settings** to update the system with your changes.

Step 5 After you complete the configuration, you must restart the server using ctlscript.sh restart under the

$PRIME_NETWORK_HOME/pentaho/ctlscript.sh restart

Figure 2-2 Sample SMTP Configuration File

```
<?xml version="1.0"?>
- <email-smtp>
        <!-- The values within <properties> are passed directly to the JavaMail API. For a list of valid properties see
        http://java.sun.com/products/javamail/javadocs/index.html -->
    - <properties>
            <!-- This is the address of your SMTP email server for sending email. e.g. smtp.pentaho.org -->
        <mail.smtp.host>144.254.72.80</mail.smtp.host>
            <!-- This is the port of your SMTP email server. Usually this is 25. For GMail this is 587 -->
        <mail.smtp.port>25</mail.smtp.port>
            <!-- The transport for accessing the email server. Usually this is smtp. For GMail this is smtps -->
        <mail.transport.protocol>smtp</mail.transport.protocol>
            <!-- Usually this is 'false'. For GMail it is 'true' -->
        <mail.smtp.starttls.enable>false</mail.smtp.starttls.enable>
            <!-- Set to true if the email server requires the sender to authenticate -->
        <mail.smtp.auth>false</mail.smtp.auth>
            <!-- This is true if the email server requires an SSL connection. Usually 'false'. For GMail this is 'true' -->
        <mail.smtp.ssl>false</mail.smtp.ssl>
            <!-- Run Email Send Test -->
        <mail.run.send.test>true</mail.run.send.test>
            <!-- Output debug information from the JavaMail API -->
        <mail.debug>false</mail.debug>
            <!-- For GMail this is 'false' -->
            <!--mail.smtp.quitwait>false</mail.smtp.quitwait-->
    </properties>
        <!-- The is the address of your POP3 email server for receiving email. e.g. pop.pentaho.org -->
        <!-- It is currently not used -->
    <mail.pop3/>
        <!-- This is the default 'from' address that emails from the Pentaho BI Platform will appear to come from e.g.
        joe.pentaho@pentaho.org -->
    <mail.from.default>pn40@cisco.com</mail.from.default>
        <!-- This is the user id used to connect to the email server for sending email It is only required if email-authenticate is set to true
        This is never sent or shown to anyone -->
    <mail.userid/>
        <!-- This is the password used to connect to the email server for sending email It is only required if email-authenticate is set to
        true This is never sent or shown to anyone -->
    <mail.password/>
</email-smtp>
```

370233

Troubleshooting Tasks

This section describes the tasks that the administrator might need to perform when troubleshooting problems with Operations Reports. It includes the following topics:

Checking the Status of AVM 44

AVM 44 is dedicated to Operations Reports. It transfers the inventory data to Infobright database. If AVM 44 is down, running inventory reports will fail. These reports would need to be rerun when AVM 44 is up again.

To check the status of AVM 44:

Step 1 Log into the gateway server as *pnuser*.

Step 2 Enter the following command to show the status of all the gateway processes, including AVM 44:

 # networkctl status.

Stopping/Starting Operations Reports Processes

As part of your troubleshooting activities, you might want to stop the Operations Reports processes. There are two processes that should be stopped: Jetty and Bootstrap.

Note When the processes are stopped, the Operations Reports Web pages are unavailable.

To stop the Operations Reports processes:

Step 1 Log into the Prime Network gateway as pnuser and make sure you are in the Prime Network home directory.

Step 2 Stop the Operations Reports script:

```
cd ~/pentaho
./ctlscript.sh stop
```

Step 3 Identify the Operations Reports Jetty process ID:

```
pn40# ps -ef | grep jetty | grep -v grep
```

The process ID is identified in the output. In the example below, it is 10174.

```
pn40    10174 10148 16 16:01 pts/3    00:00:08 /export/home/pn400/pentaho/java/bin/java -
```

Step 4 Stop the process:

```
pn40# kill -1 process-id; sleep 1
```

Step 5 Verify that the process is no longer running using the following command. You should get no response.

```
pn40# ps -ef | grep jetty | grep -v grep
```

Step 6 Identify the Operations Reports Bootstrap process ID (Tomcat process):

```
pn40# ps -ef | grep Bootstrap | grep -v grep
```

The process ID is identified in the output. In the example below, it is 20784.

```
pn40    20784     1 10 15:44 ?        00:01:30 /export/home/pn400/pentaho/java/bin/java -
```

Step 7 Stop the process:

```
pn40# kill -1 process-id; sleep 1
```

Step 8 Verify that the process is no longer running using the following command. You should get no response.

```
pn40# ps -ef | grep Bootstrap | grep -v grep
```

Before starting a new Operations Reports script from the Pentaho server, ensure that you stop the Operations Reports that is running using the following script **./ctlscript.sh stop** or the **kill** command.

To start the Operations Reports processes:

Step 1 Log into the Prime Network gateway as pnuser and make sure you are in the Prime Network home directory.

Step 2 Start the Operations Reports script:

```
cd ~/pentaho
./ctlscript.sh start
```

You must be able to launch Operations Reports successfully after starting the processes.

Resolving POODLE Vulnerability in Prime Network

POODLE vulnerability potentially allows sensitive information to be obtained from the remote host with SSL or TLS-enabled services. Prime Network supports services that can enable TLS fallback Signaling Cipher Suite Value (SCSV) mechanism and disable SSLv3 to resolve the POODLE vulnerability.

The TLS fallback SCSV mechanism prevents version rollback attacks without impacting legacy clients; however, it can only protect connections when the client and service support the mechanism. Sites that cannot disable SSLv3 immediately should enable this mechanism.This impact processes on port numbers 1311 and 8445.

To enable TLS services, perform the following steps:

Step 1 Log into the Prime Network gateway as pnuser and make sure you are in the Prime Network home directory.

Step 2 Install Pentaho in the Prime Network gateway server. The remote host supports the use of SSL ciphers that offer high-strength encryption (256 bits) SSL ciphers.

Step 3 Check if Pentaho is running using the following commands:

```
ctlscript.sh stop
ctlscript.sh start
```

Step 4 Change the browser settings:

- In Internet Explorer, choose **Setting > Internet Options > Advanced**.
  In the **Advanced** window, select the **SSL 2.0** option under **Security**.
- In Mozilla Firefox, perform the following steps to enable SSLv3:
  a. Enter **about:config** in the address bar of the Mozilla Firefox browser.
  b. Open the detailed list of configuration parameters.
  c. Check if the value for the **security.tls.version.min** option is **0** to turn on the support for SSLv3.
  d. Search for the **security.ssl** option and change the value from **false** to **true**. The value for all ciphers do not include **aes_256** in the name.
- In Mozilla Firefox, perform the following steps to enable TLS:
  a. Enter **about:config** in the address bar of the Mozilla Firefox browser.
  b. Open the detailed list of configuration parameters.
  c. Check if the value for the **security.tls.version.min** option is **1** to turn on the support for TLSv1.
  d. Search for the **security.ssl** option and change the value from **true** to **false**. The value for all ciphers do not include **aes_256** in the name.
  e. View the Prime Network monitoring graphs using the port number 1311.
  a. View the Prime Network reports using the port number 8445.

Monitoring Log Files

Table 2-1 lists the Operations Reports log files that are created and stored. You can view these files using any text editor.

Table 2-1 Gateway Server Log Files

Gateway Server Log File	Found in the Server...	Component
44.out	Gateway server	Operations Reports
From the folder, the Infobright integration zip file is extracted to: primenw_integration_<*date*>.log	Gateway server	Infobright database installation log
/usr/local/infobright/IBRestore/Log/IBRestore_*ddmmyy_hhmmss*.log	Infobright database server	Infobright database restore log
/usr/local/infobright/data/brighthouse.log /usr/local/infobright/data/bh.err	Infobright database server	Infobright database log
/usr/local/infobright/data/Old_Log	Infobright database server	Infobright database old logs

Manually Refreshing Operations Reports to Show VNE Changes

When VNE changes are made in Prime Network, for example, a new VNE is added, it takes approximately 30 minutes before the new device appears in the device selector in Operations Reports. Administrators can manually refresh the system so that the VNE changes are reflected in the Operations Reports GUI.

To manually refresh Operations Reports:

Step 1 Choose **Tools > CDA Cache Manager**.

The CDA Manager launches.

Step 2 Click on **Cached Queries**, and then click **Clear Cache**.

Step 3 Check that the new device now appears in the device selector when creating or generating a report.

Manually Reloading Report Folders to Show New Reports

If you create a new report, it will be added to the specified folder in the Browse pane. However, it will not be visible to other users on the same server unless you do a manual reload.

To manually reload report folders:

Step 1 Choose **Tools > Refresh > Repository Cache**.

Step 2 Check that your new report is visible to other users in the relevant report folder.

Managing Infobright Database Operations

The Infobright database is dedicated to Prime Network Operations Reports. It can be installed on the gateway server or on a separate server.

The following data is stored in the Infobright database under schema **pnibdb**:

- Archived fault database information—Data that has been moved to the archive partitions in the Oracle database (NETWORKEVENT, ALARM, TICKET, NEWTRAPEVENT, NEWTRAPVALUE).
- Raw events in the event archive (HP_TRAP, HP_SYSLOG).
- Inventory information used by Operations Reports. Inventory information is sent to the Infobright database by the Operations Reports engine (AVM 44), which runs on the gateway.

The data for the Network Service Reports is obtained from the Oracle database. However, unlike the Oracle database, the **pnibdb** schema does not have active and archive partitions. All the data is kept in a single table. The Infobright database is regularly backed up by Prime Network, as described in Backing Up and Recovering the Infobright Database, page 2-11. Data is purged by default after 180 days, but this setting can be changed during installation or at a later stage, as described in Changing Settings for Purging Infobright Backups, page 2-12.

Starting/Stopping the Infobright Database

Note Make sure that the Prime Network system is operational before stopping the database.

To stop/start the database:

Step 1 Log into the database as the root user.

Step 2 Enter the following command:

```
/etc/init.d/mysqld-ib stop
or
/etc/init.d/mysqld-ib start
```

Stopping the Writing of Information to the Database

To stop writing data to the database:

Step 1 Log into the Prime Network gateway as pnuser.

Step 2 Enter the following command:

```
runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/infobright/enable false
```

Step 3 Restart the Prime Network gateway:

```
# cd $ANAHOME/Main
# networkctl restart
```

Enabling/Disabling Infobright Backup Process

This setting is defined during installation.

To change the setting for writing data to backup files:

Step 1 Log into the database as pnuser.

Step 2 Enter the following command:

```
runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/infobright/enableBackupFiles
true/false
```

Step 3 Restart the Prime Network gateway:

```
# cd $ANAHOME/Main
# networkctl restart
```

Controlling When Operations Reports Data is Purged

To protect system stability and performance, Prime Network purges historical data from the Infobright database after 180 days (default). This setting can be changed during installation or at a later stage.

Note The settings in **Global Settings > Report Settings** do not affect Operations Reports.

To change the default setting that determines when data is purged from the Infobright database:

Step 1 Log into the gateway as *pnuser* and change to the Main directory.

Step 2 Check the current data purge setting for the Infobright database and change the arch_history value (in days) as required:

```
runRegTool.sh -gs 127.0.0.1 get 0.0.0.0 persistency/nodes/infobright/arch_history
180
```

Step 3 Restart AVM 11 using **networkctl**.

Monitoring Infobright Data Logs

Infobright data logs are stored on the Prime Network gateway server in the following location:

/usr/local/infobright/data
bh.err
brighthouse.log

Log files are rotated daily. Old files are kept under /usr/local/infobright/data/Old_Log, using the following naming convention:

<file_name>.YYYY-MM-DD

Old logs are kept for 30 days, then deleted.

Backing Up and Recovering the Infobright Database

This section describes the Infobright database backup and restore mechanism and provides procedures for basic tasks that the administrator might need to perform. It contains the following topics:

How the Infobright Database Backup Works

Infobright database backups are enabled during installation. Because Infobright data is continuously backed up, you cannot change the backup timing or perform a manual backup.

The working directory for creation and processing of backup files is the “ibdlp” directory in persistency.xml, located on the Prime Network gateway and units. The location of this directory is specified during installation. To identify the location of the working directory, enter:

cd Main; ./runRegTool.sh localhost get persistency/nodes/infobright/ibdlp

The following sub-directories are created under the ibdlp directory:

- DLP_Active_Files
- DLP_Input_Files
- DLP_Log_Files
- DLP_Output_Files
- DLP_Rjct_Files

The flow of the Infobright database backup is as follows:

1. Prime Network processes, such as auto-archive, AVM 25 and AVM 100, write data to the Infobright database and generate a text file in the DLP_Active_Files directory for each fault and event persistency (EP) table (HP_TRAP, HP_SYSLOG, NETWORKEVENT, ALARM, TICKET, NEWTRAPEVENT, NEWTRAPVALUE).
2. The text file is closed after one hour and moved to the DLP_Input_Files directory.

3. Every ten minutes, a cron job DLP process generates a compressed Infobright backup file from the closed text file and places it in the DLP_Output_Files directory. It also creates a log file under the DLP_Log_Files directory. If the cron job DLP process fails, a RJCT file is generated and placed in the DLP_Rjct_Files directory, and a system event is generated.
   This cronjob is configured under prime-network crontab:
   $SHEERHOME/Main/unix/cron/every_10_minutes.main.gateway.cmd DLPProcessExecuter
4. A cron job SCP process transfers the backup file from the DLP_Output_Files directory to the backup directory on the database server (ibbackup). The path to this directory is specified during installation.
   This cronjob is configured under root crontab:
   */10 * * * * /var/adm/cisco/infobright/move_dlp_files.sh >/dev/null 2>&1

Checking that Backups are Taking Place

You can verify that backups are happening by checking the backup directory specified during installation. Backups will begin appearing 1-2 hours after events noise starts on the Infobright database server.

To identify the location of the backup directory, use this command:

cd Main; ./runRegTool.sh -gs 127.0.0.1 get 0.0.0.0 persistency/nodes/infobright/ibbackup

We recommend that you maintain a copy of the backup files in an external storage device if one of the following apply:

- You prefer to retain a longer backup history than the one configured in the bkp_history parameter.
- If the backup directory cannot accommodate all the files due to lack of space.
- You prefer to save a duplicate version of the backup files to recover from a disk failure.

Changing Settings for Purging Infobright Backups

By default, backups are purged after 180 days.

To change the settings for purging Infobright backups:

Step 1 Log into the gateway and change to the Main directory.

Step 2 Check the current backup purge setting for Infobright backups:

```
./runRegTool.sh -gs 127.0.0.1 get 0.0.0.0 site/persistency/nodes/infobright/bkp_history
```

Step 3 To change the setting, specify the number of days after which backups should be purged.

```
# ./runRegTool.sh localhost 127.0.0.1 set 0.0.0.0
site/persistency/nodes/infobright/bkp_history value
```

Step 4 Log into the Infobright DB server as root user. From the following script, edit the **mtime +number_of_days** to reflect the new setting:

```
/var/adm/cisco/infobright/ibbackup_clean.sh.
```

Recovering the Infobright Database

The Infobright recovery procedure can be used to load backup files into a fresh Infobright database. To do this, prepare the following:

- A backup directory containing all your backup files.
- A fresh installation of the Infobright database.

To recover the Infobright database:

Step 1 Rename the **ibbackup** directory or save it to a different location, to prevent the directory from being removed during the uninstall procedure. For example:

```
mv ibbackup ibbackup_old.
```

Step 2 Reinstall the Infobright database. This creates a fresh database, which is required for the recovery process.

Step 3 Log into the Infobright database server as the Infobright UNIX root user, and move to the following directory:

```
# cd /var/adm/cisco/infobright
```

Step 4 Start the recovery script:

```
# ./ib_recovery.sh
```

You will be prompted for the following information:

- The Infobright schema name. Use the default (**pnibdb**).
- The path to the directory containing your backup files (see Step 1).

The script will list the available date ranges for the recovery operation.

Step 5 Choose the recovery option you require.

If you want to:	Choose the following:
Recover all the backed up data	**1. FULL LOAD**
Recover only backed up data from a specified time period	**2. PARTIAL LOAD**

Step 6 If you chose PARTIAL LOAD, enter the start and end time when prompted.

Restarting the Recovery Process After a Recovery Failure

To restart the recovery process after an interruption or recovery failure:

Step 1 Start the recovery script:

```
# ./ib_recovery.sh
```

Step 2 Choose **3. RESUME A FAILED LOAD SESSION**.

Step 3 Enter the same information (schema and backup directory) as you did for the interrupted/failed recovery.

Step 4 When the script displays the previous recovery ID, type, and start time, confirm the choice. (The script automatically uses the information from the last recovery operation.)

Step 5 Check the recovered tables for duplicates.

a. Start the recovery script:

```
# ./ib_recovery.sh
```

b. Choose **4. CHECK DUPLICATES**.

- If duplicates are found in a production system, repeat the recovery process.
- If duplicates are found in a test setup, you can rerun the recovery script with option **5. TRUNCATE FAULT AND EVENT ARCHIVE (EP) TABLES - Truncate tables in preparation for load.**" and then repeat the recovery process.

c. Enter the same information (schema and backup directory) as you did for the interrupted restoration.

Recovery Log Files

For each operation (full load, partial load, resume failed operation, truncate tables), a log is created and placed in /usr/local/infobright/IBRestore/Log/IBRestore_*ddmmyy_hhmmss*.

If the recovery process fails, you can investigate further by querying the following Infobright tables under pnibdb schema:

- PNIB_RECOVERY_INFO - Holds one row for each recovery run
- PNIB_RECOVERY_LOG - Holds many rows for each recovery run, with detailed information about the activities and files involved.

Backup and Restore Customized Reports

To backup and restore Customized reports:

Step 1 Copy the primenetwork-reports folder to your local server.

```
Example: $ANA_HOME/pentaho/server/biserver-ee/pentaho-solutions/primenetwork-reports
```

Step 2 Rename the ibbackup directory or save ibbackup directory to a different location. This is to prevent the directory from being removed during uninstallation.

```
Example: mv ibbackup ibbackup_old.
```

Step 3 Re-install the Prime Network, Oracle and Operations Reports.

Step 4 After re-installation, remove the primenetwork-reports directory.

```
For example, remove the primenetwork-reports from
$ANA_HOME/pentaho/server/biserver-ee/pentaho-solutions/primenetwork-reports.
```

Step 5 Paste the primenetwork-reports directory.

Note You must paste the primenetwork-reports directory which you already have on your local machine to the location "$ANA_HOME/pentaho/server/biserver-ee/pentaho-solutions/"

Step 6 Change the primenetwork-reports folder permission to pnusername:ana

```
For example, if pn user name is pn50, then the permission should be pn50:ana.
```

Step 7 To recover the database follow the steps mentioned in the Recovering the Infobright Database procedure.

Step 8 Launch the Prime Network Operations reports by manually entering the following URL in your browser address field:

https:// < gateway-IP >:< port-number >/ prime-network-reports/Login and view the Customized report.

If the Customized report is not available, do the following

Click **Tools > Refresh > Repository Cache** to view your customized report under the Files pane.

Note Files pane is located at the left bottom of the home page.