IT Analytics Administrator Guide

Release 11.8

VERITAS
now part of Cohesity

IT Analytics System Administrator Guide

Last updated: 2026-08-12

Legal Notice

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Chapter 1

Introduction

This chapter includes the following topics:

IT Analytics Overview

IT Analytics delivers unified backup and storage insights for heterogeneous IT environments across on-prem and cloud services. Gain actionable insights through a single pane of glass for enterprise data wherever it resides. It quickly and effectively aggregate insights from all environments through a single and unified interface.

Purpose of this document

The System Administrator's Guide describes administrative tasks required to manage the operation of IT Analytics on your system. This document is intended for a technical audience, specifically the system administrator. This guide explains administrative tasks for IT Analytics running on all platforms.

Chapter 2

Preparing for updates

This chapter includes the following topics:

About upgrades and updates

IT Analytics releases changes to this software through different types of releases:

- **Major release**. Includes all new features and enhancements.
- **Minor release**. Includes mostly fixes and some enhancements to existing features. Minor releases are packaged with an auto-installer.
- aptare.jar release. Includes updates to data collectors. This periodically released file encapsulates the main processing logic of the Data Collector.

To receive updates for a minor release, you must upgrade to the corresponding major release.

Preparing for Cloud Reports

If you have restricted access to the Internet, you will need to enable the following URL for cloud reports:

https://cloud.aptare.com/remoting/CloudReportService

Verify the connectivity by entering the URL into a browser window. The following message should display:

```
"HTTP Status 405 - HessianServiceExporter only supports POST requests"
```

Performance profiling

Enable the following URL for Performance Profiling:

https://cloud.aptare.com/remoting/CommunityService

Verify the connectivity by enter the URL into a browser window. The following message should display:

```
"HTTP Status 405 - HessianServiceExporter only supports POST requests"
```

Determine the data collector version

Two methods are available for checking the version of Data Collectors:

Data Collector Version via the Portal

In the Portal, select:

Admin > Data Collection> Collector Updates

Data Collector Version via the Command-Line Interface

To determine your current version of your Data Collector via the command-line interface (CLI)

Linux:

```
# /opt/aptare/mbs/bin/agentversion.sh
```

Windows:

```
> C:\opt\aptare\mbs\bin\agentversion.bat
```

Sample Output:

```
Version information for the Data Collector installed at /opt/aptare
 on this server DC-Centos-2
        Version: 9.0.00 08072013-1028
Version information for datarcvr, aptare.jar and Upgrade Manager at
 http://itanalyticsagent.QAtest

        datarcvr Version
                Version: 9.0.0.01
        aptare.jar Version
                Current Version: 9.0.0.03
                Build Number: 12102014-0001
        Upgrade Manager Version
```

```
                Current Version: 9.0.0.01
                Build Number: 07032013-1121
Version information for aptare.jar and Upgrade Manager at
/opt/aptare/upgrade on this server DC-Centos-2
        aptare.jar Version
                Current Version: 9.0.0.01
                Build Number: 08072013-1028
        Upgrade Manager Version
                Current Version: 9.0.0.01
                Build Number: 07032013-1122
Version information for other jars :
        aptare-dc-nbu-col.jar version is: 9.0.0.01|08072013-0003
        aptare-dc-confgen.jar version is: 9.0.0.01|08072013-0001
        aptare-legacy.jar version is: 9.0.0.01|08072013-0001
        aptare-dc-util.jar version is: 9.0.0.01|08072013-0000
        aptare-dc-hnas-com.jar version is: 9.0.0.01|06052013-0004
        aptare-dc-fw-com.jar version is: 9.0.0.01|08072013-0002
        aptare-dc-hnas-col.jar version is: 9.0.0.01|06052013-0004
        aptare-dc-fw-col.jar version is: 9.0.0.01|08072013-0002
        aptare-dc-nbu-com.jar version is: 9.0.0.01|08072013-0003
        aptare-dc-probe.jar version is: 9.0.0.01|08072013-0003
        aptare-dc-spi.jar version is: 9.0.0.01|08072013-0001
Validating aptare.jar Version Compatibility - SUCCESS
```

Data collector updates with an aptare.jar file

The Data Collector **aptare.jar** file is periodically updated. The **aptare.jar** encapsulates the main processing logic of the Data Collector.

Data collector updates can happen in the following ways:

- When data collector services are running: Data Collectors look for an update every 11 minutes. Once the portal is upgraded, the Data Collector will upgrade automatically.
- When data collector services are not running: Data Collectors can be manually updated as mentioned in the following section:
  See “Manual download of the `aptare.jar` file” on page 17.
- When data collector services are running and for some reason, such as local changes, the Data Collector must be updated again: The Portal can be used to start the update from **Admin>Data Collection>Collector Updates**.

Note: The latest distribution of aptare.jar updates is distributed through the regular patch release process.

Manual download of the `aptare.jar` file

Use the following procedure to manually update your Data Collector:

1. Use the downloadlib utility to manually download the `aptare.jar` file. Note that this utility must be run with administrative privileges.

   Windows: **`<Home>\mbs\bin\downloadlib.bat`**

   Linux: **`<Home>/mbs/bin/downloadlib.sh`**

   Note: From IT Analytics version 11.2 onwards, `aptare.jar` signing and verification enhancement has been added in Data Collector upgrade process. Data collector allows `aptare.jar` upgrade/downgrade only when it is signed using the code signing certificate. If `downloadlib.bat` (Windows) or `downloadlib.sh` (Linux) is used to downgrade the `aptare.jar` in the scenario where the Data Collector is on version 11.2 or later and the portal is on version lower than 11.2, the downgrade will fail because the portal has an unsigned jar. Hence, pass the `-skipVerification` parameter to the `downloadlib.bat` or `downloadlib.sh` utility to skip the verification.

2. If you run the downloadlib utility and you get the following message, it indicates that the restore failed or is in progress: Restore has been running for more than 10 minutes, unable to proceed with collector upgrade.

   To proceed, delete the restore.txt file and re-run downloadlib.

   Windows: **`<Home>\upgrade\restore.txt`**

   Linux: **`<Home>/upgrade/restore.txt`**

Portal updates

Portal updates contain feature enhancements and bug fixes. These updates are packaged with an auto-installer.

Once the upgrade package has been installed, perform the upgrade using:

Linux:

```
# /opt/aptare/upgrade/upgrade.sh
```

Windows:

```
> C:\opt\aptare\upgrade\upgrade.bat
```

Chapter 3

Backing up and restoring data

This chapter includes the following topics:

Best practices for disaster recovery

In the event of data loss, for whatever reason, it is critical that you have a backup of the IT Analytics system. This section lists the key files and data associated with IT Analytics that you need to protect. These should be backed up regularly to mitigate risk of data loss.

At a high level, your backup and recovery strategy will consist of:

- Oracle Backups--the most critical component of the system

- File System Backups

Each of these is discussed in detail in subsequent sections.

For steps to recover data that has been backed up refer to the following.See "Restoring the IT Analytics system" on page 26.

See "Oracle database backups" on page 20.

See "File system backups" on page 20.

Oracle database backups

- **Oracle Database: Cold Backup**: This should be done about once a month. A cold backup is required of the Oracle data file directories; by default, /data01, /data02, /data03, /data04, /data05, and /data06 for Linux and \oradata for Windows. Cold backups require Oracle to be shut down, thereby disrupting the Portal's availability. Perform a cold backup, then back up (export) the database tables on a more regular basis (nightly is recommended). In addition to monthly cold backups, a cold backup is recommended after a significant software upgrade to capture the updated database schema.
- **Oracle Database: Export**: The oracle database should be exported nightly. Exports do not require Oracle to be shut down.
  See "Oracle database: Export backups" on page 22.

File system backups

File system backups are required to protect the IT Analytics application directories, files, and OS-specific settings, such as the registry and services for Windows. The file system backup should also capture the nightly database export, which resides in the file system.

Use your company's backup method of choice to perform the file system backup. While a full system backup is recommended, we have identified the key files created by IT Analytics.

IT Analytics - Created Files and Directory Structures

The file paths are shown with Linux usage, with the Windows-specific file system noted.

Table 3-1 Files and Directory Structures

File	Comments
.bash_profile for tomcat and users	

Table 3-1 Files and Directory Structures *(continued)*

File	Comments
/etc/init.d/*aptare*	
/etc/rc3.d/*aptare*	
/etc/rc5.d/*aptare*	
/data0?	DB server (Linux) Contains the Oracle database files;must be backed up using a Cold Backup.
\oradata	DB server (Windows) Contains the Oracle database files; must be backed up using a Cold Backup.
/opt/apache*, /opt/aptare*, /opt/tomcat*	

Other Files

- /usr/java
- /etc/profile.d/java.sh

Note: When backing up the above directories, follow the symbolic links to back up the source directory. For example: /user/java is typically a symbolic link to /usr/java_versions/jdk<version>

Other Symbolic Links

These symbolic links may vary in your environment. Check the *Certified Configurations Guide* for the latest third-party and open source versions.

tomcat --> apache-tomcat-5.5.25

apache --> httpd-2.4.6

New Users and Groups

Users: aptare, tomcat (all part of /etc/passwd)

Groups: aptare, dba, tomcat (all part of /etc/group)

Note: Both /etc/passwd and /etc/group should be backed up.

Oracle database: Cold backup

Prior to deploying the Portal for operational use and periodically after installation (at least once a month is recommended), it is recommended to perform a cold backup of the Oracle database. In addition to monthly cold backups, it is also recommended to take a cold backup after a significant software upgrade to re-capture the updated database schema.

This off-line, cold backup means you physically copy or back up the Oracle database data files to another location. This cold backup will simplify the restore process, in the event of unanticipated data loss. With a cold backup, you simply have to restore the files and then import the most recent database export.

To perform a cold backup (Linux)

1. Shutdown the Oracle service:

```
/opt/aptare/bin/oracle stop
```

2. Using your organization's file system backup software, back up all the data files from:

```
$ORACLE_HOME/dbs/initscdb.ora
/data0?/*
```

To perform a cold backup (Windows)

1. Shut down the Oracle service: **OracleServicescdb** in the Windows Administrative Tools, Component Services window.
2. Backup the following files:

```
$ORACLE_HOME\dbs\iniscdb.ora
C:\oradata
```

Note: During installation, you may choose a different drive for the **oradata** install, so verify its location before backing up the data.

Oracle database: Export backups

This is the preferred method of ensuring that your database is backed up on a regular basis.

Database Export (Windows)

The database user **Aptare** must have access to the export files stored in the directory:

c:\opt\oracle\database\tools

Verify that Oracle user has read and execute privileges on these files before starting the database export.

1. Log into the Windows database server.
2. Ensure Oracle TNS Listener and Oracle services are running.
3. At the command prompt execute the script:

   `c:\opt\oracle\database\tools\expdp_database_template.bat`

4. After successful completion, the export file `aptare_scdb.exp` is saved on the Windows database server in the directory:

   `c:\opt\oracle\logs`

5. Copy the `c:\opt\aptare\datarcvrconf\aptare.ks` to `c:\opt\oracle\logs` folder. This file will be required during the database import.

   Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.

6. Copy the `c:\opt\aptare\datarcvrconf\aptare_external_password.properties` file to `c:\opt\oracle\logs` folder. This file will be required during the database import.

   Note: This step is required only if database is exported from an IT Analytics version 11.0 or above.

To schedule the export task refer to the following.

See the section called “Scheduling the Export Job (Windows)” on page 25.

Database Export (Linux)

The database user **Aptare** must have access to the export files stored in the directory:

/opt/aptare/database/tools

Verify that Oracle user has read and execute privileges on these files before starting the database export.

IT Analytics installer supports Portal and database installation on custom path. If your Portal or database is installed in a non-default location, replace `\opt` with the respective absolute installation path in the commands and procedures provided in this section.

1. Log into the Linux database server and switch to user Aptare.
2. Ensure Oracle Listener and Oracle services are running.
3. Change to the directory:

   `/opt/aptare/database/tools`
4. Execute the command:

   `./export_database_template.sh`
5. After successful completion, the export file aptare_scdb.exp is saved in the /tmp directory on the Linux database server.
6. Run the `cp /opt/aptare/datarcvrconf/aptare.ks /tmp` command. This file will be required during the database import.

   Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.
7. Run the `cp /opt/aptare/datarcvrconf/aptare_external_password.properties /tmp` command. This file will be required during the database import.

   Note: This step is required only if database is exported from an IT Analytics version 11.0 or above.

To schedule the export task refer to the following.

See the section called “Scheduling the Export (Linux)” on page 25.

Scheduling the oracle database export

To set up a scheduled Oracle database export refer to the following steps for Linux or Winodws.

Scheduling the Export (Linux)

1. Switch to user **aptare**.

   ```
   su - aptare
   ```

2. Edit or create the crontab file.

   ```
   crontab -e
   ```

3. Using the following diagram, specify the database export schedule.

   Example:

   ```
   04 15 * * * /opt/aptare/database/tools/export_database_template.sh
    2>&1 >>/tmp/database_export.log
   ```

   ```
   *     *     *   *    *        command to be executed
   -     -     -   -    -
   |     |     |   |    |
   |     |     |   |    +----- day of week (0 - 6) (Sunday=0)
   |     |     |   +------- month (1 - 12)
   |     |     +---------- day of        month (1 - 31)
   |     +------------ hour (0 - 23)
   +-------------- min (0 - 59)
   ```

4. Verify the newly scheduled job:

   ```
   crontab -l
   ```

Scheduling the Export Job (Windows)

1. **Start > Programs > Accessories > System Tools > Task Scheduler**
2. Create a new task in the Task Scheduler by specifying the schedule and the script to be run:

   ```
   c:\opt\oracle\database\tools\expdp_database_template.bat
   ```

Oracle database: On demand backup

This method is useful for cases where an immediate backup/export is required; for example, when Customer Support requests a copy of your database for troubleshooting.

To back up the database tables, IT Analytics provides an Oracle data pump utility, exdp. This utility exports the user **portal**, which contains the database tables. To create a full export file of all database objects, refer to the following section.

See "Oracle database: Export backups" on page 22.

For optimum performance, use this utility rather than your favorite backup solution's backup utility (for example, rman) because most backup solutions require archive logging. This setting is not enabled or exposed because archive logging can have a significant, negative impact on performance.

You will import this export in the event that you need to:

- Restore the entire Reporting Database.
  See "Restoring the IT Analytics system" on page 26.
- Retrieve a data table that's been corrupted or accidentally deleted. Simply drop the portal user then import the export.
  See "Import the Oracle database" on page 27.

Restoring the IT Analytics system

The following steps provide guidelines for restoring your system; however, depending on your circumstances, all of these steps may not be relevant.

Performing a Full IT Analytics System Restore

To perform a full system restore:

1. Follow your company's process to perform a system restore from the full system backup that you took, as described in the previous sections. This should restore the application directories and files, in addition to any OS level-specific files, such as the Windows registry, boot sectors, etc.
2. Restore the Oracle database.

Note: If you do not have a full system backup, it still may be possible to recover by re-installing the IT Analytics application, re-installing the Oracle binaries, and then restoring the Oracle database. Contact Cohesity Support if you need to follow this recovery method.

If your data loss is isolated to the Oracle database, it may be possible to skip a full restore and proceed with restoring the Oracle database.

Restoring the Oracle Database

IMPORTANT: Before you try to restore the Oracle database, contact Support. The following steps provide guidance for restoring your database; however, once again, depending on your circumstances, all of these steps may not be relevant.

1. Import the most recent version of the database that you exported.

   See "Import the Oracle database" on page 27.

2. If the import of the database fails it is likely that there are issues with your Oracle environment that are preventing Oracle from starting. In this case you will need to restore the database files from your Cold Backup.

   Using the backup software that you used to perform a cold backup of the database, restore the database with the most recent successful Cold Backup. If you do not have a successful Cold Backup of the Oracle database, import the most recent version of the database that you exported.

Import the Oracle database

Caution: The following process deletes your existing IT Analytics database. Before you try to import the Oracle Database, verify you have no other recovery options and you have a valid database export. Contact Cohesity Support.

Some common problems include:

- Importing unsuccessful backups. Ensure that your backups were successful before you accidentally import old data.
- Importing more than what you need. Do not restore the entire database or import all database tables if you only need to restore one database table. Import only what you need.

Note: Before restoring user objects, stop the Tomcat and Portal processes.

See "Starting and stopping portal server software" on page 44.

Database Import (Linux)

Prerequisites:

- To import the latest successful export of the Oracle database, the database user **Aptare** user must have access to the export files stored in the `/opt/aptare/database/tools` directory.

- The Oracle user must have read and execute privileges on these files before starting the database export.

1. Log into the Linux database server and switch to user Aptare.
2. Place the export file `aptare_scdb.exp` in the `/tmp` directory.

   If you have a different preferred directory (for example, `/new_directory_path`), then place `aptare_scdb.exp` in your preferred directory (`/new_directory_path`). Subsequently, change the path for the creation of directory from `/tmp` to the new directory (`new_directory_path`) in the `/opt/aptare/database/tools/drop_users_linux.sql` file.

   Note: You can also use `/opt/aptare/` or `/opt/aptare/oracle/rdbms/log/` directories.

3. Place the `aptare.ks` file in the `/tmp` directory.

   Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.

4. Place the `aptare_external_password.properties` file in the /tmp directory.

   Note: This step is required only if database is exported from an IT Analytics version 11.0 or above.

5. Stop all Oracle and APTARE services as a root user:

   ```
   /opt/aptare/bin/aptare stop
   ```

6. Copy `/opt/aptare/portalconf/master_data_keys.jceks` file from source IT Analytics virtual machine to target IT Analytics virtual machine to same path.
7. Start Oracle services as the root user:

   ```
   /opt/aptare/bin/oracle start
   ```

8. Verify the Oracle Listener is running.

   ```
   /opt/aptare/bin/aptare status
   ```

9. Run the following command:

```
chmod +x /opt/aptare/database/tools/import_database_template.sh
```

10. Run the script `import_database_template.sh` as Aptare user.

    ```
    /opt/aptare/database/tools/import_database_template.sh
    ```

11. After successful completion, the data pump export file `aptare_scdb.exp` is saved on the Linux database server in the `/tmp` directory.

The `import_database_template.sh` script, unlocks Portal user, grants privileges, and validates the packages after the completion of the import, so they are not required to be run manually. The scripts also address the compilation warnings for the packages that follow.

The import log - `import_scdb.log` is located in the `/tmp` directory.

1. Check the log file for compilation warnings for the packages:
   - view apt_v_solution_history_log
   - cmv_adaptor_pkg
   - avm_common_pkg
   - sdk_common_pkg
   - load_package
   - common_package
   - util

   These compilation warnings are addressed by the script itself and no action is required from the user.

   Note: If you are importing a database from version 10.4, upgrade the portal after the import to a 10.5 build.

   ---

2. This step is required only if you are exporting the database from IT Analytics version 10.5 or above. Run the following commands to copy the `aptare.ks` file to `datarcvrconf` folder.

   ```
   cp /tmp/aptare.ks /opt/aptare/datarcvrconf/
   chown aptare:tomcat /opt/aptare/datarcvrconf/aptare.ks
   chmod 660 /opt/aptare/datarcvrconf/aptare.ks
   ```

3. The encrypted information in `portal.properties` and `datararcvrproperties.xml` files must be replaced after the `aptare.ks` file has been copied to the portal. The information is re-encrypted when the portal services are restarted. The following modifications are required for both the files.

   Modification required for `portal.properties` file

   - Edit `/opt/aptare/portalconf/portal.properties` file
   - Remove all the characters following the first "=" on the lines containing `db.password.encrypted` and `db.ro_user_password.encrypted`
   - Specify the passwords for both users in the relevant lines

   Note: The `db.user` and `db.ro_user` information shown is below is for a default installation. Adjust the blank entries to match your environment. These will be re-encrypted when the portal services are restarted.

   ```
   #Database connection
   db.driver=oracle.jdbc.driver.OracleDriver
   db.url=jdbc:oracle:thin:@//localhost:1521/scdb
   db.user=portal
   db.password=portal
   db.password.encrypted=
   db.connection.max=75
   db.connection.min=25
   db.connection.expiration=30
   db.ro_user=aptare_ro
   db.ro_user_password=aptaresoftware123
   db.ro_user_password.encrypted=
   ```

   Modification required for `datararcvrproperties.xml` file

   - Edit the datasource section of `/opt/aptare/datarcvrconf/datrarcvrproperties.xml` file
   - Clear the encrypted values for **UserId, Password, ro_user**, and **ro_password**

   Note: The **UserId** and **ro_user** information shown is for a default installation. Clear the text entries to match your environment. These will be re-encrypted when the portal services are restarted.

```
<dataSource>
        <Driver>oracle.jdbc.driver.OracleDriver</Driver>
        <URL>jdbc:oracle:thin:@//localhost:1521/scdb</URL>
        <UserId>portal<=/UserId>
        <Password>portal</Password>
        <oracle_service_name>scdb</oracle_service_name>
        <ro_user>aptare_ro</ro_user>
        <ro_password>aptaresoftware123</ro_password>
        <MaxConnections>150</MaxConnections>
        <MinConnections>5</MinConnections>
        <ConnExpirationTime>5</ConnExpirationTime>
    </dataSource>
```

4. This step is required only if you are exporting the database from IT Analytics version 11.0 or above. Execute the following commands to copy the file `aptare_external_password.properties` to `datarcvrconf` directory.

   ```
   cp /tmp/aptare_external_password.properties
   /opt/aptare/datarcvrconf/
   chown aptare:tomcat
   /opt/aptare/datarcvrconf/aptare_external_password.properties
   chmod 660
   /opt/aptare/datarcvrconf/aptare_external_password.properties
   ```

5. Run `updateUser.sh` to change the password of the application account. For example, to change the password for the admin123 application user, run `updateUser.sh admin123 <newPassword>`
6. Restart all Oracle and APTARE services by running them from the root user:

   ```
   /opt/aptare/bin/aptare restart
   ```

7. If the Portal is deployed at a custom path (other than the default path : /opt/aptare), update the system parameter as follows:
   - Login to the Portal host as root user and run:

     ```
     su - aptare
     cd /customPath/aptare/database/tools
     ```

   - Login to SQL Plus as a Portal user and substitute pwd with your password

     ```
     sqlplus portal/pwd@<ServiceName>
     @update_system_parameter.sql
     ```

- Enter the custom installation path of the Portal, when you prompted as:

  ```
  "Enter value for input_path:"
  ```

  Example: If you have the Portal deployed under `/<CustomPath>/aptare/` directory, provide the input path as **/<CustomPath>/** while importing the database on the same host where the Portal is installed.

8. Log into the Portal using the application account.

Database Import (Windows)

To import the latest successful export of the Oracle database, the database user **Aptare** user must have access to the export files stored in the directory:

c:\opt\oracle\database\tools

Verify the Oracle user has read and execute privileges on these files before starting the database export.

1. Log into the Windows database server as a user who is a member of the ORA_DBA group.
2. Place the export file aptare_scdb.exp in the directory:

   C:\opt\oracle\logs

3. The encrypted information in `portal.properties` and `datararcvrproperties.xml` files must be replaced after the `aptare.ks` file has been copied to the portal. The information is re-encrypted when the portal services are restarted. The following modifications are required for both the files.

   Modification required for `portal.properties` file

   - Edit `c:\opt\aptare\portalconf\portal.properties` file
   - Remove all the characters following the first "=" on the lines containing `db.password.encrypted` and `db.ro_user_password.encrypted`
   - Specify the passwords for both users in the relevant lines

   Note: The `db.user` and `db.ro_user` information shown is below is for a default installation. Adjust the blank entries to match your environment. These will be re-encrypted when the portal services are restarted.

   ```
   #Database connection
   db.driver=oracle.jdbc.driver.OracleDriver
   db.url=jdbc:oracle:thin:@//localhost:1521/scdb
   db.user=portal
   ```

```
db.password=portal
db.password.encrypted=
db.connection.max=75
db.connection.min=25
db.connection.expiration=30
db.ro_user=aptare_ro
db.ro_user_password=aptaresoftware123
db.ro_user_password.encrypted=
```

Modification required for `datararcvrproperties.xml` file

- Edit the datasource section of `c:\opt\aptare\datarcvrconf\datrarcvrproperties.xml` file
- Clear the encrypted values for **UserId, Password, ro_user**, and **ro_password**

Note: The **UserId** and **ro_user** information shown is for a default installation. Clear the text entries to match your environment. These will be re-encrypted when the portal services are restarted.

```
<dataSource>
        <Driver>oracle.jdbc.driver.OracleDriver</Driver>
        <URL>jdbc:oracle:thin:@//localhost:1521/scdb</URL>
        <UserId>portal<=/UserId>
        <Password>portal</Password>
        <oracle_service_name>scdb</oracle_service_name>
        <ro_user>aptare_ro</ro_user>
        <ro_password>aptaresoftware123</ro_password>
        <MaxConnections>150</MaxConnections>
        <MinConnections>5</MinConnections>
        <ConnExpirationTime>5</ConnExpirationTime>
    </dataSource>
```

4. Place the `aptare.ks` file in the `C:\opt\oracle\logs` folder.

Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.

---

5. Place the `aptare_external_password.properties` file in the `c:\opt\oracle\logs` folder.

Note: This step is required only if database is exported from an IT Analytics version 11.0 or above.

6. Stop all Oracle and Aptare services using **stopAllServices** from the **Windows Services** tab.
7. Copy `C:/opt/aptare/portalconf/master_data_keys.jceks` file from source IT Analytics virtual machine to target IT Analytics virtual machine to same path.
8. Verify the Oracle TNS Listener is running and start **OracleServicescdb** from the **Windows Services** tab.
9. From the command prompt, run the script import_database_template.bat by executing the command:

   ```
   c:\opt\oracle\database\tools\import_database_template.bat
   ```

10. After successful completion of the import, run **Startallservices** from the **Windows Services** tab.

The `import_database_template.bat` script, unlocks Portal user, grants privileges, and validates the packages after the completion of import, so they are not required to be run manually. The scripts also address compilation warnings for the packages that follow.

The import log `import_scdb.log` is located `c:\opt\aptare\oracle\logs`.

1. Check the log file for compilation warnings for the packages:
   - view apt_v_solution_history_log
   - cmv_adaptor_pkg
   - avm_common_pkg
   - sdk_common_pkg
   - load_package
   - common_package
   - util

   These compilation warnings are addressed by the script itself and no action is required from the user.

Note: If you are importing a database from version 10.4, upgrade the portal after the import to a 10.5 build.

2. Copy the saved file `c:\opt\oracle\logs\aptare.ks` to `c:\opt\aptare\datarcvrconf` folder. Ensure that the file is owned by IT Analytics user and has appropriate Read and Write access to the copied file.

   Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.

3. Copy the saved file `c:\opt\oracle\logs\aptare_external_password.properties` to `c:\opt\aptare\datarcvrconf` folder. Ensure that the file is owned by IT Analytics user and has appropriate Read and Write access to the copied file.

   Note: This step is required only if database is exported from an IT Analytics version 11.0 or above.

4. Run `startAllServices` using the **Windows Services** tab.
5. Run `updateuser.bat` to change the password of the application account. For example, to change the password for the admin123 application user, run: `updateuser.bat admin123 newPassword`
6. Log into the portal Application using the application account.

Restoring When a Cold Backup is Not Available

If you do not have a successful cold backup of your Oracle database, take the following steps to recover your IT Analytics Oracle database. Refer to the Installation and Upgrade Guides for complete details about re-installing the binaries and the database schema. Contact Support for assistance.

1. Re-install the Oracle binaries.
2. Re-install the database schema.
3. Import the latest successful export of your Oracle database to restore the Portal user objects.

   See “Import the Oracle database” on page 27.

Manual steps for database import / export using data pump

Manual steps for Linux Data Pump Export (CDB and non-CDB environments)

Follow the steps to execute the Data Pump Export in a Linux environment

1. Login to the Linux database server and switch to user **aptare**.
2. Ensure that file `/opt/aptare/database/tools/expdp_scdb.par` is owned by **aptare** user and has 755 permissions.
3. Ensure Oracle listener and Oracle services are running.
4. Run following commands:

```
su - aptare
source <INSTALL_PATH>/aptare/bin/aptare_env.sh
sqlplus / as sysdba
alter session set container=scdb;
```

Note: The `alter session set container=scdb;` command is required for a container database. Please ignore it for a non-cdb environment.

```
CREATE OR REPLACE DIRECTORY datapump_dir AS '/tmp';
```

In case of a preferred folder such as `new_directory_path`:

```
CREATE OR REPLACE DIRECTORY datapump_dir AS '/new_directory_path';
```

5. Export the database using following command:

```
/opt/aptare/oracle/bin/expdp
parfile=/opt/aptare/database/tools/expdp_scdb.par
```

6 You can also choose to ignore the `par` file and include parameters in the `expdp` command directly. In other words, the above command can be replaced by the following command which can also be executed from **aptare** user.

```
/opt/aptare/oracle/bin/expdp
system/aptaresoftware@//localhost:1521/scdb FULL=Y
directory=datapump_dir dumpfile=aptare_scdb.exp
logfile=export_scdb.log CONTENT=ALL flashback_time=systimestamp
```

After successful completion, the data pump export file `aptare_scdb.exp` is saved in `/tmp` directory of the Linux Database server.

In case you have specified a preferred directory, `aptare_scdb.exp` is saved to that preferred location (such as `/new_directory_path`).

7 This step is required only if the database is exported from an IT Analytics version of 10.5 or above. Execute `cp /opt/aptare/datarcvrconf/aptare.ks /tmp` command to copy the aptare.ks file to /tmp folder

Manual steps for Linux Data Pump Import (CDB and non-CDB environments)

Follow the steps to execute Data Pump Import in a Linux Environment

1 Place the export file `aptare_scdb.exp` created using data pump export in `/tmp` directory.

If you have a different preferred directory (for example `/new_directory_path`), then place `aptare_scdb.exp` in your preferred directory (`/new_directory_path`).

2 Ensure the `aptare_scdb.exp` file is owned by the **aptare** user and has 755 permissions.

3 Ensure that files `/opt/aptare/database/tools/unlock_portal_linux.sql` and `/opt/aptare/database/tools/impdp_scdb.par` are owned by **aptare** user and have 755 permissions.

4 Using root user, stop all Oracle and Aptare services by running following command: `/opt/aptare/bin/aptare stop` from root user.

5 Using root user start Oracle services by running following command: `/opt/aptare/bin/oracle start`

6 Ensure Oracle listener is running. Using aptare user check for status of Listener using following command: `lsnrctl status`

7 Run following commands:

```
su - aptare
source <INSTALL_PATH>/aptare/bin/aptare_env.sh
sqlplus / as sysdba
alter session set container=scdb;
```

Note: The `alter session set container=scdb;` command is required for a container database. Please ignore it for a non-cdb environment.

Exit sqlplus and drop the existing database users by running the `drop_users_linux.sql script` as the **aptare** user:

```
sqlplus "/ as sysdba"
@/opt/aptare/database/tools/drop_users_linux.sql
```

Note: The `drop_users_linux.sql` script drops the portal and **aptare_ro** users and also drops the portal ROLE (not just the user). If the users or the role do not already exist, the script handles this gracefully and completes without error, so no manual pre-check is needed.

```
CREATE OR REPLACE DIRECTORY datapump_dir AS '/tmp';
```

In case of a preferred folder such as `new_directory_path`:

```
CREATE OR REPLACE DIRECTORY datapump_dir AS '/new_directory_path';
```

8 Run the following command using aptare user

```
/opt/aptare/oracle/bin/impdp
parfile=/opt/aptare/database/tools/impdp_scdb.par
```

9 You can also choose to ignore the par file and include parameters in the impdp command directly. In other words, the above command can be replaced by the following command which can also be executed from aptare user.

```
/opt/aptare/oracle/bin/impdp
system/aptaresoftware@//localhost:1521/scdb
schemas=portal,aptare_ro directory=datapump_dir
dumpfile=aptare_scdb.exp logfile=import_scdb.log
```

10 When import completes in a non-CDB environment, remove first command `'alter session set container = scdb;'` from the file unlock_portal_linux.sql and run following command from aptare user.

Note: The removal of the `'alter session set container = scdb;'` is required only for a non-CDB environment and no change is required if it is a Container database.

```
sqlplus / as sysdba
@/opt/aptare/database/tools/unlock_portal_linux.sql
```

11 After exiting from sqlplus, execute following command from aptare user

```
sqlplus portal/<portal_user_password>@//localhost:1521/scdb
@/opt/aptare/database/tools/validate_sp.sql
```

Note: Use the configured portal user password. The default portal/portal credential no longer applies, because the password is set from the value configured for the portal user during the drop / re-create step.

Post procedure steps

1. Go to `/tmp` directory and check the file `import_scdb.log`.

   In case you have specified a preferred directory, check for `import_scdb.log` in your preferred location.

2. Check the log file for the compilation warnings for the packages: view apt_v_solution_history_log, cmv_adaptor_pkg, avm_common_pkg, sdk_common_pkg, server_group_package, load_package, common_package, util. These compilation warnings are addressed by the script itself and no action is required from the user.

   Note: If you are importing DB from 10.4, upgrade the portal after the import to 10.5 build.

3. This step is required only if the database is exported from an IT Analytics version of 10.5 or above. Run the following commands to copy the `aptare.ks` file to `datarcvrconf` folder.

```
cp /tmp/aptare.ks /opt/aptare/datarcvrconf/
chown aptare:tomcat /opt/aptare/datarcvrconf/
chmod 664 /opt/aptare/datarcvrconf/aptare.ks
```

4. Run `updateUser.sh` to change the password of the application account. For example, to change the password for the admin123 application user, run: `updateUser.sh admin123 newPassword`
5. Restart all Oracle and Aptare services by running `/opt/aptare/bin/aptare restart` from root user.
6. Login to the portal Application using the application account.

Manual steps for Windows Data Pump Export (Both CDB and non-CDB environments)

Follow the steps for Windows Data Pump Export

1 Login to the Windows database server.

2 Ensure Oracle TNS listener and Oracle services are running.

3 Ensure Aptare user has access to the file c:\opt\oracle\database\tools\expdp_scdb_win.par

Run following commands:

```
sqlplus system/aptaresoftware@//localhost:1521/scdb

create or replace directory datapump_dir as 'c:\opt\oracle\logs';

Exit
```

4 After exiting out of sqlplus, execute the following command:
`c:\opt\oracle\bin\expdp parfile=c:\opt\oracle\database\tools\expdp_scdb_win.par`

5 You can also choose to ignore the par file and include parameters in the expdp command directly. In other words, the above command can be replaced by the following command: `c:\opt\oracle\bin\expdp system/aptaresoftware@//localhost:1521/scdb FULL=Y DIRECTORY=datapump_dir LOGFILE=export_scdb.log DUMPFILE=aptare_scdb.exp CONTENT=ALL FLASHBACK_TIME=systimestamp`

6 After successful completion, the data pump export file aptare_scdb.exp is saved in C:\opt\oracle\logs directory of the Windows Database server.

7 Copy file **c:\opt\datarcvrconf\aptare.ks** to **c:\opt\oracle\logs** folder.

Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.

Manual steps for Windows Data Pump Import (Both CDB and non-CDB environments)

Follow the steps for Windows Data Pump Import

1 Login to the Windows database server.

2 The Aptare user will already have access to import files c:\opt\oracle\database\tools\unlock_portal_win.sql and c:\opt\oracle\database\tools\impdp_scdb_win.par. In case the Oracle user does not have read and execute privileges on these files, please ensure privileges are granted before starting the import.

3 Before running `import_database_template.bat`, open `C:\opt\oracle\database\tools\create_impdp_users_win.sql` in a text editor and replace the password placeholders with the passwords you want to use:

- PORTAL_USER_PASSWORD – password for the **portal** user.
- APTARE_RO_USER_PASSWORD – password for the **aptare_ro** user.

Save the file before you start the import. If the placeholders are not updated, the import users are created with invalid passwords and the subsequent import and validation steps fail.

4 Place the export file aptare_scdb.exp in c:\opt\oracle\logs directory

5 In case the name of the export file is capitalized, please change it to lower case. For example, change the name 'APTARE_SCDB.EXP' to 'aptare_scdb.exp'

6 Stop all Oracle and Aptare services using stopAllServices from windows services tab.

7 Start OracleServicescdb from windows services tab and ensure Oracle TNS listener is running.

Run following commands:

```
Sqlplus / as sysdba
```

Alter session set container = scdb; (note this command is included only for a container database, otherwise switch to container database is not required)

```
DROP USER aptare_ro CASCADE;

DROP USER portal CASCADE;

CREATE OR REPLACE DIRECTORY datapump_dir AS 'c:\opt\oracle\logs';

EXIT;
```

8 After exiting out of sqlplus execute following command:

```
c:\opt\oracle\bin\impdp
parfile=c:\opt\oracle\database\tools\impdp_scdb_win.par
```

9 You can also choose to ignore the par file and include parameters in the impdp command directly. In other words, the above command can be replaced by the following command: `c:\opt\oracle\bin\impdp "sys/*@//localhost:1521/scdb as sysdba" SCHEMAS=portal,aptare_ro DIRECTORY=datapump_dir LOGFILE=import_scdb.log DUMPFILE=aptare_scdb.exp`

10 After import is complete, execute the following command: `sqlplus "sys/*@//localhost:1521/scdb as sysdba" @c:\opt\oracle\database\tools\unlock_portal_win.sql`

11 After exiting out of sqlplus, execute the following command:

```
sqlplus portal/<portal_user_password>@//localhost:1521/scdb

@c:\opt\oracle\database\tools\validate_sp.sql
```

Note: Use the portal password that you configured in `create_impdp_users_win.sql` (the value substituted for *PORTAL_USER_PASSWORD*). The default portal/portal credential is no longer used.

Post procedure steps

- To check for import logs, go to `c:\opt\aptare\oracle\logs` and check the file `import_scdb.log`.

- Check the log file for the compilation warnings for the packages: view apt_v_solution_history_log, cmv_adaptor_pkg, avm_common_pkg, sdk_common_pkg, server_group_package, load_package, common_package, util. These compilation warnings are addressed by the script itself and no action is required from the user.

  Note: If you are importing DB from 10.4, upgrade the portal after the import to 10.5 build

- Copy the saved file from `c:\opt\oracle\logs\aptare.ks` to `c:\opt\datarcvrconf\` folder. Ensure that the file is owned by IT Analytics user and has appropriate Read and Write access to the copied file.

  Note: This step is required only if database is exported from an IT Analytics version of 10.5 or above.

- After successful completion of the import process, run StopAllservices using service tab on Windows.
- Run `startAllServices` using service tab on Windows.
- Run `updateUser.bat` from utils directory to change the password of the application account. For example, to change the password for the admin123 application user, run: `updateUser.bat admin123 newPassword`
- Login to the portal Application using the application account.

Chapter 4

Monitoring IT Analytics

This chapter includes the following topics:

Starting and stopping portal server software

Restarting the Portal Server does not have a negative impact on the Data Collector, assuming that the Data Collector established an initial connection to the Portal Server. You do not need to restart the Data Collector after you restart the Portal Server. If the Data Collector is sending data at the time that the Portal Server becomes unavailable, the Data Collector receives an error, and then tries to send the information again. The Data Collector continues to retry and sends alerts until it can reconnect and retransmit.

When you start the Portal Server software, you will start all services that are not already running. How you start and stop the Portal Server software depends on your operating system. Choose the procedure that represents your operating system:

To start Portal Server software on Windows

1. Locate the script to start the Portal Server software:

   ```
   C:\opt\aptare\utils\startportal.bat
   ```

2. Double-click to start the software.
3. Verify the following services are running using the Windows Services Control panel:

- APTARE Agent Tomcat
- APTARE Apache
- APTARE Portal Tomcat
- Oracle Service SCDB
- OracleSCDBTNSListener

To stop Portal Server software on Windows

1. Locate the script to stop the Portal Server software:

   ```
   C:\opt\aptare\utils\stopportal.bat
   ```

2. Double-click to stop the software.
3. Verify the following services are stopped using the Windows Services Control panel:
   - APTARE Portal Tomcat

To start Portal Server software on Linux

- As user root, run one of the following commands:

  ```
  # cd /opt/aptare/bin
  ./tomcat-portal start|restart
  ```

  Individual component startup scripts can be found in the /opt/aptare/bin directory. Startup log files can be found in the /opt/aptare/logs directory.

To stop Portal Server software in Linux

- As user root, run the following commands:

  ```
  # cd /opt/aptare/bin
  ./tomcat-portal stop
  ```

  Individual component startup/shutdown scripts can be found in the /opt/aptare/bin directory. Shutdown log files can be found in the /opt/aptare/logs directory.

Starting and stopping the reporting database

Consider restarting the Reporting Database every 4-8 weeks.

To start the Reporting Database on Windows

1. Locate the scripts to start the Reporting Database:

   ```
   C:\opt\aptare\utils\
   ```

   If all components (Portal Server and Reporting Database) are on the same server and you want to start allcomponents including the Reporting Database use:

   ```
   startallservices.bat
   ```

   If all components (Portal Server and Reporting Database) are on the same server and you want to only start the Reporting Database use:

   ```
   startoracle.bat
   ```

   If the Reporting Database is on a separate server use:

   ```
   startoracle.bat
   ```

2. Verify the following services have started using the Windows Services Control panel:
   - Oracle Service SCDB
   - OracleSCDBTNSListener

To stop the Reporting Database on Windows

1. Locate the scripts to stop the Reporting Database:

   ```
   C:\opt\aptare\utils\
   ```

   If all components (Portal Server and Reporting Database) are on the same server and you want to stop allcomponents including the Reporting Database use:

   ```
   stopallservices.bat
   ```

   If all components (Portal Server and Reporting Database) are on the same server and you want to only stop the Reporting Database use:

   ```
   stoporacle.bat
   ```

   If the Reporting Database is on a separate server use:

   ```
   stoporacle.bat
   ```

2. Verify the following services have stopped using the Windows Services Control panel:
   - Oracle Service SCDB

- OracleSCDBTNSListener

To start the Reporting Database on Linux

Do one of the following

1. If all components (Portal Server and Reporting Database) are on the same server and you want to start all components including the Reporting Database, run the following command:

```
# cd /opt/aptare/bin
./aptare start|restart
```

2. If all components (Portal Server and Reporting Database) are on the same server and you want to only start the Reporting Database, run the following command:

```
# cd /opt/aptare/bin
oracle start
```

3. If the Reporting Database is on a separate server, run the following command:

```
# cd /opt/aptare/bin
oracle start
```

Starting and stopping data collectors

How you start and stop the Data Collector depends on your operating system. Also, the location of the start and stop script/service depends on your backup solution. For most backup solutions, the Data Collector does not run on the backup server (primary server).

To start the Data Collector on Windows

- Start the following services using the Windows Services Control panel:
  - APTARE Agent
  - APTARE WMIServer

To stop the Data Collector on Windows

- Stop the following services using the Windows Services Control panel:
  - APTARE Agent
  - APTARE WMIServer

To start Data Collector on Linux

- As user **root**, run the start or restart command:

```
# cd /etc/init.d
# ./aptare_agent status
# ./aptare_agent start|restart
```

To stop Data Collector on Linux

- As user **root**, run the stop command:

```
# cd /etc/init.d
# ./aptare_agent status
# ./aptare_agent stop
```

Monitoring tablespaces

The Reporting Database contains the user tablespaces outlined in the following table.

See Table 4-1 on page 48.

During your initial installation, IT Analytics created these user tablespaces and corresponding data files. These tablespaces have **AUTOEXTEND** turned on, so when a data file fills up, the tablespace increases the data file. You do not need to add any data files. However, you must add disk space to the mount point as needed; otherwise, IT Analytics cannot extend the data files.

Table 4-1 User Tablespaces

Tablespace	Data Files
aptare_tbs_data_1m	aptare_tbs_data_1m_01.dbf
aptare_tbs_idx_1m	aptare_tbs_idx_1m_01.dbf
aptare_tbs_data_20m	aptare_tbs_data_20m_01-09.dbf
aptare_tbs_idx_10m	aptare_tbs_idx_10m_01-09.dbf
aptare_tbs_data_200m	aptare_tbs_data_200m_01-09.dbf
aptare_tbs_idx_100m	aptare_tbs_idx_100m_01-09.dbf
aptare_tbs_data_200m_lob	aptare_tbs_data_200m_lob_01-09.dbf
aptare_tbs_data_200m_col	aptare_tbs_data_200m_col_01-09.dbf

Table 4-1 User Tablespaces *(continued)*

Tablespace	Data Files
aptare_undo_tbs	aptare_undo_tbs_01.dbf
aptare_temp_tbs	aptare_tbs_temp_01.dbf

Chapter 5

Accessing IT Analytics reports with the REST API

This chapter includes the following topics:

Overview

With the REST APIs, you can access report data as follows:

- You can extract data from tabular reports using pagination in JSON and XML formats.
- You can export reports as HTML, PDF, and CSV formats
- You can export custom dashboards in HTML and PDF formats.

Authentication for REST APIs

To authenticate for the REST APIs, you must generate an API key from IT Analytics Portal and then pass the key as the authentication header from Swagger for the portal for user authentication. Once authenticated, your user information is used for authorization on the portal.

If you have upgraded from 10.5 or lower versions of IT Analytics, you may see authentication error while accessing the reports, as the authentication method is changed from basic to API key. You may need to make changes in your API code/script to be able to access the reports.

Generate API key

To generate the API key from the portal:

1 Open the user menu and click **Create API Key**

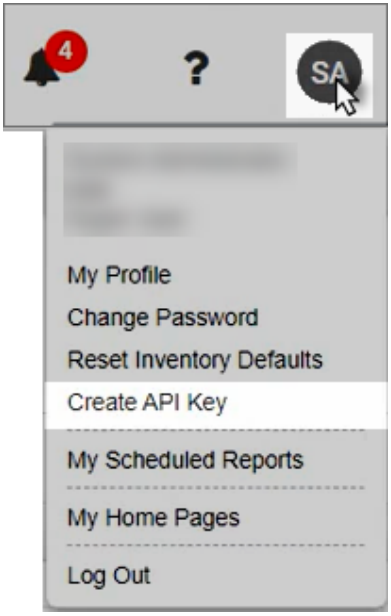

The portal generates and API key and prompts you to copy it to your system. This key is unique for each user and is generated only once, as a new key is generated at every attempt.

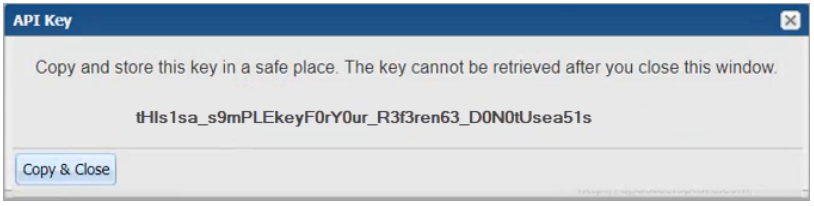

2 Click **Copy & Close**. You need to use this key to execute the REST API.

3 Save the key securely as you need to provide the API key every time you access APTARE reports using REST APIs.

Note: If a user is inactive or removed from the LDAP, update the IT Analytics Portal manually to prevent the user from using the REST API. If not done, the user automatically becomes inactive after the configured number of days.

API key properties

The API key generated using the above steps has the following properties:

- User can have only one active API key at a time.
- The key does not have an expiry date.
- Key is displayed only once when it is generated on the portal. Hence, user must copy and save it securely.
- If the key goes missing, the user can revoke the existing key from the user menu and create a new one.

Extracting data from tabular reports (with pagination)

With the REST API, you can extract data from tabular reports using pagination in JSON and XML formats.

As a prerequisite, ensure you have the API key required for user authentication on Swagger. See “Authentication for REST APIs” on page 50. for steps to generate the API key.

To extract data from tabular reports

1. In the IT Analytics portal, generate a tabular report and save it.
2. Run the saved report.

3 Press **Ctrl + Alt + T** to view the **Report Statistics** and find the **Report ID**.

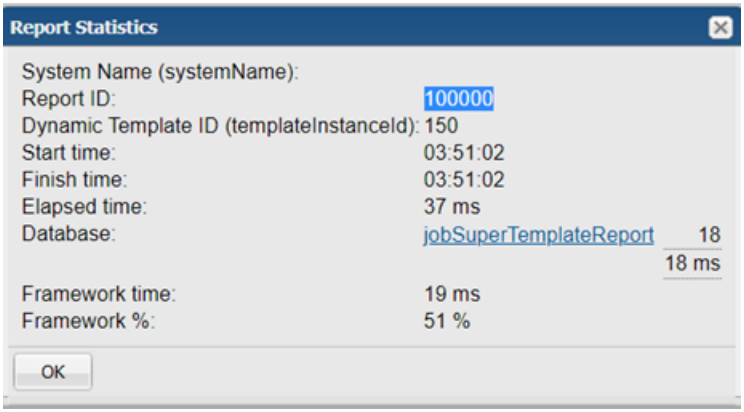

4 Access swagger (<portal_url>/api/swagger-ui.html).

5 Select option **v1** from **Select a spec**.

6 Click **Authorize** and enter the API key generated from the IT Analytics portal.

The API key is passed in the authorization header of the REST API.

7 Expand Report section.

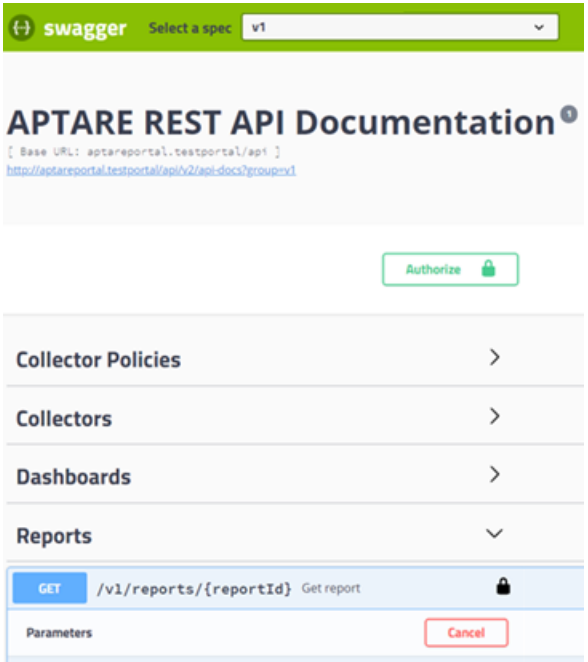

8 Expand **/v1/reports/{reportId}** and click **try it out**.

9 Enter Report Id and click **Execute**.

Exporting reports

With the REST API, you can export reports as HTML, PDF and CSV formats.

As a prerequisite, ensure you have the API key required for user authentication on Swagger. See “Authentication for REST APIs” on page 50. for steps to generate the API key.

To extract data from tabular reports

1 In the IT Analytics portal, generate a tabular report and save it.

2 Run the saved report.

3 Press **Ctrl + Alt + T** to view the **Report Statistics** and find the **Report ID**.

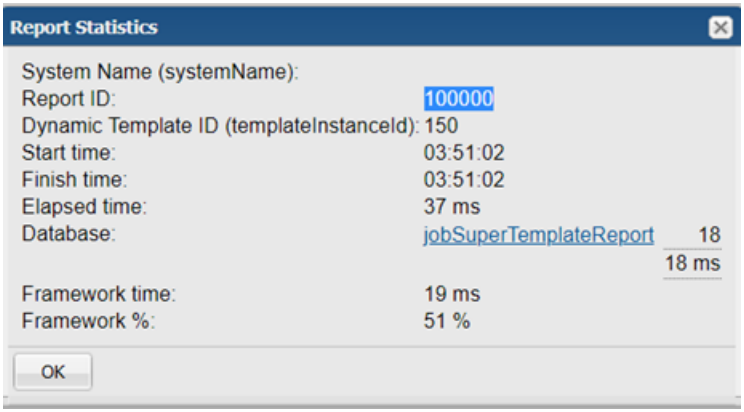

4 Access swagger (<portal_url>/api/swagger-ui.html).

5 Select option **v1** from **Select a spec**.

6 Click **Authorize** and enter the API key generated from the IT Analytics portal.

The API key is passed in the authorization header of the REST API.

7 Expand Report section.

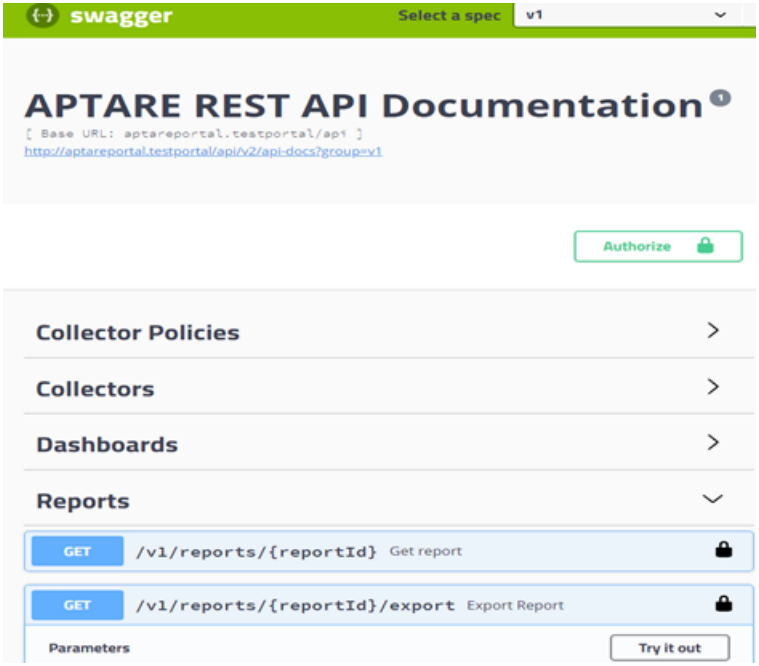

8 Expand **/v1/reports/{reportId}/export** and click **try it out**.

9 Enter Report Id and click **Execute**.

Note: To obtain the report in JSON format, use the URL `<portal_url>/api/v1/reports/{reportId}` to download the JSON file, for example using the curl command.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from IT Analytics Portal tech note for further details.

Exporting custom dashboards

With the REST API, you can export custom dashboards as HTML and PDF formats.

As a prerequisite, ensure you have the API key required for user authentication on Swagger. See “Authentication for REST APIs” on page 50. for steps to generate the API key.

To export custom dashboards

1. In the IT Analytics portal, generate a custom dashboard and save it.
2. Run the saved dashboard.
3. Press **Ctrl + Alt + T** to view the **Report Statistics** and find the **Dashboard ID**.

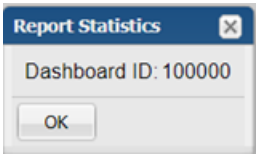

4. Access swagger (<portal_url>/api/swagger-ui.html).
5. Select option **v1** from **Select a spec**.
6. Click **Authorize** and enter the API key generated from the IT Analytics portal.

   The API key is passed in the authorization header of the REST API.
7. Expand Report section.

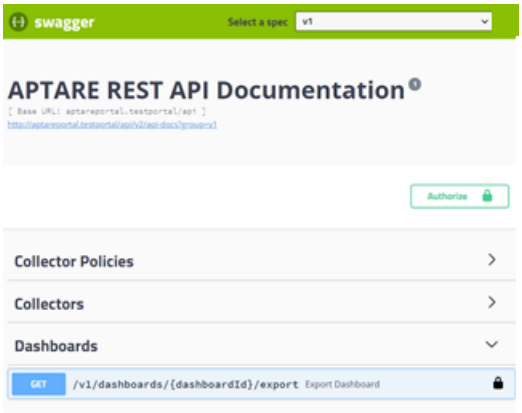

8. Expand **/v1/dashboards/{dashboardId}/export** and click **try it out**.
9. Enter dashboard Id and click **Execute**.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from IT Analytics Portal tech note for further details.

Chapter 6

Defining NetBackup estimated tape capacity

This chapter includes the following topics:

NetBackup estimated tape capacity overview

For NetBackup, some enterprise environments freeze or suspend tape media and do not allow their backup tapes to fill to capacity. In this situation, a IT Analytics database table (**apt_nbu_tape_capacity_est)** can be updated to override the standard estimated capacity calculation. In this table, values for Volume Pool ID, Media Type, and Estimated Capacity will be used to override the estimated capacity calculation.

See “Estimated capacity notes” on page 61.

When values are supplied in this table, the NetBackup estimated capacity calculations will first determine the average size of full tapes--by tape type and volume pool. If this value is less than the value provided in the **apt_nbu_tape_capacity_est** table, the value in the table will be used for the estimated capacity.

Various IT Analytics backup media reports use an estimated capacity calculation:

- Tape Media Summary
- Tape Media Detail

- Media Availability Forecast
- Media Consumption Forecast

Estimated capacity notes

Every time tape data is captured, IT Analytics recalculates the estimated tape capacity for the tape media type and volume pool combination. It takes into consideration the written KBs for the media type and volume pool for full tapes and then stores the average as the estimated tape capacity. Initially, when there is not enough data captured, this value may look low compared to the capacity stated by the vendor. Over time, the estimated capacity improves to the actual number of KBs that is being written to the tapes. Note that the amount of data that fits on the tapes differs based on the compression algorithm used and the type of data that is being backed up, which results in different compression ratios.

Updating the estimated capacity table

Use the following procedure to insert rows into the NetBackup Estimated Capacity database table.

1. Log on to the Portal Server as user **aptare**.
2. Type the following command:

   ```
   sqlplus portal/portal@//localhost:1521/scdb
   ```

3. Insert a row into the **apt_nbu_tape_capacity_est** table. The following example shows how to insert the values.

   ```
   INSERT INTO apt_nbu_tape_capacity_est (volume_pool_id, media_type,
    estimated_mbyte_capacity)
   VALUES (100304, 10, 35850);
   commit;
   ```

   In this example, both the volume_pool_id and media_type will be used to establish the estimated capacity.

4. You also can insert a row into this table using media_type only (omitting the volume_pool_id), as shown in the following example.

   ```
   INSERT INTO apt_nbu_tape_capacity_est (volume_pool_id, media_type,
    estimated_mbyte_capacity)
   VALUES (NULL, 9, 30000);
   commit;
   ```

In this example, only the media_type will be used when the calculation searches for an estimated capacity override.

5. To verify estimated capacities after updating the database table, execute the following commands, supplying the NetBackup Primary Server ID:

```
sqlplus portal/portal@//localhost:1521/scdb
execute media_package.setupTapeMediaCapacity(<primary server ID>);
```

Listing volume pool IDs and media types

Using the Report Template Designer, create a custom report using the following query to identify Volume Pool IDs and Media Type codes:

```
select DISTINCT n.vendor_media_type, t.vendor_media_type_name,
n.volume_pool_id
from apt_v_nbu_tape_media_detail n, apt_v_tape_media t
where n.tape_media_id = t.tape_media_id
and t.server_id in ($(hosts))
```

When you create this custom report via the Report Template Designer, configure the Report Designer to include the selection of a host group, enabling users to narrow the scope of the report when they generate the report.

Chapter 7

Automating host group management

This chapter includes the following topics:

- Automate NetBackup utilities
- Organize clients into groups by management server
- Set up an inactive clients group
- Set up a host group for clients in inactive policies
- Set up clients by policy
- Set up clients by policy type
- IBM Tivoli storage manager utilities
- Set up clients by policy domain
- Set up clients by IBM Tivoli storage manager instance
- Scheduling utilities to run automatically
- Host matching identification for single-domain multi-customer environments

About automating host group management

You can create, move, and organize host groups and link clients/servers to host groups through the Portal. However, you might want to automatically set up your company's host group hierarchy and membership based on unique enterprise business rules and to move or link large quantities of clients/servers in one large batch.

Note: In a Managed Service Provider (MSP) environment, the Application Administrator does not have access to the Reporting Database, so the System Administrator in a MSP environment needs to partner with the Application Administrator, who knows what changes need to be made to the host group hierarchy.

To make host group changes in bulk, use the PL/SQL utilities that IT Analytics provides. Instead of manually creating and organizing host groups through the Portal, you can run PL/SQL utilities to do the work for you.

These utilities provide the following capabilities:

- Matching. You can base your host group management on specific criteria. For example, if you want to organize backup servers by geographical location and your backup servers have a specific naming convention that indicates the servers' region, you need only specify that the SQL utilities to match on that naming convention.

- Automation. You can automate how you create and organize host groups. You can automate how you do the following:
  - Move or copy clients.
  - Move and delete host groups.
  - Organize clients into groups by management server and IBM Tivoli Storage Manager server.
  - Set up an inactive clients group.
  - Set up host group for clients in inactive policies.
  - Set up clients by policy, policy type, policy domain, and IBM Tivoli Storage Manager instance.
  - Load details of new hosts or update existing hosts.
  - Load relationships between hosts and host groups.

These utilities communicate directly with the Reporting Database to manage and manipulate the host group membership for large quantities of servers. There are two types of utilities:

- General. These utilities apply to all backup solutions.
- Product-specific. These utilities only apply to a specific backup solution.

Task overview: managing host groups in bulk

To manage host groups in large quantities, perform the following sequence of steps:

1. Learn about host groups.

   Search online documentation for Host Group.
2. Learn about bulk host group management.

   See “About automating host group management” on page 64.
3. Prepare your SQL environment so that you can run the utilities.

   See “Preparing to use PL/SQL utilities” on page 66.
4. Based on the type of update you want to make to the Reporting Database, choose the utility that can perform that action.
   - See “General utilities” on page 66.
   - See “Merge duplicate backup clients” on page 75.
   - See “Veritas NetBackup utilities” on page 89.
   - See “IBM Tivoli storage manager utilities” on page 97.

5. (Optional) Schedule utilities to run automatically.

   See “Scheduling utilities to run automatically” on page 99.

Preparing to use PL/SQL utilities

A few things you need to know about running the utilities:

- Any time that you are passing a string value as a parameter, the value must be contained within single quotes. For example: **'text'**
- Some functions require that you pass a **group_id**. To obtain this value, go to See “Identifying a host group ID” on page 70.

Also, the PL/SQL utilities have functions that need to be executed from within a PL/SQL session, so set up your environment before you begin to use the utilities.

To prepare your SQL environment

1 Start the **sqlplus** session. If any error messages appear, resolve these errors before you continue:

   - If your database server is on a Linux system, log in as the Linux user **aptare**.
   - If your database is on a Windows system, log in as a user who is a member of the **ORA_DBA** group and open a command prompt window.

```
sqlplus portal/<portal_password>@//localhost:1521/scdb
```

2 Enable server output in PL/SQL:

```
SET SERVEROUTPUT ON
```

General utilities

The utilities contained in this section apply to all host groups and hosts.

- See “Categorize host operating systems by platform and version” on page 67.
- See “Identifying a host group ID” on page 70.
- See “Move or copy clients” on page 70.
- See “Organize clients by attribute” on page 71.
- See “Move host group” on page 72.
- See “Delete host group” on page 73.
- See “Move hosts and remove host groups” on page 73.

- See “Organize clients into groups by backup server” on page 74.
- See “Merge duplicate backup clients” on page 75.

Categorize host operating systems by platform and version

Host data can be collected from various IT Analytics products, such as Capacity Manager, Backup Manager, and Virtualization Manager. Data Collectors persist values as they are collected from the subsystems. For a host’s operating system, subsystems supply values (operating system names) in a variety of formats that do not lend themselves to grouping hosts by OS for reports. For example, Red Hat Linux may be represented as RedHat Linux, rhel, or Red Hat Linux. In order to report on hosts in reasonable groupings, database processing references a set of default regular expressions to parse OS names to categorize the collected host OS data by platform and version.

The following sections provide the details for maintaining and customizing the default Host OS categorization:

- See “Use Regular Expressions to Override or Modify Default Host OS Categorization” on page 67.
- See “Host OS Categorization Default Settings” on page 68.
- See “Categorize Host Operating Systems On Demand” on page 70.

Use Regular Expressions to Override or Modify Default Host OS Categorization

IT Analytics supplies a set of regular expressions (Regex) that define the processing used to process the collected Host OS data to glean the platform and version from the text strings.

- The Host OS Categorization default settings can be modified using the utility to update these categories.
  See “Host OS Categorization Default Settings” on page 68.
  See “Utility to Update Host OS Categorizations” on page 69.
- Data collection automatically processes the Host OS regular expressions and updates the database with the normalized values. If you want to make immediate changes and not wait for the next collection cycle, you can categorize host operating systems in demand.
  See “Categorize Host Operating Systems On Demand” on page 70.
- Regex processing is case-insensitive.

Host OS Categorization Default Settings

Each row in this table represents a regular expression used to determine a common OS platform and version.

Table 7-1 Regular Expressions and Platforms

OS Platform Regex	Version Regex	Ignore	ID
FreeBSD	\d+\.?\d?+	(64-bit)|(32-bit)	1
Solaris	\d+\.?\d?+	(64-bit)|(32-bit)	2
AIX	\d+\.?\d?+	(64-bit)|(32-bit)	3
Ubuntu	6.06LTS|8.04LTS| 10.04LTS|12.04LTS| 14.04LTS|16.04LTS| \d+\.?\d?+	(64-bit)|(32-bit)	4
Data ONTAP	\d+\.?\d?+	(64-bit)|(32-bit)	5
MAC	Rhapsody Developer Release|Server 1.0|Developer Preview|Public Beta|\d+\.?\d?+	Linux|(64-bit)|(32-bit)	6
CentOS	\d+?(/?)\d+?(/?)\d?+	(64-bit)|(32-bit)	8
openSUSE|SUSE|SLES	Leap 42.1| \d+\.?\d?+	(64-bit)|(32-bit)|(x86_64)	9
Windows	(\d+?|SERVER|NT|XP|Vista) \s?\d?+	(64-bit)|(32-bit)	10
Linux		(64-bit)|(32-bit)	11
Red Hat|RedHat|rhel	\d+\.?\d?+	(64-bit)|(32-bit)	12
vmnix		-x86|x86	13
SunOS	\d+\.?\d?+	(64-bit)|(32-bit)	14
Data Domain		(64-bit)|(32-bit)	15
Fedora	\d+\.?\d?+	(64-bit)|(32-bit)	16
Debian	\d+\.?\d?+	(64-bit)|(32-bit)	17

Utility to Update Host OS Categorizations

Use this utility to insert new regular expression rules into the database table or to modify existing rules.

Usage

To insert a regular expression row into the database table, use this command:

```
execute server_group_package.insertCustomerOsNormData(null,
'os_platform_regex', 'os_platform', 'os_version_regex',
'ignore_string', priority, domain_id);
```

To update values in a regular expression row into the database table, use this command:

```
execute server_group_package.insertCustomerOsNormData
(os_normalization_id, 'os_platform_regex', 'os_platform',
'os_version_regex' ,'ignore_strin g', priority, domain_id);
```

Where:

os_normalization_id: This value is unique for each Regex row

(See "Host OS Categorization Default Settings" on page 68.)

IDs less than 100000 are system defaults and cannot be removed, but their values can be modified. When inserting a regular expression into the database table, this value must be null because the process assigns this number.

os_platform_regex: These strings are used to match a substring in the collected text to identify the platform. This field cannot be null.

os_platform: This is the value that is saved to the database when the regular expression is encountered in the collected Host OS. This platform value can never be null, however, the version derived from the version regex may be null.

os_version_regex: This is the regular expression used to match a substring in the collected text to identify the version.

ignore_string: These strings are ignored and are treated as irrelevant details when determining the platform or version.

priority: This value indicates precedence: the higher the value, the higher the priority. For example, Red Hat has a higher priority than Linux, which means that a Host OS that contains a Red Hat substring and a Linux substring will result in a Host OS of Red Hat. User-defined regular expressions must have a priority higher than 1 to override system defaults. This field cannot be null.

domain_id: The Domain ID is shipped with a null default value. In multi-tenancy environments, such as Managed Services Providers, the Domain ID can be updated to change the processing for a specific domain/customer.

Note that a Creation Date also is saved in the database table. This is the date and time that the Regex record was created in the database.

Categorize Host Operating Systems On Demand

Use this utility to process hosts to categorize their operating systems without waiting until the next data collection cycle. This utility uses the regular expressions as described in

See "Categorize host operating systems by platform and version" on page 67.

Usage

```
execute server_group_package.updateExisting
HostOsinfo(hostGroupId);
```

Where:

host_group_id: This is the numeric identifier of a host group.

See "Identifying a host group ID" on page 70.

Identifying a host group ID

Whenever you want to apply a change to a table in the Reporting Database, you must specify the host group's ID in your SQL utility. For example, you need to identify host group IDs if you intend to delete a host group through the reporting database or if you want to use PL/SQL utilities. The host group ID that you specify in your SQL statement is very important and determines how the change is applied:

- If you are not an Managed Services Provider (MSP) and you want to apply a change to a specific host group, determine the group ID for that host group. If you want to apply a change to all host groups, choose the host group ID of the top level folder, the root folder. The group ID for the root folder is always **300000**.
- If you are an MSP and you want to apply a change to a specific customer, choose the host group ID of that customer's domain. If you want to apply a change to all customers, choose the group ID of the top-level root folder (ID = 300000).

Move or copy clients

Description

This utility enables you to move or link large quantities of clients to different host groups. This tool accepts wildcards for the **client_name_mask** attribute. This utility will first validate the existence or the two group names passed as the first two parameters (see below). If the groups do not exist, the utility return an error.

Usage

```
execute server_mgmt_pkg.moveOrCopyClients('<source_host_group>',
'<destination_host_group>', '<client_name_mask>', <move_or_copy_flag>);
```

Where:

source_host_group is the full pathname to the source host group, for example **/ITAnalytics/Primary/GroupA**

destination_host_group is the full pathname to the destination host group.

client_name_mask is a string that can contain wildcards (*). For example, **abc*** indicates all clients that have an **internal_name** that starts with **abc**. To process all clients use the value **NULL** (which should not be within quotes).

move_or_copy_flag is **0**=copy and **1**=move.

Organize clients by attribute

Description

This **groupClientsbyAttributes** utility enables you to organize clients within a host group, based on client attributes. This utility will create host groups named for each attribute value, underneath a parent host group, as shown in the following example.

Example of Organizing Clients by Geographical Location:

Create a host group named Geography. This will be the destination group that will be used to organize the clients by location.

Then, create an attribute named Geography.

For a subset of a host group's clients, set their Geography attribute value to London and for another subset of clients, set their Geography attribute to New York.

Use the following **groupClientsbyAttributes**utility to organize the clients that have a Geography attribute configured.

execute server_mgmt_pkg.groupClientsbyAttributes(300000, 302398, 1, StringObjectListType(stringObjectType('Geography')));

Where **300000** is the group ID of the root group, Global; **302398**is the ID of the Geography group you just created.

Additional References:

See "Identifying a host group ID" on page 70.

Usage

```
execute server_mgmt_pkg.groupClientsbyAttributes(<source_Group_ID>,
<destination_group_ID>, <cascade_Source_Group>,
StringObjectListType(<attribute_List>));
```

Where:

source_Group_ID is the numeric identifier of the host group for which you want to group the clients.

See "Identifying a host group ID" on page 70.

for the steps for finding a group ID.

destination_group_ID is the numeric identifier of the group under which you want to group the clients.

cascade_Source_Group is a numeric flag that indicates if you want this utility to process the source host group's sub-groups and organize those clients in the destination group.

attribute_Listis a comma-separated list of attribute names, each enclosed in straight single quotes. These names are used to create the sub-groups that organize the clients underneath the source group.

Move host group

Description

This utility enables you to move a host group and all of its hierarchical contents to another host group.

Usage

```
execute server_mgmt_pkg.moveServerGroup('<source_host_group>',
'<destination_host_group>');
```

Where:

source_host_group is the full pathname to the source host group, for example **/ITAnalytics/Primary/GroupA**. Be sure to use the host group name, not the host group ID.

destination_host_group is the full pathname to the destination host group. Be sure to use the host group name, not the host group ID.

Delete host group

Description

There are some occasions where you may wish to remove the entire contents of a host group, including any sub-groups within this group and any clients under the hierarchy. This utility allows administrators to perform this type of delete operation.

Note: CAUTION: IF YOU USE THIS COMMAND TO REMOVE A HOST GROUP, AND IF CLIENTS UNDER THE HOST GROUP HIERARCHY DO NOT EXIST IN ANY OTHER GROUP, THE CLIENTS AND ALL OF THEIR ASSOCIATED CONFIGURATIONS WILL BE PERMANENTLY REMOVED FROM THE DATABASE.

Usage

```
execute server_group_package.deleteEntireGroupContents
(100, <parent_group_id>, <group_to_remove_id>);
```

Where:

parent_group_id is the group id of the parent group which contains the group to be deleted.

group_to_remove_id is the group id of the group to be deleted.

Move hosts and remove host groups

This process, often referred to as Server Group Cleanup, enables removal of backup server groups that had been created automatically in prior IT Analytics versions. In addition to cleaning up server groups, this process can also be used on other host groups.

Description

Prior versions of IT Analytics automatically created several server/host groups during backup data collection. In certain environments, these auto-generated groups may not be needed, as other host groups are more relevant. This utility can be used to clean up a Portal's host groups by moving servers/hosts and child host groups from a host group and then deleting the source host group. While this utility, by default, is intended for system-created host groups, it can be used for any host group that you want to delete, but retain its contents.

Note: Once this process completes, log out of the Portal and log back in before accessing host groups and hosts in the Inventory.

Best Practice: In multi-tenancy environments, run this command on a domain-by-domain basis, starting from the bottom of the domain hierarchy to the top IT Analytics domain. This ensures that each domain has been explicitly processed with log messages that confirm the actions taken.

Note: CAUTION: THIS COMMAND MOVES CLIENTS TO THE DOMAIN'S HOME HOST GROUP AND THEN PERMANENTLY REMOVES THE SPECIFIED HOST GROUPS FROM THE DATABASE. RUN THIS COMMAND IN VALIDATE MODE FIRST TO VERIFY THAT THE ACTIONS REPRESENT THE INTENDED RESULTS.

Usage

```
server_mgmt_pkg.serverGroupCleanup(<processingMode>,
'<domain_name>', (<server_group_names_list>),
'<log_file_path_name>', '<log_file_name>');
```

Where:

processing_mode is either 1 = Validate or 2 = Execute. Run this command in Validate mode first to understand what hosts will be moved and what host groups will be deleted.

domain_name, enclosed in single straight quotes, is the case-insensitive name of the IT Analytics domain for the group to be deleted. See the Best Practice listed above.

server_group_names_list is a comma-separated list of host group names to remove, in single straight quotes. This list must be enclosed in parentheses and prefaced with **stringListType**. If NULL is specified, the utility will process these system-created host groups: NetBackup Policy Types, NetBackup Policies, Inactive Policy Clients, and Policy Domains.

log_file_path_name, enclosed in single straight quotes, is the location of the log file for this process.

log_file_name, enclosed in single straight quotes, is the name of the log file.

Example of Validate Mode:

```
execute server_mgmt_pkg.serverGroupCleanup
(1, 'EMEAfinance', stringListType
('NetBackup Policy Types','NetBackup Policies',
'Inactive Policy Clients', 'Policy Domains'),
'/tmp', 'serverGrpCleanup.log');
```

Example of Execute Mode:

```
execute server_mgmt_pkg.serverGroupCleanup
(2, 'EMEAfinance', stringListType
('NetBackup Policy Types','NetBackup Policies',
'Inactive Policy Clients', 'Policy Domains'),
'/tmp', 'serverGrpCleanup.log');
```

Example of Validate without a list of Host Group Names:

```
exec server_mgmt_pkg.serverGroupCleanup
(1, 'EMEAfinance', NULL, '/tmp',
'serverGrpCleanup.log');
```

Organize clients into groups by backup server

This utility can be used for any backup product, such as IBM Tivoli Storage Manager, Veritas Backup Exec, or HP Data Protector.

Description

This utility enables you to create a hierarchy of servers and links all clients that are members of a server into the respective host group.

For example, in an IBM Tivoli Storage Manager environment if you have two IBM Tivoli Storage Manager servers called **TSM1**, **TSM2**, this utility creates two host groups, **TSM1** and **TSM2**, and links the IBM Tivoli Storage Manager server's clients into the corresponding IBM Tivoli Storage Manager host group.

Usage

```
execute common_package.moveClientsIntoServerGroups
( <source_group_id>,<destination_group_id>,
<move_or_copy_flag>, <latest_server_only>);
```

Example:

```
exec common_package.moveClientsIntoServerGroups
(300000, 300010, 1, 1), ;
```

Where:

source_group_id is the internal group ID of the group hierarchy to traverse.

destination_group_id is the group ID in which host groups by management server will be created. Create a host group under **source_group_id** called <vendor_name> Servers and use the group ID of this new host group for the second parameter.

See "Identifying a host group ID" on page 70.

When you organize by server, if a host group exists anywhere under the source group hierarchy with the name of that server, the routine associates the clients with that folder and does not create a new folder under the destination folder. This association occurs whether you explicitly specify the destination folder or if the destination is NULL. However, if you pass a source folder that is at a lower level, the routine only checks for a folder under that hierarchy. If you specify NULL as the destination folder, the routine creates a host group under the source_group_id called Servers.

move_or_copy_flag can be set to **0**=Link (copy) clients or **1**=Move clients. If set to **0**, the utility links the clients to their respective host groups and keeps the clients in their original group location. If set to **1**, the utility moves all clients from the source host group and to their respective host groups.

The utility processes and organizes all clients of the source group hierarchy into the target server grouping. However, if the **move_or_copy flag** is set to **1**, the utility removes only clients in the top level **source_group_id** group--and does not remove those already organized in lower-level sub-groups.

latest_server_only, when set to 1, indicates the last server to back up the client; otherwise, set this flag to 0.

Merge duplicate backup clients

ADVANCED USE ONLY - Due to the nature of this utility's processing, if executed incorrectly, backup clients could be incorrectly moved. This process is particularly

risky in a multi-tenancy environment, where multiple IT Analytics domains could have hosts/servers with the same name.

This merge utility can be used for clients that have been collected from the following backup products: Veritas NetBackup and Commvault Simpana.

Description

Under certain circumstances, backup clients may have duplicate entries in the IT Analytics database. This utility enables you to merge the data of clients that appear more than once in the database.

In most cases, it is not necessary to shut down the data receiver while the client records are being merged. Although not required, it is recommended that you shut down the data receiver before executing this utility so that data will not continue to be collected for the hosts that are being merged.

Merging of NetBackup primary servers is not supported.

Usage

```
execute
duplicate_package.mergeDuplicateServers(<'host_grp'>,<host_name_type>);
```

Example:

```
exec duplicate_package.mergeDuplicateServers('/Corp',1);
```

Where:

host_grp is the explicit host group path and name.

host_name_typeindicates whether to use only the host's base name while finding duplicates, or use the fully qualified name. 0 = fully qualified host name, 1 = host base name.

Example of a host base name: **QAhost1**

Example of a fully qualified host name: **QAhost1.yourcompany.com**

Merge duplicate hosts

To merge duplicate hosts, you must have a local CSV copy of the Duplicate Host Analysis report. This report serves as an input to the duplicate host merging script.

Recommendations:

Follow these recommendations before you merge duplicate hosts:

- Carefully set the report scope and generate the Duplicate Host Analysis report, as its CSV export copy serves as an input for the host merge script.
- Use a copy of the original CSV export as an input for the merge duplicate hosts script. The original CSV can serve as a reference in future.
- Since the host merge process is irreversible, it must be executed by an administrator with comprehensive knowledge of backup solutions.

- Back up the database before performing the host merge since the process is irreversible.

Merge duplicate hosts

Since merge duplicate hosts is an advanced process, make sure you have followed all the recommendations suggested above.

The merging duplicate hosts is performed in the following order:

1. Generate the Duplicate Host Analysis report on the IT Analytics export it in CSV format.
2. Edit the CSV copy of the Duplicate Host Analysis report for the host merge script.
3. Run the host merge script using the CSV input.

Step-1: Generate and export Duplicate Host Analysis report

1 Access the Duplicate Host Analysis report from **Reports** tab > **System Administration** reports.

2 Click the **Duplicate Host Analysis** report name. Use the descriptions in the table below to set the required report scope and generate the report.

Field name	Description
Host Group	Allow you to select the host groups, backup servers, or hosts for the report scope. Your selection narrows down the scope and helps find duplicates more efficiently by targeting specific host groups or host name.
Find Duplicates Using: ■ **Host Name** ■ **Display Name**	You can choose between **Host Name** (default) and **Display Name**. Both searches are case-sensitive. ■ For **Host Name**, the system compares the internal host names to find duplicates. This is the default criterion in the legacy host merge option. ■ For **Display Name**, the system uses the display or external names of the hosts to find duplicates.

Field name	Description
Host Type for the Duplicate Host ■ **Clients Only** ■ **All**	■ **Clients Only** allows you to find duplicates only for hosts that are identified as Clients (hosts backed up by any backup system). ■ **All** detects duplicates for all types of hosts.
Surviving host Selection Criteria ■ **Highest Job Count** ■ **Most Recently Updated**	Allows you to specify the criteria to select the surviving host among the duplicates when performing a host merge. ■ **Highest Job Count**: Selects the host with most associated jobs as the surviving host. This is the default criterion of the legacy host merge option, as a higher job count suggests that the host has more data associated with it. ■ **Most Recently Updated**: Selects the most recently updated host as the surviving host. Use this option when the duplicate hosts found are no longer actively collecting new data, as it helps to retain the most current host.
Cascade into sub-groups	The scope selector default is to cascade to all child sub-groups when generating the report. If you prefer to report ONLY on the host group you selected, then uncheck Cascade into sub-groups.
Filter by Common Attributes	Select this checkbox to have the report scope display attributes using "AND" logic. By selecting this feature, the report will display those results with the intersection of the selected criteria. If this checkbox is not selected, the report will display attributes using "OR" logic. For example, if you select attribute values, Campbell, Engineering, Cost Center 1 and select Filter by Common Attributes, the report will display only the results that contain all 3 attribute values. If you do not select Filter by Common Attributes, the report will display all results with attributes Campbell, Engineering, or Cost Center 1.

Field name	Description
Apply Attributes to Backup Servers	Select this checkbox to apply the attributes only to the backup servers, instead of hosts.

3. After generating the report, export the report in CSV format on your system.
4. Create a copy of the CSV report and prepare the copy for the host merge script as described in the next step.

Step-2: Edit the CSV copy of the Duplicate Host Analysis report

This Duplicate Host Analysis report displays one row each for each suspected duplicate pair. If multiple duplicate hosts are detected, the report displays one row for each duplicate pair. For example, if host A has three potential duplicates, the report displays three rows - one for each duplicate.

Update the values of the following columns in the CSV copy as suggested below:

1. **Surviving Host**: Default value of this report column is **Main**, which indicates that the duplicates will be merged into the Main host. To change the surviving host, change its value to **Duplicate**. This way, all hosts are merged into the duplicate host. **Main** and **Duplicate** are the only acceptable values in this column.
2. **Is Duplicate Host's Merge supported**: This column supports only **Yes** and **No** as values. Delete all the rows containing the value **No** from the report CSV that you plan to use as input for the host merge process.

None other than the above modifications must be made to the report CSV that you plan to use for the host merge process. Your report CSV is now ready to serve as an input for the host merge script.

Step-3: Run the host merge script using the report CSV

The host merge script has a provision to perform a pre-assessment during which it evaluates errors in the CSV and suggests corrections before proceeding further. You must ensure a successful pre-assessment and only then proceed to merge the hosts. Any error in the report CSV will result in the script aborting the process. You must provide the report CSV path along with the file name, log file path, and log file name when you run the script.

Caution: As the host merge process is irreversible, you must back up your database and follow all the recommendations suggested above before you proceed.

To merge duplicate hosts:

1. Run the `host_merge.sql` script from the `../database/tools` directory.

   You can run the script from `SQL*Plus` or `SQL Developer` as **portal** or equivalent user which has access to all the schema tables.

2. Enter the following details when requested by the script:

   - **Enter 1 or 2:** Enter **1** to run the pre-assessment when you run the script for the first time. You can enter **2** when the pre-assessment is successful.
   - **Enter Duplicate Host Analysis CSV file name with full path:** Enter the report CSV file path including the file name.
   - **Enter log file path:** Enter the location for log file (without the file name)
   - **Enter log file name:** Enter the name of the log file.

   After your pre-assessment is successful, repeat this step with option **2** to complete the host merge.

Merge duplicate hosts script run samples

Figure 7-1 Sample of how `host_merge.sql` script is run

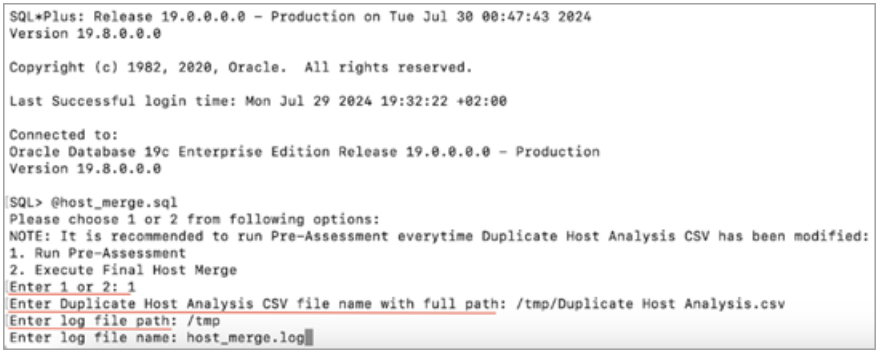

Figure 7-2 Sample-1of an unsuccessful pre-assessment

```
SQL> @host_merge.sql
Please choose 1 or 2 from following options:
NOTE: It is recommended to run Pre-Assessment everytime Duplicate Host Analysis CSV has been modified:
1. Run Pre-Assessment
2. Execute Final Host Merge
Enter 1 or 2: 1
Enter Duplicate Host Analysis CSV file name with full path: /tmp/Duplicate Host Analysis.csv
Enter log file path: /tmp
Enter log file name: host_merge.log
================================================================================

PRE-ASSESSMENT SUMMARY:
================================================================================
 File Name: /tmp/Duplicate Host Analysis.csv
 Log File path: /tmp
 Log file name: host_merge.log
 Total Records Processed: 24
 Number of Unique Surviving Hosts: 22
 Number of Unique Duplicate Hosts: 22
 Total Jobs to be merged from duplicate hosts: 0
 Total lines with Main as the surviving host: 19
 Total lines with Duplicate as the surviving host: 3

================================================================================

ISSUES TO BE RESOLVED IN /tmp/Duplicate Host Analysis.csv BEFORE PROCEEDING FOR HOST MERGE:
================================================================================

INCORRECT SURVIVING HOST:
--------------------------------------------------------------------------------
 Total lines with 'Unknown' surviving host(allowed values are: Main/Duplicate): 2
 List of Sequence Numbers with 'Unknown' surviving host: 180,216
 List of values identified as incorrect surviving host indicator: Maine,Dup
RESOLUTION: Please select surviving host as either 'Main' or 'Duplicate' on the mentioned sequenece numbers.
--------------------------------------------------------------------------------
```

Figure 7-3 Sample-2 of an unsuccessful pre-assessment

```
SQL> @host_merge.sql
Please choose 1 or 2 from following options:
NOTE: It is recommended to run Pre-Assessment everytime Duplicate Host Analysis CSV has been modified:
1. Run Pre-Assessment
2. Execute Final Host Merge
Enter 1 or 2: 1
Enter Duplicate Host Analysis CSV file name with full path: /tmp/Duplicate Host Analysis.csv
Enter log file path: /tmp
Enter log file name: host_merge.sql
================================================================================

PRE-ASSESSMENT SUMMARY:
================================================================================
 File Name: /tmp/Duplicate Host Analysis.csv
 Log File path: /tmp
 Log file name: host_merge.sql
 Total Records Processed: 24
 Number of Unique Surviving Hosts: 23
 Number of Unique Duplicate Hosts: 23
 Total Jobs to be merged from duplicate hosts: 0
 Total lines with Main as the surviving host: 19
 Total lines with Duplicate as the surviving host: 4

================================================================================

ISSUES TO BE RESOLVED IN /tmp/Duplicate Host Analysis.csv BEFORE PROCEEDING FOR HOST MERGE:
================================================================================

SURVIVING HOST SELECTED AS DUPLICATE OR VICE-VERSA:
--------------------------------------------------------------------------------
 Total hosts where Surviving host is marked as Duplicate in other line : 1
 List of Sequence Numbers where Surviving host is marked as Duplicate: 191
 Display name of host(s) selected as surviving as well as duplicate: auvimel000s930-bckup.downergroup.internal
 Internal host id of the host(s) selected as surviving as well as duplicate: 322015
RESOLUTION: Please make sure that a surviving host is not marked duplicate to maintain data consistency post merge as expect
--------------------------------------------------------------------------------

SOME HOST ID IN THE FILE DO NOT EXIST IN DATABASE:
--------------------------------------------------------------------------------
 Total hosts where Surviving host is not present in DB : 1
 List of Sequence Numbers where Surviving host is not present in DB: 217
RESOLUTION: Please re-run Duplicate Host Analysis report OR fix the host ids for the mentioned sequence numbers.
```

Figure 7-4 Successful pre-assessment

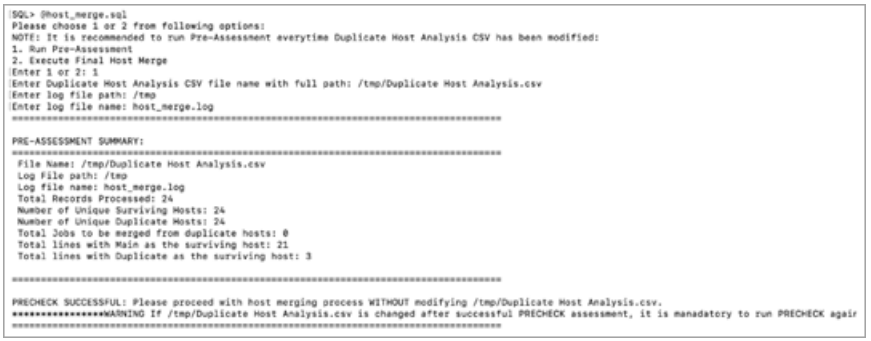

```
SQL> @host_merge.sql
Please choose 1 or 2 from following options:
NOTE: It is recommended to run Pre-Assessment everytime Duplicate Host Analysis CSV has been modified:
1. Run Pre-Assessment
2. Execute Final Host Merge
Enter 1 or 2: 1
Enter Duplicate Host Analysis CSV file name with full path: /tmp/Duplicate Host Analysis.csv
Enter log file path: /tmp
Enter log file name: host_merge.log
================================================================================

PRE-ASSESSMENT SUMMARY:
================================================================================
 File Name: /tmp/Duplicate Host Analysis.csv
 Log File path: /tmp
 Log file name: host_merge.log
 Total Records Processed: 24
 Number of Unique Surviving Hosts: 24
 Number of Unique Duplicate Hosts: 24
 Total Jobs to be merged from duplicate hosts: 0
 Total lines with Main as the surviving host: 21
 Total lines with Duplicate as the surviving host: 3

================================================================================

PRECHECK SUCCESSFUL: Please proceed with host merging process WITHOUT modifying /tmp/Duplicate Host Analysis.csv.
**********WARNING If /tmp/Duplicate Host Analysis.csv is changed after successful PRECHECK assessment, it is manadatory to run PRECHECK agair
================================================================================
```

Figure 7-5 Host merge sample

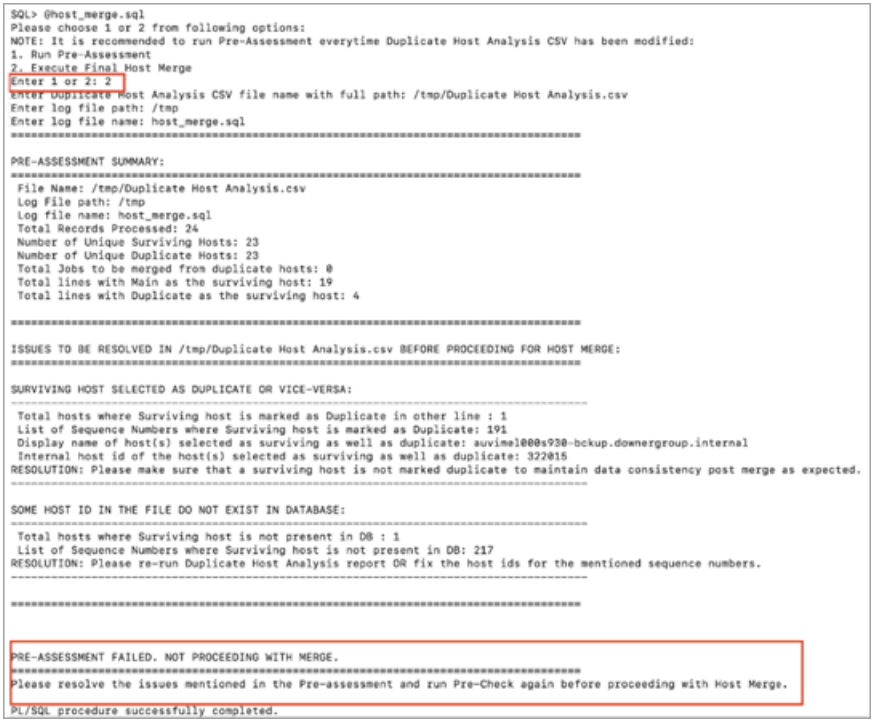

```
SQL> @host_merge.sql
Please choose 1 or 2 from following options:
NOTE: It is recommended to run Pre-Assessment everytime Duplicate Host Analysis CSV has been modified:
1. Run Pre-Assessment
2. Execute Final Host Merge
Enter 1 or 2: 2
Enter Duplicate Host Analysis CSV file name with full path: /tmp/Duplicate Host Analysis.csv
Enter log file path: /tmp
Enter log file name: host_merge.sql
================================================================================

PRE-ASSESSMENT SUMMARY:
================================================================================
 File Name: /tmp/Duplicate Host Analysis.csv
 Log File path: /tmp
 Log file name: host_merge.sql
 Total Records Processed: 24
 Number of Unique Surviving Hosts: 23
 Number of Unique Duplicate Hosts: 23
 Total Jobs to be merged from duplicate hosts: 0
 Total lines with Main as the surviving host: 19
 Total lines with Duplicate as the surviving host: 4

================================================================================

ISSUES TO BE RESOLVED IN /tmp/Duplicate Host Analysis.csv BEFORE PROCEEDING FOR HOST MERGE:
================================================================================

SURVIVING HOST SELECTED AS DUPLICATE OR VICE-VERSA:
--------------------------------------------------------------------------------
 Total hosts where Surviving host is marked as Duplicate in other line : 1
 List of Sequence Numbers where Surviving host is marked as Duplicate: 191
 Display name of host(s) selected as surviving as well as duplicate: auvimel000s930-bckup.downergroup.internal
 Internal host id of the host(s) selected as surviving as well as duplicate: 322015
RESOLUTION: Please make sure that a surviving host is not marked duplicate to maintain data consistency post merge as expected.
--------------------------------------------------------------------------------

SOME HOST ID IN THE FILE DO NOT EXIST IN DATABASE:
--------------------------------------------------------------------------------
 Total hosts where Surviving host is not present in DB : 1
 List of Sequence Numbers where Surviving host is not present in DB: 217
RESOLUTION: Please re-run Duplicate Host Analysis report OR fix the host ids for the mentioned sequence numbers.
--------------------------------------------------------------------------------

================================================================================


PRE-ASSESSMENT FAILED. NOT PROCEEDING WITH MERGE.
================================================================================
Please resolve the issues mentioned in the Pre-assessment and run Pre-Check again before proceeding with Host Merge.

PL/SQL procedure successfully completed.
```

Bulk load utilities

The utilities contained in this section load hosts and their relationships to host groups into the Reporting Database as a batch process using a comma-delimited file. These utilities load new hosts as well as update existing hosts.

Bulk Load utilities must be run in SQLPLUS as user **aptare**.

1. Log in to the Portal server as the user **aptare** for Linux, or an administrator user for Windows.

   su - aptare

2. Open a command-line window.
3. Change the directory to the stored_procedures directory.

   ```
   /opt/aptare/database/stored_procedures (Linux)
   \opt\oracle\database\stored_procedures (Windows)
   ```

4. Execute the command: `sqlplus portal/portal@//localhost:1521/scdb`
5. Run the desired command as listed in the following tables.

This section contains the following topics:

- See "Load host aliases" on page 83.
- See "Load details of new hosts or update existing hosts" on page 84.
- See "Load relationships between hosts and host group" on page 86.

Load host aliases

Description

Sets up host aliases from a comma-delimited file containing a list of hosts and host aliases.

for a description of the logic used to process hosts and aliases.

Usage

```
execute load_package.loadHostAliasFile('<file_name>',
'<log_path_name>','<log_file_name>');
```

Where file_name is the fully qualified path to the **csv** file that contains the aliases to be loaded.

Example of the file_name specification:

'/opt/aptare/database/HostAliases.csv'

Example of the command execution:

execute load_package.loadHostAliasFile('/opt/aptare/database HostAliases.csv',
'/tmp','loadHostAlias.log');

Load File Specification

The specification for the comma-delimited file is as follows:

```
<domain>, <hostname>, <alias_hostname>
```

Example: **Enterprise, whitney,182.16.1.101**

The detailed specification for each field follows:

```
domain                CHAR(128)     NOT NULL
hostname              CHAR(64)      NOT NULL
alias_hostname        CHAR(64)
```

The alias_hostname can be either a host name (up to 64 characters) or an IP address.

Data Constraints

Field values cannot contain embedded commas.

The second field, **hostname**, is the external name, as defined in the Portal database.

The csv file must exist and be readable.

The csv filename must be specified within single quotes.

Logic Conditions

If the host alias already exists, no updates take place.

If the host alias does not already exist in the Reporting Database, the utility adds it.

The utility applies case differences in the input file as updates to preexisting rows.

Logging

The utility logs all additions, updates, warnings and errors to the specified log file. Logging strings are typically in the format: **Date -> Time -> load_package:sub_routine -> Action**

sub_routine is the sub-routine that is being executed (e.g., loadHostLine).

Load details of new hosts or update existing hosts

Description

Imports host details from a comma-delimited file containing a list of hosts and host attributes.

Usage

```
execute load_package.loadServerFile('<file_name>'
[,'<source_name>']);
```

file_name is the fully qualified path to the **csv** file. Example:

```
/opt/aptare/database/hosts.csv
```

source_name is an optional, case-insensitive string, up to 100 characters, representing the source of the host details; for example, CMDB might be relevant for a change management database. This source information is retained for historical purposes, to track how the host was added to the database. If nothing or NULL is provided for this parameter, CSV Load will be inserted as the source into the reporting database.

Load File Specification

The specification for the comma-delimited file is as follows:

```
path_to_host_group, internal_name, external_name,
description, location, IP_address, make, model, OS
```

Example:

```
/APTARE/Test,testhost01,testhost01,description,location,
172.20.16.1,Sun,E450,Solaris 10
/APTARE/Test,testhost02,testhost02,,location,
172.20.16.2,Sun,,Solaris
```

The detailed specification for each field follows:

```
internal_name              CHAR(128)    NOT NULL
external_name              CHAR(128)    NOT NULL
description                CHAR(256)
location                   CHAR(64)
ip_address                 CHAR(40)
make                       CHAR(64)
model                      CHAR(64)
os_version                 CHAR(128)
```

Data Constraints

Field values cannot contain embedded commas.

The first field, **path_to_host_group**, must be the full path to an existing host group otherwise the host will not be inserted.

The csv file must exist and be readable.

The csv filename must be specified within single quotes.

Logic Conditions

If the host already exists in the specified host group, the utility updates its details.

If the host does not already exist in the Reporting Database, the utility adds the host to the specified host group.

If a host attribute field has a **NULL** value in the input file, the corresponding field in the database will not be updated for a pre-existing row.

The utility applies case differences in the input file as updates to preexisting rows.

Since the primary key to the record is the **internal_name**, the **internal_name** for a host cannot be updated via this utility.

If the number of parameters passed in a row exceeds 9, the utility skips the row.

Logging

The utility logs all additions, updates, warnings and errors to the file **scon.log**, which is located under **/tmp** by default on Linux systems and `C:\opt\oracle\logs` on Windows systems. Logging strings are typically in the following format:

Date -> Time -> Level -> load_package:sub_routine -> Action

Where:

Level is DBG, INFO,WARN, or ERR.

sub_routine is the sub-routine that is being executed (e.g. loadHostLine).

Action is the action that was being reported on.

Example:

```
14-MAY 21:56:08 INFO : updating host: z0001-web0600-s
in /APTARE/Infrastructure
```

Load relationships between hosts and host group

Description

Imports host-to-host-group relationships from a comma-delimited file. You can choose to audit the host movement, which records the details when a host is removed, added, or moved.

See “Sample Audit File (output from load_package.loadGroupMemberFile)” on page 89.

Usage

```
execute load_package.loadGroupMemberFile('<file_name>',
'<recycle_group>', <remove_old_entries>,'<audit_pathname>',
'<audit_output_file>', <do_log>);
```

Where:

file_name is the fully qualified path to the csv file. For example:

```
/opt/aptare/database/hosts.csv
```

recycle_group is the full path to the group into which deleted hosts will be moved (i.e., the 'recycle bin').

remove_old_entries enables you to remove relationships in the Reporting Database that are not in the file. If set to **1** and where there are hosts with a previous relationship to a host group and where that relationship is no longer represented within the file, the utility moves those hosts to the recycle group. If set to **0**, the utility does not remove those hosts.

audit_pathnameis the full path to the audit file, not including the filename.

audit_output_fileis the name of the audit file where the audit results will be stored.

do_logenables you to turn on the auditing function so that all host movements are logged in the audit_output_file. Enter a numeric: 0 or 1, where 0 = No, 1 = Yes.

Example command:

execute load_package.loadgroupmemberfile ('/opt/aptare/database/movehosts.csv','/Global1/Recycle',1,'/opt/aptare/database','movehosts.out',1);

Load File Specification

The specification for the comma-delimited file is as follows:

```
path_to_host_group, internal_name1, internal_name2,
internal_name3, etc.
```

Where path_to_host_group is the fully qualified path to the host group into which the hosts should be added, and **internal_name1** is the internal name of a host within the existing host group hierarchy.

Example:

```
/APTARE/Test, testhost01, testhost02
/APTARE/Infrastructure, testhost02, testhost03
```

Detailed specification for each field follows:

```
internal_name              CHAR(64)    NOT NULL
```

Data Constraints

The first field, **path_to_host_group**, must be the full path to an existing host group. If any host groups in the **path_to_host_group** field value do not exist, the utility creates them.

Field values cannot contain embedded commas.

The **csv** file must exist and be readable.

The recycle group folder must exist.

Each row must have at least one host specified, otherwise the row will not be processed.

Logic Conditions

If you list hosts after the **path_to_host_group** field and those hosts are located in the existing host group hierarchy, the utility adds those host groups to the specified host group.

If a host with the specified internal name does not exist in the hierarchy, the relationship will not be added. The host must already be configured in the reporting database.

If any host groups in the **path_to_host_group** field value do not exist, the utility creates them.

If the **removeOldEntries** parameter is set to **1**, the utility assumes that this file will contain all the required relationships. In other words, for all the host groups that you specify in the file, only those hosts will be in that group after you run this utility. If the host group previously contained other host(s) that are now no longer listed in the file, the utility removes those host(s) from the host group and moves them to the recycle folder.

The utility does not delete host groups from the Reporting Database; it only removes members of a host group.

If a host group in the Reporting Database is not listed in the file, the utility does not take any processing action against that host group.

Host groups with many hosts can be split into multiple lines for ease of file maintenance--for example, the host group and some of the hosts appear on the first line, then the same host group and other hosts appear on subsequent lines.

Logging

The utility logs all additions, updates, warnings, and errors to the **scon.log** file, which is located under **/tmp** by default on Linux systems and `C:\opt\oracle\logs` on Windows systems. Logging strings are typically in the following format:

Date -> Time -> Level -> load_package:sub_routine -> Action

Where:

Level is DBG, INFO,WARN, or ERR.

sub_routine is the sub routine that is being executed (e.g. **loadServerLine**).

Action is the action that was being reported on.

Example:

```
14-MAY 19:00:06 ERR  load_package:loadServerGroupMembers:
Host group /APTARE/Business Views/Regional Offices/
Connecticut does not exist on line 6 of the data load file
```

Sample Audit File (output from load_package.loadGroupMemberFile)

```
27-FEB 14:18:35 Start processing host group membership, filePathname:
 /opt/mycompany/database/movehosts.csv recycleGroup: /Global/Recycle
 , removeOldEntries: 1)
27-FEB 14:18:35 Adding host: Whitney (104637) to group: /Global/Corp
 (102572)
27-FEB 14:18:35 Adding host: K2(104638) to group: /Global/Corp
(102572)
27-FEB 14:18:35 Adding host: Everest (102573) to group:
/Global/Bangalore (104538)
27-FEB 14:18:35 Adding host: McKinley (104637) to group:
/Global/United States/Northwest(104639)
27-FEB 14:18:35 Moving host: testhost (102573) from group: testgroup
 (102572) to group: /Global/Recycle (104541)
27-FEB 14:18:35 Completed processing host group membership,
filePathname: /opt/mycompany/database/movehosts.csv recycleGroup:
/Global/Recycle , removeOldEntries: 1)
```

Veritas NetBackup utilities

The following utilities apply only to NetBackup environments.

See “Automate NetBackup utilities” on page 89.

See “Organize clients into groups by management server” on page 91.

See “Set up an inactive clients group” on page 94.

See “Set up a host group for clients in inactive policies” on page 94.

See “Set up clients by policy” on page 95.

See “Set up clients by policy type” on page 96.

Automate NetBackup utilities

The Cohesity NetBackup utilities listed in this section can be set up to run automatically. A stored procedure can be edited to customize the list of utilities to be run as a background job on a particular schedule.

Windows:

```
C:\opt\oracle\database\stored_procedures\nbu\setup_nbu_jobs_manual.sql
```

Linux:
`opt/aptare/database/stored_procedures/nbu/setup_nbu_jobs_manual.sql`

All five are included in this file. To omit a particular utility from the scheduled job, use the following syntax before and after the block of code.

See "Veritas NetBackup utilities" on page 89.

- Before the block of code to be omitted, use: **/***
- After the block of code to be omitted, use: ***/**

Example of a Scheduled NetBackup Utility

In a text editor, open the **setup_nbu_jobs_manual.sql** file and modify the schedule to meet your needs. The following example illustrates how to edit syntax to customize the schedule.

```
-------------------------------------------------------------------------------------------

    -- Move clients that are in inactive policies
    -- Frequency: Every day at 02:30

-------------------------------------------------------------------------------------------

    jobName := dba_package.getSchedulerJobName('setupInactivePolicyClients');

    IF (jobName IS NOT NULL AND LOWER(jobName) <> LOWER('setupInactivePolicyClients'))
 THEN
        DBMS_OUTPUT.PUT_LINE('setupInactivePolicyClients exists with default name '||
jobName ||
' hence will be removed and recreated.');
        DBMS_SCHEDULER.DROP_JOB(job_name => jobName);
        jobName := NULL;
    END IF;

    IF jobName IS NULL THEN
        DBMS_SCHEDULER.CREATE_JOB(
            job_name         => 'setupInactivePolicyClients',
            job_type         => 'PLSQL_BLOCK',
           job_action       => 'server_mgmt_pkg.setupInactivePolicyClients(NULL, NULL,
 0, 0);',  -- What to run
            start_date       => SYSDATE + (5/48),           -- First run is 150 mins
from initial installation
          repeat_interval  => 'TRUNC(SYSDATE+1,''DD'') + (5/48)',  -- Next run is 2:30
```

```
each subsequent day
           enabled          => TRUE);
   ELSE
      DBMS_OUTPUT.PUT_LINE('setupInactivePolicyClients exists and will be altered with
updated version.');
      DBMS_SCHEDULER.SET_ATTRIBUTE(
           name         => jobName,
           attribute    => 'job_type',
           value        => 'PLSQL_BLOCK'
           );
      DBMS_SCHEDULER.SET_ATTRIBUTE(
           name         => jobName,
           attribute    => 'job_action',
           value        =>  'server_mgmt_pkg.setupInactivePolicyClients(NULL, NULL, 0,
0);'
           );
      DBMS_SCHEDULER.SET_ATTRIBUTE(
           name         => jobName,
           attribute    => 'repeat_interval',
           value        => 'TRUNC(SYSDATE+1,''DD'') + (5/48)'
           );
   END IF;
   DBMS_OUTPUT.put_line('setupInactivePolicyClients set to run at 2:30 every day');
```

Scheduling a NetBackup Utility Job to Run Automatically

Execute the SQL file, as **aptare** user on a Linux system or on a Windows system, as an Administrator who is a member of the **ORA_DBA** group:

```
sqlplus portal/<portal_password>@//localhost:1521/
scdb@setup_nbu_jobs_manual.sql
```

Organize clients into groups by management server

This utility enables you to create a hierarchy of management servers and links all clients that are members of the management server into the respective host group. For example, in a NetBackup environment if you have two primary servers called **primary 1 and primary 2**, this utility creates host groups named **primary 1** and **primary 2** and links the primary server's clients into the corresponding group. Two versions of this utility are available:

- Move clients into a Primary Server Group:

See “Organize clients into groups by management server” on page 91.

- Move clients into a Primary Server Group & Exclude the Policy Client & Cascade to Sub-groups:
  See “Organize clients into groups by management server” on page 91.

Basic Usage with 4 Parameters

```
execute nbu_adaptor_pkg.moveClientsIntoPrimaryGroups
( <source_group_id>,<destination_group_id>,
<move_clients>, <latest_primary_only> );
```

Example: **exec nbu_adaptor_pkg.moveClientsIntoPrimaryGroups(300000, 300010, 1, 1);**

source_group_id is the internal group ID of the host group hierarchy to traverse.

destination_group_id is the group ID in which the host group for your primary servers groups will be created. Create a host group under source_group_id called Primaries or Management Servers and use the group ID of this new host group for the second parameter.

See “Identifying a host group ID” on page 70.

When you organize by primary server, if a host group exists anywhere under the source group hierarchy with the name of the primary server, the routine associates the clients with that folder and does not create a new folder under the destination folder. This association occurs whether you explicitly specify the destination folder or if the destination is NULL. However, if you pass a source folder that is at a lower level, the routine only checks for a folder under that hierarchy. If you specify NULL as the destination, the routine will create (if it does not exist already) a group called “NetBackup” under the Source group ID. It then creates a host group called “Primary Servers” under the “NetBackup” group.

move_clients If set to **0**, the clients link into the respective host group and remain in their original host group location. If set to **1**, all the clients move from the source host group and into the respective host groups.

The utility processes and organizes all clients of the source group hierarchy into the target primary server grouping. However, if the **move_clients** flag is set to **1**, the utility removes only clients in the top level **source_group_id** group--and those already organized in lower level sub-groups remain.

latest_primary_only defaults to 0, but can be set to 1, indicating organization by the latest primary server. If a client is backed up by two primary servers, or if a client was backed up by primary server A in the past, but is now backed up by primary server B, setting this flag to true will result in the client being organized by the latest primary server.

Usage with 6 Parameters

```
execute nbu_adaptor_pkg.moveClientsIntoPrimaryGroups
( <source_group_id>,<destination_group_id>,
<cascade_source_group>, <move_clients>,
<latest_primary_only>, <exclude_policy_client>);
```

Example: **exec moveClientsIntoPrimaryGroups(300000, 300010, 1, 1, 1, 0);**

source_group_id is the internal group ID of the host group hierarchy to traverse.

destination_group_id is the group ID in which the new host group for your primary servers will be created. Create a host group under source_group_id called Primaries or Management Servers and use the group ID of this new host group for the second parameter.

See "Identifying a host group ID" on page 70.

When you organize by primary server, if a host group exists anywhere under the source group hierarchy with the name of the primary server, the routine associates the clients with that folder and does not create a new folder under the destination folder. This association occurs whether you explicitly specify the destination folder or if the destination is NULL. However, if you pass a source folder that is at a lower level, the routine only checks for a folder under that hierarchy. If you specify NULL as the destination, the routine will create (if it does not exist already) a host group called "NetBackup" under the Source group ID. It then creates a host group called "Primary Servers" under the "NetBackup" host group.

cascade_source_groupcan be set to **0=** Do not include sub-groups, **1=** Include sub-groups. Use cascade_source_group to find and re-sort all of the defined host groups that are under the source group. Use this parameter so that you do not have to move all of your clients to the top before re-sorting.

move_clients If set to **0**, the clients link into the respective host group and remain in their original host group location. If set to **1**, all the clients move from the source group and into the respective management server host groups.

The utility processes and organizes all clients of the source group hierarchy into the target primary server grouping. However, if the move_clients flag is set to 1, the utility removes only clients in the top-level source_group_id group--and those already organized in lower-level sub-groups remain.

latest_primary_only **defaults to 0**, but can be **set to 1**, indicating organization by the latest primary server. If a client is backed up by two primary servers, or if a client was backed up by primary server A in the past, but is now backed up by primary server B, setting this flag to true will result in the client being organized by the latest primary server.

exclude_policy_client **defaults to 0**, but can be **set to 1**, indicating that you want to organize the clients based on backups and exclude policy-based clients. If this flag is set to 0, the utility finds the clients that are backed up by the primary server and also clients that are in the policy that is controlled by the primary server.

Set up an inactive clients group

Description

This utility enables you to automatically generate a list of all clients that are not a member of any policy and move or link the clients into a user-definable host group.

Usage

```
execute server_mgmt_pkg.setupInactiveClientsGroup
('<host_group_to_traverse>','<inactive_clients_group>',
<move_or_copy_flag>);
```

Example: **exec setupInactiveClientsGroup('/Global/Corp',NULL,0);**

Where:

host_group_to_traverse is the full pathname to the host group hierarchy to traverse looking for inactive clients, for example **/Aptare/hostgroup1**.

inactive_clients_group is the full pathname to the host group into which the inactive clients will be moved or linked. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically creates a host group called **Clients Not In Policy** within **host_group_to_traverse**.

move_or_copy_flag can be set to **0**=Link (copy) clients or **1**=Move clients. If set to **0**, the utility links the clients to the **inactive_clients_group** and keeps the clients in their original host group location. If set to **1**, the utility moves all the inactive clients from their current host group location and consolidates them into the **inactive_clients_group**.

Set up a host group for clients in inactive policies

Description

This utility enables you to automatically generate a list of all clients that are members of a policy, but the policy is NOT active. These clients are then linked or copied into a user-definable host group.

Usage

```
execute server_mgmt_pkg.setupInactivePolicyClients
('<host_group_to_traverse>','<inactive_clients_group>',
<move_or_copy_flag>, <include_deleted_flag>);
```

Example: **exec setupInactivePolicyClients('/Global/Corp',NULL,1,0);**

Where:

host_group_to_traverse is the full pathname to the host group hierarchy to traverse looking for inactive policies, for example **/Aptare/hostgroup1**.

inactive_clients_group is the full pathname to the host group into which the clients in an inactive policy will be moved or linked. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically creates a host group called **Inactive Policy Clients** within **host_group_to_traverse**.

move_or_copy_flagcan be set to **0**=Link (copy) clients or **1**=Move clients. If set to **0**, the utility links the client to the **inactive_clients_group** and keeps the client in the original host group location. If set to **1**, the utility moves all the clients in inactive policies from their current host group location and consolidates them into the **inactive_clients_group**.

include_deleted_flag can be used in conjunction with the move_or_copy_flag can be set to **1** to include policies deleted from NetBackup while organizing inactive policy clients into the **inactive_clients_group**.

Set up clients by policy

Description

This utility enables you to automatically organize clients by the Cohesity NetBackup Policy (or Policies) to which they belong. The utility automatically creates host groups for each Cohesity NetBackup Policy and links clients that are members of these policies to the host group(s) accordingly.

Usage

```
execute nbu_adaptor_pkg.setupClientsByPolicy
('<source_host_group>','<destination_host_group>');
```

Example: **exec setupClientsByPolicy('/Global/Corp', NULL);**

Where:

source_host_group is the full pathname to the host group hierarchy to traverse for clients, for example **/Aptare/hostgroup1**. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically locates the highest level host group to traverse.

destination_host_group is the full pathname to the host group under which the new groups by policy name will be automatically created. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically creates a host group called **NetBackup Policies** within **source_server_group**.

If a client is removed from a Cohesity NetBackup policy, added to a new policy and the utility is subsequently run again, the client will appear in the new policy group but will not be deleted from the old policy group. To remove the client from the old policy group and completely re-synchronize the grouping structure, simply delete the Policy grouping hierarchy via the **deleteEntireGroupContents** utility, referenced in

See "Delete host group" on page 73.

, and then run the **setupClientsByPolicy** utility again.

Set up clients by policy type

Description

This utility enables you to automatically organize clients by the type of the Cohesity NetBackup policy to which they belong; for example, Standard, NDMP, ORACLE. The utility automatically creates host groups for each Cohesity NetBackup policy type and links clients that are members of these policy types to the host group(s) accordingly.

Usage

```
execute nbu_adaptor_pkg.setupClientsByPolicyType
('<source_host_group>','<destination_host_group>');
```

Example: **exec setupClientsByPolicyType('/Global/Corp', NULL);**

Where:

source_host_group is the full pathname to the server group hierarchy to traverse for clients, for example **/Aptare/hostgroup1**. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically locates the highest level host group to traverse.

destination_server_group is the full pathname to the server group under which the new groups by Policy type will be automatically created. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL** , the utility automatically creates a host group called **NetBackup Policy Types** within **source_host_group**.

If a client is removed from one Cohesity NetBackup policy type, added to a new policy type and the utility is subsequently run again, the client will appear in the new policy type host group but will not be deleted from the old policy group. To remove the client from the old policy group and completely re-synchronize the grouping structure, simply delete the Policy Type grouping hierarchy via the **deleteEntireGroupContents** utility, referenced in

See "Delete host group" on page 73.

, and then run the **setupClientsByPolicy** utility again.

IBM Tivoli storage manager utilities

The utilities contained in this section apply only to IBM Tivoli Storage Manager environments and clients that have been backed up by IBM Tivoli Storage Manager.

- See "Set up clients by policy domain" on page 97.
- See "Set up clients by IBM Tivoli storage manager instance" on page 98.

Set up clients by policy domain

Description

This utility enables you to automatically organize clients by the IBM Tivoli Storage Manager policy domain(s) to which they belong. The utility automatically creates host groups for each policy domain and links clients that are members of these Policy domain(s) to the host group(s) accordingly.

Usage

```
execute tsm_common_pkg.TSMsetupClientsByPolicyDomain
('<source_host_group>', '<destination_host_group>');
```

Example: **exec tsm_common_pkg.TSMsetupClientsByPolicyDomain('/Global', NULL);**

Where:

source_host_group is the full pathname to the server group hierarchy to traverse for clients, for example **/Aptare/hostgroup1**. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically locates the highest level host group to traverse.

destination_host_group is the full pathname to the host group under which the new groups by policy domain name will be automatically created. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL** the utility automatically creates a host group called **Policy Domains** within **source_host_group** .

If a client is removed from a IBM Tivoli Storage Manager policy domain, added to a new policy domain and the utility is subsequently run again, the client will appear in the new policy group but will not be deleted from the old policy group. To remove the client from the old policy group and completely re-synchronize the grouping structure, simply delete the Policy domain grouping hierarchy via the **deleteEntireGroupContents** utility, referenced in

See “Delete host group” on page 73.

, and then run the **TSMsetupClientsByPolicyDomain** utility again.

Set up clients by IBM Tivoli storage manager instance

Description

This utility enables you to automatically organize clients by the IBM Tivoli Storage Manager instance to which they belong. The utility automatically create host groups for each instance, then links clients that are members of the instance into the host group(s) accordingly.

Usage

```
execute tsm_common_pkg.TSMsetupClientsByInstance
('<source_host_group>', '<destination_host_group>',
<move_or_copy_flag>);
```

Example: **exec tsm_common_pkg.TSMsetupClientsByInstance('/Global', NULL, 1);**

Where

source_host_group is the full pathname to the host group hierarchy to traverse for clients, for example **/Aptare/hostgroup1**. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL**, the utility automatically locates the highest level host group to traverse.

destination_host_group is the full pathname to the host group under which new groups by instance name will be automatically created. The default value for this parameter is **NULL** (which should not be within quotes). If set to **NULL** the utility automatically creates a host group called IBM Tivoli Storage Manager instances within **source_host_group**.

move_or_copy_flag can be set to **0**=Link (copy) clients or **1**=Move clients. If set to **0**, the utility links the clients to their respective host groups and keeps the clients in their original host group location. If set to **1**, the utility moves all clients from the source host group and to their respective host groups.

Scheduling utilities to run automatically

The utilities can be run on a one-time basis, or scheduled to run every day to automatically keep your hosts and host groups up to date. Scheduling can be accomplished by creating an Oracle job. IT Analytics already makes use of Oracle jobs to run many background tasks such as purging old data and rebuilding indices.

The sample SQL file in Example-Scheduling Utilities to Run Automatically sets up an Oracle job to run every day at 5:00 a.m. and call the **moveOrCopyClients** utility to move clients from one folder to another. This example can be used as a template for other automatic jobs you need to set up. Simply customize the text in bold for your particular requirements.

To see the Oracle jobs that are automatically configured as part of a new install, review the following files:

- `<database_home>/stored_procedures/setup_base_jobs.plb`
- `<database_home>/stored_procedures/nbu/setup_nbu_jobs.sql`
- `<database_home>/stored_procedures/nbu/setup_nbu_jobs.sql`
- `<database_home>/stored_procedures/tsm/setup_leg_jobs.plb`

Where database_home is **/opt/aptare/database** for Linux servers and `C:\opt\oracle\database` for Windows servers.

Example-Scheduling Utilities to Run Automatically

In order to execute the following sample SQL file, as user **aptare** on a Linux system or on a Windows system as an Administrator who is a member of the **ORA_DBA** group execute the following:

```
sqlplus portal/<portal_password>@//localhost:1521/scdb@setup_ora_job.sql
```

The following example uses two methods:

moveOracleClients	The name you want to assign to this job you are defining.
MoveOrCopyClients	The utility you are calling along with the parameters you are passing.

Note: The parameters passed to the **moveOrCopyClients** method which must be quoted actually have two single quotes. The two single quotes is the standard Oracle syntax to incorporate a literal quote within an already quoted string.

Sample .sql file (setup_ora_job.sql) to set up an automatic job

```
DECLARE
   jobName   user_scheduler_jobs.job_name%TYPE := NULL;
BEGIN
   ----------------------------------------------------------------------
   -- Move new clients whose server name ends with 'ORA' into the 'database' host group

   -- Frequency: Every day at 5am (Portal Time)
   ----------------------------------------------------------------------
   jobName := dba_package.getSchedulerJobName('moveOracleClients');

    IF (jobName IS NOT NULL AND LOWER(jobName) <> LOWER('moveOracleClients')) THEN
        DBMS_OUTPUT.PUT_LINE('setupInactivePolicyClients exists with default name
'|| jobName ||' hence will be removed and recreated.');
        DBMS_SCHEDULER.DROP_JOB(job_name => jobName);
        jobName := NULL;
    END IF;

    IF jobName IS NULL THEN
        DBMS_SCHEDULER.CREATE_JOB(
            job_name         => 'moveOracleClients',
            job_type         => 'PLSQL_BLOCK',
```

```
            job_action       => 'server_mgmt_pkg.moveOrCopyClients(''/Aptare'',
''/Aptare/database'',''*ORA'', 1);',  -- What to run
            start_date       => SYSDATE + (5/24),
-- First run is 5am server time
            repeat_interval  => 'TRUNC(SYSDATE+1,''DD'') + (5/24)',
 -- Next run is 5am each subsequent day
            enabled          => TRUE);
    ELSE
       DBMS_OUTPUT.PUT_LINE('setupInactivePolicyClients exists and will be
altered with updated version.');
       DBMS_SCHEDULER.SET_ATTRIBUTE(
            name         => jobName,
            attribute    => 'job_type',
            value        => 'PLSQL_BLOCK'
            );
       DBMS_SCHEDULER.SET_ATTRIBUTE(
            name         => jobName,
            attribute    => 'job_action',
            value        =>  'server_mgmt_pkg.moveOrCopyClients(''/Aptare'',
''/Aptare/database'',''*ORA'', 1);'
            );
       DBMS_SCHEDULER.SET_ATTRIBUTE(
            name         => jobName,
            attribute    => 'repeat_interval',
            value        => 'TRUNC(SYSDATE+1,''DD'') + (5/24)'
            );
    END IF;
    DBMS_OUTPUT.put_line('moveOracleClients job set to run at 5am every day');
END;
/
quit;
```

Host matching identification for single-domain multi-customer environments

Host matching identification using a Host Matching Identifier is intended for deployments having a single IT Analytics domain that handles multiple hosts. Each host within the domain is agnostic of other hosts and is an organization by itself. It has its distinct set of users or user groups. If two or more hosts have identical names, you can assign a host matching identifier to the Data Collector or the policy that collects data from each host to differentiate the data collected from each host. Host Matching Identifier assigned to a Data Collector is inherited only by those

policies added to the collector after the identifier was assigned. For existing policies within the collector, you must manually assign the Host Matching Identifier. Once a Host Matching Identifier is assigned to a policy that collects host data, it becomes an integral part of the host matching process. Changing the Host Matching Identifier after the data collection results in creation of new hosts with the new host matching identifier.

Note: When the same physical host is collected by more than one data-collector policy, for example, a NetBackup policy and a VMware policy collecting the same VM, and only one of those policies is configured with a host matching identifier (HMI) while the other is not, the Portal may create duplicate host records as the alias/IP addresses may also match which are used in host lookup.

Given the expertise required to use this feature and its potential to create duplicate hosts if not used responsibly, the feature is disabled by default. Also, this feature access is restricted to Super User or Administrator, who can grant access to specific users or user groups on the Portal.

Enable host matching identification

You can enable host matching identification with the following steps:

Step-1: Enable Portal custom parameter
`portal.hostMatchingIdentifier.enabled`:

1. On the IT Analytics Portal, go to **Admin** tab > **Advanced** > **System Configuration** > **Custom Parameters**.
2. Click **Add**.
3. Enter **Custom Parameter Name** as **portal.hostMatchingIdentifier.enabled** and **Custom Parameter Value** as **True**.
4. Click **Save** and also click **Save and Apply**.

   This introduces the **Host Matching Identifier** column and **Edit Host Matching Identifier** button on the **Collector Administration** view. You may need to logout and login again on the Portal to view the changes.

Step-2: Add Host Matching Identifier to the Data Collector or the collector policy

1. On the IT Analytics Portal, go to **Admin** tab > **Collector Administration**.
2. Select a Data Collector and click **Edit Host Matching Identifier**.
3. Enter **Host Matching Identifier** label of your choice and click **OK**.

   The label you assign appears under the **Host Matching Identifier** column on the **Collector Administration** view.

   You can assign an identifier to a policy by repeating the above steps only after you have assigned it to its Data Collector. While assigning to a policy, you need to select the Host Matching Identifier from a drop-down list.

Chapter 8

Attribute management

This chapter includes the following topics:

Attribute bulk load utilities

The Attributes feature provides the ability to define a scope for a report based on specific characteristics, such as operating system or host criticality. In effect, when you generate a report, you can select a set of objects that share the same attribute or characteristic.

Often, large enterprise environments need to configure attributes for many objects, such as hosts, arrays, and switches. Bulk load utilities assign attributes to objects. While the Portal has capabilities for assigning attributes to certain objects, the utilities described in this section fulfill the large-scale requirement for assigning attributes to a large number of objects.

To facilitate bulk loading and configuration of attributes, several utilities are provided. These utilities load attributes and values into the IT Analytics database from comma-separated-values (CSV) files.

Note: Currently, there are no utilities available for the bulk load of Datastore attributes.

For Managed Service Providers (MSPs)

Attributes are IT Analytics domain-specific; therefore, you can configure different object attributes for different clients without impacting the reports and environments of your other clients.

See "Load host attributes and values" on page 107.

See "Load attributes and values and assign to hosts" on page 108.

Attribute naming rules

Adhere to the following rules when creating attribute names. Attributes are validated against these rules so that there are no conflicts in the database, such as duplicates or the use of Oracle reserved words.

- Limit the length to 30 characters.
- Begin the name with an alphabetic character.
- Use only ASCII, alpha, numeric, or underscore characters in the name. Spaces and special characters other than underscores are not allowed in attribute names, although they are allowed in the list of values (LOV) for an attribute.
- Names are not case-sensitive.
- Do not use Oracle reserved words. See http://docs.oracle.com/cd/E15817_01/appdev.111/b31231/appb.htm. To list the Oracle reserved words, use this SQLPlus query at the command line:

  ```
  SQL> SELECT * from v$reserved_words;
  ```

- Attribute names within a domain hierarchy must be unique.

Rename attributes before upgrading

Beginning with Release Version 10, all attributes are multi-object attributes--that is, a single attribute is defined and that attribute, with its values, can be used for multiple object types. For example, prior to Release Version 10, you could have a Location attribute for a host and a separate Location attribute for an array. During the upgrade, a system attribute named Location is added to the database and this single attribute can be used for multiple object types--in this example, it would be used for both hosts and arrays.

During the Portal upgrade, if the names of existing attributes match the name of a system attribute introduced with the upgrade, you may want to rename existing attributes so that their values do not get merged into a single attribute. Renaming of attributes, before the upgrade is completed, must be performed using SQL at the command line.

Note: After you rename an attribute, any report templates that used these attributes must be updated via the Portal SQL Template Designer.

To rename existing attributes so that their values do not get merged into a single attribute, take the following steps.

1. Log in to the Portal server.
2. At the command line:

   ```
   su - aptare
   ```

3. At the command line, launch sqlplus:

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb>
   ```

   - Example: `sqlplus portal/portal@//localhost:1521/scdb`

4. At the command line, execute the following at the **SQL** prompt:

   ```
   UPDATE apt_attribute
   SET attribute_name - '<NewAttributeName>'
   WHERE attribute_id = <ExistingAttributeID>;
   Commit;
   ```

   Where <NewAttributeName> is the new name you supply and <ExistingAttributeID> is the ID listed during the Portal upgrade process.

Load host attributes and values

Function: This utility provides an efficient method for creating multiple attributes, along with a list of possible values for each attribute. Note that the result of this process is simply an inventory of attributes with an associated list of values (LOV). These attributes need to be applied to hosts. This can be accomplished via the Portal or by using other attribute load utilities.

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

Create a CSV File

In preparation for loading host attributes, enter the information into a spreadsheet from which you will create a comma-separated values file.

The table in the spreadsheet should be in the following format:

- first column - list the attribute name
- subsequent columns - list the possible values for the attribute

Table 8-1 Table 1: Example of a Table of Attributes and Values

Attribute Name	Attribute Value	Attribute Value
Operating System	Windows	Win
Criticality	Mission Critical	Low

Execute the Load Utility

To load attributes and values for a domain

1. Create a table in a spreadsheet, as shown in the above example. Save the table as a comma-separated file in a temporary directory.
   - Windows example: `C:\temp\attributes.csv`
   - Linux example: `/tmp/attributes.csv`
2. Log in to the Portal server.
3. At the command line:

   ```
   su - aptare
   ```

4 At the command line, launch sqlplus:

```
sqlplus <pwd>/<pwd>@//localhost:1521/scdb
```

- Example: `sqlplus portal/portal@//localhost:1521/scdb`

5 At the command line, execute the following at the **SQL** prompt:

```
SQL> Execute
load_package.loadAttributeFile('pathname_and_filename',
'domain_name');
```

where:

'pathname_and_filename'	full path + filename (enclosed in single straight quotes) of the CSV file that you created
'domain_name'	name of the domain in which the hosts reside (enclosed in single straight quotes).

Example:

```
Execute load_package.loadAttributeFile('c:\temp\attributes.csv',
 'APTARE');
```

6 Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Finding the Domain Name

To identify the domain, in the Portal: **Admin > Domains > Domains**

Load attributes and values and assign to hosts

Function: This utility provides an efficient method of assigning attributes to a large number of hosts.

To create an inventory of a large number of host attributes and associated values using a bulk load utility refer to the following:

See “Rename attributes before upgrading” on page 106.

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

Take the following steps to load attributes and values and assign those attributes to hosts:

1. Create a CSV File of Hosts, Attributes, and Values
2. Execute the Load Host Attribute Utility
3. Verify the Host Attributes Load
4. Create a report template using a Report Template Designer.

   Once attribute values are assigned to hosts, a report can query the database to report on hosts, filtered by the attributes that you've created to categorize them.

Create a CSV File of Hosts, Attributes, and Values

In preparation for loading host attributes, you will need to enter information into a spreadsheet from which you will create a comma-separated values (CSV) file. The table in the spreadsheet should be in the following format:

Columns

- One column lists the hosts, which must already exist in the IT Analytics database.
- Each additional column lists attributes and values that will be applied to the host.

Rows

- First (Header) Row - Enter the object type--in this case, Host Name--followed by attribute names. Note that any column may be used for the list of host names. When you run the utility, you'll indicate which column contains the host names. The header row is information only and is not processed as a data row.
- Subsequent rows list host names, followed by the attribute values that you are assigning to each host.

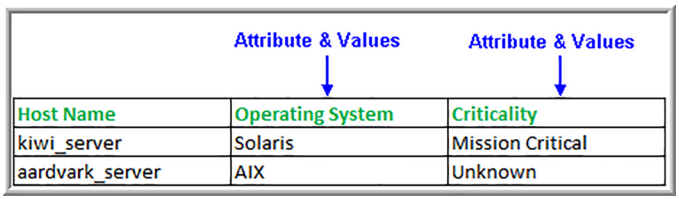

Rules for Attributes and Values in the CSV File

- The hosts listed in the CSV must already exist in the Portal database.
- A host should be listed only once. If a host name is listed in multiple rows, only the attributes from the last row with the same host name will be saved.
- The Host Name cannot begin with #. Any line that begins with # will be ignored.
- Follow the Attribute Naming Rules for the attribute name.
- The maximum line size--that is, characters in a row--is 8192 characters.
- Every column must have a value for each row. Those columns that do not have any actual values can be filled with **N/A** or **Unknown** or a period (**.**). Attribute values of “N/A”, “Unknown”, “.” will be ignored.
- A single attribute **value** cannot include commas. For example: LastName, FirstNameUse spaces instead of commas to separate words. For example: LastName FirstName

Execute the Load Host Attribute Utility

Note: This utility can be used to load new data as well as to update previously loaded data. To revise existing data, simply run the utility with an updated CSV file.

To assign attributes to hosts

1. Create a table in a spreadsheet.
2. Save the table as a comma-separated file (for example, HostAttributes.csv).
3. Log in to the Portal server.
4. At the command line:

   ```
   su - aptare
   ```

5. At the command line, launch sqlplus:

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb
   ```

   - Example: `sqlplus portal/portal@//localhost:1521/scdb`

6 Execute the following at the **SQL** prompt:

```
SQL> Execute load_package.loadServerAttributeFile
('pathname_and_filename','domain_name',
host_name_column_num,'log_path_name',
'log_file_name','check_valid_value');
```

Where:

'pathname_and_filename'	Full path + filename (enclosed in single straight quotes) of the CSV file that you created.
'domain_name'	Name (enclosed in single straight quotes) of the IT Analytics domain in which the host groups and hosts reside.
host_name_column_num	Column number in the csv file where the host names are listed. These hosts must already exist in the IT Analytics database. Typically, this would be column 1.
'log_path_name'	Full path (enclosed in single straight quotes) where the log file will be created/updated. Verify that you have write access to this directory. Example: 'C:\tmp' Optional: If you do not specify a path and log file name, only error messages will be written to the scon.err file. To omit this parameter, enter: ''
'log_file_name'	Filename of the log where execution status and errors messages are written. Example: 'HostAttributeLoad.log' Optional: If you do not specify a path and log file name, only error messages will be written to the scon.err file. To omit this parameter, enter: ''
'check_valid_value'	'**Y**' or '**N**'Indicates if you want the utility to check if the values provided in this file are among the existing possible values for the attributes. Y or N must be enclosed in single straight quotes.

Example:

```
Execute load_package.loadServerAttributeFile
('C:\myfiles\HostAttributes.csv','QA_Portal',
1,'C:\tmp','HostAttributeLoad.log','Y');
```

7 Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Verify the Host Attributes Load

To verify that the attribute load took effect:

1. In the Portal, go to **Reports**.

2. Select a blue user folder.
3. Select **New SQL Template**.
4. With the SQL Template Designer open, click the **Query** tab.
5. Enter the following query in the SQL Template Designer to verify Host attributes:

```
select * from apt_v_server_attribute
```

Load array attributes and values and assign to arrays

Function: The Load Array Attributes utility provides an efficient method of assigning attributes to a large number of arrays.

Take the following steps to load array attributes and values:

1. Create a CSV File of Arrays, Attributes and Values.
2. Execute the Load Array Attribute Utility.
3. Verify the Array Attributes Load.
4. Create a report template using a Report Template Designer.

Once attribute values are assigned to hosts, a Report Template Designer report can query the database to report on arrays, filtered by the attributes that you've created to categorize them.

Create a CSV File of Arrays, Attributes, and Values

In preparation for loading array attributes, you will need to enter information into a spreadsheet from which you will create a comma-separated values (CSV) file. The table in the spreadsheet should be in the following format:

Columns

- One column lists the arrays, which must already exist in the IT Analytics database.
- Each additional column lists attributes and values.

Rows

- First (Header) Row - Enter the object type--in this case, Array Name--followed by attribute names. Note that any column may be used for the list of array names. When you run the utility, you'll indicate which column contains the array names. The header row is information only and is not processed as a data row.

- Subsequent rows list arrays, followed by the attribute values that you are assigning to each array.

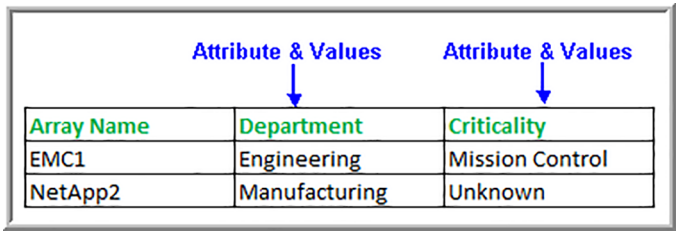

Execute the Load Array Attribute Utility

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

To assign attributes to arrays

Note: This utility only assigns attributes to active arrays. If an array exists in the system, but it is inactive, the log will indicate that no attribute was assigned.

---

1. Create a table in a spreadsheet.
2. Save the table as a comma-separated file (for example, ArrayAttributes.csv).
3. Log in to the Portal server.
4. At the command line:

   ```
   su - aptare
   ```

5. At the command line, launch sqlplus:

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb
   ```

   - Example: `sqlplus portal/portal@//localhost:1521/scdb`

6 Execute the following at the **SQL** prompt:

```
SQL> Execute load_package.loadArrayAttributeFile
('pathname_and_filename','domain_name',
array_name_column_num,'log_path_name',
'log_file_name','check_valid_value');
```

Where:

'pathname_and_filename'	Full path + filename (enclosed in single straight quotes) of the CSV file you created.
'domain_name'	Name (enclosed in single straight quotes) of the IT Analytics domain in which the arrays reside.
array_name_column_num	Column number in the csv file where the array names are listed. These arrays must already exist in the IT Analytics database. Typically, this would be column 1.
'log_path_name'	Full path (enclosed in single straight quotes) where the log file will be created/updated. Verify that you have write access to this directory. Example: 'C:\tmp' Optional: If you do not specify a path and log file name, only error messages will be written to the scon.err file. To omit this parameter, enter: ''
'log_file_name'	Filename of the log where execution status and errors messages are written. Example: 'ArrayAttributeLoad.log' Optional: If you do not specify a path and log file name, only error messages will be written to the scon.err file. To omit this parameter, enter: ''
'check_valid_value'	'**Y**' or '**N**' Indicates if you want the utility to check if the values provided in this file are among the existing possible values for the attributes. Y or N must be enclosed in single straight quotes.

Example:

```
Execute load_package.loadArrayAttributeFile
('C:\myfiles\ArrayAttributes.csv','QA_Portal',
1,'C:\tmp','ArrayAttributeLoad.log','Y');
```

7 Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Verify the Array Attributes Load

To verify that the attribute load took effect:

1. In the Portal, go to **Reports**.

2. Select a blue user folder.
3. Select **New SQL Template.**
4. With the SQL Template Designer open, click the **Query** tab.
5. Enter the following query in the SQL Template Designer to verify Array attributes:

```
select * from aps_v_storage_array_attribute
```

Overview of application attributes and values

Applications consume storage on a host and therefore it is often necessary to account for who is using that storage. Since several applications may reside on a host, additional criteria can be used to classify a specific application.

Application attributes provide the mechanism for identifying a host's applications--an Application Database or MS Exchange Organization--enabling reports that:

- Identify the group within your organization that is using an application, to be able to determine accountability.
  Example: Which applications are being used by the Engineering department so that I can charge them for capacity usage?
- Identify applications by type/function.
  Example: Which applications are used for production and which are development applications?

Note: Currently, application attributes can be used only in reports created with the SQL Template Designer.

Load application database attributes and values

Function: The Load Application Database Attributes utility provides an efficient method of assigning attributes to a large number of application databases.

Take the following steps to load application database attributes and values:

1. Create a CSV File of Application Database Objects and Attributes.
2. Execute the Load Application Database Attribute Utility.
3. Verify the Application Database Attributes Load.
4. Create a report template using the SQL Template Designer.

Once attribute values are assigned to application databases, a SQL Template Designer report can query the database to report on the application databases.

Create a CSV File of Application Database Objects and Attributes

The LoadDBAttribute utility assigns application attribute values to a host's database application. This utility takes as input a comma-separated values (CSV) file.

Note: This CSV file becomes the primary document of record for Application Database Attributes and therefore should be preserved in a working directory for future updates.

---

1. Create a spreadsheet table, in the format shown in the following example, and save it as a CSV file in a working directory. This file is specific to application databases.

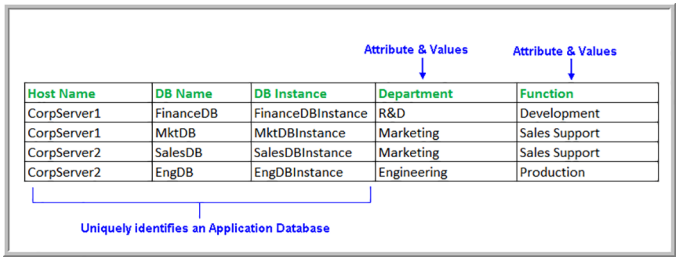

In the above example:

- The first 3 columns comprise the unique identifier for an Application Database--in this example, CorpServer1, FinanceDB, and FinanceDBInstance.
- Subsequent columns list the attributes and values.

Columns

- Columns list the objects that uniquely identify an application. For an Application Database, the required columns are: Host Name, DB Name, DB Instance.
- Each additional column lists attributes and values.

Rows

- First (Header) Row - Contains the fields that uniquely identify an application, followed by the attribute names. The header row is information only and is not processed as a data row.

- Subsequent rows list the objects that uniquely identify an application database, followed by the attribute values that you are assigning to each application database.

Execute the Load Application Database Attribute Utility

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

To assign attributes to application databases

1. Create a table in a spreadsheet.
2. Save the table as a comma-separated file (for example, DBAttributes.csv).
3. Log in to the Portal server.
4. At the command line:

   ```
   su - aptare
   ```

5. At the command line, launch sqlplus:

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb
   ```

   - Example: `sqlplus portal/portal@//localhost:1521/scdb`
   - Execute the following at the **SQL** prompt:

     ```
     SQL> Execute load_package.loadDBAttributeFile
     ('pathname_and_filename','domain_name',
     db_name_column_num,db_instance_column_num,host_name_column_num,
     'log_path_name','log_file_name','check_valid_value');
     ```

   Where:

'pathname_and_filename'	Full path + filename (enclosed in single straight quotes) of the CSV file Windows Example: '`c:\config\DBAttributes.csv`' Linux Example: '`/config/DBAttributes.csv`'
'domain_name'	Name (enclosed in single straight quotes) of the domain in which the host groups and hosts reside; Example: '**DomainEMEA**'

db_name_column_num	Column number in the csv file where the DB Name is listed; Example: **2**
db_instance_column_num	Column number in the csv file where the DB Instance is listed; Example: **3**
host_name_column_num	Column number in the csv file where the Host Name is listed; Example: **1**
'log_path_name'	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to scon.log and scon.err. To omit this parameter, enter: '' Example: '`c:\configs`'
'log_file_name'	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to scon.log and scon.err. To omit this parameter, enter: '' Example: '**DBAttributes.log**'
'check_valid_value'	'**Y**' or '**N**' enclosed in single straight quotes. **Y** - Checks if the attribute value exists. If the utility determines that the attribute value is not valid, it skips this row and does not assign the attribute value to the application database. **N** - Updates without checking that the attribute value exists. This option is seldom chosen, but it is available for certain customer environments where attributes may have been created without values (with scripts that bypass the user interface). Example: `SQL> Execute load_package.loadDBAttributeFile('/config/DBAttributes.csv', 'DomainEMEA', 2, 3, 1,'/config/logs','DBAttributes.log','Y');`

6. Check the log file for status and errors.
7. Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Verify the Application Database Attributes Load

To verify that the attribute load took effect:

1. In the Portal, go to **Reports**.
2. Select a blue user folder.
3. Select **New SQL Template**.
4. With the SQL Template Designer open, click the **Query** tab.

5. Enter the following query in the SQL Template Designer to verify Application Database attributes:

```
select * from aps_v_database_attribute
```

Load MS Exchange organization attributes and values

The Load MS Exchange Organization Attributes utility provides an efficient method of assigning attributes to a large number of Exchange Organizations.

Take the following steps to load Exchange Organization attributes and values:

1. Create a CSV File of Exchange Organization Objects and Attributes.
2. Load MS Exchange Organization Attributes and Values.
3. Verify the MS Exchange Organization Attributes Load.
4. Create a report template using the SQL Template Designer.

Once attribute values are assigned to MS Exchange Organizations, a Report Template Designer report can query the database to report on metrics such as capacity usage for chargebacks.

Create a CSV File of Exchange Organization Objects and Attributes

The Load Exchange Organization Attribute utility assigns application attribute values to a host's Microsoft Exchange Organizations. This utility takes as input a comma-separated values (CSV) file.

Note: This CSV file becomes the primary document of record for MS Exchange Organization Attributes and therefore should be preserved in a working directory for future updates.

---

1. Create a spreadsheet table, in the format shown in the following example, and save it as a CSV file in a working directory. This file is specific to MS Exchange Organizations.

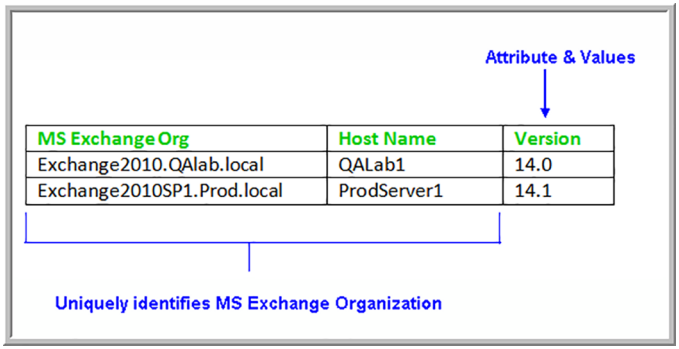

In the above example:

- The first 2 columns comprise the unique identifier for a Microsoft Exchange Organization--in this example, Exchange2010.QAlab.local, QALab1.
- Subsequent columns list the attributes and values--in this example, Version.

Columns

- Columns list the objects that uniquely identify an application. For MS Exchange, the required columns are: MS Exchange Organization and Host Name.

Rows

- First (Header) Row - Names the fields that uniquely identify an application, followed by the attribute names.
- Subsequent rows list the objects that uniquely identify an MS Exchange Organization--in this case, MS Exchange Organization and Host Name--followed by the attribute values that you are assigning to each MS Exchange Organization.

Execute the Load MS Exchange Organization Attribute Utility

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

To assign attributes to Microsoft Exchange Organizations

1. Log in to the Portal server.
2. At the command line:

   ```
   su - aptare
   ```

3. At the command line, launch sqlplus:

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb
   ```

   - Example: `sqlplus portal/portal@//localhost:1521/scdb`

4 Execute the following at the **SQL** prompt:

```
SQL> Execute
load_package.loadExchOrgAttributeFile('pathname_and_filename',
'domain_name',exchange_org_column_num,host_name_column_num,
'log_path_name','log_file_name','check_valid_value');
```

Where:

'pathname_and_filename'	Full path + filename (enclosed in single straight quotes) of the CSV file Windows Example: '`c:\config\MSExchangeAttributes.csv`' Linux Example: '`/config/MSExchangeAttributes.csv`'
'domain_name'	Name (enclosed in single straight quotes) of the IT Analytics Domain in which the host groups and hosts reside; Example: '**DomainEMEA**'
exchange_org_column_num	Column number in the csv file where the MS Exchange Organization is listed; Example: **1**
host_name_column_num	Column number in the csv file where the Host Name is listed; Example: **2**
'log_path_name'	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to scon.log and scon.err. Example: '`c:\configs`'
'log_file_name'	Name of the log file enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to scon.log and scon.err. Example: '**MSExchangeAttributes.log**'
'check_valid_value'	'**Y**' or '**N**'enclosed in single straight quotes. Y - Checks if the attribute value exists. If the utility determines that the attribute value is not valid, it skips this row and does not assign the attribute value to the Exchange Organization. **N** - Updates without checking that the attribute value exists. This option is seldom chosen, but is available for certain customer environments where attributes may have been created without values (with scripts that bypass the user interface).

Example:

```
SQL> Execute
load_package.loadExchOrgAttributeFile('/config/MSExchangeAttributes.csv',
 'DomainEMEA',1,2,'/config/logs','MSExchangeAttributes.log','Y');
```

5. Check the log file for status and errors.
6. Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Verify the MS Exchange Organization Attributes Load

To verify that the attribute load took effect:

1. In the Portal, go to **Reports**.
2. **Select a blue user folder.**
3. **Select New SQL Template.**
4. With the SQL Template Designer open, click the Query tab.
5. Enter use the following query in the SQL Template Designer to verify the MS Exchange Organization attributes:
6. **select * from aps_v_exch_org_attribute**

Load LUN attributes and values

Function: The Load LUN Attributes utility provides an efficient method of assigning attributes to a large number of storage array logical units (LUNs).

To load switch attributes and values

1. Create a CSV File of LUN Objects and Attributes.
2. Execute the Load LUN Attribute Utility.
3. Verify the LUN Attributes Load.
4. Create a report template using the SQL Template Designer.
5. Once attribute values are assigned to application databases, a SQL Template Designer report can query the database to report on the application databases.

Create a CSV File of LUN Objects and Attributes

The loadLunAttributeFile utility assigns attribute values to a list of LUNs. This utility takes as input a comma-separated values (CSV) file.

Note: This CSV file becomes the primary document of record for LUN Attributes and therefore should be preserved in a working directory for future updates.

---

1. Create a spreadsheet table, in the format shown in the following example, and save it as a CSV file in a working directory. This file is specific to loading LUN attributes.

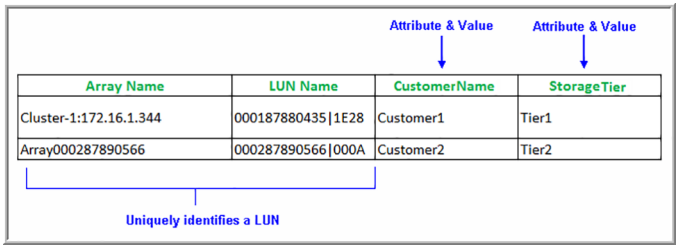

Columns

- The first column lists the Array Name.
- The second column lists the LUN Name.
- Each additional column lists attributes and values that will be applied to the LUN. Multiple attributes can be assigned to a single LUN object.

Rows

- First (Header) Row - Contains the fields that uniquely identify the LUN (array and LUN names), followed by Attribute names. The header row is information only and is not processed as a data row.
- Subsequent rows list the Array name and LUN name, followed by the attribute values that you are assigning to each LUN.

Execute the Load LUN Attribute Utility

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

To assign attributes to application databases

1 Create a table in a spreadsheet, as shown in

2 Save the table as a comma-separated file (for example, SwitchAttributes.csv).

3 Log in to the Portal server.

4 At the command line:

```
su - aptare
```

5 At the command line, launch sqlplus:

```
sqlplus <pwd>/<pwd>@//localhost:1521/scdb
```

- Example: `sqlplus portal/portal@//localhost:1521/scdb`

6 Execute the following at the **SQL** prompt:

```
SQL> Execute
load_package.loadLunAttributeFile('pathname_and_filename',
'domain_name',array_name_column_num, lun_name_column_num
,'log_path_name','log_file_name','check_valid_value');
```

Where:

'pathname_and_filename'	Full path + filename (enclosed in single straight quotes) of the CSV file Windows Example: '`c:\config\SwitchAttributes.csv`' Linux Example: '`/config/SwitchAttributes.csv`'
'domain_name'	Name (enclosed in single straight quotes) of the domain in which the host groups and hosts reside; Example: '**DomainEMEA**'
array_name_column_num	Column number in the csv file where the Array Name is listed; Example: **1** Note that the Array Name and the LUN Name can be either column 1 or 2 of the CSV. This parameter tells the utility in which column the Array Name will be found.
lun_name_column_num	Column number in the csv file where the LUN Name is listed; Example: **2**
'log_path_name'	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to scon.log and scon.err. To omit this parameter, enter: '' Example: '`c:\config`'
'log_file_name'	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to scon.log and scon.err. To omit this parameter, enter: '' Example: '**SwitchAttributes.log**'
'check_valid_value'	'**Y**' or '**N**'enclosed in single straight quotes. **Y** - Checks if the attribute value exists. If the utility determines that the attribute value is not valid, it skips this row and does not assign the attribute value to the switch object. **N** - Updates without checking that the attribute value exists. This option is seldom chosen, but it is available for certain customer environments where attributes may have been created without values (with scripts that bypass the user interface).

Example:

```
SQL> Execute
load_package.loadLunAttributeFile('/config/LUNAttributes.csv',
'DomainEMEA', 1, 2,'/config/logs','LUNAttributes.log','Y');
```

7 Check the log file for status and errors.

8 Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Verify the LUN Attributes Load

To verify that the attribute load was successful:

1. In the Portal, go to **Reports**.
2. Select a blue user folder.
3. Select **New SQL Template**.
4. With the SQL Template Designer open, click the **Query** tab.
5. Enter the following query in the SQL Template Designer to verify Switch attributes:

```
select * from aps_v_logical_unit_attribute
```

Load switch attributes and values

Function: The Load Switch Attributes utility provides an efficient method of assigning attributes to a large number of switches. Please note Fabric Manager must be installed or the loading will fail.

To load switch attributes and values

1 Create a CSV File of Switches, Attributes and Values.

2 Execute the Load Switch Attribute Utility.

3 Verify the Switch Attributes Load.

4 Create a report template using the SQL Template Designer.

Once attribute values are assigned to application databases, a SQL Template Designer report can query the database to report on the application databases.

Create a CSV File of Switches, Attributes, and Values

The loadSwitchAttributeFile utility assigns attribute values to a list of switches. This utility takes as input a comma-separated values (CSV) file.

Note: This CSV file becomes the primary document of record for Switch Attributes and therefore should be preserved in a working directory for future updates.

1. Create a spreadsheet table, in the format shown in the following example, and save it as a CSV file in a working directory. This file is specific to loading switch attributes.

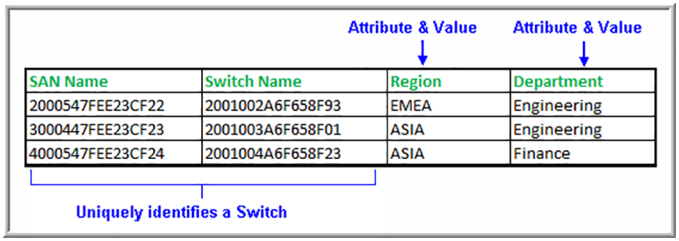

Columns

- The first column lists the SAN Name.
- The second column lists the Switch Name.
- Each additional column lists attributes and values that will be applied to the switch. Multiple attributes can be assigned to a single switch object.

Rows

- First (Header) Row - Contains the fields that uniquely identify the SAN and Switch names, followed by Attribute names. The header row is information only and is not processed as a data row.
- Subsequent rows list the SAN Name and Switch Name, followed by the attribute values that you are assigning to each switch.

Execute the Load Switch Attribute Utility

Before You Begin

Bulk Load utilities must be run in SQLPLUS as user APTARE. The load_package utility is located in:

```
/opt/aptare/database/stored_procedures (Linux)
\opt\oracle\database\stored_procedures (Windows)
```

To assign attributes to application databases

1 Create a table in a spreadsheet.

2 Save the table as a comma-separated file (for example, SwitchAttributes.csv).

3 Log in to the Portal server.

4 At the command line:

```
su - aptare
```

5 At the command line, launch sqlplus:

```
sqlplus <pwd>/<pwd>@//localhost:1521/scdb
```

- Example: `sqlplus portal/portal@//localhost:1521/scdb`

6 Execute the following at the **SQL** prompt:

```
SQL> Execute load_package.loadSwitchAttributeFile
('pathname_and_filename','domain_name',
san_name_col_num,switch_name_col_num,'log_path_name',
'log_file_name','check_valid_value');
```

Where:

'pathname_and_filename'	Full path + filename (enclosed in single straight quotes) of the CSV file Windows Example: '`c:\config\SwitchAttributes.csv`' Linux Example: '`/config/SwitchAttributes.csv`'
'domain_name'	Name (enclosed in single straight quotes) of the domain in which the host groups and hosts reside; Example: '**DomainEMEA**'
san_name_column_num	Column number in the csv file where the SAN Name is listed; Example: **1** Note that the SAN Name and the Switch Name can be either column 1 or 2 of the CSV. This parameter tells the utility in which column the SAN Name will be found.
switch_name_column_num	Column number in the csv file where the Switch Name is listed; Example: **2**
'log_path_name'	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to scon.log and scon.err. To omit this parameter, enter: '' Example: '`c:\config`'
'log_file_name'	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to scon.log and scon.err. To omit this parameter, enter: '' Example: '**SwitchAttributes.log**'
'check_valid_value'	'**Y**' or '**N**'enclosed in single straight quotes. **Y** - Checks if the attribute value exists. If the utility determines that the attribute value is not valid, it skips this row and does not assign the attribute value to the switch object. **N** - Updates without checking that the attribute value exists. This option is seldom chosen, but it is available for certain customer environments where attributes may have been created without values (with scripts that bypass the user interface).

Example:

```
SQL> Execute
load_package.loadSwitchAttributeFile('/config/SwitchAttributes.csv',
 'DomainEMEA', 1, 2,'/config/logs','SwitchAttributes.log','Y');
```

7 Check the log file for status and errors.

8 Restart the Portal services so that the newly added attributes become available in the Dynamic Template Designer.

Verify the Switch Attributes Load

To verify that the attribute load was successful:

1. In the Portal, go to **Reports**.
2. Select a blue user folder.
3. Select **New SQL Template**.
4. With the SQL Template Designer open, click the **Query** tab.
5. Enter the following query in the SQL Template Designer to verify Switch attributes:

```
select * from aps_v_switch_attribute
```

Load port attributes and values

Function: The Load Port Attributes utility provides an efficient method of assigning attributes to a large number of ports. Please note Fabric Manager must be installed or the loading will fail.

To load port attributes and values:

1 Create a CSV file of Ports, Attributes and Values.

2 Execute the Load Port Attribute Utility.

3 Verify the Port Attributes Load.

4 Create a report template using the SQL Template Designer.

Once attribute values are assigned to application databases, a SQL Template Designer report can query the database to report on the application databases.

Create a CSV file of Ports, Attributes, and Values

The loadPortAttributeFile utility assigns attribute values to a list of switches. This utility takes as input a comma-separated values (CSV) file.

Note: This CSV file becomes the master document of record for Port Attributes and therefore must be preserved in a working directory for future updates.

- Create a spreadsheet table, in the format shown in the following example, and save it as a CSV file in a working directory. This file is specific to loading port attributes.

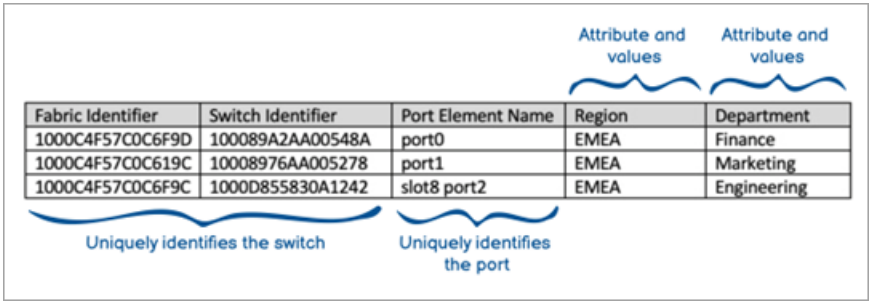

Columns:

- The first column lists the Fabric Identifier.
- The second column lists the Switch Identifier.
- The third column lists the Port Element Name.
- Each additional column lists attributes and values that will be applied to the port. Multiple attributes can be assigned to a single port object.

Rows:

- First (Header) Row - Contains the fields that uniquely identify the Fabric Identifier Name, Switch Identifier, Port element name followed by Attribute names. The header row is information only and is not processed as a data row.
- Subsequent rows list the Fabric Identifier, Switch Identifier, Port element name followed by the attribute values that you assign to each port.

Execute the Load Ports Attribute Utility

Before you begin, Bulk Load utility must be run in SQLPLUS as APTARE user.

The load_package utility is located in:

- Linux: `/opt/aptare/database/stored_procedures`
- Windows: `\opt\oracle\database\stored_procedures`

To assign attributes to application databases

1. Create a table in a spreadsheet.
2. Save the table as a comma-separated file (for example, `PortAttributes.csv`).

3 Log in to the portal server.

4 At the command line:

```
su -aptare
```

5 At the command line, launch sqlplus:

```
sqlplus <pwd>/<pwd>@//localhost:1521/scdb
```

Example:

```
sqlplus portal/portal@//localhost:1521/scdb
```

6 Execute the following at the SQL prompt:

```
SQL> Execute load_package.loadPortAttributeFile
('pathname_and_filename','domain_name',
Fabric_identifier_col_num, switch_identifier_col_num,
port_ele_name_col_num,  'log_path_name',
'log_file_name','check_valid_value');
```

Example:

```
SQL> Execute
load_package.loadPortAttributeFile('/tmp/portAttributes.csv',
'DomainEMEA', 1, 2,3,'/tmp/logs','portAttributes.log','Y');
```

Where:

pathname_and_filename	Full path + filename (enclosed in single straight quotes) of the CSV file. Windows example: 'c:\temp\PortAttributes.csv' Linux example: '/tmp/PortAttributes.csv'
domain_name	Name (enclosed in single straight quotes) of the domain in which the ports reside. Example: 'DomainEMEA'
Fabric_identifier_col_num	Column number in the CSV file where the Fabric Identifier is listed; Example: 1
switch_identifier_col_num	Column number in the CSV file where the Switch Identifier is listed; Example: 2
port_ele_name_col_num	Column number in the CSV file where the Port Element Name is listed; Example: 3
log_path_name	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to scon.log and scon.err. To omit this parameter, enter: Example: 'c:\temp' or '/tmp'
log_file_name	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to scon.log and scon.err. To omit this parameter, enter: '' Example: 'PortAttributes.log'

check_valid_value	'Y' or 'N' - enclosed in single straight quotes. Y - Checks if the attribute value exists. If the utility determines that the attribute value is not valid, it skips this row and does not assign the attribute value to the switch object. N - Updates without checking that the attribute value exists. This option is seldom chosen, but it is available for certain customer environments where attributes may have been created without values (with scripts that bypass the user interface).

7 Check the log file for status and errors.

8 Restart the portal services so that the newly added attributes become available in the product.

Verify the Port Attributes Load

To verify that the attribute load was successful:

1 In the portal, go to Reports.

2 Select a blue user folder.

3 Select **New SQL Template**.

4 With the **SQL Template Designer** open, click the **Query** tab.

5 Enter the following query in the **SQL Template Designer** to verify Switch attributes:

```
select * from aps_v_swi_port_attribute
```

Load Subscription attributes and values

Function: The Load Subscription Attributes utility provides an efficient method of assigning attributes to a large number of subscriptions.

To load subscription attributes and values:

1 Create a CSV file of Subscriptions, Attributes, and Values.

2 Execute the Load Subscription Attribute Utility.

3 Verify the Subscription Attributes Load.

4 Create a report template using the SQL Template Designer.

Once attribute values are assigned to application databases, a SQL Template Designer report can query the database to report on the application databases.

Create a CSV file of Subscription, Attributes, and Values

The loadSubscriptionAttrFile utility assigns attribute values to a list of subscriptions. This utility takes a comma-separated values (CSV) file as input.

Note: This CSV file becomes the master document of record for Subscription Attributes and hence must be preserved in a working directory for future updates.

Create a spreadsheet table, in the format shown in the following example, and save it as a CSV file in a working directory. This file is specific to loading Subscription attributes.

SubscriptionID	Level1	Level2	Level3
e8e71c52-9613247febe-4e28-467d-a65c	Charlie	Buster	Curly
e8e71c52-9612-43c1-89a6-a728379218d0	Charlie	Buster	Larry
1f0d0ca6-e405-4f13-84ff-691fd10d41bd	Charlie	Buster	Curly
2c2b34a3-1122-49cb-adca-4c1d1c413567	Charlie	Buster	Moe

Cloud Account Subscription ID | Attributes and Their Values

Columns:

- The first column lists the Subscription Identifier.
- Each additional column lists attributes and values that will be applied to the subscription.
  Multiple attributes can be assigned to a single subscription object.

Rows:

- First (Header) Row - Contains the fields that uniquely identify the Subscription Identifier followed by Attribute names. The header row is information only and is not processed as a data row.
- Subsequent rows list the Subscription followed by the attribute values that you are assigning to each subscription.

Execute the Load Subscription Attribute Utility

Before you begin, Bulk Load utilities must be run in SQLPLUS as user APTARE.

The load_package utility is located in:

- Linux: `/opt/aptare/database/stored_procedures`
- Windows: `\opt\oracle\database\stored_procedures`

To assign attributes to application databases:

1. Create a table in a spreadsheet.
2. Save the table as a comma-separated file (for example, `subscription.csv`).

3 Log in to the portal server.

4 At the command line:

```
su -aptare
```

5 At the command line, launch sqlplus:

```
sqlplus <pwd>/<pwd>@//localhost:1521/scdb
```

Example:

```
sqlplus portal/portal@//localhost:1521/scdb
```

6 Execute the following at the SQL prompt:

```
SQL> Execute load_package.loadSubscriptionAttrFile
('pathname_and_filename','domain_name',
subscription_identifier_col_num, 'log_path_name',
'log_file_name','check_valid_value');
```

Example:

```
Execute load_package.loadSubscriptionAttrFile
('/tmp/subscription.csv', 'DomainEMEA',1,'/tmp',
'subscription.log',’Y');
```

Where:

pathname_and_filename	Full path + filename (enclosed in single straight quotes) of the CSV file. Windows example: '`c:\temp\subscription.csv`' Linux example: '`/tmp/subscription.csv`'
domain_name	Name (enclosed in single straight quotes) of the domain in which the host groups and hosts reside. Example: 'DomainEMEA'
subscription_identifier_col_num	Column number in the CSV file where the Subscription Identifier is listed; Example: 1
log_path_name	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to scon.log and scon.err. To omit this parameter, enter: Example: 'c:\temp’ or ‘/tmp’

log_file_name	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to `scon.log` and `scon.err`. To omit this parameter, enter: '' Example: '`subscription.log`'
check_valid_value	'Y' or 'N' - enclosed in single straight quotes. Y - Checks if the attribute value exists. If the utility determines that the attribute value is not valid, it skips this row and does not assign the attribute value to the switch object. N - Updates without checking that the attribute value exists. This option is seldom chosen, but it is available for certain customer environments where attributes may have been created without values (with scripts that bypass the user interface).

7 Check the log file for status and errors.

8 Restart the portal services so that the newly added attributes become available in the product.

Verify the Subscription Attributes Load

To verify that the attribute load was successful:

1 In the portal, go to Reports.

2 Select a blue user folder.

3 Select **New SQL Template**.

4 With the **SQL Template Designer** open, click the **Query** tab.

5 Enter the following query in the **SQL Template Designer** to verify Switch attributes:

```
select * from aps_v_swi_subscription_attribute
```

Chapter 9

Importing generic backup data

This chapter includes the following topics:

About generic backup data collection

Backup Manager can report on data from backup products that are not native to IT Analytics--such as PureDisk, BakBone, and BrightStor. Using the backup vendor's export feature, create a comma-separated values (CSV) file. The IT Analytics Data Collection process will import the data into the Portal database, to be included in IT Analytics reports, such as the Job Summary report. The data can be scheduled for regular collection intervals.

Note: In addition to the regularly scheduled data collection, the CSV file also can be imported manually.

See "Manually loading the CSV file" on page 143.

Considerations

- Files can be imported more than once. Importing will not result in duplicate entries.

- Data is job-centric only--that is, no tape media or tape library information is imported.
- When checking Data Collection Status, the indicator may display a red error status, which for generic backup data collection, may not be a true error condition. The Generic Backup Data Collector checks the timestamp of the CSV file and if it is the same as the last collection, it does not attempt to re-import the data. In this regard, the Generic Backup data collection process differs from other collectors, as it expects to have data provided via the CSV file.
- Data is stored securely and can be used for historical tracking and trending.

Configuring generic backup data collection

The following tasks must be performed as part of the Generic Backup Data Collection setup:

- In the IT Analytics Portal, create a host of type: Generic Backup Server. This is the server that is managing the backups of the clients for which you will be importing backup data.
- In the IT Analytics Portal, add a Data Collector Policy of Product Type: Generic Backup. In this policy, supply the path to the CSV file.
- Create a comma-separated file of the backup/restore data--typically, a file that has been exported using the backup software utilities.
  See "CSV Format Specification" on page 141.

CSV Format Specification

Using the backup software, create a comma-separated file that contains the following 15 data elements from the backup/restore job(s). Note that each field must have an entry, even if it is a null entry within the commas. Field values cannot contain embedded commas. All string fields must be enclosed within single straight quotes.

Note: The CSV file must be UTF-8 encoded, however be sure to remove any UTF-8 BOMs (Byte Order Marks). The CSV cannot be properly parsed with these additional characters.

Table 9-1 CSV file with data elements.

Data Elements	Data Type	Value
VendorName	STRING	The name of the backup application used to perform the backup, enclosed in single straight quotes.
ClientName	STRING	The host name of the machine being backed up, enclosed in single straight quotes.
ClientIPAddress	NUMBER	The IP address of the machine being backed up. If an IP address is not available, simply use two single straight quotes (") or 'null' to indicate a blank/missing value.
VendorJobType	STRING	Valid values include: BACKUP or RESTORE--enclosed in single straight quotes.
StartDateString	DATE	The start date and time of the backup job in the format: YYYY-MM-DD HH:MI:SS (enclosed in single straight quotes). **Note:** Adhere to the specific date format--number of digits and special characters--as shown above.
FinishDateString	DATE	The end date and time of the backup job in the format: YYYY-MM-DD HH:MI:SS (enclosed in single straight quotes). **Note:** Adhere to the specific date format--number of digits and special characters--as shown above.
BackupKilobytes	NUMBER	The numeric size of the backup in kilobytes (otherwise use 0). Remember IT Analytics uses 1024 for a KiB.
NbrOfFiles	NUMBER	The number of files that were backed up (otherwise use 0).
MediaType	STRING	The type of media that was used: T for Tape or D for Disk, enclosed within single straight quotes.
VendorStatus	NUMBER	A numeric job status: 0=Successful, 1=Partially Successful, or 2=Failed.
VendorJobId	STRING	Vendor job ID, enclosed in single straight quotes.
VendorPolicyName	STRING	Vendor policy name, enclosed in single straight quotes.
JobLevel	STRING	Job level, enclosed in single straight quotes. Example: Incremental, Full.
TargetName	STRING	File system backed up by the managed backup system (MBS), enclosed in single straight quotes.

Table 9-1 CSV file with data elements. *(continued)*

Data Elements	Data Type	Value
ScheduleName	STRING	Name of the backup schedule, enclosed in single straight quotes.

EXAMPLE: genericBackupJobs.csv

```
'Mainframe Backup','mainframe_name','10.10.10.10','BACKUP','2008-03-24
 10:25:00', '2008-03-24
11:50:00',3713,45221,'D',0,'413824','Retail_s01002030','Incremental','/I:/Shared/','Daily'
'UNIX tar backup','host_xyz.anyco.com','null','BACKUP','2008-03-24
10:22:00','2008-03-24
12:50:00',1713,45221,'T',1,'5201','HQ_Finance','Full','/D:/Backups/','Daily'
'ArcServe','host_123.anyco.com','null','RESTORE','2008-03-24
8:22:00','2008-03-24
9:12:00',0,0,'T',0,'2300','Retail_s03442012','Incremental','/I:/Shared/','EOM'
```

Manually loading the CSV file

Use the following procedure to manually load the Generic Backup CSV file into the Portal database.

Prerequisites:

- These scripts must be run on the **Collector Server**.
- The **checkinstall** script must be run first to register the event collector ID.

1. List the Data Collectors to get the **Event Collector ID** and the **Server ID**, which will be used in step 2.

   Windows:

   ```
   C:\opt\APTARE\mbs\bin\listcollectors.bat
   ```

   Linux:

   ```
   /opt/aptare/mbs/bin/listcollectors.sh
   ```

   In the output, look for the Event Collectors section associated with the Software Home--the location of the CSV file (the path that was specified when the Data Collector Policy was created). Find the **Event Collector ID** and **Server ID**.

   ```
   ==== Event Collectors ===
   ```

Event Collector Id: EVENT_1029161_9

```
Active: true
Active: true
Software Home: C:\gkgenericBackup.csv
Server Address: 102961
Domain: gkdomain
Group Id: 102961
Sub-system/Server Instance/Device Manager Id: 102961
Schedule: */10 * * * *
```

2. Use the following commands to load the data from the CSV file into the Portal database.

   Windows:

   ```
   C:\opt\APTARE\mbs\bin\loadGenericBackupData.bat <EventCollectorID>
    <ServerID> [verbose]
   ```

   Linux:

   ```
   /opt/aptare/mbs/bin/loadGenericBackupData.sh <EventCollectorID>
   <ServerID> [verbose]
   ```

Note: If you run the command with no parameters, it will display the syntax.

The load script will check to see if the backup server and client already exist; if not, they will be added to the database. The script then checks for a backup job with the exact same backup server, client, start date and finish date. If no matches are found, the job will be added; otherwise, it will be ignored. This prevents duplicate entries and allows the import of the script to be repeated, if it has not been updated. Once the load is complete, these clients and jobs will be visible via the IT Analytics Portal and the data will be available for reporting.

Chapter 10

Backup job overrides

This chapter includes the following topics:

Overview

In some backup environments, it is desirable to treat backup warning status messages as successful backups. A configuration modification can change the default behavior of IT Analytics reports. You may want to override other backup statuses as well. IT Analytics supports job overrides for all supported backup products.

Use the following procedure to update the job override configuration.

Note: For the purpose of simplicity, only NetWorker and NetBackup job override steps are shown. Similar configuration changes can be done for other backup products.

Configure a backup job override

1. At the command line, log in to the Portal Server as user **aptare**.
2. Make a copy of the following files. Note that in this step, only the NetWorker (leg) override is shown, however, other backup product job overrides can be customized using their relevant .plb files.

   Linux:

   ```
   /opt/aptare/database/stored_procedures/job_override.sql
   /opt/aptare/database/stored_procedures/leg/leg_adaptor_pkg.plb
   ```

Windows:

```
C:\opt\aptare\database\stored_procedures\job_override.sql
C:\opt\aptare\database\stored_procedures\leg\leg_adaptor_pkg.plb
```

3. On the Portal server, shut down the Data Receiver.
   - Linux: `/opt/aptare/bin/tomcat-agent stop`
   - Windows: `C:\opt\aptare\utils\stopagent.bat`
4. Edit **job_override.sql** and find the following section.

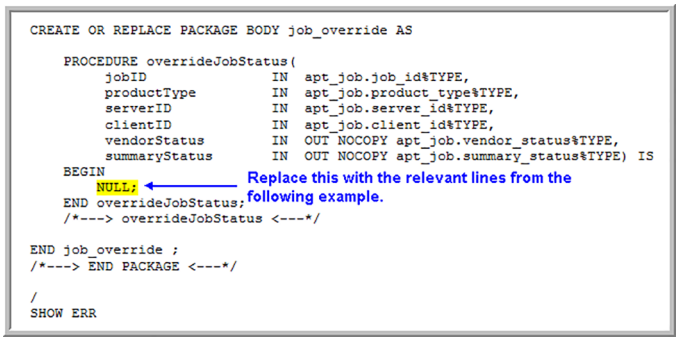

5. Update the job_override.sql file with code to customize how IT Analytics will override the backup status. The following example illustrates updates for NetBackup or NetWorker, enabling a warning to be treated as a successful backup. Note that overrides are supported for all backup products.

Example of Job Overrides for NetWorker and NetBackup

```
CREATE OR REPLACE PACKAGE BODY job_override AS
  PROCEDURE overrideJobStatus(
            jobID           IN  apt_job.job_id%TYPE,
            productType     IN  apt_job.product_type%TYPE,
            serverID        IN  apt_job.server_id%TYPE,
            clientID        IN  apt_job.client_id%TYPE,
            vendorStatus    IN  OUT NOCOPY apt_job.vendor_status%TYPE,
            summaryStatus   IN  OUT NOCOPY apt_job.summary_status%TYPE) IS
  BEGIN
    IF productType = constant.PRODUCT_LEGATO_NW THEN
      IF summaryStatus = constant.JOB_STATUS_WARNING THEN
        summaryStatus := constant.JOB_STATUS_OK;
        vendorStatus  := constant.JOB_STATUS_OK;
      END IF;
    END IF;
    IF productType = constant.PRODUCT_VERITAS_NBU THEN
      IF summaryStatus = constant.JOB_STATUS_WARNING THEN
        summaryStatus := constant.JOB_STATUS_OK;
        vendorStatus  := constant.JOB_STATUS_OK;
      END IF;
    END IF;
  END overrideJobStatus;
  /*---> overrideJobStatus <---*/
END job_override ;
/*---> END PACKAGE <---*/
/
SHOW ERR
```

Note: The above example is for illustration purposes only. You may choose to customize job overrides for other backup vendor job statuses.

6. Go to:
   - Linux: `/opt/aptare/database/stored_procedures/`
   - Windows: `C:\opt\aptare\database\stored_procedures\`
7. Compile the SQL binary (.sql file):

   ```
   sqlplus portal/portal@//localhost:1521/scdb @ job_override.sql
   ```

8. Go to:
   - Linux: `/opt/aptare/database/stored_procedures/leg/`
   - Windows: `C:\opt\aptare\database\stored_procedures\leg`
9. Compile the PL/SQL binary (.plb file):

sqlplus portal/portal@//localhost:1521/scdb @ leg_adaptor_pkg.plb

10. Restart the Data Receiver.
    - Linux: `/opt/aptare/bin/tomcat-agent start`
    - Windows: `C:\opt\aptare\utils\startagent.bat`

Chapter 11

Managing host data collection

This chapter includes the following topics:

Identifying hosts by WWN to avoid duplicates

By default, IT Analytics data collection finds hosts based on IP address or host name. Often hosts are collected from multiple sources and these sources have different names for the same host. In such environments, host name or IP address are not sufficient for uniquely identifying a host. In order to prevent duplicate hosts from being created in the database, IT Analytics can use the host's port WWN to uniquely identify a host.

A typical scenario that warrants WWN matching for unique host identification is described in the following example. Host data can be collected in multiple ways; for example, via a manual CSV load, as well as from Virtualization Manager and from Capacity Manager HP 3PAR collection. In this example, all three sources provide different names for the host, which would cause duplicates to be saved in the IT Analytics database. Therefore, in this case, matching on a host port WWN

ensures unique hosts. With WWN matching, if different host names are encountered, a host alias record is also created in anticipation of future host data collection.

By default, WWN matching is turned off. A system parameter can be configured to turn on WWN matching prior to data collection.

To turn on host WWN matching, type the following command at the command line.

```
update apt_system_parameter set param_value = '1'
  where param_name = 'SEARCH_HOST_BY_WWN_IS_ENABLED';
COMMIT;
```

A value of 1 turns on WWN matching and 0 turns it off (the default).

Setting a host's priority

IT Analytics can collect host data from multiple vendor products (subsystems), such as Veritas NetBackup and IBM XIV. When host data is collected from more than one subsystem, host reports will display the data from the primary subsystem. IT Analytics provides a default ranking for subsystems. When a host is collected, that rank order is referenced to determine if the collected host is coming from the primary subsystem.

An Administrator can override that default ranking and configure a different source subsystem as primary by using the following instructions to customize the ranking for your enterprise.

1. Log on to the Portal Server as user **aptare**.
2. At the command prompt, type: **sqlplus <pwd>/<pwd>@//localhost:1521/scdb**
3. Execute the following at the **SQL** prompt to view the default host ranking. In this table, the Product Types translate to: 1 = backup, 2 = capacity, 4 = virtualization, 8 = replication, 16 = fabric, 32 = File Analytics.

   ```
   SQL> SELECT * from apt_host_source_rank;
   ```

aptare_product_type	product_vendor	priority	product_vendor_name
100	1	202	Veritas NetBackup
1	4	203	Tivoli Storage Manager
1	3	204	EMC NetWorker
1	5	205	CommVault Simpana

1	6	206	HP Data Protector
1	2	207	Veritas Backup Exec
1	9	208	Generic Backup
1	7	209	EMC Avamar
4	51	301	VMware
4	52	302	IBM VIO
2	41	201	Host Resource
2	21	401	Hitachi Data Systems
2	211	402	Hitachi NAS
2	22	403	EMC
2	221	404	EMC CLARiiON
2	222	405	EMC Symmetrix
2	223	406	EMC VNX (Celerra)
2	225	407	EMC Isilon
2	231	408	NetApp Cluster-Mode
2	23	409	NetApp
2	24	410	HP
2	241	411	HP 3PAR
2	25	412	IBM
2	26	413	NetApp E-Series
2	27	414	IBM SVC
2	28	415	HP EVA
2	254	416	IBM XIV
2	29	417	Dell Compellent
8	61	501	NetApp SnapMirror
8	62	502	NetApp SnapVault

16	701	601	Brocade Switch
16	703	603	Cisco Switch
32	801	700	File Analytics
1	32	701	EMC Data Domain
		800	HBA CSV Load
		801	CSV Load

4. Execute the following to customize the host source subsystem ranking for your enterprise. This command can be repeated for as many vendor products (subsystems) as needed in your environment. It updates a custom host source ranking table, which is specific to your environment.

   See "Determining host ranking" on page 154.

   ```
   SQL> INSERT INTO apt_host_user_source_rank (domain_id,
   aptare_product_type, product_vendor, priority,
   product_vendor_name)    VALUES (<domain_id_value>,
   <aptare_product_type_value>, <product_vendor_value>,
   <priority_value>, '<product_vendor_name_value>');
   SQL> Commit;
   ```

   where:

   <domain_id_value>

   Most environments have only one Domain ID, however, Managed Service Providers (MSPs) will have a different Domain ID for each of their customers.

   To list the currently configured Domain IDs, use the following SQL SELECT statement:

   SQL> SELECT * from apt_domain;

   <aptare_product_type_value>

   The product type is a number that represents the IT Analytics product, such as Capacity Manager.

   1 = backup, 2 = capacity, 4 = virtualization, 8 = replication, 16 = fabric, 32 = File Analytics

<product_vendor_value>	This number represents the vendor and subsystem from which the host data is collected.
<priority_value>	This number sets the priority ranking for the host. Priority numbers used by customers should be between 1 and 99.
'**<product_vendor_name_value>**'	This name corresponds to the vendor_number; for example, EMC Avamar. This must be entered exactly as it is listed in the Default Host Ranking Table. **Note:** The product vendor name is not mandatory when the product vendor (number) is available. In this case, a null value within single quotes can be used for the product vendor name value.

Example:

```
INSERT INTO apt_host_user_source_rank (domain_id,
aptare_product_type, product_vendor, priority,
product_vendor_name) VALUES (100396, 1, 1, 88, 'Veritas
NetBackup');
Commit;
```

5. Execute the following to view the host ranking that you customized for your enterprise:

```
SQL> SELECT * from apt_host_user_source_rank;
```

6. To update a rank that you have customized, use the following steps. Refer to the

   for column names.

```
SQL> update apt_host_user_source_rank set  <column_name> = <value>
 where <column_name> = <Value>;
```

Example:

```
update apt_host_user_source_rank set priority = 91 where
product_vendor = 1;
```

Determining host ranking

The user-specific host ranking table differs from the default ranking table, as it has a Domain ID column that enables Domain-specific ranking for hosts. You can override host ranking with a Null Domain ID; this becomes a system-wide override.

When determining the priority of a host, IT Analytics checks the ranking in the following order of priority:

1. User host source rank with a Domain ID populated
2. User host source rank with a Null Domain ID
3. Default host source rank, as defined in the

Loading host and WWN relationships

In environments with firewall restrictions, Capacity Manager Host Resources data cannot be collected. IT Analytics provides a utility for manually loading host and Worldwide Name (WWN) relationships into the Portal database so that this limited information can be included in the standard Capacity Manager reports.

This utility takes a comma-separated values (CSV) file as input and populates database tables so that the Host Bus Adapter (HBA) data is available in reports.

Loading the host HBA port data

This utility provides an bulk-load method for populating the database with host data.

Create a CSV file

In preparation for loading host data, enter the information into a spreadsheet from which you will create a comma-separated file.

The table in the spreadsheet should be in the following format:

- first column - host name
- second column - node WWN
- third column - port WWN

Example of a host WWN CSV file

```
test_host, 50:06:0E:80:05:63:B7:20, 50:06:0E:80:05:63:B7:21
```

If a host has multiple HBAs, the CSV should contain a row for every HBA so that all HBAs for the host will be loaded. For example:

```
host123, 50:06:0E:80:05:63:B7:20, 50:06:0E:80:05:63:B7:21
host123, 50:06:0E:80:05:63:B7:21, 50:06:0E:80:05:63:B7:24
host123, 50:06:0E:80:05:63:B7:25, 50:06:0E:80:05:63:B7:26
```

Execute the script

To load host HBA port data

Note: When running this script, pay attention to the value you supply for the isIncremental parameter. When you specify '**N**' your existing host data is deleted. When you specify '**Y**' your host data is added without removing existing records.

---

1. Create a table in a spreadsheet, as shown in the above example. Save the table as a comma-separated file (for example, hostWWN.csv).
2. Log on to the Portal Server as user **aptare**.
3. At the command prompt, type: **sqlplus <pwd>/<pwd>@//localhost:1521/scdb**

4 Execute the following at the **SQL** prompt:

```
SQL> Execute srm_load_pkg.loadHBAPortFile('<
domainName>', '<isIncremental>', '<CSVfile>', '<logPathname>',
'<logFilename>',[,'<source_name>']);
```

where:

'domainName'	IT Analytics domain name (enclosed in single straight quotes)

'isIncremental'	'**Y**' **or** '**N**' (enclosed in single straight quotes) to indicate if it is an incremental load.
	
	`If 'Y',`
	`an HBA port record will be created if none exists.`
	`If 'N',`
	`old HBA port records will be deleted first`
	`and then new records created. Take care`
	`when choosing this option, as it will`
	`remove existing host data from the database.`

'CSVfile'	CSV file path and name (enclosed in single straight quotes)
'logPathname'	Log path name (enclosed in single straight quotes). The audit log file is created only if errors occur. Other status is logged in scon.log.
'logFilename'	Log file name (enclosed in single straight quotes). This audit log file is created only if errors occur. Other status is logged in scon.log.
'source_name'	**source_name** is an optional, case-insensitive string, up to 100 characters, representing the source of the host details; for example, CMDB might be relevant for a change management database. This source information is retained for historical purposes, to track how the host was added to the database. If nothing or NULL is provided for this parameter, **HBA CSV Load** will be inserted as the source into the reporting database.

Example:

```
Execute srm_load_pkg.loadHBAPortFile('corpHost1', 'Y',
'/tmp/hba_port_data.txt','/tmp','hba_port_data.log','CMDB');
```

5 IMPORTANT: If you created a new **source_name**, you need to insert it into the custom host source ranking table using the instructions provided in

See “Setting a host’s priority” on page 150.

Chapter 12

System configuration in the Portal

This chapter includes the following topics:

System configuration in the Portal

Configure a number of the components in your system directly from the Portal. Using the System Configuration feature available to Super Users, you can modify default values established during installation for everything from data retention period to how email is setup.

Functional Areas

Functional areas are divided into separate tabs as follows:

- Data Collection - Set values for all collection, product-based collection and vendor-based.
  See "Data collection: Capacity chargeback" on page 162.
- Data Retention - Modify default retention periods for systems that are collected by traditional Data Collectors to determine when data is purged from the database. Purging is required to maintain reasonable table sizes. Data types include historical and performance data. Fields are displayed based on what has been installed and collected.
  For systems collected by Data Collectors deployed via the SDK, use the procedure described in:
  See "Data retention periods for SDK database objects" on page 280.
- Database Administration - Set values to configure the structure of the database.
  See "Database administration: database" on page 162.
- Host Discovery - Enable rules for host matching when the system is discovering new hosts/clients.
  See "Host discovery: EMC Avamar" on page 163.
  See "Host discovery: Host" on page 164.
- Inventory - Modify the database polling frequency for Inventory objects.
- Portal - Modify default values for a variety of Portal properties including:
  - Host attribute import parameters
  - Maximum number of open tabs
  - Security settings such as time out values and allowed login attempts
  - Custom headers and footers for reports and dashboards.

    Note: Requires Portal Tomcat restart.

  - Event audit logging. See "Events captured for audit" on page 165.
- Custom Parameters - Add, edit and delete custom system parameters, Portal properties and their associated values. This area allows free form entry for name/value pairs.
  See "Custom parameters" on page 167.

System configuration: functions

Buttons available at the top of the tabbed space, apply across all functional tabs.

Table 12-1 Function description

Function buttons	Description
Save and Apply	Before saving and applying changes, a dialog is displayed to show old values and new values to verify the update. Some changes require a Portal restart. If a restart is required, this is displayed in the confirmation dialog. See "Starting and stopping portal server software" on page 44. for information about restarting systems.
Undo Changes	Cancels changes and resets to the last value across all tabs within the System Configuration area. Use the field level refresh icon to reset values field by field.
Restore All Defaults	Resets default values for all parameters across all tabs within the System Configuration area. Rollover the icons to display the parameters default value.
Download	Click to download a text file of all your system setting values. This includes any custom parameter values.
Expand/Collapse All	Click to expand or collapse all categories on within the System Configuration area. Categories can be expanded or collapsed individually using the icon beside the category title.

Navigation overview

This self-service portal makes it easy to quickly determine what parameters you are setting. The following graphic outlines some of the built-in features.

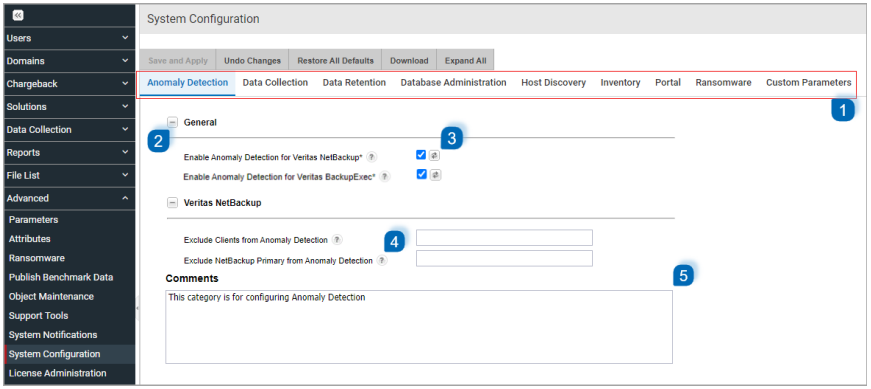

1. Functional areas organized into tabs
2. Expand and collapse categories individually
3. Field-level restore icons reset default values per parameter. Default values are displayed in rollover text.
4. Field-level help displays short parameter descriptions for use-case clarification.
5. Free form comments are allowed to enter notes about updates.

System configuration parameter descriptions: Additional info

Some configuration settings are straight forward and do not require more explanation than is available in the Portal through the field-level help. Other settings require more information. The section covers the additional information not displayed in the Portal.

Anomaly detection

Anomaly detection helps to detect suspicious activity in backup operation through analysis of various backup attributes such as:

- Backup image size
- Number of files backed up
- Kilobytes transferred
- Deduplication rate
- Backup job completion time

Any significant change in the above parameters detected during a backup job are reported as anomalies. The **General** section lists all the data protection tools for which you can enable anomaly detection.

You can configure anomaly detection from NetBackup Web UI (see *Configure anomaly detection settings* section of the *NetBackup Security and Encryption Guide*). The anomalies detected by NetBackup are capture by the Security Details probe in the NetBackup Policy for the primary server.

Data collection: Capacity chargeback

- Drive capacity chargeback range in GB: When configuring Capacity Chargeback Policies, a Drive Capacity (GB) policy type can be defined. This value is treated as a range of capacity, with the default set to: plus or minus 10 GB. Often this default is not sufficient for an environment's chargeback policies. Using system parameters, the range for the Drive Capacity policy type can be modified. The value configured in the policy will be treated as a range of values--that is, the Policy's Drive Capacity plus or minus the Drive Capacity Rangethat is configured in the system parameters.
- Enable chargeback logging: Enable or disable capacity chargeback logging. This allows data to be collected and then custom report templates developed with one of the report template designers can query this data to generate chargeback reports. Capacity chargeback logging is enabled by default. If this is not a requirement for your environment, disable it.

Database administration: database

- Large index number leaf blocks: In large environments, to improve the performance of index builds, a system parameter can be configured to define the number of leaf blocks for a large index. The default value is 10,000 leaf blocks. This parameter configures the number of leaf blocks in a database index.
- Maximum number of large indexes for rebuild: To improve the performance of index recreation, modify this parameter to change the number of large database indexes that will be processed in a single run. The default value is 10.
- Rebuild indexes schedule (days): In large environments, if report generation performance begins to degrade, database indexes can be built more frequently. The default is to rebuild indexes every 60 days.
- Maximum time in minutes for large index rebuilds: To improve the performance of index recreation, configure this parameter to define the number of minutes the index rebuild should run. The default value is 10 minutes. If the rebuild takes

longer than this time, the job will stop. In very large environments, it may be necessary to increase this time to accommodate large indexes.

Host discovery: EMC Avamar

In a specific circumstance, EMC Avamar data collection can persist duplicate clients in the Portal database. This occurs in the following case:

- Multiple enterprise domains are configured.
- The same host name is used in multiple domains, but for different hosts.

If your Portal has the above configuration, configure the following to prevent the creation of duplicate Avamar clients. And, logging can be configured to identify how a host is determined and persisted in the Portal database.

Note: Another parameter, Enable IP address matching for Host search, is also used by the Avamar host-matching algorithm. If your environment has already enabled this parameter, it will be honored by the host-matching algorithm.

- Enable short name matching for Avamar Host Search: This parameter is used to enable comparisons of a client's base name. During data collection, the data persistence logic will compare the short name retrieved by data collection and compare it to what exists in the Portal database. This parameter currently is used only while searching for a host in Avamar data. For example, the host name in the database might be **xyz.aptare.com**, but the collected host name is **xyz.apt.com**. If this parameter is enabled, the host-matching algorithm will find the host with the name, **xyz.aptare.com**, based on matching the short name, **xyz**, thereby preventing the creation of a duplicate host.
- Remove patterns in host matching: This parameter will enable the stripping of unwanted suffixes while searching for hosts based on host name. This parameter is currently used only while searching for a host in Avamar data.
  - **Prerequisite**: Any suffix that needs to be ignored must first be inserted into the **apt_host_name_excld_suffix** database table, as described in the following procedure. When the parameter is enabled, the host-matching algorithm searches this table for suffixes that should be ignored.

  Add suffixes to the database table:

```
INSERT INTO apt_host_name_excld_suffix (excld_suffix,
suffix_length, priority) VALUES
(<<excludeSuffixInitials>>,<<totalSuffixLength -
lengthOfexcludeSuffixInitials>>, <<priority>>);
COMMIT;
```

Examples

The data searching logic used when this system parameter is enabled is described in the following examples.

- Host name in the database is **xyz** and the collected host name is **xyz_UCMAAZWIR6kihhBHN5R8iA**. The host-matching algorithm will find the host with the name **xyz** and **_UCMAAZWIR6kihhBHN5R8iA** will be removed while searching.
- Host name in the database is **xyz** and the collected host name is **xyz_UA3rT06VdULrQyVilxEFuQ2011.07.22.16.05.49**. The host-matching algorithm will find the host with the name **xyz** and **_UA3rT06VdULrQyVilxEFuQ2011.07.22.16.05.49** will be removed while searching. The time portion, **2011.07.22.16.05.49**, is automatically removed if the parameter is enabled.
- Host name in the database is **xyz** and the collected host name is **xyz2011.07.22.16.05.49**. The host-matching algorithm will find the host with the name **xyz** and **2011.07.22.16.05.49** will be removed while searching. The time portion, **2011.07.22.16.05.49**, is automatically removed if the parameter is enabled.
- Host name in database is **xyz** and the collected host name is **xyz2011.07.22.16.05.49_UA3rT06VdULrQyVilxEFuQ**. The host-matching algorithm will find the host with the name **xyz** and **2011.07.22.16.05.49_UA3rT06VdULrQyVilxEFuQ** will be removed while searching. The time portion, **2011.07.22.16.05.49**, is automatically removed if the parameter is enabled.

Host discovery: Host

Enable IP address matching for Host search: In certain environments, where hosts or VMs are frequently provisioned and/or decommissioned causing IP addresses to be re-used, these duplicate IP addresses can result in multiple aliases for a single host. Use this parameter to activate/deactivate IP address matching. When the parameter is disabled, a collected host with an IP address that matches a host in the database, but that has a different host name, will result in the creation of a new host in the database. By default, this value is disabled and its value is preserved even after upgrades.

You can enable this parameter only when you are certain that the IP addresses will never get reassigned and reused by different hosts.

The options to enable base name (**Enable base name matching for Host search**) and short name (**Enable short name matching for Host search**) matching for host search provide administrators with better control to find hosts in IT Analytics.

The short name refers to the shortest possible matching name when the FQDN is parsed into its constituent tokens. The base name matches the left-most token in the FQDN of a host, also known as the host name, when it is parsed into its constituent tokens.

Events captured for audit

Event audit logging is always enabled by default in IT Analytics and can be edited by an admin user from **Admin** > **Advanced** > **System Configuration** > **Portal** tab. When enabled, the events captured can be viewed on the Audit Events Summary and Audit Event Details reports. You can trace the user actions and object value changes using this feature.

Events that fit the following criteria are captured:

- Only selected audit events (or user actions) are audited.
- For template events (create/modify/delete) attribute changes are not captured.
- Attribute changes (such as create/modify/delete) must not be for template events
- Attribute changes (such as create/modify/delete) must not belong to Veritas Backup Exec policy.

Event audit logging captures the following events:

Table 12-2 Events captured

Event group	Event	Event description
User Management	Add user	A new user is added.
	Delete user	An existing user is deleted.
	Modify user details	Existing user details are modified.
User Group Management	Add user	A user is added to a user group.
	Remove user	A user is removed from a user group.

Table 12-2 Events captured *(continued)*

Event group	Event	Event description
Security	User login succeeded	A successful user login.
	User login failed	User login has failed.
	User logged out	User has successfully logged out.
	User switched ID	User has successfully switched the ID.
	Unauthorized action	User has performed an unauthorized action.
	API key created	API key generated successfully.

Table 12-2 Events captured *(continued)*

Event group	Event	Event description
User Actions	Collector added	User has added a collector.
	Collector details modified	User has modified collector details, such as host name, servers, and password.
	Collector key file generated	Collector key file was successfully generated.
	Collector deleted	User has deleted a collector.
	Policy added	User has added a policy.
	Policy details modified	User has modified the policy details.
	Policy deleted	User has deleted a policy.
	Policy portability	User has moved a policy from one collector to other.
	Report template created	User has created a report template.
	Report template modified	User has modified a report template.
	Report template deleted	User has deleted a report template.
	Report instance created	User has created a report instance.
	Report instance modified	User has modified a report instance.
	Report instance deleted	User has deleted a report instance.

Custom parameters

Customizations to the Portal extend beyond what is available in the System Configuration. When working with Services and Cohesity Support, you may be required to add or edit custom parameters to address a particular issue. The Custom Parameters tab enables free-form key value pairs to further customize IT Analytics.

Note: Prior to version 10.3, customizations to the Portal were made using a file, portal.properties. Not all of those settings are displayed in the System Configuration feature. If you upgrade from a version prior to 10.3, those properties are displayed and automatically populated in the Custom Parameters.

Adding/editing a custom parameter

1. Navigate to Admin>Advanced>System Configuration>Custom Parameters.
2. Select a custom parameter if editing.
3. Click Add/Edit. The Add Custom Parameters dialog is displayed.
4. Enter the Custom Parameter Name and Custom Parameter Value.
5. Click Save. The parameter is added to the list and available for you to revise in future sessions.

Portal customizations

This section covers customizations for the portal that are not available through the user interface. Use Custom Parameters to add/edit and delete these properties.

- See "Configuring global default inventory object selection" on page 168.
- See "Restricting user IDs to single sessions" on page 169.
- See "Customizing date format in the report scope selector" on page 169.
- See "Customizing the maximum number of lines for exported reports" on page 169.
- See "Customizing the total label display in tabular reports" on page 170.
- See "Customizing the host management page size" on page 170.
- See "Customizing the path and directory for File Analytics database" on page 170.
- See "Configuring badge expiration" on page 171.
- See "Configuring the maximum cache size in memory" on page 171.
- See "Configuring the cache time for reports" on page 172.

Configuring global default inventory object selection

To globally configure the selection of default Inventory objects for users, modify the portal.properties files. This is useful for filtering environments with large volumes of data that may be impacted by browser limitations. For new users who have never logged into the portal, the objects defined with this setting are shown selected when

they log in. For existing users, this property can be used to reset a users environment when large volumes of data can present issues with certain browsers.

- Use the following:

```
portal.ocn.defaultVisibleObjectType=
HOST,ARRAY,SWITCH,BACKUPSERVER,VM_SERVER,VM_GUEST,
DEDUPLICATION_APPLIANCE,DATASTORE,EC2_INSTANCE,
S3_BUCKET,AZURE_STORAGE_ACCOUNT,AZURE_VIRTUAL_MACHINE
```

Restricting user IDs to single sessions

To restrict a user ID from signing on multiple times using different browsers on the same machine or the same browser on different machines, modify the portal.properties file. The last browser session with the user ID to login will have access to the portal. Other browser sessions with same user ID will be logged out.

- Use the following:

```
portal.security.allowUserToLoginMultipleTimes=false
```

Customizing date format in the report scope selector

To customize the date format displayed in the report scope selector for all Portal users, you can modify the portal.properties file. For example, you can set the date to display: dd/MM/yyyy or MM/dd/yyyy

- Use the following:

```
#Formatters that define specific presentations of numbers and dates
formatter.decimalPlaces=2
fileSize.base2=true
formatter.number=###,###,##0
formatter.date=MMM dd, yyyy hh:mm:ssa
formatter.dateZone=MMM dd, yyyy hh:mm:ssa z
formatter.yearMonth=MMM dd
formatter.groupByDate=MMM dd
formatter.designerDate=MM/dd/yyyy
formatter.currency=$ ###,###,##0.00
```

Customizing the maximum number of lines for exported reports

When you export or email a large report, IT Analytics limits the maximum number of lines to 20,000. The report truncates when that value is exceeded. The report

can still be exported or emailed, but will contain a message that the report has been truncated.

- Use the following:

  ```
  portal.report.maxRowsExported=<enter new limit value here>
  ```

  Where the <new limit value> is the number of rows greater than 20,000 that your report export requires. For example, if your report has 36,000 rows enter a number greater than 36000. Note that the new limit value cannot contain commas or decimal points. Keep in mind that Portal server performance can degrade considerably for very large reports. For very large reports, you may want to segment the scope into multiple reports.

Customizing the total label display in tabular reports

To customize the minimum number of records needed to display the Total label in a report, you can modify the portal.properties file. The default value is 10.

- Use the following:

  ```
  portal.rowCountDisplayMinimum = <enter numeric value>
  ```

Customizing the host management page size

In the Portal, the Host Management page (Admin > Advanced > Host Management) displays 200 rows by default. You can change the default value by modifying the portal.properties file. System performance will be impacted if you increase the number of rows past the 200 value.

- Use the following:

  ```
  portal.hostManagementPageSize=xxxx
  ```

Customizing the path and directory for File Analytics database

You can customize the location of the File Analytics database. The default paths are:

- Linux: `/opt/aptare/fa`
- Windows: `C:\opt\aptare\fa`

Use the following to revise the path:

```
fa.root=<preferred folder location>
```

For example:

Linux:

```
fa.root=/opt/aptare/fa_db
```

Windows:

```
fa.root= D:\opt\aptare\fa_db (forward slash recommended)
```

Note: Specified preferred folder location should have a folder named **raw** . If the **raw** folder is not specified, IT Analytics will display an error message.

For example: After overriding the preferred location `D:/opt/aptare/fa_db`, the new preferred folder location should be as follows:

- Windows: `fa.root= D:\opt\aptare\fa_db/raw`
- Linux: `fa.root=/opt/aptare/fa_db/raw`

Configuring badge expiration

Configure the expiration of NEW badges in the Home section of the Reports tab. By default, NEW badges will no longer display after 14 days.

- Use the following:

  ```
  cloudTemplateNewBadgeExpireInDays = 14
  ```

Configuring the maximum cache size in memory

The cache can retain up to 0.5 GB of reporting data and if it reaches capacity, it frees up space for new reports by purging the data for the least frequently used reports.

- Use the following:

  ```
  portal.reports.cache.maxSizeInMemory
  ```

  The unit of measure for the cache maxSizeInMemory value is bytes.
  Example: **portal.reports.cache.maxSizeInMemory=536870912**

Configuring the cache time for reports

The cache retains reporting data and if it reaches capacity, it frees up space for new reports by purging the data for the least frequently used reports. Purging also occurs when a cached report is more than 24 hours old.

1. Use the following:

   ```
   portal.reports.cache.timeOut
   ```

   The unit of measure for the cache timeOut value is seconds.

   Example: **portal.reports.cache.timeOut=86400**

Configure SMTP Authentication

You can send email notifications to designated recipients when the scheduled reports are generated and alerts are raised in IT Analytics by configuring SMTP authentication.

You can configure SMTP authentication from the **Admin** tab of the portal. To perform the email configuration for SMTP authentication, go to **Admin** tab > **Advanced** > **Email Configuration** and add details as described in the table below.

Note that **SMTP password** and **Repeat SMTP password** fields are displayed only if you set the **SMTP authentication type** as **Basic Authentication**.

Table 12-3 Email configuration for SMTP authentication

Field name	Description
Email SMTP host IP address	IP address of a functional email server that will be used for authentication.
Email enable Transport Layer Security (TLS)	Enables the SMTP transport layer security (TLS) encryption and authentication protocol. The data channel used to communicate between the client and server will be encrypted.
SMTP Port	SMTP communication port number.
SMTP authentication type	Authentication method used by the client to prove its identity.
SMTP User	User name to be used for authentication on the above-mentioned email server.
SMTP password	The password associated with the SMTP User name.

Table 12-3 Email configuration for SMTP authentication *(continued)*

Field name	Description
Repeat SMTP Password	Repeat the secret associated with the SMTP User name password supplied above.
Email from address	Email address to be used for reply back.
Email from name	The name associated with the email address to be used for reply back.
Email debug mode	Select to enable Cohesity to address mail transmission issues.

Chapter 13

Webhook Configuration

This chapter includes the following topics:

Overview

Webhook is a notification mechanism that pushes event data in real time. It is more prompt in notifying events compared to polling-based or log-based mechanisms.

In IT Analytics, you can add a Webhook configuration using its name, URL, and API key on the portal and configure it as a notification option within your alert policy, in addition to the other notification options. When a Webhook detects an event, it immediately sends an HTTP request with structured data to its host to initiate a notification. The mechanism hence is very useful for sending alerts.

Add a Webhook

To get real-time notifications of system events on the IT Analytics Portal, you must add a Webhook configuration. Before you add a Webhook on the portal, make sure you have a Webhook created on its server and get its URL and API key.

To add a Webhook configuration:

1. Login to the IT Analytics Portal and go to **Admin** > **Advanced** > **Webhook**.
2. Click **Add** and enter the following details on the pop-up:
   - **Domain**: Select the domain in which you want to add the Webhook.
   - **Webhook Name**: Enter a name of your choice.
   - **Webhook URL**: Enter the Webhook URL.

- **API Key for Webhook**: Enter the Webhook API key.
- **Allow insecure connection**: Select if the HTTPS URL provides an untrusted certificate.
- **Template library**: Click to view and select a template.
  The slected template appears in **Payload Template** area. You can either edit the existing template or add a new one. The templates are in `.json` format. The keys within the template are always static strings whereas the values can be static or dynamic. The dynamic values are of the format {{data.fieldName}}, {{data.fieldName1.fieldName2}}, and so on. There can be any level of nesting.
- Click **Validate template**.
- Click **Test Transformation** to transform the template into test event message which will be shown in the **Preview text** area.

3 Click **Test** to check whether event notifications are sent to the Webhook host.

4 Click **OK** to save the configuration.

The configuration is saved to the Webhook page.

Manage Webhook Configurations

Every Webhook configuration is displayed on the Webhook page. You can select any configuration and **Edit**, **Delete**, **Refresh**, **Enable** and **Disable** a configuration.

Chapter 14

Performance profile schedule customization

This chapter includes the following topics:

Overview

Array Performance Profiling enables you to monitor performance over time and to compare your enterprise-specific performance with the performance found in a broader community. You can customize the time of day when your environment's profiling job will run.

Customize the performance profile schedule

To customize the time period for profiling the collected performance data, take the following steps.

1. On the Portal server, go to the database procedures directory.

   Windows: `C:\opt\oracle\database\stored_procedures\srm`

   Linux: `/opt/aptare/database/stored_procedures/srm`

2. Edit the script: **setup_srm_jobs.plb**.
3. Note the parameters shown in red and modify them accordingly.

```
jobName := dba_package.getSchedulerJobName('recalIntPerformanceProfile');

    IF (jobName IS NOT NULL AND LOWER(jobName) <> LOWER('recalIntPerformanceProfile'))
```

```
THEN
        DBMS_OUTPUT.PUT_LINE('recalIntPerformanceProfile exists with default name
'|| jobName ||' hence will be removed and recreated.');
        DBMS_SCHEDULER.DROP_JOB(job_name => jobName);
        jobName := NULL;
    END IF;

    IF jobName IS NULL THEN
        DBMS_SCHEDULER.CREATE_JOB(
            job_name         => 'recalIntPerformanceProfile',
            job_type         => 'PLSQL_BLOCK',
            job_action       => 'srm_array_perf_report_pkg.recalIntPerformanceProfile
(dateRangeType(null,null,null,SYSDATE-2/24, SYSDATE, null, 0));',  -- What to run
            start_date       => SYSDATE + (3/24),
-- First run is 3 hour after upgrade or install
            repeat_interval  => 'TRUNC(SYSDATE+1,''DD'') + (10/24)',
-- Next run is 10am each subsequent day
            enabled          => TRUE);
    ELSE
       DBMS_OUTPUT.PUT_LINE('recalIntPerformanceProfile exists and will be altered
with updated version.');
       DBMS_SCHEDULER.SET_ATTRIBUTE(
            name         => jobName,
            attribute    => 'job_type',
            value        => 'PLSQL_BLOCK'
            );
       DBMS_SCHEDULER.SET_ATTRIBUTE(
            name         => jobName,
            attribute    => 'job_action',
            value        => 'srm_array_perf_report_pkg.recalIntPerformanceProfile
(dateRangeType(null,null,null,SYSDATE-2/24, SYSDATE, null, 0));'  -- What to run
            );
       DBMS_SCHEDULER.SET_ATTRIBUTE(
            name         => jobName,
            attribute    => 'repeat_interval',
            value        => 'TRUNC(SYSDATE+1,''DD'') + (10/24)'
-- Next run is 10am each subsequent day
            );
    END IF;

    DBMS_OUTPUT.put_line('recalIntPerformanceProfile set to run on daily at 10am');
```

- Three hours after a Portal Installation or Upgrade, this job runs for the first time. See the parameter: **(SYSDATE + (3/24))**
- After the first run, this job will run at 10:00 a.m. every day. See the parameter: **(TRUNC(SYSDATE+1, "DD ") + (10/24)**
- This Performance Profiler will calculate the last two hours of statistics. See the parameter: **SYSDATE-2/24**

4. Execute the following command to activate your new schedule:

```
su - aptare
sqlplus portal/portal@//localhost:1521/scdb @setup_srm_jobs.plb
```

Chapter 15

LDAP and SSO authentication for Portal access

This chapter includes the following topics:

Overview

IT Analytics supports user authentication and authorization using:

- Active Directory (AD) or Lightweight Directory Access Protocol (LDAP): IT Analytics supports user authentication and optionally supports authorization using Active Directory (AD) or Lightweight Directory Access Protocol (LDAP).
- Single Sign On (SSO) for a standard unified login: IT Analytics supports Single Sign On (SSO) for a standard unified login. User authentication is performed through an external Identity Management Server allowing for an increased level of security for user passwords and identity details.

IT Analytics can support both authentication types simultaneously or individually as required. User's login experience changes based on the authentication type set for the user or the user group from the Portal.

Portal login options

Based on the options enabled for SSO and LDAP from the Portal, the following options appear on the login screen:

- For Active Directory (AD) or Lightweight Directory Access Protocol (LDAP), the login page displays **LDAP** and **Local** as options in the **Authentication Type** drop-down list.
- For SSO authentication, the login page displays an additional login option of **SIGN IN WITH SINGLE SIGN ON**. Users can click **SIGN IN WITH SINGLE SIGN ON** to login to the Portal.
- For local authentication, the login page displays an option for local authentication. Users can select **Local authentication** to login using local credentials.

Note: If the Portal was upgraded from a lower version, you may have to clear the browser cache for the authentication type and SSO options to appear on the login screen.

Configure authentication for Portal access

Configuration of AD/LDAP authentication and authorization is driven through the LDAP configuration screen.

The prerequisites and steps to enable user authentication are described under the following sections:

- See “Configure AD/LDAP” on page 182.
- See “Configure single sign-on (SSO)” on page 196.
- See “Enable local authentication” on page 201.

Active directory tools

Several tools are available for identifying the Active Directory details. If these tools are not present on the Active Directory server, download them from the Microsoft web site and use the following links to access the documentation.

- Remote Server Administration Tools (RSAT) for Windows
- Microsoft’s Active Directory Interface Editor (adsiedit.msc)

Using LDP to find the base DN

Use the following procedure to search the Active Directory hierarchy.

1. Execute **ldp.exe** to log in to the Active Directory server.

2. Choose **Connection -> Connect** and enter the Server and Port number.

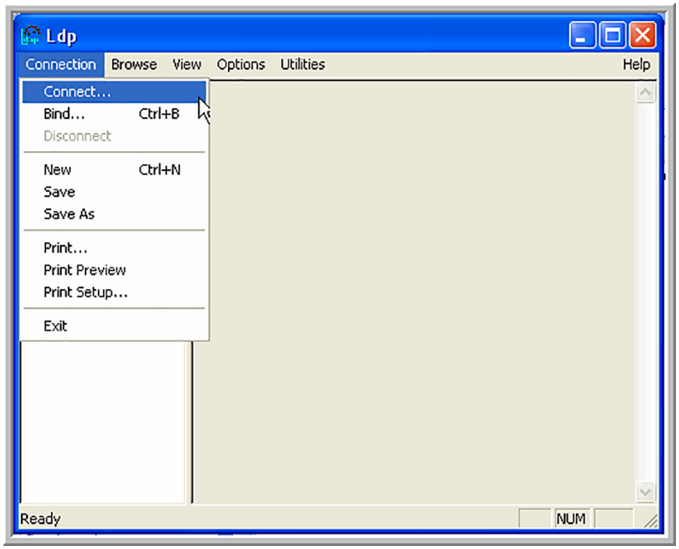

3. Choose **Connection -> Bind** and enter the Administrator for the User ID and then the password to authenticate the user access.
4. Choose **View -> Tree** to browse the Active Directory tree.
5. The Tree View window expects a **BaseDN** entry. The Tree View displays a tree hierarchy with the settings of the users under the Base DN. Most environments have Exchange Objects located in:

   ```
   CN=Services, CN=Configuration, DC=<yourdomain>, DC=com
   ```

   Alternatively, try one of the following:

   - **DC=<yourdomain>, DC=com**
   - **DC=<yourdomain>, DC=local**
   - **CN=Users, DC=<yourdomain>, DC=com**

   Example

   If your domain is **support.aptare.com**, your Base DN would be:

   ```
   DC=support, DC=aptare, DC=com
   ```

Using LDP to search active directory

To confirm the Base DN, use the LDP Search option.

1. In LDP, choose **Browse -> Search**.
2. Use the **Sub-tree Scope** for all searches.
3. Leave the BASE DN entry untouched and enter the following **Filter**:

   ```
   (&(objectClass=msExchExchangeServer)(cn=<serverShortName>))
   ```

   where <serverShortName> is the name before the dot (.) of a fully qualified domain name

   The filtered attributes of interest are: legacyExchangeDN and serialNumber
4. If the filter in the previous step does not result in what you need, try the following **Filter**:

   ```
   (&(objectClass=msExchExchangeServer)(cn=<serverName>))
   ```

   where <serverName> is the fully qualified domain name of the server
5. A Group search will display the DN.

   ```
   (objectClass=msExchStorageGroup)
   ```

6. A Stores search will display the DN.

   ```
   (objectClass=msExchMDB)
   ```

Configure AD/LDAP

IT Analytics supports user authentication and optionally supports authorization using Active Directory (AD) or Lightweight Directory Access Protocol (LDAP). Configuration of AD/LDAP authentication and authorization is driven through the configuration in **Admin** > **Authentication** > **LDAP**.

This section covers the configuration steps for the following scenarios:

- AD/LDAP configuration for authentication - describes the procedure to configure AD/LDAP for user authentication only.
- AD/LDAP Configuration for authentication and authorization - describes the procedure to configure AD/LDAP for user authentication and authorization.
- Migrate portal users when AD/LDAP authentication is configured - describes the configuration required to authenticate using AD/LDAP for users previously using database for authentication.
- Migrate portal users with LDAP authentication and authorization configured - describes the configuration required for authentication and authorization using AD/LDAP for users previously using database for authentication.

AD/LDAP configuration for authentication

To configure AD/LDAP only for user authentication, the portal administrator must create users in the IT Analytics Portal similar to those present in AD/LDAP database.

1. Login to the IT Analytics Portal as a super user and go to **Admin** > **Users**.
2. Create a new user with the same user name present in AD/LDAP.
3. Assign appropriate privileges to the new user.
4. If the portal server fails to resolve the AD/LDAP hostname, add the hostname entry in the `/etc/hosts` file on the portal server.

   ```
   <AD_IP_Address> <AD_Domain_Name>
   ```

   For example:

   ```
   192.168.2.90 ad.gold
   ```

5. Go to **Admin** > **Authentication** > **LDAP**.

6 Enable authentication, authorization, and enter the configuration as suggested in the table below:

Field name	Description
Enabled	Select to enable AD/LDAP authentication
Authorisation	You can skip enabling this as you are only enabling authentication. When selected, Portal authorizes the user against AD groups. At least one AD group of which the new user is a member must be configured as a User Group in the Portal. **Note:** If the AD group is not mapped with the User Group in the Portal, then authentication fails during login with the error: "No user group mapping present for external LDAP user."
LDAP Domain Name	This field is deprecated. If this field appears in your Portal, enter **LDAP** as its value. Example: `LDAP Domain Name: LDAP`

LDAP URL	Set to the host and port of your AD. Note that this URL value has a prefix **ldap:**. If using SSL, change the prefix to **ldaps:**.
	If you are using Active Directory for your external LDAP configuration, you may want to use the global catalog port of 3268 instead of port 389.
	If using SSL, you may want to use the secure global catalog port of 3269 or 636 for standard LDAPs.
	Example:
	`ldap://example.company.com:389`
	or
	`ldaps://example.company.com:636`

Field name	Description
Search Base	Set the location from where the search will be performed to locate users in the authentication directory. Often referred to as the Active Directory (AD) Search Base, this is the starting point in the Active Directory tree for searching for LD AP users. This search base, in LDAP distinguished name format, contains a fully qualified domain name. IT Analytics supports only one Search Base. Example: `dc=example,dc=company,dc=com`
DN	Set to the ID of a user or a service account that has permission to search the search base. This user or service account must be able to search all LDAP directory servers. IT Analytics requires a user or a service account name that has privileges to search under the Base DN (Distinguished Name) within the Active Directory structure. This must be an account that has administrative privileges, typically an Administrator. It can be the Administrator account that was created when Active Directory was installed, or it can be an account that was created and either was given administrative privileges or was placed into a group with administrative privileges. If you use Active Directory, specify this setting because Active Directory services do not allow anonymous binds. Microsoft Active Directory requires the username and password of a user that has enough privileges to search the LDAP directory. `CN=Admin,CN=Users,DC=example,DC=company,DC=com` In this case the `Admin` user in the common name `Users` has permissions to search the search bases: `DC=example,DC=company,DC=com`. **Note:** Special characters #, >, <, ;, (,), and = are supported, but / and \ are not supported in DN and CN.
DN Password	Set to the password of the user who is used in the **DN** field.

Field name	Description
Certificate	If LDAP is configured with SSL, you need an LDAP server certificate for a secure connection with the LDAP server. You can get this certificate from the LDAP administrator or an Active Directory administrator for upload. **Note:** If you have a certificate chain, you must isolate the LDAP server certificate from the chain and upload it separately as a single certificate. Also, the certificate is expected to contain plain text anchor lines like **BEGIN CERTIFICATE** and **END CERTIFICATE**.
Login Attribute	Enter the login attribute used for authentication. This is the attribute name in Active Directory that specifies the username, such as uid or sAMAccountName. Example: `sAMAccountName`
New User Domain	Enter the domain name on which the user needs to be authorized. Get the domain name from **Admin** > **Domains** > **Domain Name**. Example: `example.company.com`
Disable User Attribute Name	Enter the value of the AD attribute that indicates whether the user is active or inactive. During Portal authentication via AD, the REST API uses the AD attribute assigned to this property to check whether the user is still an active AD user. For example, if `ad.user.active` is the AD attribute that indicates whether a user is active or disabled, then `ad.user.active` must be assigned as the value of this field.

Field name	Description
Disable User Attribute Value	Enter the same value as that of the AD attribute (specified in **Disable User Attribute Name**, which indicates the AD user is disabled. For example: If `ad.user.active` is the attribute for user status in AD, it may have several values such as `live, inactive, joined`, and so on. If the value `inactive` indicates the user is disabled in AD, then inactive must be set as value for this property (`inactive`). REST API matches this value with the value of the AD attribute specified in this field. If the values match, the user is disabled on the IT Analytics Portal. **Note:** A Portal super user must explicitly activate the user that was deactivated in both AD and Portal in the past but is again activated only in AD. A Portal administrator with adequate privileges can also activate such a user. Without user activation, Portal access will be restricted.

7 Click **Test Connection**. Make the required changes if the test fails.

8 Click **Save**.

Enabling LDAP authentication is complete.

Note: If you are unable to save the configuration, check if the JDK truststore password was changed before the last upgrade and ensure the updated password is assigned to the `portal.jdk.trustStore.password` parameter from **Admin** > **System Configuration** > **Custom** page of the Portal. The JDK truststore locations for Windows and Linux are `<portal_installation_path>\jdk\lib\security\cacerts` and `/usr/java/lib/security/cacerts` respectively.

---

9 To change the existing superuser LDAP_ID to map to the AD username, update the existing record on the Oracle database server.

For example: If the login attribute is user_name and actual value is Admin, update the existing record as below:

```
# sqlplus portal/<portal_password>@scdb
# UPDATE ptl_user SET ldap_id = 'Admin' WHERE user_id = 100000;
# commit;
```

Use this updated username to login to the external directory, instead of **aptare**. Since the user account **aptare** (user_id=100), is an internal bootstrap user, it is required to maintain referential integrity among database tables and therefore you must avoid using **aptare** for external LDAP integration.

Note: The user_id = 100000 is always the default user_id for the super user account.

---

10 Login to the portal using any user name common across AD/LDAP and the IT Analytics Portal.

If the Portal was upgraded from a lower version, you may have to clear the browser cache for the authentication type and SSO options to appear on the login screen.

AD/LDAP Configuration for authentication and authorization

To configure AD/LDAP for user authentication as well as authorization, Portal Administrator must create at least one User Group in portal which is also present in AD/LDAP as a UserGroup.

1 Login to Portal as a SuperUser, navigate to **Admin** > **User Groups**.

2 Create a new user group with the same group name present in AD/LDAP. Only members of this user group can access the portal when **LDAP Authroization** is enabled on the Portal.

3 Assign appropriate privileges to the newly created user group. See *Setting user group privileges* section in the *User Guide*.

4 Record the domain name where new users will be created.

To find domain name, navigate to **Admin** > **Domains** > **Domain Name**.

You are required to specify this in the **LDAP Domain Name** field when you enable LDAP authentication.

5 Go to **Admin** > **Authentication** > **LDAP**.

6 Enable authentication, authorization, and enter the configuration as suggested in the table below:

Field name	Description
Enabled	Select to enable AD/LDAP authentication
Authorisation	Select to enable AD/LDAP authorisation When selected, Portal authorizes the user against AD groups. At least one AD group of which the new user is a member must be configured as a User Group in the Portal. **Note:** If the AD group is not mapped with the User Group in the Portal, then authentication fails during login with the error: "No user group mapping present for external LDAP user."
LDAP Domain Name	Enter the Portal domain name where the new user gets created. It is used provided ldap.authorization is set to true. To find domain name in portal, go to **Admin** > **Domains** > **Domain Name**. Example: `example.company.com`

LDAP URL	Set to the host and port of your AD. Note that this URL value has a prefix **ldap:**. If using SSL, change the prefix to **ldaps**.
	If you are using Active Directory for your external LDAP configuration, you may want to use the global catalog port of 3268 instead of port 389.
	If using SSL, you may want to use the secure global catalog port of 3269 or 636 for standard LDAPs.
	Example:
	`ldap://example.company.com:389`
	or
	`ldaps://example.company.com:636`

Field name	Description
Search Base	Set the location from where the search will be performed to locate users in the authentication directory. Often referred to as the Active Directory (AD) Search Base, this is the starting point in the Active Directory tree for searching for LD AP users. This search base, in LDAP distinguished name format, contains a fully qualified domain name. IT Analytics supports only one Search Base. Example: `example, company, com`
DN	Set to the ID of a user who has permission to search the SEARCHBASE. This user must be able to search all LDAP directory servers. IT Analytics requires a user that has privileges to search under the Base DN (Distinguished Name) within the Active Directory structure. This must be an account that has administrative privileges, typically an Administrator. It can be the Administrator account that was created when Active Directory was installed, or it can be an account that was created, and either was given administrative privileges or was placed into a group with administrative privileges. If you use Active Directory, specify this setting because Active Directory services do not allow anonymous binds. Microsoft Active Directory requires the username and password of a user that has enough privileges to search the LDAP directory. Example: `CN=Admin,CN=Users,DC=example,DC=company,DC=com` **Note:** Special characters #, >, <, ;, (,), and = are supported, but / and \ are not supported in DN and CN.
DN Password	Set to the password of the user who is used in the **DN** field.
Certificate	Navigate to the keystore path location and select the AD certificate.

Field name	Description
Login Attribute	Enter the login attribute used for authentication. This is the attribute name in Active Directory that specifies the username, such as uid or sAMAccountName. Example: `sAMAccountName`
New User Domain	Enter the Portal domain name where new user gets created. It is used only if **Authorisation** is enabled. To find domain name in portal, navigate to **Admin** > **Domains** > **Domain Name**. Example: `example.company.com`
Disable User Attribute Name	Enter the value of the AD attribute that indicates whether the user is active or inactive. During Portal authentication via AD, the REST API uses the AD attribute assigned to this property to check whether the user is still an active AD user. For example, if `ad.user.active` is the AD attribute that indicates whether a user is active or disabled, then **ad.user.active** must be assigned as the value of this field. **Note:** If a user account is disabled in LDAP, subsequent API calls for that user will result in the user being disabled in IT Analytics.

Field name	Description
Disable User Attribute Value	Enter the same value as that of the AD attribute (specified in **Disable User Attribute Name**, which indicates the AD user is disabled. For example: If `ad.user.active` is the attribute for user status in AD, it may have several values such as **live**, **inactive**, **joined**, and so on. If the value **inactive** indicates the user is disabled in AD, then **inactive** must be set as value for this field. REST API matches this value with the value of the AD attribute specified in this field. If the values match, the user is disabled on the IT Analytics Portal. **Note:** A Portal super user must explicitly activate the user that was deactivated in both AD and Portal in the past but is again activated only in AD. A Portal administrator with adequate privileges can also activate such a user. Without user activation, Portal access will be restricted.

7 Click **Test Connection**. Make the required changes if the test fails.

8 Click **Save**.

Enabling LDAP authentication and authorization is complete.

Note: If you are unable to save the configuration, check if the JDK truststore password was changed before the last upgrade and ensure the updated password is assigned to the `portal.jdk.trustStore.password` parameter from **Admin** > **System Configuration** > **Custom** page of the Portal. The JDK truststore locations for Windows and Linux are `<portal_installation_path>\jdk\lib\security\cacerts` and `/usr/java/lib/security/cacerts` respectively.

---

9 To change the existing superuser LDAP_ID to map to the AD username, update the existing record on the Oracle database server.

For example: If the login attribute is user_name and actual value is Admin, update the existing record as below:

```
# sqlplus portal/<portal_password>@scdb
# UPDATE ptl_user SET ldap_id = 'Admin' WHERE user_id = 100000;
# commit;
```

Use this updated username to login to the external directory, instead of **aptare**. Since the user account **aptare** (user_id=100), is an internal bootstrap user, it is required to maintain referential integrity among database tables and therefore you must avoid using **aptare** for external LDAP integration.

Note: The user_id = 100000 is always the default user_id for the super user account.

10 Login to the portal using any user present in the Active Directory and part of the group created in step 2.

If the Portal was upgraded from a lower version, you may have to clear the browser cache for the authentication type and SSO options to appear on the login screen.

Note that to automatically create a user in the portal, these attributes must be set for each user in AD/LDAP:

- givenName: Mandatory. It is used as the first name of the user.
- telephoneNumber: Optional
- mobile: Optional
- mail: Mandatory

Note: If for any reason the LDAP configuration is disabled from the portal, the portal administrator must set the password for all the AD/LDAP users in portal.

Migrate portal users when AD/LDAP authentication is configured

To migrate portal users from using portal database for authentication to using AD/LDAP, you must ensure the LDAP_ID of the portal user present in the PTL_USER table matches the ldap.loginAttribute value in AD/LDAP. If the values fail to match, you must update the user name in the portal database.

For example: Assume Joe has joe.smith as LDAP_ID in the portal database. If ldap.loginAttribute is set to sAMAccountName on the LDAP screen, and on AD/LDAP, the value of sAMAccountName must be joe.smith for the user to login successfully. If the value of sAMAccountName is other than joe.smith, you must change the LDAP_ID of the user in the PTL_USER table of the portal database to joe.smith to match the user name present in AD/LDAP.

To update the LDAP_ID in the portal database:

1. Login to the Oracle database server of the IT Analytics Portal.
   - On Linux: Login as **aptare** user. If you have already logged in as root, use `su -aptare`.
   - On Windows: Login using user credentials of the ORA_DBA group.
2. Identify the USER_ID that corresponds to the LDAP_ID using:

   ```
   select LDAP_ID, USER_ID from PTL_USER where LDAP_ID='<user_name>';
   ```

   For example:

   ```
   select LDAP_ID, USER_ID from PTL_USER where LDAP_ID='joe.smith';
   ```

3. Update the LDAP_ID with the user ID obtained from the above step.

   ```
   update PTL_USER set LDAP_ID = '<user_name>' where USER_ID =
   <user_ID_from_previous_step>;
   ```

   For example:

   ```
   update PTL_USER set LDAP_ID = 'smith.joe'  where USER_ID =
   <user_ID_from_previous_step>;
   ```

4. Repeat steps 2 and 3 for all the users having a mismatch in their IDs.

Note: Use LDAP_ID mentioned in step 3 (above) to login to AD/LDAP. Avoid using user name **aptare** as aptare (user_id=100) is an internal bootstrap user required to maintain referential integrity amongst the database tables. Hence you must not change aptare or use it for external LDAP integration.

Migrate portal users with LDAP authentication and authorization configured

To enable authentication as well as authorization of portal users using AD/LDAP, the portal user's LDAP_ID present in the PTL_USER table must match with the value of the ldap.loginAttribute for the user in AD/ldap. If the values fail to match,

you must update the user ID in the portal database. Also for user authorization, you must create user groups in the portal which match with at least one AD group that includes the user name.

For example: Assume Joe has joe.smith as LDAP_ID in the portal database. If ldap.loginAttribute is set to sAMAccountName on the LDAP screen and on AD/LDAP, the value of sAMAccountName must be joe.smith for the user to login successfully. If the value of sAMAccountName is other than joe.smith, you must change the LDAP_ID of the user in the PTL_USER table of the portal database to joe.smith to match the user name present in AD/LDAP.

To update the LDAP_ID in the portal database:

1. Login to the IT Analytics Portal before configuraing AD for authentication and create the required user groups with appropriate privileges. The user group name must match with that of the AD/LDAP group name. This user group is used to authorize the user once AD/LDAP is configured.
2. Login to the Oracle database server of the IT Analytics Portal.
   - On Linux: Login as **aptare** user. If you have already logged in as root, use `su -aptare`.
   - On Windows: Login using user credentials of the ORA_DBA group.
3. Identify the USER_ID that corresponds to the LDAP_ID using:

   ```
   select LDAP_ID, USER_ID from PTL_USER where LDAP_ID='<user_name>';
   ```

   For example:

   ```
   select LDAP_ID, USER_ID from PTL_USER where LDAP_ID='joe.smith';
   ```

4. Update the LDAP_ID with the user ID obtained from the above step.

   ```
   update PTL_USER set LDAP_ID = '<user_name>' where USER_ID =
   <user_ID_from_previous_step>;
   ```

   For example:

   ```
   update PTL_USER set LDAP_ID = 'smith.joe'  where USER_ID =
   <user_ID_from_previous_step>;
   ```

5. Repeat steps 3 and 4 for all the users having mismatch in their user IDs.

Note: Use LDAP_ID mentioned in step 4 (above) to login to AD/LDAP. Avoid using user name **aptare** as aptare (user_id=100) is an internal bootstrap user required to maintain referential integrity amongst the database tables. Hence you must not change aptare or use it for external LDAP integration.

Configure single sign-on (SSO)

IT Analytics supports Single Sign On (SSO) for a standard unified login. User authentication is performed through an external Identity Management Server allowing for an increased level of security for user passwords and identity details. To use Single Sign On, setup is required on the IT Analytics Portal, an external Identity Provider (IDP) and an external LDAP directory.

Single sign-on (SSO) prerequisites

- IT AnalyticsPortal must be SSL enabled (https protocol) using SSL certificates with the following properties:
  - Signature algorithm name: SHA256 with RSA
  - Subject public key algorithm: 2048-bit RSA key
- An external Identity Provider (IDP) that supports SAML 2.0
- SSL certificate must be added to the Portal Keystore using the Keystore Utility (deployCert)

Setting up the external Identity Provider (IDP) server

For the IDP to communicate with the IT Analytics Portal, an LDAP directory must be configured on the external server for user management. Certain attributes must be populated for each user who will be logging into the Portal. Users must also belong to at least one group.

Users and groups in the external LDAP directory

When using an external authentication service, there are two areas that require setup for a synchronization between the two systems to occur:

- User profiles in the external directory must have specific attributes set
- Group names in the external directory must match User Group names in the IT Analytics Portal for privilege inheritance

Set the following attributes for each user in the external LDAP directory. For each attribute, the properties **name** and **friendlyName** must be present and have values

populated. These attributes must be exposed by both the external LDAP directory and the IDP server. The names of attributes are as follows:

- displayName: <first_name> <last_name> For example Jane Smith
- email: email address
- mobile: cell phone or mobile number
- telephoneNumber: work phone or home phone number
- sAMAccountName: the unique user name that is used as a login
- memberOf: List of group names to which the user belongs, supporting with or without domain prefixed for Azure IDP. This attribute requires customization for a Microsoft Azure IDP. It is recommended to set **Groups Assigned to the application** instead of **All groups** or **Security groups** for "memberOf" attribute. Click here for more details.

  The memberOf attribute must be in the below supported formats:

  - DOMAIN_NAME\userGroupName
  - CN=userGroupName,CN=Users,DC=aptareadfs,DC=com (for non-AZURE IDPs)

Before an external user can use SSO to log into the Portal, they must belong to one external directory group that also exists as a User Group in the IT Analytics Portal. If the setup criteria is met, when the user logs into the Portal for the first time, their user profile will be synchronized from the external directory. They will also inherit all privileges assigned to the User Group.

Registering with the IDP server

The registration process occurs by exchanging metadata XML files between the IT Analytics Portal and the IDP server. On the Portal side, once SSO is configured and the Portal Tomcat service restarted, you can download the metadata XML file and provide it to the IDP server. This file contains the SSL certificate and identifies the IT Analytics as a service provider for SSO. A similar metadata XML file must be downloaded from the IDP server and provided to the Portal.

Activate single Sign-on (SSO) in the portal

Set up SSO properties by navigating to **Admin** > **Authentication** > **SSO**. Access to **Authentication** settings are restricted to users with a Super User role. In addition to the SSO properties displayed in the **Authentication**, a valid SSL certificate must added to the portal keystore. Field descriptions for each of the properties are available through rollover text.

You must have downloaded the SAML metadata XML file from the external Identity Provider (IdP).

Note: If you are unable to save the configuration, check if the JDK truststore password was changed before the last upgrade and ensure the updated password is assigned to the `portal.jdk.trustStore.password` parameter from **Admin** > **System Configuration** > **Custom** page of the Portal. The JDK truststore locations for Windows and Linux are `<portal_installation_path>\jdk\lib\security\cacerts` and `/usr/java/lib/security/cacerts` respectively.

To activate single sign on in the Portal

1. Navigate to **Admin** > **Authentication** > **SSO**.
2. On the **SSO using SAML** page, select **Enable SSO** to activate SSO.

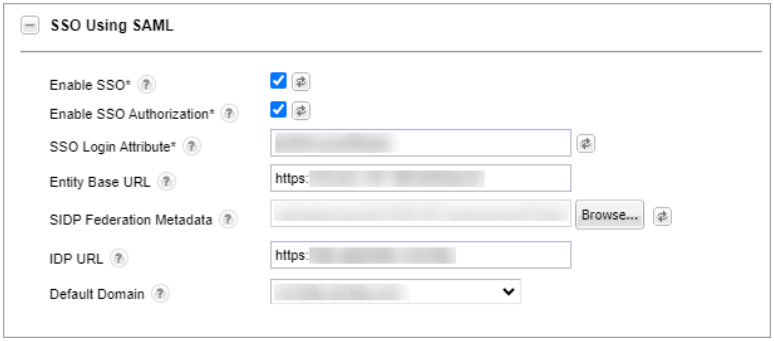

3. Select **Enable SSO Authorization** if required. Selecting this option activates group memberships from the external LDAP.

   The group name derived from the memberOf list must match the name of the User Group.
4. Enter AD attribute name in **SSO Login Attribute** - the value of this attribute serves as an access credential for users during the SSO login.

   Note: After activating SSO, a user performing a local login will have to use this attribute value as a credential to access the Portal.

   ---
5. Enter the URL for the portal application. This should be an https URL with a trailing '/' at the end of URL.

6. Browse to the metadata.xml file that was downloaded from the external Identity Provider (IdP).
7. Enter the URL for the external IdP server. The entityId must match the value of entityId as listed in the IdP metadata XML file.
8. Enter the domain to be assigned to the SSO user when the Portal creates it automatically.
9. Restart the Portal Tomcat service.

   Note: If there are issues with the configuration, your Portal may not restart. There is a utility (resetSSoConfig) available to reset the parameters so the Portal can be restarted.

10. Download the metadata XML file using the following URL:

    https://aptareportal.domain_name/saml/metadata

    If the XML is displayed in the browser, you can select “save target as” to save the XML as a file. This file will be required on the external Identity Provider (IdP) server to identify IT Analytics as a service provider for SSO and complete the registration process.

    Note: The XML file is only available once the configuration settings have been saved and the Portal Tomcat service has been restarted.

11. After the registration process is complete, open the Portal login screen and try to login with Single Sign On. If there are issues with the configuration, your Portal may not restart. Use the resetSSoConfig utility to reset the parameters so the Portal can be restarted.
12. Navigate to **Admin** > **Advanced** > **System Configuration** > **Portal**.

    See “Add a Certificate into the Portal Keystore” on page 264.

SSO troubleshooting and maintenance

The following covers common scenarios and possible solutions.

Table 15-1 Scenarios and its solutions.

Scenario	Solution
SSL Certificate Revisions	If there are changes to the SSL certificates perform the following steps: **1** Use the Keystore Utility (deployCert) to add (replace) the existing SSL certificate. See "Add a Certificate into the Portal Keystore" on page 264. **2** Verify the new SSL certificate has the mandatory prerequisite algorithm settings. See "Single sign-on (SSO) prerequisites " on page 196. **3** Restart the Portal Tomcat service. **4** Log into the Portal after restarting and navigate to **Admin>Advanced>System Configuration>Portal** . **5** In the section **SSO using SAML**, click **Download** in the **SAML Metadata** field to retrieve the metadata XML file. This file contains the content of the new SSL certificate. To re-register, the file must be resent to the external Identity Provider (IDP) server to identify IT Analytics as a service provider for SSO.
Identity Provider Server (IDP) Revisions	If there are changes to the IDP server, the entire SSO configuration and registration process must be redone.
Identity Provider (IDP) Login Screen Not Displaying	■ For this issue, check if both the IDP server and the IT Analytics Portal are able to resolve host names. ■ A second solution is try to ping the IDP server by hostname from the IT Analytics Portal.
Message Security Error	This error is displayed when SSL certificates are not as expected with regards to the metadata XML files. Verify the SSL certificates satisfy all the requirements on both the IT Analytics Portal and the IDP server.
Security Provider Not Registered	This error indicates the registration process between the IDP server andIT Analytics Portal was not completed successfully. Verify the exchange of both metadata XML files was done correctly.
IDPSSODescriptor Not Found	This error indicates the Entity Base URL and the path to the IDP server was set incorrectly. Verify there is no '/' at the end of the given path.
Stale Request	■ This error may be caused by using the browser window back button from the IDP login screen. For this issue, clear your browser cache and retry the SSO login. ■ This error may be caused by time zone discrepancies between the IDP server and the IT Analytics Portal.

Table 15-1 Scenarios and its solutions. *(continued)*

Scenario	Solution

Login Issues/Reset Utility	If Single Sign On (SSO) is not properly set up in the **Admin>Advanced>System Configuration**, after restarting you may not be able to log into the Portal. This utility resets the Single Sign On (SSO) parameters to provide Portal access. Run the following scripts from the command prompt:
	Linux
	`cd /opt/aptare/utils`
	`./resetSSOConfig.sh`
	Windows
	`cd C:\opt\aptare\utils`
	`resetSSOConfig.bat`
	Restart the Portal services after running the OS specific script.

Enable local authentication

Local authentication allows users to login using local credentials. As an administrator you can also manage this authentication through the command line, as described in this section below.

To enable local authentication:

1. Go to **Admin** > **Authentication** > **Local**.
2. Select the **Enable Local Authentication** checkbox.
3. Click **Save**.

Manage local authentication from command line

Use `localAuth.sh` (on Linux) or `localAuth.bat` (on Windows) to manage (enable/disable) the local user authentication mechanism from command line and is available under `/opt/aptare/utils` which is localAuth.(sh/bat). The following table explains the various options available to manage the authentication.

Table 15-2 Options available for `localAuth` script

Option	Example	Description
`disable`	Linux: `./localAuth.sh disable` Windows: `localAuth.bat disable`	Immediately disables local authentication.
`enable`	Linux: `./localAuth.sh enable` Windows: `localAuth.bat enable`	Immediately enables local authentication for the user
`status`	Linux: `./localAuth.sh status` Windows: `localAuth.bat status`	Displays the current status of local authentication (disabled or enabled). It is enabled by default.

Chapter 16

Change Oracle database user passwords

This chapter includes the following topics:

Overview

These instructions are for modifying the Oracle database user passwords for access to the IT Analytics database. You can modify the user passwords, but do not modify the user names without the assistance of Professional Services.

Database connection properties

The following table summarizes the portal.properties values for the Oracle users and passwords that are used by the portal.

Table 16-1 Portal properties with description.

Portal Property	Description
db.driver	This value is customized by the Portal installer and should not be modified.

Table 16-1 Portal properties with description. *(continued)*

Portal Property	Description
db.url	This is the address where the IT Analytics database resides. Depends on what was entered during the installation. This may need to be modified when there is a host name change.
db.user	The application DB user ID. Default value is portal and cannot be changed.
db.password db.password.encrypted=	Enter a password to be used with the DB user. The default value is **portal**. The password initially is stored in clear text, but after the restart of the Tomcat Portal services, the password is saved in the encrypted format and the clear text password is removed from portal.properties.
db.connection.max	Use this property to specify the maximum database connections allowed. The default value is 25.
db.connection.min	Use this property to specify the minimum number of database connections that the Portal can have. The default value is 25.
db.connection.expiration	When a Portal report initiates a long-running database query, this value (in minutes) establishes when the report will time out if the query takes too long to complete. The default value is 5.
db.ro_user_password db.ro_user_password.encrypted=	The Oracle database read-only user password for the IT Analytics database tables. The preset value is aptaresoftware123. The password initially is stored in clear text, but after the restart of the Tomcat Portal services, the password is saved in the encrypted format and the clear text password is removed from `portal.properties`.
db.sysdba_user	The Oracle database System DBA for the IT Analytics database tables. The preset value is system.

Modify the Oracle database user passwords

Use the following utilities to modify passwords for the Oracle database users **portal** and **aptare_ro**. These instructions apply only to user **portal** and **aptare_ro**.

- Linux: /opt/aptare/utils/changeDBPassword.sh
- Windows: C:\opt\aptare\utils\changeDBPassword.bat

The new password must not have the following characters:

- Double Quotes "".
- Back Slash '/'.
- Blank space ' '.

- Back Tick `'

Note: If CyberArk feature is enabled, do not modify Oracle database password.

See *"Integrate with CyberArk "* on page 208.

Complete these steps to modify passwords for the Oracle database user. These instructions apply to aptare_ro and portal users.

1. Login with root access on Linux or with admin access on Windows.
2. Stop the portal and agent Tomcat services.
3. Change the user password:

   On Linux:

   ```
   /opt/aptare/utils/changeDBPassword.sh -user <user_name> <password>
   ```

   On Windows:

   ```
   C:\opt\aptare\utils\changeDBPassword.bat -user <user_name>
   <password>
   ```

   This updates the specified user's password in Oracle configuration as well as the properties files like `portal.properties` and `datrarcvrproperties.xml`.
4. Restart the File Analytics services immediately after changing the password.

Modify the Oracle database user passwords for split architecture

The following are the steps to modify the Oracle password in a split architecture for Linux.

Steps to be performed on portal server

1. Log in with the root access.
2. Stop the portal and data receiver (Tomcat) services.

Steps to perform on the Oracle server at the command prompt

- su - aptare
  source <INSTALL_PATH>/aptare/bin/aptare_env.sh
- sqlplus / as sysdba
- alter session set container=scdb;

Note: scdb is the oracle service name.

- alter user aptare_ro identified by **aptare123456789**;

Note: aptare123456789 is the new password.

- alter user portal identified by **portal123456789**;

Note: portal123456789 is the new password.

- commit
- exit

Steps to perform on the portal server

1. Change the user password:

   ```
   /<install path>/aptare/utils/changeDBPassword.sh -user <user_name>
   <password>
   ```

2. The utility will ask if the passwords are updated on the oracle server.

   ```
   Is the password updated on Oracle Server (y/n)?
   ```

   Enter 'y' if the password is updated on the oracle server, as mentioned in step 1, else enter 'n' to exit.

   Note: This updates the specified user's password in the properties files like `portal.properties` and `datrarcvrproperties.xml`.

3. Restart the portal and data receiver (Tomcat) services.

Determine if Oracle is using the default login password

You can determine if the DBMS is using the default password.

Run the following SQL statement to determine if the PORTAL user uses the default password:

```
SELECT * FROM SYS.DBA_USERS_WITH_DEFPWD where username = 'PORTAL';
```

If the row is returned, it means the PORTAL user uses the default password.

See “Modify the Oracle database user passwords” on page 204.

Chapter 17

Integrate with CyberArk

This chapter includes the following topics:

Introduction

CyberArk, password vault application, is designed to discover, secure, rotate, and control access to privileged account passwords used to access systems throughout the enterprise IT environment. This integration enables IT Analytics to automate the process of fetching latest Oracle user accounts passwords from CyberArk, which changes passwords periodically according to the organizational policy.

To facilitate this integration, CyberArk exposes a REST API to fetch the passwords from the Agentless AAM (Central Credential Provider). IT Analytics fetches the latest oracle user account password via this REST API. The following instructions do not cover how to setup CyberArk, but they do call out the information required from the CyberArk setup.

CyberArk setup prerequisites

The following constructs must be in place within the CyberArk application before setting up IT Analytics:

- An application object must be created to uniquely identify IT Analytics.
- In the Password safe, provision the Oracle database user account or accounts for IT Analytics that will be subject to automatic password change.

- Add the CyberArk Application Password Providers and application as members of the Safes where the Oracle user account passwords are stored.

Core IT Analytics Oracle Database User Accounts

IT Analytics uses the following user accounts for the Oracle Database.

Note: At least one account must be configured in CyberArk to enable this feature.

- Portal: This user is created during installation by the IT Analytics application.
- aptare_ro: This read-only user is created during installation by the IT Analytics application.

Once at least one of these accounts are added in CyberArk, they should be configured in IT Analytics application to fetch the latest password from CyberArk.

Setting up the portal to integrate with CyberArk

Setting up the Portal to work with CyberArk requires a properties file with specific values and executing a script to enable the feature. Most values required in the properties file are derived from entries in the CyberArk application. There are two steps to setup the portal; Setting up properties and running the utility.

Set Up the Properties File on the IT Analytics Portal

1 Create a properties file in the tmp directory. For example:

/tmp/dbvaultconnection.properties or C:\tmp\dbvaultconnection.properties

2 Configure the properties file with the following information:

Table 17-1 Properties file configuration

Field Name	Value
vault_vendor_name	CyberArk **Note:** Use only CyberArk as a value in this field.
host	IP Address or hostname of the machine where Agentless AAM(Central Credential Provider) web service is running.
port	port # of Agentless AAM(Central Credential Provider) web service.

Table 17-1 Properties file configuration *(continued)*

Field Name	Value
https	Set this value to true if Agentless AAM(Central Credential Provider) is running as https service otherwise false
schedule_frequency	The value of this field is in hours. Defines the frequency that IT Analytics will poll CyberArk for password updates. This field is optional. The default is every hour.
app_id	Name/ID of the application as defined in CyberArk. For example: IT Analytics.
user_safe_id	Name of the CyberArk Safe, where 'Portal' user account password is stored.
user_password_folder_name	Name of the CyberArk folder, where the 'Portal' user account password is stored. This field is optional.
user_password_object	Name of the 'Portal' user account password object in CyberArk
For each user account repeat the following. This sample describes the aptare_ro user.	
ro_safe_id	Name of the CyberArk Safe, where 'aptare_ro' user account password is stored.
ro_password_folder_name	Name of the CyberArk folder where the 'aptare_ro' user account password is stored. This field is optional.
ro_password_object	Name of the 'aptare_ro' user account password object in CyberArk

Properties file examples

Sample of `dbvaultconnection.properties` containing 'portal' user account configured:

```
vault_vendor_name=CyberArk
host=10.x.x.x
port=443
https=true
schedule_frequency=2
app_id=testappid
```

```
user_safe_id=safe1
user_password_object=portal_account
```

Sample of `dbvaultconnection.properties` containing 'portal' user account configured:

```
vault_vendor_name=CyberArk
host=10.x.x.x
port=443
https=true
schedule_frequency=2
app_id=testappid

user_safe_id=safe1
user_password_object=portal_account

ro_safe_id=safe2
ro_password_object=ro_account
```

Run the Utility to enable the CyberArk Integration.

1. Navigate to the OS-specific utility on the Portal.

   ```
   <APTARE_HOME>/utils/configure-db-vault-connection-info.sh
   ```

   ```
   <APTARE_HOME/utils/configure-db-vault-connection-info.bat
   ```

2. Execute the utility as a root/tomcat user.

   On Linux:

   <APTARE_HOME>/utils/configure-db-vault-connection-info.sh /tmp/dbvaultconnection.properties

   On Windows

   <APTARE_HOME>/utils/configure-db-vault-connection-info.bat

   /tmp/dbvaultconnection.properties

   The utility validates the connection parameter by invoking the REST API for each user configured in the properties file. If the validation is successful, the properties file is copied to the <HOME>/datarcvrconf/passwordvault/ folder. If the validation fails, a message is displayed and the CyberArk integration is not enabled.

3. If required, navigate to the log file to troubleshoot further:

   ```
   <APTARE_HOME>\logs\passwordVaultValidator.log
   ```

Chapter 18

Tuning IT Analytics

This chapter includes the following topics:

Before you begin tuning

You should rarely need to tune any factory-default settings. If you determine that degraded system performance warrants an examination of certain configurations such as memory, take the following steps.

Note: If you encounter any issues following these directions contact Cohesity Support for further guidance.

---

1. Before modifying your configuration, make a copy of all files you plan to edit.
2. Consider tuning to be a process--that is, increase/decrease a number slightly, then monitor system performance. If your modification results in improvement, you may consider additional adjustments later.
3. Whenever you undertake this tuning process, consider the potential negative impact of settings that are either too high or too low, within the resource constraints of your environment.

Tuning the portal database

Note: Only 64-bit Operating Systems are supported for the Oracle Database.

Database Cache and Shared Pool Size Recommendation:

- Comment out the following lines in initscdb.ora:
  - #db_cache_size = 400M
  - #shared_pool_size = 256M
    Windows: `C:\opt\oracle\database\initscdb.ora`
    Linux: `/opt/aptare/oracle/dbs/initscdb.ora`

Memory Recommendation: If your database server has sufficient memory, you may consider making the changes listed below.

- Increase the values for the following fields in initscdb.ora:
  - WithinBullet>pga_aggregate_target from 1000 MB to 1500 MB
  - sga_target from 1228 MB to 2048 MB
    Windows: `C:\opt\oracle\database\initscdb.ora`
    Linux: `/opt/aptare/oracle/dbs/initscdb.ora`

Number of Connections Recommendation: On a Windows server, the number of Oracle connections is specified in the script: **setupservices.bat**. To modify the number of connections, take the steps listed below.

1. Edit/view the script that sets up the number of connections:

   ```
   C:\opt\aptare\utils\setupservices.bat
   ```

2. Note that in this file, the following commands specify the number of connections:

   ```
   C:\opt\oracle\bin\oradim -new -sid scdb -maxusers 60 -startmode
   auto -pfile
   C:\opt\oracle\database\initscdb.ora
   ```

   To change the number of Oracle connections, you must remove the service and then re-add it.

3. To remove the service:

   ```
   C:\opt\oracle\bin\oradim -delete -sid scdb
   ```

4. To re-add the service, execute the following command, substituting the new connection values:

```
C:\opt\oracle\bin\oradim -new -sid scdb -maxusers 60 -startmode
auto -pfile
```

Performance recommendations

To optimize database performance:

- Use your fastest disk storage for the portal database. When choosing the device for your Oracle database and if you have a choice between RAID1 - RAID5, choose RAID1.
- Minimize I/O wait time. Use the **top** command to determine the I/O wait time.

Enable SQL queries baseline capture

Consistent query execution is the backbone of the performance of any database application. Even though the Oracle Optimizer is self-sufficient to identify the optimal plan without any user intervention, the execution plan of a SQL statement can get impacted by reasons such as regathering optimizer statistics and changes to the optimizer parameters or schema/metadata definitions. SQL Plan Management (SPM) provides a similar framework and enables complete controlled plan evolution.

To prevent unexpected performance issues, Cohesity recommends to use 'AUTOMATIC PLAN CAPTURE' which is controlled using Oracle parameter OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES. You can find more details in SQL Plan Management in Oracle Database 19c white paper.

To enable baseline captures when the performance is stable:

1. sqlplus / as sysdba alter session set container = scdb;
2. Run following statement to find if the parameter is currently enabled/disabled (TRUE/FALSE):

   ```
   show parameter OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES;
   ```

3. If disabled (FALSE), enable sql baseline capture using following:

   ```
   alter system set OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES = TRUE;
   ```

After about 24 to 48 hours or when all the collections have completed at least once, disable the baselines capture using following steps:

1. sqlplus / as sysdba alter session set container = scdb;
2. Disable sql baseline capture using following: alter system set OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES = FALSE;

Reclaiming free space from Oracle

You may occasionally need to reclaim space from Oracle before additional storage can be provisioned.

A new monthly Oracle job `reclaimLOBTablespaces` has been incorporated which reclaims unused table space for `APTARE_TBS_DATA_20M` and `APTARE_TBS_DATA_1M`.

To complete the process successfully,

- the portal user should have ALTER DATABASE privilege.
  OR
- the job should be disabled.

The following are the steps to enable the privileges:

1. `su - aptare`
2. `sqlplus as sysdba`
3. `alter session set container = scdb;`
4. `grant alter database to PORTAL;`

The following is the command to disable the job.

```
exec dbms_scheduler.disable('PORTAL.RECLAIMLOBTABLESPACES');
```

You can run the following script at any time to reclaim space. It examines every Oracle database file (DBF) for “white space” at the end of the file. If the script discovers more than 256 MB of white space, it re-sizes the DBF file to remove the trailing space. This white space is a result of many insertions and deletions; in addition, white space can occur if you have truncated tables or purged a lot of data.

1. Log in to the database server as aptare.
2. Go to the tools directory:
   - Linux: `cd /opt/aptare/database/tools`
   - Windows: **cd C:\opt\oracle\database\tools**

```
sqlplus / as sysdba
@ reclaim_aptare_tablespace
commit;
exit
```

Portal / Data receiver Java memory settings

Use Java memory settings to improve performance by controlling the amount of memory Java programs use.

- Use -Xmx to specify the maximum heap size
- Use -Xms to specify the initial Java heap size
- Use -Xss to set the Java thread stack size

Linux Locations

- /opt/aptare/portalconf/tomcat/java-settings.sh
- /opt/aptare/datarcvrconf/tomcat/java-settings.sh

Windows Locations

- C:\opt\aptare\portalconf\tomcat\java-settings.bat
- C:\opt\aptare\ datarcvrconf \tomcat\java-settings.bat

Chapter 19

Working with log files

This chapter includes the following topics:

About debugging IT Analytics

IT Analytics logs exceptions. Often these exceptions do not indicate a problem with IT Analytics. However, if you experience system problems, Cohesity Support wants to help you troubleshoot your problem and interpret information in the log files. To speed up troubleshooting, provide the Cohesity Support with the appropriate log files. These log files are often specific to your operating system.

Log files are managed by a logging subsystem, which manages log file size and rolls and deletes old files. Log files are used only for audit trail or troubleshooting purposes and can be safely deleted.

Note: When the portal/receiver or collector is upgraded, the logging configuration will be reset to the default values.

Turn on debugging

When you turn on debugging, additional entries are logged to provide troubleshooting details.

1. In the Portal, within a report window, enter the following key combination:

   `Ctrl+Alt+D`

   This turns on debugging for the current report and it logs messages to both of the following log files:

   Linux: `/tmp/scon.log and /opt/tomcat/logs/portal.log`

   Windows: `C:\tmp\scon.log and C:\opt\tomcat\logs\portal.log`

2. See “Portal log files” on page 233.

   See “Database log files” on page 236.

Database logging

The **/tmp/scon.log** file (on Linux systems) or `C:\opt\oracle\logs\scon.log` (on Windows systems) contains a database audit trail and troubleshooting messages. You can control database logging by editing the following file, which contains instructions on what to modify in the file.

Linux: `/opt/aptare/database/stored_procedures/config.sql`

Windows: `C:\opt\oracle\database\stored_procedures\config.sql`

Portal and data collector log files - reduce logging

To manage the maximum file size and threshold parameter, edit the file:

Linux: `/opt/aptare/mbs/conf/systemlogger.xml`

Windows: `C:\opt\aptare\mbs\conf\systemlogger.xml`

```
<param name="MaxFileSize" value="10MB" />
<param name="MaxBackupIndex" value="10" />
<!--The Threshold param can either be debug/info/warn/error/fatal.-->
<param name="Threshold" value="debug"/>
```

Portal Log Files

For **Portal Tomcat**, edit this file:

Linux: `/opt/aptare/portalconf/systemlogger.xml`

Windows: `C:\opt\aptare\portalconf\systemlogger.xml`

```
<param name="MaxFileSize" value="10MB" />
<param name="MaxBackupIndex" value="10" />
```

For the **Data Receiver Tomcat**, edit this file:

Linux: `/opt/aptare/datarcvrconf/systemlogger.xml`

Windows: `C:\opt\aptare\datarcvrconf\systemlogger.xml`

```
<rollingPolicy
class="ch.qos.logback.core.rolling.FixedWindowRollingPolicy">
<fileNamePattern>/opt/tomcat/logs/datarcvr_%i.log</fileNamePattern>
 <minIndex>1</minIndex>
 <maxIndex>10</maxIndex>
</rollingPolicy>
<triggeringPolicy
class="com.aptare.dc.util.LogbackSizeBasedTriggeringPolicy">
            <maxFileSize>20MB</maxFileSize>
</triggeringPolicy>
<!--The Threshold param can either be debug/info/warn/error/fatal.-->
<param name="Threshold" value="debug"/>
```

Data Collector Log Files

For the **Data Collector** the file:

Linux: `/opt/aptare/mbs/conf/metadatalogger.xml`

Windows: `C:\Program Files\Aptare\mbs\conf\metadatalogger.xml`

```
<rollingPolicy
class="ch.qos.logback.core.rolling.FixedWindowRollingPolicy">
<fileNamePattern>/opt/aptare/agent_version/
DemoDC/mbs/logs/metadata${mdc_key}.%i.log</fileNamePattern>
                    <minIndex>1</minIndex>
                    <maxIndex>20</maxIndex>
</rollingPolicy>
<triggeringPolicy
class="com.aptare.dc.util.LogbackSizeBasedTriggeringPolicy">
                    <maxFileSize>50MB</maxFileSize>
</triggeringPolicy>
```

Database SCON logging - reduce logging

To minimize output to the scon.log file, you can prune the content using this procedure:

1. Log into the database server and switch to aptare user.
2. Edit the **config.sql** file.

   Linux: `/opt/aptare/database/stored_procedures/config.sql`

   Windows: `C:\opt\oracle\database\stored_procedures\config.sql`
3. Change the output setting to **LOW, as shown in BOLD in the following example**.

```
set Echo     Off
set Feedback OffCREATE OR REPLACE PACKAGE config AS

    -- Valid APTARE StorageConsole Logging levels are as follows:

    -- constant.DEBUG_OFF
    -- constant.DEBUG_LOW
    -- constant.DEBUG_MEDIUM
    -- constant.DEBUG_HIGH    -- To change the global APTARE
StorageConsole logging level, change the following constant
    globalDebugLevel        PLS_INTEGER := constant.DEBUG_LOW;
```

```
    reportTransLongerThan   FLOAT := 0.85;      -- Transactions
that take longer than this number of seconds will be reported in
 aptare-trans.log

   --UNCOMMENT BELOW TO ENABLE SCON.LOG writing. SAME parameter
is available under System Configurations> Database Administration
 > Enable scon.log real time logging (minutes)
       --UPDATE apt_system_parameter SET param_value = 10 WHERE
param_name = 'ENABLE_SCON_REAL_LOG_MINS'; COMMIT;

    -- The following directory will be used to store the APTARE
StorageConsole
    -- database logfiles scon.log and scon.err. On a Windows
portal server,
    -- this will default to C:\opt\oracle\logs
    LOGDIRECTORY CONSTANT VARCHAR2(64) := '/tmp' ;

   -- Following parameters are used in datacoll_exec_tracking_pkg

    NO_OF_DC_SP_LOG_ROWS                CONSTANT NUMBER(5)    :=
 100;--Defines the number of Data Collection log rows. This log
contains execution tracking.
    NO_OF_DC_QUERY_LOG_ROWS             CONSTANT NUMBER(5)    :=
 50;--Defines the number of rows for the Data Collection query
log. This log contains execution tracking.
    DATACOLL_LOGGING_ENABLED            CONSTANT NUMBER(1)    :=
 1;--Enables/disables data persistance transaction logging during
 data collection. This logging helps to isolate collection
performance issues. If collection performance is degraded, you
can temporarily disable the process.

END config;
/
SHOW ERRORS;
```

4. Apply and validate new settings with the following utilities:

   Linux:

```
sqlplus
portal/<portal_password>@//localhost:1521/scdb@/opt/aptare/database/
stored_procedures/config.sql
sqlplus
portal/<portal_password>@//localhost:1521/scdb@/opt/aptare/database/tools/validate_sp
```

Windows:

```
sqlplus
portal/<portal_password>@//localhost:1521/scdb@C:\opt\oracle\database\
stored_procedures\config.sql
sqlplus
portal/<portal_password>@//localhost:1521/scdb@C:\opt\oracle\database\tools\
validate_sp
```

Refreshing the database SCON log

Even when the amount of data written to the scon.log file is minimal, over time this file can reach a limit that causes processing to cease. The following instructions provide the steps for a utility that copies the existing log file and then empties scon.log, without impacting IT Analytics processing.

The utility to refresh scon.log executes automatically according to the following rules:

- Utility executes monthly to refresh the scon.log file
- First run executes 30 days after portal installation
- Production run is scheduled for the first Tuesday of every month

The utility searches for the following directory paths until an scon.log file is found:

```
'C:\opt\oracle\logs','C:\opt\aptare\oracle\logs','C:\opt\aptare\oracle\log','/tmp'
```

To refresh the scon.log file, use the following utility:

1. Log in to SQLPLUS, as shown below.

   Linux & Windows:

   ```
   sqlplus portal/<portal_password>@//localhost:1521/scdb
   ```

2. Execute the log cleanup utility,

   Linux & Windows:

   ```
   exec logfile_cleanup_pkg.cleanupLog('Y','Y');
   ```

   - Two parameters are required, as described in the following table:

Backup Flag	Save in Backup File	Description
Y	Y	scon.log file emptied backup in scon_<month>.log file
Y	N	scon.log file emptied backup in scon1.log file; all backups will overwrite scon1.log
N	Y	scon.log file emptied no backup of the scon.log file
N	N	scon.log file emptied no backup of the scon.log file

Note: This utility is intended to be run no more than once a month. If you plan to run it more than once in a month, be aware of the naming convention for the backup scon.log file, as shown with the parameters in the above table.

Logging user activity in audit.log

By default, the audit.log captures:

- User login
- User impersonate

Modify the logging level of the systemlogger.xml file to provide additional information about a user's activity in the Portal. You can set the level to info to capture only what a user deletes OR set it to debug to capture all the user activity (including all deletes).

Logging only what a user deletes

To log all deletes made by a user in audit.log, edit the following file:

Linux: `/opt/aptare/portalconf/systemlogger.xml`

Windows: `C:\opt\aptare\portalconf\systemlogger.xml`

```
<logger name="com.aptare.sc.gwt.shared.server.GwtSpringAdapter"
```

```
additivity="false">
    <level value="info"/>
   <appender-ref ref="SECURITY" />
   </logger>

<logger name="com.aptare.sc.presentation.filter.AuthorizationFilter"
 additivity="false">
    <level value="info"/>
   <appender-ref ref="SECURITY" />
  </logger>
```

Logging all user activity

To log all the activity of a user (including what they delete) in audit.log, edit the following file:

Linux: `/opt/aptare/portalconf/systemlogger.xml`

Windows: `C:\opt\aptare\portalconf\systemlogger.xml`

```
   <logger name="com.aptare.sc.gwt.shared.server.GwtSpringAdapter"
additivity="false">
    <level value="debug"/>
   <appender-ref ref="SECURITY" />
  </logger>

<logger name="com.aptare.sc.presentation.filter.AuthorizationFilter"
 additivity="false">
    <level value="debug"/>
   <appender-ref ref="SECURITY" />
  </logger>
```

Data collector log files

Before resorting to examining Data Collector logs, use the following System Administration reports to check collection status:

- Collection Message Summary
- Data Collection Schedule Summary
- Data Collector Status Summary
- File Analytics Collection Status

You can also view live collection status using the Collection Status page on the Admin tab.

The data collection process logs activity to provide additional processing details that support troubleshooting. Use the **Admin > Advanced > Support Tools** feature to request a support package that contains Data Collector Server files. To identify specific files, it may be helpful to have some knowledge of the logging structure and naming conventions.

Refer to the following sections to understand the logging structure.

- See "Data collector log file organization" on page 225.
- See "Data collector log file naming conventions" on page 225.
- See "General data collector log files" on page 232.

Data collector log file organization

Data Collector logs are organized into two directory hierarchies:

Checkinstall and Validation Probes

- `/opt/aptare/mbs/logs/validation/`
- `C:\Program Files\Aptare\mbs\logs\validation\`

Scheduled Probes (running as a service)

- `/opt/aptare/mbs/logs/scheduled/`
- `C:\Program Files\Aptare\mbs\logs\scheduled\`

Within this directory structure, there is a **framework** sub-directory. The **framework** sub-directory is used in the beginning of the collection process, before the Data Collector type has been identified. Once IT Analytics knows the type of collector (for example, EMC Isilon), logging is recorded in the main **scheduled** directory using the naming convention described in the following section.

See "Data collector log file naming conventions" on page 225.

Data collector log file naming conventions

Within the directory structure described in *Data Collector log file organization* section, log files have the following naming convention:

```
<vendor.product>/<subsystem>#META_<ID>/Probe.log
```

For example, an EMC Isilon probe from checkinstall would result in a file name similar to:

```
/opt/aptare/mbs/logs/validation/emc.isilon
/alphpeifr023#META_EA1BA380E95F73C72A72B3B0792111E5
/IsilonClusterDetailProbe.log
```

Some collectors may have a period of time when they are not processing a specific subsystem. For those periods, logging will occur in an aggregate log file similar to:

```
/opt/aptare/mbs/logs/validation/emc.isilon
/#META_EA1BA380E95F73C72A72B3B0792111E5
/IsilonClusterDetailProbe.log
```

See "Data collector log file organization" on page 225.

Sample Vendor.Product Naming Convention

Examples of vendor.product folder names within this directory structure include:

```
cisco.cisco
commvault.simpana
dell.compellent
emc.avamar
emc.clariion
generic.host (Valid for a host resources discovery policy)
hp.3par
veritas.bue
```

Additionally, each Java Virtual Machine (JVM) creates its own logging file(s) when starting up. This is necessary because multiple processes logging to the same file could overwrite each other's log messages. These log files can be found in the framework sub-directory.

See "Data collector log file organization" on page 225.

See "Checkinstall Log" on page 227.

Log File Names Based on Data Collector Generation

Throughout the IT Analytics life cycle, some collectors have been upgraded to include new functionality and new Data Collectors have been designed with an improved architecture. The log file naming convention depends on the generation of the Data Collector that you are using. You do not need to know the generation of the Data Collector to find relevant log files.

When running a scheduled data collection, the following log files are created, depending on the generation of the Data Collector. For technical reasons, the following naming conventions are used for data collection logs:

- older-generation collectors follow the convention, **{scheduled,validation}/vendor.product/#META_** or **{scheduled,validation}/vendor.product/EVENT_,**where **EVENT_** or **META_** is prepended to the collector policy ID; for example, **META_CA6EC7685A9E6330EC3BBFC0DD4811E4**.
- newer-generation collectors share a single file named with the main collector ID, for example, **{scheduled,validation}/vendor.product/#<Collector_PolicyID>**

Several log file names include a specific ID. This ID can be found in a System Administration report, Data Collection Schedule Summary.

See "Find the event / meta collector ID" on page 233.

Examples

- `/opt/aptare/mbs/logs/scheduled/dell.compellent /#META_EA1BA380E95F73C72A72B3B0792111E5 /META_EA1BA380E95F73C72A72B3B0792111E5.log`
- `/opt/aptare/mbs/logs/scheduled/emc.avamar /#HQBackupCollector/HQBackupCollector.log`

Checkinstall Log

The **checkinstall** process produces its own log file, but in most cases, there is very little to report in this log.

For example, the checkinstall creates:

`/opt/aptare/mbs/logs/validation/framework/#checkinstall/checkinstall.log`

Test Connection Log

When you initiate a Test Connection action from within a Data Collector policy, a **TestConnection.log**file captures the steps and their status.

- `/opt/aptare/mbs/logs/validation/<vendor.product>/#TestConnection/TestConnection.log`
- `C:\Program Files\Aptare\mbs\logs\validation\<vendor.product>\#TestConnection\TestConnection.log`

Log file naming convention by collected system

The log file names will have one of the following prefixes substituted for the **<policyID>**:

- `#EVENT_<policyID>`
- `#META_<policyID>`
- `#<policyID>`

Example:
scheduled\legato.nw\#META_D922ACBCCFFA2933A301A530A0E011E4

Note: Some collectors may have both a **#META_**and an **#EVENT_**log file.

Table 19-1 Log file locations for IT Analytics products

IT Analytics Product	Collected System	Where to find the logs (Linux syntax)
Backup Manager	Commvault Simpana	`scheduled/commvault.simpana/#<policyID>` `validation/commvault.simpana/#<policyID>`
	EMC Avamar	`scheduled/emc.avamar/#<policyID>` `validation/emc.avamar/#<policyID>`
	EMC Data Domain Backup	`scheduled/emc.datadomain/#<policyID>` `validation/emc.datadomain/#<policyID>`
	EMC NetWorker	`scheduled/legato.nw/#<policyID>` `validation/legato.nw/#<policyID>`
	HP Data Protector	`scheduled/hp.dp/#<policyID>` `validation/hp.dp/#<policyID>`
	IBM Spectrum Protect (TSM)	`scheduled/ibm.tsm/#<policyID>` `validation/ibm.tsm/#<policyID>`

	Oracle Recovery Manager (RMAN)	`scheduled/oracle.rman/#<policyID>`
		`validation/oracle.rman/#<policyID>`
		
		Note: In addition to the **hostname#<productId>**/ directories, RMAN also creates directories for each INSTANCE.SCHEMA from which it collects.

Table 19-1 Log file locations for IT Analytics products *(continued)*

IT Analytics Product	Collected System	Where to find the logs (Linux syntax)
	Veeam Backup & Replication	`scheduled/veeam.backupandreplication /#<policyID>` `validation/veeam.backupandreplication /#<policyID>`
	Veritas Backup Exec	`scheduled/veritas.bue/#<policyID>` `validation/veritas.bue/#<policyID>`
	Veritas NetBackup	`scheduled/veritas.netbackup/#<policyID>` `validation/veritas.netbackup/#<policyID>`
Capacity Manager	Dell Compellent	`scheduled/dell.compellent/#<policyID>` `validation/dell.compellent/#<policyID>`
	EMC Data Domain Storage	`scheduled/emc.datadomain/#<policyID>` `validation/emc.datadomain/#<policyID>`
	EMC Isilon	`scheduled/emc.isilon/#<policyID>` `validation/emc.isilon/#<policyID>`
	EMC Symmetrix	`scheduled/emc.symmetrix/#<policyID>` `validation/emc.symmetrix/#<policyID>`
	EMC VNX (Celerra)	`scheduled/emc.celerra/#<policyID>` `validation/emc.celerra/#<policyID>`
	EMC VNX (CLARiiON)	`scheduled/emc.clariion/#<policyID>` `validation/emc.clariion/#<policyID>`
	EMC VPLEX	`scheduled/emc.vplex/#<policyID>` `validation/emc.vplex/#<policyID>`
	EMC XtremIO	`scheduled/emc.xtremio/#<policyID>` `validation/emc.xtremio/#<policyID>`
	HP 3PAR	`scheduled/hp.3par/#<policyID>` `validation/hp.3par/#<policyID>`

Table 19-1 Log file locations for IT Analytics products *(continued)*

IT Analytics Product	Collected System	Where to find the logs (Linux syntax)
	HP EVA	`scheduled/hp.eva/#<policyID>` `validation/hp.eva/#<policyID>`
	Hitachi Block	`scheduled/hds.hds/#<policyID>` `validation/hds.hds/#<policyID>`
	Hitachi HCP	`scheduled/hds.hcp/#<policyID>` `validation/hds.hcp/#<policyID>`
	Hitachi NAS	`scheduled/hitachi.hnas/#<policyID>` `validation/hitachi.hnas/#<policyID>`
	IBM Enterprise	`scheduled/ibm.ent/#<policyID>` `validation/ibm.ent/#<policyID>`
	IBM SVC	`scheduled/ibm.svc/#<policyID>` `validation/ibm.svc/#<policyID>`
	IBM XIV	`scheduled/ibm.xiv/#<policyID>` `validation/ibm.xiv/#<policyID>`
	NetApp 7-Mode	`scheduled/netapp.netapp/#<policyID>` `validation/netapp.netapp/#<policyID>`
	NetApp Cluster	`scheduled/netapp.netapp/#<policyID>` `validation/netapp.netapp/#<policyID>`
	NetApp E-Series	`scheduled/netapp.netapp/#<policyID>` `validation/netapp.netapp/#<policyID>`
	Pure Storage FlashArray	`scheduled/purestorage.flasharray/#<policyID>` `validation/purestorage.flasharray/#<policyID>`
Cloud	Amazon Web Services	`scheduled/amazon.webservices/#<policyID>` `validation/amazon.webservices/#<policyID>`

Table 19-1 Log file locations for IT Analytics products *(continued)*

IT Analytics Product	Collected System	Where to find the logs (Linux syntax)
	Microsoft Azure	scheduled/microsoft.azure/#<policyID> validation/microsoft.azure/#<policyID>
	OpenStack Ceilometer	scheduled/openstack.ceilometer/#<policyID> validation/openstack.ceilometer/#<policyID>
	OpenStack Swift	scheduled/openstack.swift/#<policyID> validation/openstack.swift/#<policyID>
Fabric Manager	Brocade	scheduled/brocade.brocadeswitch/#<policyID> validation/brocade.brocadeswitch/#<policyID>
	Brocade Zone Alias	scheduled/brocade.brocade/#<policyID> validation/brocade.brocade/#<policyID>
	Cisco	scheduled/cisco.ciscoswitch/#<policyID> validation/cisco.ciscoswitch/#<policyID>
	Cisco Zone Alias	scheduled/cisco.cisco/#<policyID> validation/cisco.cisco/#<policyID>
File Analytics	Hosts	scheduled/generic.fa/#<policyID> validation/generic.fa/#<policyID>
Host Collection	Hosts	scheduled/generic.host/#<policyID> validation/generic.host/#<policyID>
Virtualization Manager	IBM VIO	scheduled/cisco.cisco/#<policyID> validation/cisco.cisco/#<policyID>
	VMware	scheduled/vmware.esx/ validation/vmware.esx/

General data collector log files

Locations in this table represent the default locations, but these may have been modified for your environment.

Table 19-2 General Data Collector Logs

Log File Name	Default Location	Description	Component
start_watchdog.log	`/opt/aptare/mbs/logs`	Logging for the high-level management of the Watchdog component. Management includes startup, shutdown, and initialization.	Watchdog
wrapper.log	`C:\Program Files\Aptare\mbs\logs`		
watchdog.log	`C:\Program Files\Aptare\mbs\logs` `/opt/aptare/mbs/logs`	Detailed logging for the Watchdog component. Logs all the functions Watchdog performs including downloading new configuration files.	
upgradeMgr.log upgrade_<version>.log	`C:\Program Files\Aptare\upgrade\upgradeManager\logs` `/opt/aptare/upgrade/upgradeManager/logs`	Detailed logging for the Data Collector Upgrade Manager. <version> refers to the version of the aptare.jar to which the Data Collector is being upgraded. **Note**: The logs directory will not exist for a new installation of the Data Collector. The logs directory will only be created once the collector goes through an aptare.jar upgrade cycle.	Upgrade Manager

Find the event / meta collector ID

This ID is displayed in the Data Collection Schedule Summary, but you can also use the following script to list the IDs.

1. List the Data Collectors to get the Event/Metadata Collector ID.

   Windows:

   ```
   C:\Program Files\APTARE\mbs\bin\listcollectors.bat
   ```

   Linux:

   ```
   /opt/aptare/mbs/bin/listcollectors.sh
   ```

   In the output, look for the Event Collectors section associated with the Software Home--the path that was specified when the Data Collector Policy was created.

   ```
   ==== Event Collectors ===
   Event Collector Id: EF14CEE486DF781F312E5D40411C11E5
   Active: true
   Active: true
     Software Home: C:\Program Files\EMC NetWorker\nsr\bin
     Server Address: networker3
   Domain: 100000
   Sub-system/Server Instance/Device Manager Id: 110050
       Schedule: 10
   ```

Portal log files

Key log files are managed by a logging subsystem, which manages the log size, and rolls and deletes old files. Log files are used only as an audit trail for troubleshooting, so they can be deleted safely.

Locations in this table represent the default locations, but these may have been modified for your environment.

Table 19-3 Portal Log Files

Log File Name	Default Location	Description
access.log	`C:\opt\apache\logs` `/opt/apache/logs`	Standard Web server access log.Use this log to analyze request processing time, request method, page hit count, and session activity.

Table 19-3 Portal Log Files *(continued)*

Log File Name	Default Location	Description
itanalyticsagent-access*.log	`C:\opt\apache\logs` `/opt/apache/logs`	Standard Web Server access log. Logs all http transactions between the Data Collector and the Web Server.

itanalyticsagent-error*.log	`C:\opt\apache\logs`
	`/opt/apache/logs`
	`/opt/tomcat/logs`
		Standard Web Server error log file. Logs http transaction errors between the Data Collector and the Web Server.

itanalyticsportal-access*.log	`C:\opt\apache\logs` `/opt/apache/logs`	Standard Web Server access log. Logs all http transactions between the browser based Portal application and the Web Server.
itanalyticsportal-error*.log	`C:\opt\apache\logs` `/opt/apache/logs`	Standard Web Server error log file. Logs http transaction errors between the browser-based Portal application and the Web Server.
itanalyticsStartup.log	`/opt/aptare/logs`	Contains startup and shutdown information for the Portal services (e.g., from running /opt/aptare/bin/xxx start|stop), If a service fails to start, check this log file for details.
audit.log	`C:\opt\tomcat\logs\` `/opt/tomcat/logs/`	By default, logs portal login requests and user impersonations in the Portal web browser window. Modifications can be made to log additional user activity. See "Logging user activity in audit.log" on page 223.
catalina.out	`C:\opt\tomcat\aptare-instances\portal\logs` `/opt/tomcat/aptare-instances/portal/logs`	For Tomcat, this is the standard destination log file for System.out and System.err console messaging.
datarcvr*.log	`C:\opt\tomcat\logs` `/opt/tomcat/logs`	Detailed logging for the data receiver - the servlet that receives data from the Data Collectors on the Primary Server(s). Logs Data Collector connection requests/issues, data received, and database interaction. See "Portal and data collector log files - reduce logging" on page 219.

Table 19-3 Portal Log Files *(continued)*

Log File Name	Default Location	Description
error.log	`C:\opt\apache\logs` `/opt/apache/logs`	Standard Web Server error log file.
portal*.log	`C:\opt\tomcat\logs\` `/opt/tomcat/logs/`	Detailed logging for the Portal servlet. Logs portal login requests, user impersonations, portal reports that are run - basically all actions in the Portal web browser window. Database problems are displayed as SQL exceptions and often list the associated Oracle error number (ORA nnn). See "Portal and data collector log files - reduce logging" on page 219.
stderr.log	`C:\opt\tomcat\logs\` `/opt/tomcat/logs/`	High-level log messaging for Tomcat.
Tomcat Standard Log Files (Portal)		
admin*.log	`C:\opt\tomcat\aptare-instances\portal\logs` `/opt/tomcat/aptare-instances/portal/logs`	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to admin in the name of the file.
catalina*.log	`C:\opt\tomcat\aptare-instances\portal\logs` `/opt/tomcat/aptare-instances/portal/logs`	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to catalina in the name of the file.
host-manager*.log	`C:\opt\tomcat\aptare-instances\portal\logs` `/opt/tomcat/aptare-instances/portal/logs`	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to host-manager in the name of the file.
manager*.log	`C:\opt\tomcat\aptare-instances\portal\logs` `/opt/tomcat/aptare-instances/portal/logs`	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to manager in the name of the file.
Tomcat Standard Log Files (Data Receiver)		

Table 19-3 Portal Log Files *(continued)*

Log File Name	Default Location	Description
admin*.log	C:\opt\tomcat\aptare-instances\agent\logs /opt/tomcat/aptare-instances/agent/logs	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to admin in the name of the file.
catalina*.log	C:\opt\tomcat\aptare-instances\agent\logs /opt/tomcat/aptare-instances/agent/logs	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to catalina in the name of the file.
host-manager*.log	C:\opt\tomcat\aptare-instances\agent\logs /opt/tomcat/aptare-instances/agent/logs	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to host-manager in the name of the file.
manager*.log	C:\opt\tomcat\aptare-instances\agent\logs /opt/tomcat/aptare-instances/agent/logs	Apache Tomcat provides several file handlers that write messages to application-specific log files. The date is appended to manager in the name of the file.

Managing Apache Log Files

Even if you set up Apache logs for a regularly scheduled rotation, space issues can still occur. You must purge older files if maintaining space is more important than preserving logs. Logs are not automatically deleted.

Set up a cron job to remove older apache logs. By default, these logs rotate, but they are not removed. Create scheduled jobs to delete logs that are <XX> days old.

Database log files

Locations in this table represent the default locations, but these may have been modified for your environment.

Table 19-4 Reporting Database Log Files

Log File Name	Location	Description
aptare-trans.log	`C:\tmp` `/tmp`	Records long-running transactions.
scon.err	`C:\tmp` `/tmp`	Error messages from the database stored procedures. Oracle errors are prefixed with ORA-nnnn, where nnnn is the APTARE or Oracle error number. On Windows, if the directory `C:\tmp` exists, the file will be written there. Otherwise, it will be written to `C:\opt\oracle\logs`.
scon.log	`C:\tmp` `/tmp`	Detailed logging from database stored procedures. Shows input and output values, and timing of queries. Oracle errors are prefixed with ORA-nnnn, where nnnn is the APTARE or Oracle error number. On Windows, if the directory C:\tmp exists, the file will be written there. Otherwise, it will be written to `C:\opt\oracle\logs`.
sqlnet.log	`C:\opt\oracle\network\log` `/opt/aptare/oracle/network/log`	Oracle listener log file. Logs information about connection requests made to the Reporting Database.
alert_scdb.log	`C:\opt\oracle\rdbms\log\diag\rdbms\scdb\scdb\trace` `/opt/aptare/oracle/rdbms/log/diag/rdbms/scdb/scdb/trace`	Oracle alert log file. Logs information on startup and shutdown of oracle instance, and also critical Oracle problems (such as disk I/O failures, out of disk space, system resource issues etc.)

Installation / Upgrade log files

Table 19-5 Installation Log Files

Log File Name	Location	Description
aptare_installer*.log	`C:\opt\oracle\logs` `/tmp`	Database installation log files.
install.log	`C:\opt\installlogs` `/opt/aptare/installlogs`	Log of high-level installation tasks.
installer_debug.txt	`C:\opt\installlogs` `/opt/aptare/installlogs`	Log of install tasks.
upgrade.log	`C:\opt\aptare\upgrade\logs` `/opt/aptare/upgrade/logs`	Log of upgrade tasks performed by the upgrader

Chapter 20

Defining report metrics

This chapter includes the following topics:

Changing backup success percentage

By default the Backup Status Report defines the **Success Percentage** metric as **85%**. Although this percentage is typical, your SLA might require a different percentage. You can change this percentage for specific host groups or for all host groups.

To change the success percentage metric:

1. Determine the host group's ID.
2. Log on to the Portal Server as user **aptare**.
3. Type the following command:

```
sqlplus portal/<portal_password>@//localhost:1521/scdb
```

4. Insert a row into the **ptl_sla_group_policy** table. The following example assumes that you want to change the success percentage to **95%** for all host groups.

```
INSERT INTO ptl_sla_group_policy (group_id,
successful_backups_objective, successful_restores_objective)
VALUES (300000, 95.0, 95.0);
commit;
```

Changing job status

By default, the Job Summary Report defines the job status as follows:

Table 20-1 Table 1 Job Summary Metrics

Metric	Default Value
LONG_JOB_HOURS_DEFAULT	12 hours
SLOW_JOB_DEFAULT	<200 kilobytes per second
STALLED_JOB_DEFAULT	1800 seconds (30 minutes)

Although these values are typical, your SLA might require a different values. You can change these metrics for specific host groups or for all host groups.

To change the job status:

1. Determine the host group's ID.
2. Log on to the Portal Server as user **aptare**.
3. Type the following command:

```
sqlplus portal/<portal_password>@//localhost:1521/scdb
```

4. Insert a row into the **ptl_sla_group_policy** table. The following example assumes that you want to change the metric to 1MB per second for all host groups.

```
INSERT INTO ptl_group_policy (group_id, policy_name,display_name,
 numeric_value)
VALUES (300000,'SLOW_JOB_KB_PER_SEC', 'SlowRunning Job', 1000);
commit;
```

Chapter 21

SNMP trap alerting

This chapter includes the following topics:

Overview

Tabular Reports can be configured to alert users via a number of options:

- Email
- Script
- SNMP
- Native log
- Syslog

Typically, notifications are expected when a report's content indicates a threshold crossing or an error state. For example, a Job Summary report can be configured to trigger an alert when failed backup events are reported.

Note: To receive SNMP traps and native logs from IT Analytics, the recipient system must be able to connect to the IT Analytics Portal. This is applicable to on premise and SaaS-based Solutions of IT Analytics. For SaaS-based solutions like Alta, ensure that Alta Reports can connect to the SNMP server or the native log server via public IP, secured VPN, or any another appropriate network channel and required ports are open. It is highly recommended to configure appropriate networking security measures when opening a internal system to internet.

SNMP configurations

The SNMP traps can be issued from:

1. Any saved tabular report instance, including custom reports that have been created via the Report Template Designer.

   Traps have the following characteristics:

   - A trap is sent for each row in a report; therefore, if a table is empty, no traps are sent.
   - The name of the report is included in the trap.
   - The trap includes data for each column in a row.

2. Policy Based rules, which are created by choosing from a list of predefined policy rules in Alert Policy Administration.

   Traps have the following characteristics:

   - A trap is sent for each time the rules defined in policy are triggered. Traps are triggered only after rules is set to active and rule condition is met.

SNMP traps

This section explains the content of an SNMP trap that is sent from IT Analytics.

MIB files for report based alert can be found at `<APTARE_HOME>/goodies/`. For example if <APTARE_HOME> is `/opt/aptare`, then MIB files will be present at `/opt/aptare/goodies`.

See "SNMP traps for system alerts" on page 242.

See "SNMP traps for report-based alerts" on page 244.

SNMP traps for system alerts

Each IT Analytics trap contains 2 standard object identifiers and 13 IT Analytics-specific object identifiers. An object identifier (or OID) is a numeric string that is used to uniquely identify an object.

The following table shows the contents of a trap that is sent from IT Analytics. A total of 15 bindings (or 15 name-value pairs) are present in each trap that is sent from IT Analytics. Each binding associates a particular Management Information Base (MIB) object instance with its current value.

IT Analytics trap binding shows the name-value pairs that the traps pass to the SNMP manager if SNMP payload format is selected as **Enhanced**.

Table 21-1 IT Analytics system alerts trap binding

Name	Value
1.3.6.1.2.1.1.3.0	This field is the time (in hundredths of a second) between when IT Analytics server service starts and the IT Analytics trap is sent. See Request for Comment (RFC) 1905 and 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1173792454
1.3.6.1.6.3.1.1.4.1.0	This field is the unique identifier for this trap. See RFC 1905 and RFC 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1.3.6.1.4.1.15622.3.12.10.2.0.4
1.3.6.1.4.1.15622.3.12.10.1.1 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch. alertApplicationName)	This field is the Application Name. Example: IT Analytics
1.3.6.1.4.1.15622.3.12.10.1.2 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertDomainId)	This value specifies the Domain Id of the alert. Example: 100000
1.3.6.1.4.1.15622.3.12.10.1.3 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertDescription)	This field is the alert description.
1.3.6.1.4.1.15622.3.12.10.1.4 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertPolicyName)	This field is the alert policy name.
1.3.6.1.4.1.15622.3.12.10.1.5 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertTypeName)	Defines the object type for which the alert was set. Object type dictates which drill down report is displayed.
1.3.6.1.4.1.15622.3.12.10.1.6 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertRootObjectName)	Displays the object on which the alert was issued.

Table 21-1 IT Analytics system alerts trap binding *(continued)*

Name	Value
1.3.6.1.4.1.15622.3.12.10.1.7 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertParentTypeName)	If a hierarchy structure exists, displays the object type of the parent for the alert object.
1.3.6.1.4.1.15622.3.12.10.1.8 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertParentObjectName)	If a hierarchy structure exists, displays the name of parent of the alert object.
1.3.6.1.4.1.15622.3.12.10.1.9 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertLastUpdated) (iso.org.dod.internet.private.enterprises.it-analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertLastUpdated)	Displays the date the alert condition occurred.
1.3.6.1.4.1.15622.3.12.10.1.10 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertCount)	The number of occurrences of alert.
1.3.6.1.4.1.15622.3.12.10.1.11 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertSeverity)	Indicates the severity of the alert.
1.3.6.1.4.1.15622.3.12.10.1.12 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchAlertLastAlertDate)	Displays the latest date the alert condition occurred.
1.3.6.1.4.1.15622.3.12.10.1.13 (iso.org.dod.internet.private.enterprises.analytics.products.portal.itaSystemAlert.itaSystemAlertTrapVars.branchTicketNumber)	ServiceNow ticket number.

The first two OIDs listed in the table are standard SNMP OIDs. The other OIDs starting from 1.3.6.1.4.1.15622.3.12.10.1.1 to 1.3.6.1.4.1.15622.3.12.10.1.12 are IT Analytics OIDs. As per SNMPv2c trap definition, the two standard SNMP OIDs must be present as part of every trap.

All the 12 IT Analytics OIDs are defined in the IT Analytics MIB files. However, the two standard OIDs are not defined in the IT Analytics MIB files.

SNMP traps for report-based alerts

Each IT Analytics report-based alert trap contains two standard object identifiers and 2N IT Analytics-specific object identifiers where value of N depends on number of columns present in each row in the report. An object identifier (or OID) is a numeric string that is used to uniquely identify an object.

The following table shows the contents of a trap that is sent from IT Analytics for report-based alerts. Four bindings (or 4 name-value pairs) are present in each trap that is sent from IT Analytics and 2N bindings are present for each column present in report. For each column, there are 2 bindings, one for column name and other for its value. Each binding associates a particular Management Information Base (MIB) object instance with its current value.

IT Analytics trap binding shows the name-value pairs that the traps pass to the SNMP manager if SNMP payload format is selected as Enhanced.

Table 21-2 IT Analytics report-based alert trap binding

Name	Value
1.3.6.1.2.1.1.3.0	This field is the time (in hundredths of a second) between when IT Analytics server service starts and the IT Analytics trap is sent. See Request for Comment (RFC) 1905 and 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1173792454
1.3.6.1.6.3.1.1.4.1.0	This field is the unique identifier for this trap. See RFC 1905 and RFC 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1.3.6.1.4.1.15622.3.12.11.2.0.4
1.3.6.1.4.1.15622.3.12.11.1.1 (iso.org.dod.internet.private.enterprises.ianalytics.products.itaportal.itaReport.itaReportAlert.itaReportAlertTrap.vasBranchreportName)	It displays the report name
.1.3.6.1.4.1.15622.3.12.11.1.2 (iso.org.dod.internet.private.enterprises.ianalytics.products.itaportal.itaReport.itaReportAlert.itaReportAlertTrap.vasBranchreportDataCount)	It displays number of columns present in report. This value can be used to iterate over bindings which contains column name and column value. For example, if the value is 2 then it means 2 columns are present in the report and there are 4 bindings for column name and value.

Table 21-2 IT Analytics report-based alert trap binding *(continued)*

Name	Value
1.3.6.1.4.1.15622.3.12.11.1.3.1.1.1 (.iso.org.dod.internet.private.enterprises.analytics.productData.dataReport.dataReportAlertTrap.vasBranchReportDataReportEntry.columnName.1)	It displays name of column 1 of the report.
1.3.6.1.4.1.15622.3.12.11.1.3.1.2.1 (.iso.org.dod.internet.private.enterprises.analytics.productData.dataReport.dataReportAlertTrap.vasBranchReportDataReportEntry.columnValue.1)	It displays value of column 1 of the report.
1.3.6.1.4.1.15622.3.12.11.1.3.1.1.N (.iso.org.dod.internet.private.enterprises.analytics.productData.dataReport.dataReportAlertTrap.vasBranchReportDataReportEntry.columnName.N)	It displays name of column N of the report.
1.3.6.1.4.1.15622.3.12.11.1.3.1.2.N (.iso.org.dod.internet.private.enterprises.analytics.productData.dataReport.dataReportAlertTrap.vasBranchReportDataReportEntry.columnName.N)	It displays value of column N of the report.

Note: The name-value pair in the report-based alert trap binding will be from 1 to N, where N is number of columns present in the report.

Chapter 22

SSL certificate configuration

This chapter includes the following topics:

SSL certificate configuration

The following sections provide the steps for implementing the Secure Socket Layer (SSL) protocol within IT Analytics. SSL can be used when communicating with the Portal and for data collection communication. In addition to SSL configuration details, this section provides instructions to create a self-signed certificate and to add a virtual interface to a Linux server. A sample of a default Apache SSL configuration file is also provided.

These instructions have been validated, but there are many variations in the method used to implement SSL. This document is meant only as a guide to one implementation approach and it may not be applicable in all situations. Implementing SSL requires knowledge of the underlying technology.

Note: While these steps include directions for generating a temporary, self-signed certificate, you should obtain the certificate from a third-party provider rather than using the self-signed certificate.

SSL implementation overview

The Secure Socket Layer (SSL) protocol enables corporations to leverage standards-based security to protect and encrypt traffic between the IT Analytics Portal, the Data Collector, and the client browser. The following diagram illustrates how SSL is implemented for IT Analytics. The Apache Web Server typically resides on the Portal Server.

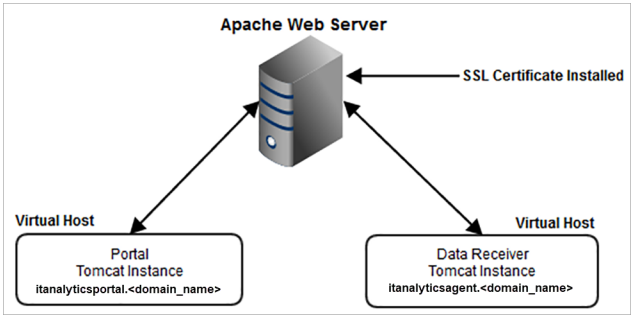

Note: The actual SSL certificates get installed and configured within the Apache Web Server, however, in cases where the issuing certificate authority (CA) is not automatically trusted (for example, self-signed or a one-off domain reseller), the certificates will need to be imported and configured to be trusted on the Data Collector Server. In this case, follow the process to import certificates into the keystore for both the Data Collector and the Upgrade Manager:

See "Configure the Data Collector to trust the certificate" on page 261.

Implementing SSL involves these main tasks:

- See "Obtain an SSL certificate" on page 249.
- See "Update the web server configuration to enable SSL" on page 251.
- See "Enable / Disable SSL for a Data Collector" on page 258.
- See "Enable / Disable SSL for emailed reports" on page 259.

Obtain an SSL certificate

Obtain a third-party certificate from a certificate authority (CA) such as VeriSign, Thawte, or GeoTrust. The methods for obtaining a certificate vary. Therefore, refer to the vendor's web site for specific instructions.

You may, for testing purposes or as a permanent solution, use a self-signed certificate. This is not recommended as it makes the implementation slightly more complex and may limit access to IT Analytics to some of your users.

The following outlines the process for creating a Subject Alternative Name (the certificate covers more than one hostname under a single certificate) self-signed certificate on a Linux operating system. Steps will be similar on Windows. This certificate secures communication for both the portal and data receiver web instances.

```
cd /tmp

vi san.cnf
```

Sample `san.cnf` file – use this faile as a template and modify this for your environment. The `san.cnf` file will be an input parameter during certificate generation. Note the use of an example domain name of `example.com`; change this to own your environment's domain name.

Under the v3 section, in addition to the portal name, also provision the data receiver under this same certificate.

```
[ req ]
default_bits = 4096
prompt = no
default_md = sha256
distinguished_name = req_distinguished_name
x509_extensions = v3_req

[ req_distinguished_name ]
C = US
ST = New York
L = New York City
O = Veritas
OU = ITA
emailAddress = aReal.emailaddress@yourdomain.com
CN = itanalyticsportal.example.com

[ v3_req ]
subjectAltName = @alternate_names

[alternate_names]
DNS.1 = itanalyticsportal.example.com
DNS.2 = itanalyticsagent.example.com
```

Generate Certificate using the `san.cnf` file created above

The following command results in the private key name of server.key, and certificate name of server.crt. These names will be used through the remainder of this chapter. You are free to use different names for the certificate and private key files if desired. With this command, we are also creating a self-signed certificate for 3650 days, or 10 years.

```
openssl req -x509 -newkey rsa:4096 -sha256 -days 3650 -nodes -keyout
 server.key -out server.crt -config /tmp/san.cnf

Generating a RSA private key
.....................................++++
.....................................++++
writing new private key to 'server.key'
-----
 tmp]# ll
total 276
-rwxrwxrwx 1 root   root      513 Dec 11 01:03 san.cnf
```

```
-rw-r--r-- 1 root   root     2187 Dec 11 01:25 server.crt
-rw------ 1 root   root     3272 Dec 11 01:25 server.key
```

See “Create a self-signed SSL certificate” on page 260.

Update the web server configuration to enable SSL

These instructions apply to Apache version 2.4.xx and the procedure must be executed on the designated Web server.

1. Copy the certificate files, typically generated via a certificate authority (CA), to a folder in the Web server’s Apache configuration folder.

   Note: Configuration files shipped with IT Analytics licensed modules may use path names with recommended folder names. To use folders with different names, be sure to update all references to the recommended name in the default configuration files.

   Linux

   ```
   /opt/apache/conf/ssl_cert
   ```

   Windows

   ```
   C:\opt\apache\conf\ssl_cert
   ```

2. Stop the Apache and Tomcat services. From a terminal console, enter the following commands.

   Linux

   ```
   /opt/aptare/bin/tomcat-agent stop
   /opt/aptare/bin/tomcat-portal stop
   /opt/aptare/bin/apache stop
   ```

   Windows

   ```
   C:\opt\aptare\utils\stopagent.bat
   C:\opt\aptare\utils\stopportal.bat
   C:\opt\aptare\utils\stopapache.bat
   ```

3. Update the Apache configuration file to enable SSL.

   Linux: `/opt/apache/conf/httpd.conf`

Windows: `C:\opt\apache\conf\httpd.conf`

Un-comment the following lines by removing the # character.

Linux

```
#LoadModule ssl_module modules/mod_ssl.so
#Include conf/extra/httpd-ssl.conf
```

Windows

```
#LoadModule ssl_module modules/mod_ssl.so
#Include conf/extra/httpd-ssl.conf
```

4. On an SSL enabled Portal server, it is recommended to either disable http or redirect http protocol to https.
   - To disable http protocol, edit `httpd.conf` file and remove VirtualHost section of portal configuration
   - To redirect http protocol to https, edit `httpd.conf` file, remove all entries of VirtualHost section of portal configuration and add following lines in same VirtualHost:

     ```
     ServerName   itanalyticsportal.<hostname>
     Redirect permanent / https://itanalyticsportal.<hostname>/
     ```

5. To ensure a secure web server, remove any port 80 VirtualHost sections from the `/opt/apache/conf/httpd.conf` file.

   This prevents the HTTP message headers from getting unencrypted if one end of the communication is using non-HTTPS protocols.

6. If a Virtual Host declaration is missing from the default Apache SSL configuration file, add the missing virtual host declaration to the configuration file. See the relevant section for instructions.
   - See “Configure virtual hosts for portal and / or data collection SSL” on page 254.
   - See “SSL Implementation for the Portal Only” on page 255.
   - See “SSL Implementation for Data Collection Only” on page 255.
   - See “SSL Implementation for Both the Portal and Data Collection” on page 256.

7. For each active virtual host section in the Apache SSL configuration file (httpd-ssl.conf), ensure that declaration lines beginning with the following are un-commented (they do not have a # at the beginning of the line):

```
SSLEngine
SSLCertificateFile (update certificate file details)
SSLCertificateKeyFile (update certificate key file details)
```

8. Run the deployCert utility as root user on the Portal server to save the ssl certificates configured with Apache in java keystore itanalytics.jks .

   This will be used while configuring SingleSignOn and Syslog over SSL.

   - Linux: /opt/aptare/utils/deployCert.sh update
   - Windows: C:\opt\aptare\utils>deployCert.bat update

9. Verify the Apache configuration is valid:, for Linux only

   ```
   # export LD_LIBRARY_PATH=/opt/apache/ssl/lib:$LD_LIBRARY_PATH
   (If https is enabled)

   # /opt/apache/bin/apachectl -t
   ```

10. If SSL is enabled, change the applicationUrl in portal.properties to https instead of http.
11. Restart Apache and both Tomcat (Portal and Data Collector) services.

    Linux

    ```
    /opt/aptare/bin/apache start
    /opt/aptare/bin/tomcat-portal start
    /opt/aptare/bin/tomcat-agent start
    ```

    Windows

    ```
    C:\opt\aptare\utils\startapache.bat
    C:\opt\aptare\utils\startagent.bat
    C:\opt\aptare\utils\startportal.bat
    ```

Update the default CipherSuite provided in the Apache Httpd configuration

IT Analytics Portal supports only one CipherSuite ECDHE-RSA-AES256-GCM-SHA384 by default if configured with https.

To update the Apache httpd setting with a different value, follow the steps described below.

1. Stop Apache httpd and Tomcat services.

   Linux

   ```
   /opt/aptare/bin/tomcat-agent stop
   /opt/aptare/bin/tomcat-portal stop
   /opt/aptare/bin/apache stop
   ```

   Windows

   ```
   C:\opt\aptare\utils\stopagent.bat
   C:\opt\aptare\utils\stopportal.bat
   C:\opt\aptare\utils\stopapache.bat
   ```

2. Update the `httpd-ssl-override.conf` file with required values for SSLCipherSuite.

   Linux: **/opt/apache/conf/extra/httpd-ssl-override.conf**

   Windows: **C:\opt\apache\conf\extra\httpd-ssl-override.conf**

   ```
   SSLCipherSuite <new list of Cipher Suites separated by colon>
   ```

3. Start Apache httpd and Tomcat services

   Linux

   ```
   /opt/aptare/bin/apache start
   /opt/aptare/bin/tomcat-portal start
   /opt/aptare/bin/tomcat-agent star
   ```

   Windows

   ```
   C:\opt\aptare\utils\startapache.bat
   C:\opt\aptare\utils\startagent.bat
   C:\opt\aptare\utils\startportal.bat
   ```

Configure virtual hosts for portal and / or data collection SSL

Refer to the following sections that are relevant for your environment.

- See “Update the web server configuration to enable SSL” on page 251.
- See “SSL Implementation for the Portal Only” on page 255.
- See “SSL Implementation for Data Collection Only” on page 255.

- See "SSL Implementation for Both the Portal and Data Collection" on page 256.

SSL Implementation for the Portal Only

1. Verify that there is a VirtualHost section with the **IP address assigned to the Portal host**. This section starts with the following lines. These lines must be present and enabled.

   ```
   <VirtualHost IP_ADDRESS_PORTAL:443>
   ServerName aptareportal.domainname:443
   Document Root /opt/aptare/portal
   ```

2. In the VirtualHost declaration, replace **IP_ADDRESS_PORTAL** with the IP address assigned to the Portal server.
3. If the Portal VirtualHost section is not found, the configuration for the Portal VirtualHost must be added.
4. If there is a configuration section for the Data Collection virtual host, ensure that this section is disabled by adding a # to the beginning of each line in the section, as shown below.

   ```
   #<VirtualHost itanalyticsagent.domainname:443>
   ```

5. Set the Document Root path to a valid path for the Web Server's OS.

   Linux

   ```
   /opt/aptare/portal
   ```

   Windows

   ```
   C:\opt\aptare\portal
   ```

SSL Implementation for Data Collection Only

1. Verify that there is a VirtualHost section for the data collection with the **IP address of the Data Receiver**. This section starts with the following lines. These lines must be present and enabled.

   ```
   <VirtualHost IP_ADDRESS_DATARCVR:443>
   ServerName aptareagent.domainname:443
   DocumentRoot /opt/aptare/datarcvr
   ```

2. Replace **IP_ADDRESS_DATARCVR** in the VirtualHost declaration with the IP address assigned to the Data Receiver.
3. If the data collection VirtualHost section is not found, the configuration for the data collection VirtualHost must be added.
4. If there is a configuration section for the Portal virtual host, ensure that this section is disabled by added a # to the beginning of each line in the section, as shown below.

   ```
   #<VirtualHost itanalyticsportal.domainname:443>
   ```

5. Set the **Document Root**path to a valid path for the Web Server's OS.

   Linux

   ```
   /opt/aptare/datarcvr
   ```

   Windows

   ```
   C:\opt\aptare\datarcvr
   ```

6. If implementing SSL for Data Collection, complete the following steps:

   Ensure that the Data Collector global properties file does not have the protocol (http or https) in the URL specified in that file.

   Linux

   ```
   /opt/aptare/datarcvrconf/
   collectorConfig.global.properties
   ```

   Windows

   ```
   C:\opt\aptare\datarcvrconf\
   collectorConfig.global.properties
   ```

SSL Implementation for Both the Portal and Data Collection

To implement SSL for both the Portal and Data Collection, the Portal server must be configured with two IP addresses, one for the Portal and one for Data Collection. The two required IP addresses may be implemented using two NICs. If only a single NIC is available, a virtual interface may be added for the second IP address.

See "Add a virtual interface to a Linux server" on page 266.

See "Add a virtual / secondary IP address on Windows" on page 267.

Verify the following from `/opt/apache/conf/httpd-ssl.conf` **on Linux host or from** `C:\opt\apache\conf\httpd-ssl.conf` **in case of Windows host:**

1. Verify there is a VirtualHost section with the **IP address assigned to the Portal host**. This section starts with the following lines. These lines must be present and enabled.

```
<VirtualHost IP_ADDRESS_PORTAL:443>
ServerName itanalyticsportal.domainname:443
Document Root /opt/aptare/portal
...
</VirtualHost>
```

2. Replace **IP_ADDRESS_PORTAL** in the VirtualHost declaration with the IP address assigned to the Portal server.
3. Verify there is a VirtualHost section with the **Data Receiver IP address**. This section starts with the following lines. These lines must be present and enabled.

```
<VirtualHost IP_ADDRESS_DATARCVR:443>
ServerName itanalyticsagent.domainname:443
DocumentRoot /opt/aptare/datarcvr
...
</VirtualHost>
```

4. Replace **IP_ADDRESS_DATARCVR** in the Virtual Host declaration with the IP address assigned to the Data Receiver.
5. Set the Document Root paths to valid paths for the Web Server's OS.

   Linux

```
/opt/aptare/portal
/opt/aptare/datarcvr
```

   Windows

```
C:\opt\aptare\portal
C:\opt\aptare\datarcvr
```

Enable / Disable SSL for a Data Collector

Once you globally configure SSL, you can change the settings for individual Data Collectors. This provides the capability of supporting a mix of both http and https among your Data Collector servers.

A system property **-Dhttps.protocols=TLSv1.2 and TLSv1.3** are included in the startup and checkinstall scripts (aptaredc.bat/sh, checkinstall.bat/sh) to support TLS 1.2 and TLS 1.3.

To enable and disable SSL for a specific Data Collector:

1. In the IT Analytics Portal, navigate to **Admin** > **Data Collection** > **Collector Administration**.
2. Double-click a Data Collector to view the existing settings or click **Add** to add a Data Collector.

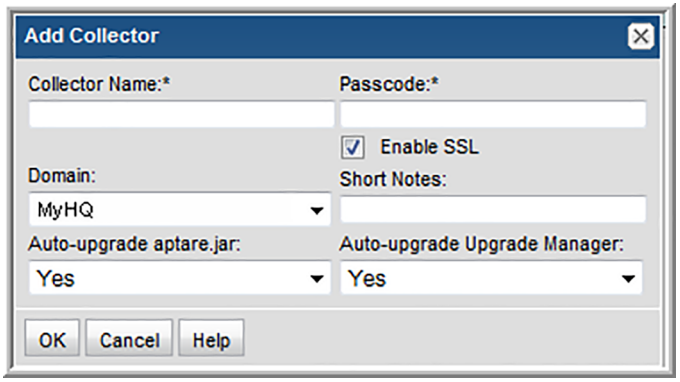

3. Check the **Enable SSL** checkbox.

   Both secure (SSL) and non-secure Data Collectors can send data to the same Portal. Check this box to select the secure communication protocol (https) that the Data Collector will use.

   This check box will not appear in the dialog box if SSL is not enabled in your environment. The Portal data receiver must be listening for https traffic; for example: https://agent.mycollector.com

4. Click **OK** to save the setting.

See “Configure the Data Collector to trust the certificate” on page 261.

Enable / Disable SSL for emailed reports

When emailing reports, an Add a Live Link option provides the capability of having a hyperlink (View this report in the Portal) in the email to take the user directly to the Portal. In environments where SSL is enabled, a configuration change is required in the portal.properties file to ensure that this link is secure.

Linux: `/opt/aptare/portalconf/portal.properties`

Windows: `C:\opt\aptare\portalconf\portal.properties`

1. In the `portal.properties` file, find the following section and update the value portal.applicationUrl to replace http with https.

   Example:

   ```
   #The Portal environment
   portal.sessionTimeout=3600
   portal.applicationUrl=https://aptareportal.mycompany.com
   ```

2. Restart the Portal service.

   See "Starting and stopping portal server software" on page 44.

Test and troubleshoot SSL configurations

The following sections cover common configuration for testing and troubleshooting.

Test if SSL is set up for the Portal

1. Enter `https://itanalyticsportal.domain.com/` in a browser.

   The Portal login page should display.

Test if SSL is set up for Data Collection

1. Enter `https://itanalyticsagent.domain.com/` in a browser. The following should display: Cohesity IT Analytics Data Receiver.
2. Enter `https://itanalyticsagent.domain.com/servlet/util/`in a browser.

   The error message, **GET not SUPPORTED. Illegal Operation!!!**, should display.

Create a self-signed SSL certificate

Use OpenSSL open source software to create your self-signed certificate. For more information on creating self-signed certificates using OpenSSL, refer to the FAQs and documentation on the OpenSSL site at www.openssl.org.

The instructions and examples in this section are applicable for the Linux operating system. OpenSSL also may be used with the Windows operating system. Check the OpenSSL web site for specific instructions. Note that the certificate is independent of the operating system under which it was created. A self-signed certificate created on a Linux computer may be installed on a Windows web server.

You can create a self-signed certificate with multiple options depending on how you want to configure your certificate. Use the following OpenSSL command to create a self-signed certificate. The command creates two files: **server.key** and **server.crt**. You must install these files on the IT Analytics web server.

```
openssl req -newkey rsa:2048 -nodes -keyout server.key
-x509 -days 365 -out server.crt
```

where

-x509 is used to create a certificate as opposed to a certificate request that is sent to a certificate authority

-days determines the number of days that the certificate is valid

-newkey rsa:2048 sets the key as 2048-bit RSA

-nodes specifies that no passkey will be used

-keyout specifies the name of the key file

-out specified the name of the certificate file

Example:

```
openssl req -newkey rsa:2048 -nodes -
keyout server.key -x509 -days 365 -out server.cert
```

Specify the following parameters for the command:

- Country Name (2 letter code) [XX]: :
- State or Province Name (full name) []:
- Locality Name (eg, city) [Default City]:
- Organization Name (eg, company) [Default Company Ltd]:
- Organizational Unit Name (eg, section) []:

- Common Name (eg, your name or your server's hostname) []:
- Email Address []

Note: The use of the **-nodes** option in the previous example creates a certificate that does not require a pass phrase. This makes it easier to install and use the certificate, but weakens the security of the certificate. If the certificate is created with a pass phrase, it must be entered when the certificate is installed and used.

The actual certificates get installed and configured on the Apache web server, however, in cases where the issuing certificate authority (CA) is not automatically trusted (such as self-signed certificates), the certificates need to be imported and trusted on the Data Collector server.

Once the self-signed certificates have been created, configure the Data Collector to trust the certificate.

Configure the Data Collector to trust the certificate

In cases where the certificate authority (CA) is not trusted, as may be the case when using a self-signed or unknown certificate, both the Data Collector and the Upgrade Manager will need to have the certificate imported into the keystore to ensure that the Data Collector can communicate using SSL.

See “Keystore file locations on the Data Collector server” on page 261.

See “Import a certificate into the Data Collector Java keystore” on page 262.

Keystore file locations on the Data Collector server

If you are not running the Data Collector installer from the default collector location for an upgrade (`/opt/aptare` or `C:\Program Files\Aptare`) or for a fresh installation (`/usr/openv/analyticscollector/` or `C:\Program Files\Veritas\AnalyticsCollector`), substitute the appropriate path for <APTARE_HOME> in the command path in the following commands:

- Linux Data Collector: `<APTARE_HOME>/java/lib/security/cacerts`
- Windows Data Collector: `<APTARE_HOME>\java\lib\security\cacerts`
- Linux Upgrade Manager: `<APTARE_HOME>/upgrade/upgradeManager/jre/lib/security/cacerts`

- Windows Upgrade Manager:
  `<APTARE_HOME>\upgrade\upgradeManager\jre\lib\security\cacerts`

See "Import a certificate into the Data Collector Java keystore" on page 262.

Import a certificate into the Data Collector Java keystore

Use the following steps to add an SSL certificate to the Java keystore for a Data Collector. Some servers, such as vSphere, require a certificate for connection while communicating with SSL.

See "Configure the Data Collector to trust the certificate" on page 261.

See "Keystore file locations on the Data Collector server" on page 261.

1. Copy the certificate file (server.crt file) to the Data Collector.
2. If you are not running the Data Collector installer from the default collector location for an upgrade (`/opt/aptare` or `C:\Program Files\Aptare`) or for a fresh installation (`/usr/openv/analyticscollector/` or `C:\Program Files\Veritas\AnalyticsCollector`), substitute the appropriate path for <APTARE_HOME> in the command path in the following commands:

   Linux:

   ```
   <APTARE_HOME>/java/bin/keytool -importcert -alias "somealias"
   -file server.crt -keystore <APTARE_HOME>/java/lib/security/cacerts
   <APTARE_HOME>/java/bin/keytool -import -alias "somealias" -file
   server.crt -keystore
   <APTARE_HOME>/upgrade/upgradeManager/jre/lib/security/cacerts
   ```

   Windows:

   ```
   "<APTARE_HOME>\java\bin\keytool" -importcert -alias "somealias"
   -file server.crt -keystore
   "<APTARE_HOME>\java\lib\security\cacerts"
   "<APTARE_HOME>\java\bin\keytool" -import -alias "somealias" -file
    server.crt -keystore
   "<APTARE_HOME>\upgrade\upgradeManager\jre\lib\security\cacerts"
   ```

3. When prompted, enter the default password to the keystore:

   ```
   changeit
   ```

   The results will be similar to the following example:

```
Enter keystore password:
......
Certificate Shown here
......
Trust this certificate? [no]: yes
```

4. Once completed, run the following **keytool** command to view a list of certificates from the keystore and confirm that the certificate was successfully added. The certificate fingerprint line displays with the alias name used during the import.

   Linux:

   ```
   <APTARE_HOME>/java/bin/keytool -list -keystore
   <APTARE_HOME>/java/lib/security/cacerts
   ```

   Windows:

   ```
   "<APTARE_HOME>\java\bin\keytool" -list -keystore
   "<APTARE_HOME>\java\lib\security\cacerts"
   ```

   Sample Linux Output

   ```
   Enter keystore password:
   Keystore type: JKS
   Keystore provider: SUN
   Your keystore contains 79 entries
   digicertassuredidrootca, Apr 16, 2008, trustedCertEntry,
   Certificate fingerprint (SHA1):
   05:63:B8:63:0D:62:D7:5A:BB:C8:AB:1E:4B:DF:B5:A8:99:B2:4D:43
   trustcenterclass2caii, Apr 29, 2008, trustedCertEntry,
   Certificate fingerprint (SHA1):
   AE:50:83:ED:7C:F4:5C:BC:8F:61:C6:21:FE:68:5D:79:42:21:15:6E
   .....
   ```

Keystore on the portal server

A separate keystore is used on the Portal Server to store and manage certificates when SSL is enabled. Use the Keystore Utility (deployCert) to manage your SSL certificate on the Portal Server. This utility is only needed when the portal is configured for https. Using the utility you can:

- Add the certificate
- Update the certificate
- Download the certificate from the keystore.

The utility is located here:

- linux: /opt/aptare/utils/deployCert.sh
- windows: C:\opt\aptare\utils\deployCert.bat

The keystore is located here:

/opt/aptare/portalconf/itanalytics.jks

After running the Keystore Utility (deployCert), restart Portal Services.

Features that Require the SSL Certificate

Depending on your installation scenario, you may need to manually add the SSL certificate to the Portal Keystore to use Security Assertion Markup Language (SAML)/Single-Sign On (SSO) or syslog with SSL operations .

Add a Certificate into the Portal Keystore

Use the Utility (deployCert) to add an SSL certificate to the Portal keystore. After configuring https, run the **add** command to add certificates into the keystore. Execute the following from the command window:

Note: Root user can execute deployCert utility.

Linux

```
cd /opt/aptare/utils

./deployCert add
```

If the certificate key is password protected, use:

```
./deployCert.sh add -p <passphrase>
```

Windows

```
cd c:\opt\aptare\utils

deployCert.bat add
```

If the certificate key is password protected, use:

```
deployCert.bat add -p <passphrase>
```

Update a Certificate in the Portal Keystore

Use the Utility (deployCert) to update an SSL certificate to the Portal keystore. Run the **update** command to update certificates in the keystore. Execute the following from the command window:

Linux

```
cd /opt/aptare/utils

# ./deployCert.sh update
```

If the certificate key is password protected use:

```
# ./deployCert.sh update -p <passphrase>
```

Windows

```
cd c:\opt\aptare\utils

deployCert.bat update
```

If the certificate key is password protected use:

```
deployCert.bat update -p <passphrase>
```

Download a Certificate from the Portal Keystore

Use the Utility (deployCert) to download an SSL certificate from the Portal keystore. Before configuring SSL for syslog in the Portal, you must add the portal certificate into the syslog server as a trusted certificate. Use the download command to retrieve the certificate for the syslog server. Run the **download** command to download certificates from the keystore. The certificate will be downloaded into the current working directory. Execute the following from the command window:

Linux

```
cd /opt/aptare/utils

./deployCert download
```

Windows

```
cd c:\opt\aptare\utils

deployCert.bat download
```

Add a virtual interface to a Linux server

The standard APTARE server configuration uses two virtual hosts on the server. One host, identified by the sub-domain **itanalyticsportal**, handles Portal requests to deliver IT Analytics administration and reporting functionality. The second host, identified by the sub-domain **itanalyticsagent**, handles data collection functionality between the data collection agent and the various devices that report to the agent. These virtual hosts are defined in the Apache configuration file; the sub-domain names are used to identify the each host.

When using SSL, unique IP addresses must be assigned to each virtual host. Therefore, if SSL is to be enabled for both the Portal and Data Collection, two IP addresses are required. Two IP addresses can be assigned using two NICs or, on a Linux server, a virtual interface can be created to assign two IP addresses to a single NIC.

See "Configure virtual hosts for portal and / or data collection SSL" on page 254.

1. To verify the number of IP addresses assigned to a Linux server, use the following command:

   ```
   ifconfig -a
   ```

   Example result of the **ifconfig -a** command:

   ```
   eth0      Link encap:Ethernet  HWaddr 08:00:27:71:44:C4
             inet addr:10.0.2.15  Bcast:10.0.2.255
   Mask:255.255.255.0
             UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
             RX packets:310 errors:0 dropped:0 overruns:0 frame:0
             TX packets:372 errors:0 dropped:0 overruns:0 carrier:0

             collisions:0 txqueuelen:1000
             RX bytes:63235 (61.7 KiB)  TX bytes:28143 (27.4 KiB)
   lo        Link encap:Local Loopback
             inet addr:127.0.0.1  Mask:255.0.0.0
             UP LOOPBACK RUNNING  MTU:16436  Metric:1
             RX packets:8762 errors:0 dropped:0 overruns:0 frame:0
            TX packets:8762 errors:0 dropped:0 overruns:0 carrier:0

             collisions:0 txqueuelen:0
             RX bytes:5422509 (5.1 MiB)  TX bytes:5422509 (5.1 MiB)
   ```

2. You must have two Ethernet connections, identified by the **eth0** label. To add a virtual interface on a Linux server, with a second IP address, to the existing Ethernet interface, use the following command:

   ```
   ifconfig eth0:0 111.222.333.444
   ```

   where

   111.222.333.444 is the new IP address for the virtual interface.

3. You must add a file to the network scripts to recreate the virtual interface when the server is rebooted. If the IP address assigned to the **eth0** interface is static, make a copy of the **ifcfg-eth0** file in **/etc/sysconfig/network-scripts** and name it **ifcfg-eth0:0**.
4. Update the IP address in **ifcfg-eth0:0** to be the new IP address assigned to the virtual interface.
5. If the IP address in the **eth0** interface is dynamically assigned, as indicated by the line **BOOTPROTO=dhcp** in the **ifcfg-eth0** file, create a file named **ifcfg-eth0:0** with the following lines:

   ```
   DEVICE=eth0:0
   IPADDR=111.222.333.444
   ```

6. Finally, update your DNS server so that the new IP address is mapped to the data collection URL (for example, **itanalyticsportal.<domainname>**).

Add a virtual / secondary IP address on Windows

To add a Virtual IP Address on Windows, go to:

Control Panel > Network and Internet > Network and Sharing Center > Change adapter settings

Right-click on a Network connection and select **Properties**.

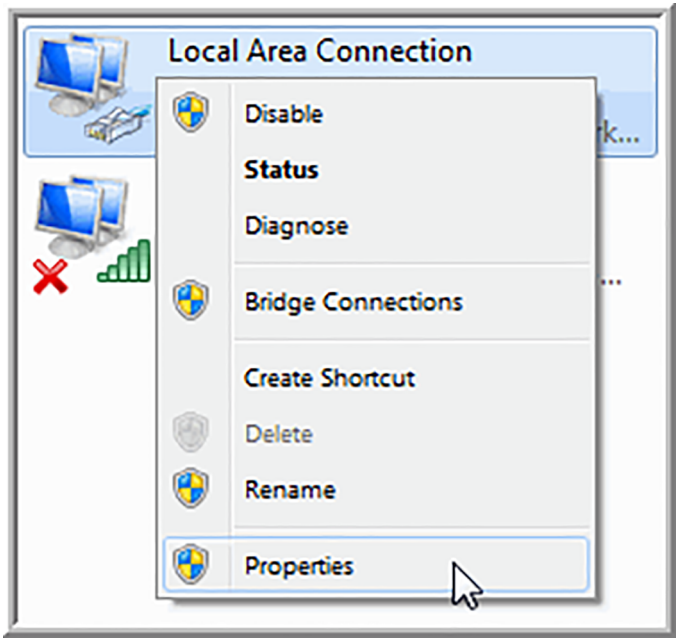

Take the following steps to configure a secondary IP address.

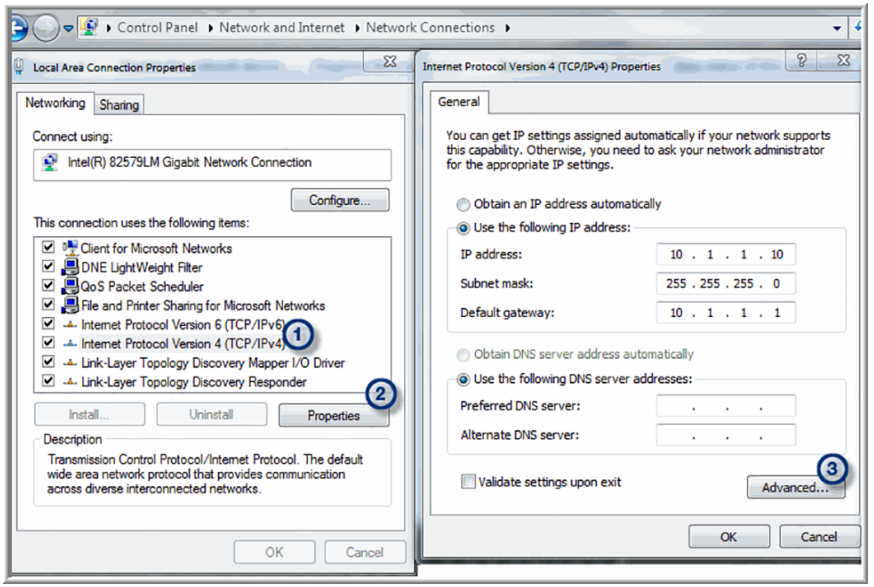

1. Select the TCP/IP connection.

2. Click **Properties**.
3. For the configured IP address, click **Advanced**.

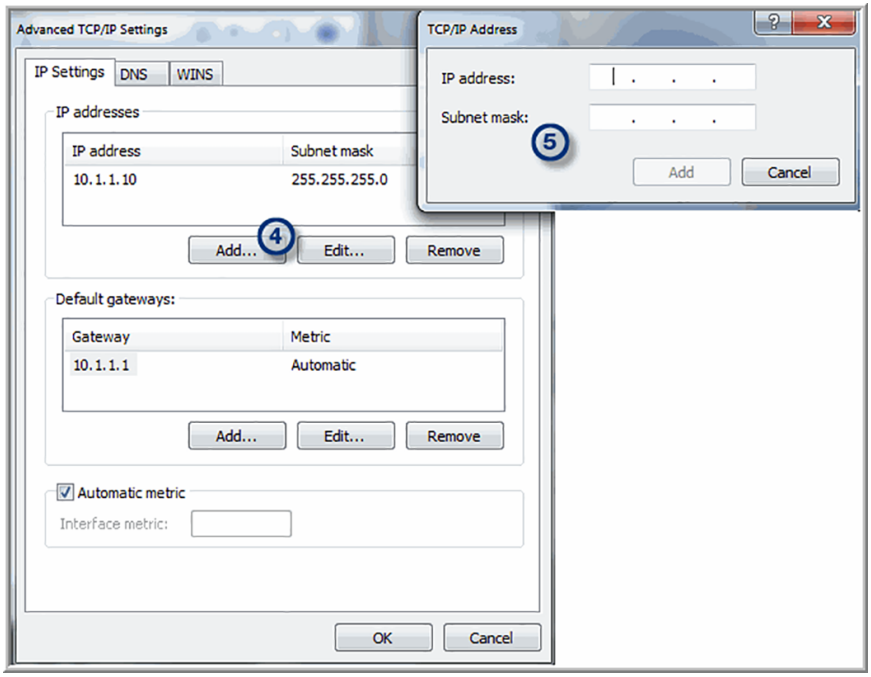

4. In the Advanced TCP/IP Settings window, click **Add**.
5. Enter the **IP address** and **Subnet mask** and click **Add**.

Chapter 23

Portal properties: Format and portal customizations

This chapter includes the following topics:

Introduction

This section covers customizations for the portal that are not available through the user interface. Use Custom Parameters to add/edit and delete these properties.

Note: Prior to version 10.3, customizations to the Portal were made using a file, portal.properties. Not all of those settings are displayed in the System Configuration feature. If you upgrade from a version prior to 10.3, those properties are displayed and automatically populated in the Custom Parameters.

Configuring global default inventory object selection

To globally configure the selection of default Inventory objects for users, modify the portal.properties files. This is useful for filtering environments with large volumes of data that may be impacted by browser limitations. For new users who have never logged into the portal, the objects defined with this setting are shown selected when they log in. For existing users, this property can be used to reset a users environment when large volumes of data can present issues with certain browsers.

The portal.properties file is located here:

Linux:

```
/opt/aptare/portalconf/portal.properties
```

Windows:

```
C:\opt\aptare\portalconf\portal.properties
```

1. Update the following section:

```
portal.ocn.defaultVisibleObjectType=HOST,ARRAY,SWITCH,BACKUPSERVER,
VM_SERVER,VM_GUEST,
DEDUPLICATION_APPLIANCE,DATASTORE,EC2_INSTANCE,S3_BUCKET,
AZURE_STORAGE_ACCOUNT,AZURE_VIRTUAL_MACHINE
```

2. Restart the Tomcat Portal services after making your modification.

Restricting user IDs to single sessions

To restrict a user ID from signing on multiple times using different browsers on the same machine or the same browser on different machines, modify the portal.properties file. The last browser session with the user ID to login will have access to the portal. Other browser sessions with same user ID will be logged out.

The portal.properties file is located here:

Linux:

```
/opt/aptare/portalconf/portal.properties
```

Windows:

```
C:\opt\aptare\portalconf\portal.properties
```

1. Update the following section:

   ```
   portal.security.allowUserToLoginMultipleTimes=false
   ```

2. Restart the Tomcat Portal services after making your modification.

Customizing date format in the report scope selector

To customize the date format displayed in the report scope selector for all Portal users, you can modify the portal.properties file. For example, you can set the date to display: dd/MM/yyyy or MM/dd/yyyy.

The portal.properties file is located here:

Linux:

```
/opt/aptare/portalconf/portal.properties
```

Windows:

```
C:\opt\aptare\portalconf\portal.properties
```

1. Update the following section:

   ```
   #Formatters that define specific presentations of numbers and dates
   formatter.decimalPlaces=2
   fileSize.base2=true
   formatter.number=###,###,##0
   formatter.date=MMM dd, yyyy hh:mm:ssa
   formatter.dateZone=MMM dd, yyyy hh:mm:ssa z
   formatter.yearMonth=MMM dd
   formatter.groupByDate=MMM dd
   formatter.designerDate=MM/dd/yyyy
   formatter.currency=$ ###,###,##0.00
   ```

2. Restart the Tomcat Portal services after making your modification.

Customizing the maximum number of lines for exported reports

When you export or email a large report, APTARE IT Analytics limits the maximum number of lines to 20,000. The report truncates when that value is exceeded. The report can still be exported or emailed, but will contain a message that the report has been truncated.

To change the 20,000 line limit for all Portal users, modify the portal.properties file:

Linux:

```
/opt/aptare/portalconf/portal.properties
```

Windows:

```
C:\opt\aptare\portalconf\portal.properties
```

1. Add the following line:

   ```
   portal.report.maxRowsExported=<enter new limit value here>
   ```

   Where the <new limit value> is the number of rows greater than 20,000 that your report export requires. For example, if your report has 36,000 rows enter a number greater than 36000. Note that the new limit value cannot contain commas or decimal points. Keep in mind that Portal server performance can degrade considerably for very large reports. For very large reports, you may want to segment the scope into multiple reports.
2. Restart the Tomcat Portal services after making your modification.

Customizing the total label display in tabular reports

To customize the minimum number of records needed to display the Total label in a report, you can modify the portal.properties file. The default value is 10.

The portal.properties file is located here:

Linux:

```
/opt/aptare/portalconf/portal.properties
```

Windows:

```
C:\opt\aptare\portalconf\portal.properties
```

1. Add the following line:

   ```
   portal.rowCountDisplayMinimum = <enter numeric value>
   ```

2. Restart the Tomcat Portal services after making your modification.

Customizing the host management page size

In the Portal, the Host Management page (Admin > Advanced > Host Management) displays 200 rows by default. You can change the default value by modifying the portal.properties file. System performance will be impacted if you increase the number of rows past the 200 value.

1. To customize the number of rows displayed on the Host Management page access the portal.properties file:

   Linux:

   ```
   /opt/aptare/portalconf/portal.properties
   ```

   Windows:

   ```
   C:\opt\aptare\portalconf\portal.properties
   ```

2. Add the following line and set the page size value:

   ```
   portal.hostManagementPageSize=xxxx
   ```

   For example, to only display 50 rows on the page, it would appear as follows:

   ```
   portal.hostManagementPageSize=50
   ```

Customizing the path and directory for file analytics database

You can customize the location of the File Analytics database. The default paths are as follows:

Linux:

```
/opt/aptare/fa
```

Windows:

```
C:\opt\aptare\fa
```

1. To customize the path for the location of the File Analytics database access the portal.properties file:

   Linux:

   ```
   /opt/aptare/portalconf/portal.properties
   ```

   Windows:

   ```
   C:\opt\aptare\portalconf\portal.properties
   ```

2. Add the following line and enter the new path:

   ```
   fa.root=/opt/aptare/fa
   ```

   Example

   Linux:

   ```
   fa.root=/opt/aptare/fa_db
   ```

   Windows:

   ```
   fa.root=D:\opt\aptare\fa
   ```

3. Restart the Tomcat Portal services.

Configuring badge expiration

Configure the expiration of **NEW** badges in the **Home**section of the **Reports** tab. By default, **NEW**badges will no longer display after 14 days. You can change that value by adding a line to the portal.properties file.

1. Revise the default expiration period for **NEW**badges on the **Reports**tab using the portal.properties file:

   Linux:

   ```
   /opt/aptare/portalconf/portal.properties
   ```

   Windows:

   ```
   C:\opt\aptare\portalconf\portal.properties
   ```

2. Add the following line:

   ```
   cloudTemplateNewBadgeExpireInDays = <enter numeric value>
   ```

3. Restart the Tomcat Portal Services.

Configuring the maximum cache size in memory

The cache can retain up to 0.5 GB of reporting data and if it reaches capacity, it frees up space for new reports by purging the data for the least frequently used reports. You can change the retention value by adding a line to the portal.properties file.

1. Revise the maximum cache value for reporting data using the portal.properties file:

   Linux:

   ```
   /opt/aptare/portalconf/portal.properties
   ```

   Windows:

   ```
   C:\opt\aptare\portalconf\portal.properties
   ```

2. Add the following line:

   ```
   portal.reports.cache.maxSizeInMemory
   ```

   The unit of measure is bytes.

   Example:

   ```
   portal.reports.cache.maxSizeInMemory=536870912
   ```

3. Restart the Tomcat Portal services.

Configuring the cache time for reports

The cache retains reporting data and if it reaches capacity, it frees up space for new reports by purging the data for the least frequently used reports. Purging also occurs when a cached report is more than 24 hours old. You can change the value by adding a line to the portal.properties file.

1. Revise the cache purge time for reporting data using the portal.properties file:

   Linux:

   ```
   /opt/aptare/portalconf/portal.properties
   ```

   Windows:

   ```
   C:\opt\aptare\portalconf\portal.properties
   ```

2. Add the following line:

   ```
   portal.reports.cache.timeOut
   ```

   The unit of measure is seconds.

3. Restart the Tomcat Portal services.

Configuring LDAP to use active directory (AD) for user group privileges

This attribute enables administrators to control access to privileges through user creation and group membership assignments in their enterprise LDAP. Create AD user groups with the same name in the Portal, add the following parameters, then when new users are added to AD, they will be automatically added to the corresponding portal-based user groups and inherit the group privileges.

Prerequisites and Assumptions

- External LDAP may be used.
- User Groups with the same names must exist in both Active Directory and the portal.
- If a user is assigned to a User Group they acquire all associated privileges of that User Group.

LDAP authorization is enabled by adding two lines to the portal.properties:

- ldap.authorization - Indicates if Active Directory user groups should be used for authorization. When set to true, the user permissions from a group perspective are derived solely from the AD groups they belong to that have a corresponding group defined withinIT Analytics. If set to false (the default), user permissions are derived fromIT Analytics.
- ldap.newUserDomain - Must be a domain that exists in the APTARE IT Analytics portal. This is assigned during user creation and is based on the domain of the user's home group.

Note: For **ldap.authorization=false**: AD users must be manually created in portal so that they can login using AD credentials. For **ldap.authorization=true**: AD users will be automatically created in the portal.

---

1. Enable LDAP authorization using the portal.properties file:

   Linux:

```
/opt/aptare/portalconf/portal.properties
```

Windows:

```
C;\opt\aptare\portalconf\portal.properties
```

2. Add the following lines:

```
ldap.authorization=true

ldap.newUserDomain=<string>
```

3. Restart the Tomcat Portal services.

Chapter 24

Data retention periods for SDK database objects

This chapter includes the following topics:

Data retention periods for SDK database objects

This section covers the procedure for modifying data retention periods for systems collected by Data Collectors deployed via the SDK.

Because data collection continuously adds data to the database, management of the database size becomes an essential task. **For objects created via the SDK**, a data retention period can be defined along with the database schema, to align with your IT Analytics reporting requirements. Data retention (purging rules) described in this section apply only to user-defined objects defined by an SDK project; API and database objects are handled by system parameters.

Purging rules (data retention periods) are specific to an object, such as array disk performance, and they are used to maintain corresponding database tables. These rules are configured at the beginning of the object's definition in the schema template json file.

Best Practice: Configure retention days for objects that accumulate historical data over time, such as array performance. Other objects should not have data retention days.

In multi-tenancy environments, where IT Analytics domains partition client data, data retention periods can be configured for each domain. The data retention period for a domain applies to that domain and all of its sub-domains. For systems collected by traditional Data Collectors, use the **Data Retention** tab located: **Admin>Advanced>System Configuration**.

To enable purging for a user-defined object for a specific domain, take the following steps.

1. Log in to the Portal Server as user **aptare**.
2. Type the following command:

   ```
   sqlplus <portal_user>/<portal_password>@//localhost:1521/scdb
   where <portal_user> and <portal_password> are portal credentials
    to connect to the database.
   ```

3. At the command line, execute the following SQL statement, substituting relevant values for the variables shown < > in the syntax.

   See "Find the domain ID and database table names" on page 285.

   See the section called "Retention Period Update for Multi-Tenancy Environments Example" on page 293.

   See "Retention period update for SDK user-defined objects example" on page 285.

```
INSERT INTO apt_purge_rules (table_name, product_type,
table_description, date_column_name, retention_days,
default_retention_days, domain_id, creation_date, last_updated)
SELECT table_name, product_type, table_description,
date_column_name, <NewRetentionDays>, default_retention_days,
<domain_id_value>, SYSDATE, SYSDATE
  FROM apt_purge_rules
 WHERE domain_id IS NULL
   AND table_name = <tableName>;
Commit;
```

Data aggregation

Overview

Storage, Fabric, and Virtualization are just a few of the subsystems that IT Analytics supports in terms of performance metric gathering.

When regulating the growth of performance metrics data, IT Analytics purges old data periodically using customizable retention parameters. But if the retention periods need to be longer than the defaults, this data can eventually take up a lot of space on the disc where the database is stored.

Additionally, when the data in the underlying tables expand exponentially, data retrieval becomes difficult.

With the release of version 11.2.02 for **Capacity** and **Switches** sub-systems, IT Analytics now provides **Data Aggregation** for performance metrics data to manage the aforementioned scenarios. In the subsequent releases, it will also be expanded to include Virtualization and Cloud subsystems.

About

Higher data retention durations are made possible by data aggregation, which aggregates data for performance measures without increasing the database's required disc space.

The existing data in the table is aggregated as per the preselected aggregation metrics (average, max, min, count, sum, standard deviation etc.) for a specific time resulting in reduced number of records.

For instance, IT Analytics allows the collection of performance statistics for logical units within 30 seconds. This can lead to 28 million records each day for 10,000 LUNs. Without aggregation, this table's retention duration of one year will produce 10 billion records. While for the same retention, this data can be reduced by ten times with aggregation enabled.

Advantages

The following are the advantages of data aggregation.

- **Disk space**: Depending on the frequency of data collection, data aggregation can potentially shrink the quantity of raw data by up to 10 times, thus lowering the impact of the enormous performance metrics data.
- **Quick report response**: Report queries against the database processes more quickly since aggregated values have already been persisted.
- **Data retrieval**: Data retrieval is made simpler by lowering the size of the data.

Pre-requisites

The following are the pre-requisites to enable data aggregation in IT Analytics:

Enable data aggregation

To enable the data aggregation, follow the steps as below:

1. Navigate to *Administration >> System Configuration*. **System Configuration** dialog box is displayed.
2. Click **Database Administration** tab.
3. Select the **Enable data aggregation** check box option.

Table partitioning

The database tables must be partitioned in order to aggregate millions of records (composite partitioning Interval-List). The existing data in the specified non-partitioned tables will be transferred to the partitioned schema.

Meta data entry

In order for Data Aggregation to know which columns require what types of aggregation and levels of aggregation, a one-time script execution is required to load the meta data for these chosen tables.

Data aggregation will begin with the nightly purge cycle after all of the above pre-requisite steps have been fulfilled. If the table is selected for aggregation by partitioning and metadata entry, the data for that table will first be aggregated before being purged, based on level-wise retention parameters.

Note: Cohesity advise performing an environment assessment before enabling data aggregation. After the assessments, Cohesity includes separate scripts to divide any current tables that will benefit from data aggregation. Contact Veritas Support.

Data aggregation and retention levels

Data aggregation is a multi-level process. Following are the levels supported by IT Analytics.

Data aggregation levels

- **Level 0 (Zero)**: This is the level where the data is persisted. This can be at any interval of time as per probes schedules.
- **Level 1**: This is the first aggregation level which is set to 5 minutes. Any persistence level data which has been collected at a time interval lower than 5 minutes, that data will be aggregated in this level.

  For example: If LUN performance data is currently being collection at 30 seconds interval, it will get aggregated to 5 minutes aggregating 10 records to 1 record. But if the persistence collection frequency is more than 5 minutes, the same data will be inserted in the table and marked as Level 1.
- **Level 2**: This is the second aggregation level which is set to 15 minutes. Any level 1 data which has been collected at a time interval lower than 15 minutes, will get aggregated to this level.
- **Level 3**: This is the third aggregation level which is set to 30 minutes. Any level 2 data which has been collected at a time interval lower than 30 minutes will get aggregated to this level.
- **Level 4**: This is the fourth aggregation level which is set to 60 minutes. Any level 3 data which has been collected at a time interval lower than 60 minutes will get aggregated to this level.

Data retention levels

As mentioned above, IT Analytics supports data aggregation at 5 levels. Further, each level has their own data retention levels for each sub-systems.

- **Level 0 (Zero) retention**: The level is set to 45 days. The data, collected from day One to 45th day, will be persisted at level zero. However, the 46th day data collected at frequency less than 5 minutes will get aggregated. If the frequency is equal to or greater than 5 minutes, the same data will be re-inserted without aggregation and will be marked as level 1.
- **Level 1 retention**: The level is set to 90 day from the 46th day of post level 1 aggregation till 90th day. These records will be kept at level 1.
- **Level 2 retention**: The level is set to 180 days from 91st day to post level 2 aggregation till 180th day. These records will be kept at level 2.
- **Level 3 retention**: The level is set to 360 days from 181st day to post level 3 aggregation till 360th day. These records will be kept at level 3.

- **Level 4 retention**: The Level 4 retention depends on the retention period of the table.

  For example, for the LUN performance metrics table, if the retention is set to 15 months, then records at level 4 aggregation will be maintained from 361st day to 450th day.

  If it is kept at 365 days, then these records will be retained for 5 days only and then purged.

  If the table's final retention period is kept lower than the level 4 retention period, then no data will be purged till level 4 is achieved, making both retention periods equal.

Note: The following is an example to enable the data aggregation option.

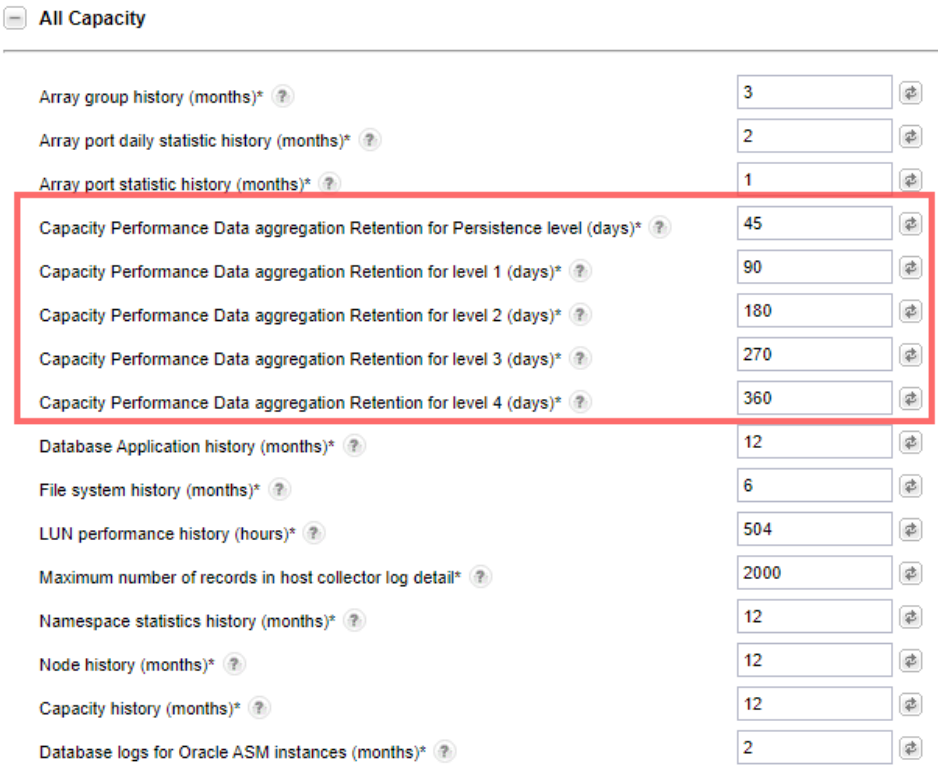

Find the domain ID and database table names

<domain_id_value>

Most environments have only one Domain ID; however, multi-tenancy environments, such as Managed Service Providers (MSPs), will have a different Domain ID for each of their customers.

To list the currently configured Domain IDs, use the following SQL SELECT statement:

SQL> SELECT * from apt_domain;

<tableName>

Find the relevant database table name for the retention period you want to change

See "SDK user-defined database objects" on page 285.

Retention period update for SDK user-defined objects example

The following example illustrates an update for a Pure Storage performance database table.

```
INSERT INTO apt_purge_rules (table_name, product_type,
table_description, date_column_name, retention_days,
default_retention_days, domain_id, creation_date, last_updated)
SELECT table_name, product_type, table_description, date_column_name,
 50, default_retention_days, 100007, SYSDATE, SYSDATE
  FROM apt_purge_rules
 WHERE domain_id IS NULL
   AND table_name = 'SDK_PURE_STORAGEARRAY_PERF';
Commit;
```

SDK user-defined database objects

The data retention period for various database tables often can be modified to maintain reasonable table sizes. Data retention periods can be modified for historical and performance data. Reference the following sections for details:

- See "Capacity: default retention for basic database tables" on page 286.
- See "Capacity: default retention for EMC Symmetrix enhanced performance" on page 287.
- See "Capacity: Default retention for EMC XtremIO" on page 288.

- See “Capacity: Default retention for Dell EMC Elastic Cloud Storage (ECS)” on page 288.
- See “Capacity: Default retention for Windows file server” on page 289.
- See “Capacity: Default retention for Pure Storage FlashArray” on page 290.
- See “Cloud: Default retention for Amazon Web Services (AWS)” on page 290.
- See “Cloud: Default retention for Microsoft Azure” on page 291.
- See “Cloud: Default retention for OpenStack Ceilometer” on page 291.

Capacity: default retention for basic database tables

Table 24-1 Default retention for basic database tables

Table Description	Table Name	Default Retention	Notes
File system log	aps_file_system_log	6 months	KEEP_FILE_SYSTEM_LOG_MONTHS
LUN performance log	apt_lun_perform_log	504 hours	KEEP_LUN_RAW_PERFORM_HOURS
Array group log	aps_array_group_log	12 months	KEEP_ARRAY_GROUP_LOG_MONTHS
Array port statistics log	aps_array_port_stats_log	1 day	KEEP_ARRAY_PORT_STAT_LOG_DAYS
Array port statistics daily log	aps_array_port_stats_daily_log	2 months	KEEP_ARRAY_PORT_STAT_DAILY_LOG_MONTHS
Capacity chargeback log	aps_chargeback_log	24 months	KEEP_SRM_CHARGEBACK_LOG_MONTHS
Host processing log	apt_host_process_log	15	KEEP_HOST_PROCESS_LOG_DAYS

Capacity: default retention for EMC Symmetrix enhanced performance

Table 24-2

Table Description	Table Name	Default Retention	Notes
Array Performance	SDK_esym_ARRAY_PERF	21 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
Back-end Director Performance	SDK_esym_be_director_PERF	21 days	
Database Performance	SDK_esym_database_PERF	21 days	
Device Group Performance	SDK_esym_device_group_PERF	21 days	
Disk Group Performance	SDK_esym_disk_group_PERF	21 days	
Disk Performance	SDK_esym_disk_PERF	21 days	
Disk Tech Pool Performance	SDK_esym_disk_tchpool_PERF	21 days	
Front-end Director Performance	SDK_esym_fe_director_PERF	21 days	
Front-end Port Performance	SDK_esym_fe_port_PERF	21 days	
Storage Group Performance	SDK_esym_storage_grp_PERF	21 days	
Storage Tier Performance	SDK_esym_storage_tier_PERF	21 days	
Thin Pool Performance	SDK_esym_thin_pool_PERF	21 days	
Thin Tier Performance	SDK_esym_thin_tier_PERF	21 days	

Table 24-2 *(continued)*

Table Description	Table Name	Default Retention	Notes
Array's Cache Usage Performance	sdk_esym_array_cache_usage	21 days	

Capacity: Default retention for EMC XtremIO

Table 24-3

Table Description	Table Name	Default Retention	Notes
XtremIO Data Reduction Rate	sdk_exio_datareductionrate	15	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Capacity: Default retention for Dell EMC Elastic Cloud Storage (ECS)

Table 24-4

Table Description	Table Name	Default Retention	Notes
ECS Storage Array Performance	sdk_decs_ecsstg_ary_stats	21	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
ECS Replication Group Performance	sdk_decs_ecsrep_grp_stats	21	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
ECS Node Performance	sdk_decs_ecsnode_stats	21	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Table 24-4 *(continued)*

Table Description	Table Name	Default Retention	Notes
ECS Node's Disk Performance	sdk_decs_ecsnodedisk_stats	21	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
ECS Storage Pools Performance	sdk_decs_ecsstg_pls_stats	21	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Capacity: Default retention for Windows file server

Table 24-5

Table Description	Table Name	Default Retention	Notes
History and Performance Table for Windows File Systems	sdk_msws_filesystem_perf	21 months	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
History and Performance Table for Windows Cifs Shares	sdk_msws_cifsshare_perf	21 months	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
History and Performance Table for Windows NFS Server	sdk_msws_nfs_statistics	21 months	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
History and Performance Table for Windows Event Log Details	sdk_msws_eventlog_details	30 months	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Capacity: Default retention for Pure Storage FlashArray

Table 24-6

Table Description	Table Name	Default Retention	Notes
Array Performance	SDK_PURE_STORAGEARRAY_PERF	21 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Cloud: Default retention for Amazon Web Services (AWS)

Table 24-7

Table Description	Table Name	Default Retention	Notes
Billing Record Tag	SDK_aws_billing_rec_tag	366 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
Mapping from any resource ID or name to one of many different entities (or none at all)	sdk_aws_resource_map	999999 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
S3 bucket usage	sdk_aws_s3_bucket_usage	999999 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Table 24-7 *(continued)*

Table Description	Table Name	Default Retention	Notes
Billing Record	sdk_aws_billing_record	367 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280. **Note:** The retention period for the billing record data must be one day more than the billing record tag data to ensure that the dependent tag data is removed before purging the billing record.

Cloud: Default retention for Microsoft Azure

Table 24-8

Table Description	Table Name	Default Retention	Notes
Resource Mapping from any resource ID or name to one of many different entities (Virtual Machines or Storage Accounts)	sdk_msaz_resourcemap	999999 days	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Cloud: Default retention for OpenStack Ceilometer

Table 24-9

Table Description	Table Name	Default Retention	Notes
Instance Metrics	sdk_oscm_instance_metrics	15	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Table 24-9 *(continued)*

Table Description	Table Name	Default Retention	Notes
Network Metrics	sdk_oscm_network_metrics	15	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.
Disk Metrics	sdk_oscm_disk_metrics	15	To modify the retention period, See "Data retention periods for SDK database objects" on page 280.

Configure multi-tenancy data purging retention periods

In multi-tenancy environments, where IT Analytics domains partition client data, data retention periods can be configured for each domain. The data retention period for a domain applies to that domain and all of its sub-domains. For systems collected by traditional Data Collectors, use the procedure described in the following section:

See "System configuration: functions" on page 160.

To enable purging for a user-defined object for a specific domain, take the following steps.

1. Log in to the Portal Server as user **aptare**.
2. Type the following command:

```
sqlplus portal/<portal_password>@//localhost:1521/scdb
```

3. At the command line, execute the following SQL statement, substituting relevant values for the variables shown < > in the syntax.

```
INSERT INTO apt_purge_rules (table_name, product_type,
table_description, date_column_name, retention_days,
default_retention_days, domain_id, creation_date, last_updated)
SELECT table_name, product_type, table_description,
date_column_name, <NewRetentionDays>, default_retention_days,
<domain_id_value>, SYSDATE, SYSDATE
  FROM apt_purge_rules
 WHERE domain_id IS NULL
```

```
   AND table_name = <tableName>;
Commit;
```

See “Find the domain ID and database table names” on page 285.

See the section called “Retention Period Update for Multi-Tenancy Environments Example” on page 293.

Retention Period Update for Multi-Tenancy Environments Example

The following example illustrates an update for a Pure Storage performance database table.

```
INSERT INTO apt_purge_rules (table_name, product_type,
table_description, date_column_name, retention_days,
default_retention_days, domain_id, creation_date, last_updated)
SELECT table_name, product_type, table_description, date_column_name,
 50, default_retention_days, 100007, SYSDATE, SYSDATE
  FROM apt_purge_rules
 WHERE domain_id IS NULL
   AND table_name = 'SDK_PURE_STORAGEARRAY_PERF';
Commit;
```

Chapter 25

Troubleshooting

This chapter includes the following topics:

Troubleshooting user login problems

To troubleshoot LDAP/login issues, use the **findUser** tool.

1. Launch the tool's usage instructions.

   Linux:

   ```
   # cd /opt/aptare/utils/
   # ./findUser.sh
   ```

   Windows:

   ```
   C:\opt\aptare\utils\finduser.bat
   ```

2. Use any of the following commands:

Add User — Linux: `./addUser.sh <userId> <lastName> <password> <restoreWizPassword>`

Windows: `adduser.bat <userId> <lastName> <password> <restoreWizPassword>`

Delete User Linux: `./delUser.sh <userId>`

Windows: `deluser.bat <userId>`

Find User Linux: `./findUser.sh <userId>`

Windows: `finduser.bat <userId>`

Modify User Linux: `./updateUser.sh <currentUserId> <modUserId> <modLastName> <modPassword> <modRestoreWizPassword>`

Windows: `updateuser.bat <currentUserId> <modUserId> <modLastName> <modPassword> <modRestoreWizPassword>`

See "Forgotten password procedure" on page 295.

Forgotten password procedure

If a user forgets a user ID or password, they can be reset using the following command, run with root/administrator privileges:

Linux:

```
cd /opt/aptare/utils
./updateUser.sh <currentUserId> <modLastName> <modPassword>
<modRestoreWizPassword>
```

Windows:

```
C:\opt\aptare\utils
updateuser.bat <currentUserId> <modLastName> <modPassword>
<modRestoreWizPassword>
```

For example:

```
./updateUser.sh admin@domain.com adminguy@domain.com Administrator
newpwd newrestpwd
```

Login issues

The following sections highlight common login issues and their possible solutions.

Portal Login Errors

- File system is out of disk space.
- Fully qualified URL incorrectly set up

- URL not in the local hosts file or in DNS.
- Domain incorrectly specified
- Values in Tomcat Application Server and Apache Web Server do not match what's in DNS.

Cannot Log On To Portal

When a user can't log in to the Portal, the reasons are usually one of the following:

- User forgot password. Change the user's password.
- LDAP service is not running.
- There is a port conflict. Another program is listening on port 80, resulting in a port conflict. Apache Web Server needs port 80. Use the **netstat** command to determine the other application that's listening on port 80, then assign that application a different port.

Single Sign On (SSO) Configuration Reset Utility (resetSSOConfig)

If Single Sign On (SSO) is not properly set up in the **Admin>Advanced>System Configuration**, after restarting you may not be able to log into the Portal. This utility resets the Single Sign On (SSO) parameters to provide Portal access. Run the following scripts from the command prompt:

Linux

```
cd /opt/aptare/utils
./resetSSOConfig.sh
```

Windows

```
cd C:\opt\aptare\utils
resetSSOConfig.bat
```

Restart the Portal services after running the OS specific script.

Connectivity issues

No Connectivity

A number of conditions can disrupt connectivity, including:

- Firewall issues can prevent connectivity.
- A network change occurred. Typically a DNS, system domain, or hostname change is the culprit.

- Oracle service is not running.
- Network timeouts

Action Recommendations

The following list suggests actions you can take to determine what is causing the issue.

- Ping **itanalyticsagent.mydomain.com** from the primary server.
- Check if you can connect to **itanalyticsagent.mydomain.com** telnet port 80.

Connection failure between portal server and database server

Portal server fails to connect to database server and the following error is logged in `portal.log`.

```
Error code: APT-UI-1556759233
The Network Adapter could not establish the connection
```

Possible root causes:

- Oracle service is not running.
- Firewall is enabled and database port is not configured properly.
- SELinux is set to enforce.
- Required memory is not available to run Oracle.

Resolution:

1. Check if Oracle service is running:

   ```
   # /opt/aptare/bin/oracle status
   ```

2. Check for errors in the respective log files:

   Linux: `/opt/aptare/oracle/rdbms/log/startup.log`

   Windows: `C:\opt\oracle\rdbms\log\startup.log`

3. Verify whether Firewall is disabled.
4. If Firewall is enabled, ensure the database port (typically 1521) is allowed for inbound connections.

5 Ensure SELinux status is either disabled or permissive.

```
# getenforce
```

6 Connect to database using sqlplus on portal server and run basic `sql` commands like `select * from ptl_users`.

Data Collector and database issues

Data Collector: Changing the Name and Passcode

The Data Collector uses a name and a passcode to identify itself with the Data Receiver. If this name or passcode is changed on the Portal, it needs to be changed on the Data Collector side.

To find and update name and passcode information on the Data Collector server, follow these steps:

Changing the Name and Passcode on a Windows Data Collector Server

1 Edit the following file:

```
$APTARE_HOME/mbs/conf/wrapper.conf
```

by modifying the following entries accordingly:

```
wrapper.app.parameter.2="$COLLECTOR_NAME$"
wrapper.app.parameter.3="$COLLECTOR_PASSWORD$"
```

2 Edit the following file:

```
$APTARE_HOME/mbs/bin/updateconfig.bat
```

to modify the name and passcode. They are the two parameters immediately following "com.storage.mbs.watchdog.ConfigFileMonitorThread".

Changing the Name and Passcode on a Linux Data Collector Server

- Edit the following files:

```
$APTARE_HOME/mbs/bin/updateconfig.sh
$APTARE_HOME/mbs/bin/startup.sh
```

  - The name and the passcode will be passed as program arguments to the Java program in the above two scripts.
  - In **updateconfig.sh**, the name and passcode are the two parameters immediately following "com.storage.mbs.watchdog.ConfigFileMonitorThread".

- In **startup.sh**, the name and passcode follow "com.storage.mbs.watchdog.WatchDog".

Insufficient Privileges

When creating the database, you may get an insufficient privileges message when the Windows user is not local, or is not a member of the **ORA_DBA** group. In this case, you receive an **insufficient privileges** error when you create the database.

Remove an Inactive Hitachi Array from the Database

In Capacity Manager, if an array has been removed from Hitachi Device Manager (HDvM), it remains in the IT Analytics database as an inactive array and it will continue to be included in reports. Take the following command-line steps to delete the array.

1. Log in to SQLPlus and set the user to aptare.
2. Execute the query:

```
select storage_array_id from aps_storage_array where
array_name='<ARRAY NAME>';
```

3. Using the storage_array_id from the above query to execute this code:

```
Begin
  srm_common_pkg.deleteStorageArray(<STORAGEARRAYID>);
End;
/
```

4. Verify that the array was deleted successfully using this query (should return 0).

```
select count(*)
FROM aps_storage_array
WHERE storage_array_id = <STORAGEARRAYID>;
```

Report Emails are not Being Sent

If Report Emails are not being sent:

The SMTP_HOST value does not have a running mailer.

You can also set the authentication mode under **Admin** tab > **Advanced** > **Email Configuration**

As shown in the following table, several properties can be customized.

Table 25-1 Email configuration for basic SMTP authentication type

Field names	Description
Email SMTP authentication	Set to true to enable email authentication.
Email SMTP host IP address	Make sure this references a running mail server.
Email debug mode	Used for product developers to debug mail transmission issues.
Email enable Transport Layer Security (TLS)	The SMTP Transport Layer Security extension--the encryption and authentication protocol. When a client and server that support TLS talk to each other, they can encrypt the data channel to guard against eavesdroppers.
Email from address	The email address for a reply back.
Email from name	The name associated with the reply back email address.
SMTP User	User name used for authentication on the email server.
SMTP password	Password required by email server.

Table 25-2 Email configuration for OAuth SMTP authentication type

Field names	Description
Email SMTP host IP address	IP address of a functional authorization server that will be used for authentication.
Email enable Transport Layer Security (TLS)	Enables the SMTP transport layer security (TLS) encryption and authentication protocol. The data channel used to communicate between the client and server will be encrypted.
SMTP Port	SMTP communication port number.
SMTP authentication type	Authentication method used by the client to prove its identity. **OAUTH** in this case.

Table 25-2 Email configuration for OAuth SMTP authentication type *(continued)*

Field names	Description
Client Id	The ID that IT Analytics will use to request the access token from the authorization server.
Client Secret	The confidential string between IT Analytics and the authorization server.
Auth url	URL of the authorization endpoint on the authorization server.
Auth Token url	The token used by IT Analytics to access the SMTP resource server.
Token Info url	The URL to the token server to be used to regenerate tokens.

Email SMTP authentication	Set to true to enable email authentication.
Email SMTP host IP address	Make sure this references a running mail server.
Email debug mode	Used for product developers to debug mail transmission issues.
Email enable Transport Layer Security (TLS)	The SMTP Transport Layer Security extension--the encryption and authentication protocol. When a client and server that support TLS talk to each other, they can encrypt the data channel to guard against eavesdroppers.
Email from address	The email address for a reply back.
Email from name	The name associated with the reply back email address.
SMTP User	User name used for authentication on the email server.
SMTP password	Password required by email server.

Additional Email Troubleshooting Recommendations

To determine the cause of the issue:

- Telnet to port 25 to test connectivity to the email server.

```
Typed text is in bold.
$ telnet localhost 25
```

```
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
220 www.aptare.com ESMTP Sendmail 8.11.6/8.11.6; Wed, 26 Jan 2008
 10:16:10 -0800
hello aptare
250 www.aptare.com Hello localhost.localdomain [127.0.0.1], pleased
 to meet you
mail from: gthom@aptare.com
250 2.1.0 gthom@aptare.com... Sender ok
rcpt to: gthom@aptare.com
250 2.1.5 gthom@aptare.com... Recipient ok
data
354 Enter mail, end with "." on a line by itself
this is the email body
.
250 2.0.0 j0QIMUQ10566 Message accepted for delivery
$
```

General Reporting Issues

Reports with Graphs Fail to Load

Certain reports that include graphs may result in a report failing to load due to an “out-of-memory: Java heap space” error. This may occur when generating a report for a large set of data or exporting it to a PDF. For example, the following reports are known to occasionally encounter this condition: Tape Library and Drive Utilization for one year grouped by hours, Job Volume Summary for three years of daily data, and Data Domain Filesystem Capacity & Forecast for two months grouped by hours.

To work around this condition, try these changes in the Report Designer:

- reduce the number of selected objects (for example, servers or arrays)
- limit the time period to a shorter timeframe
- select a less granular “group by” option (for example, by days instead of by hours)

Missing backups from clients

Typically, this is a Data Collector issue.

Charts Do Not Display

File system is out of disk space.

Performance Issues

Performance can be impacted by a number of issues. Use the following checklist to help you isolate problems.

- Check the number of backup jobs that have been processed in the last 24-hour period.
- Determine the level of database logging that has been enabled in **scon.log**. If DBG level messages have been enabled, this can negatively impact performance. INFO and WARNING messages have negligible impact.
- Check if anything else is running on the server.
- Note if performance suffers at specific times of the day and determine which processes are running during those times.
- Verify the server's configuration: memory and CPU size
- Check if Oracle has been tuned or if the default file (**initscdb.ora**) is in use.
  - On Windows, by default, Oracle can use only 1 GB of memory. If the system has more memory available, use the /3GB option to tell windows to allocate 2 GB to the Oracle process.
- Determine the top running processes--run **top**.
  - What are the top running processes and what is the average CPU Utilization?
  - Is a lot of Virtual Memory being consumed?
  - Is there a high I/O Wait? (This implies a disk bottleneck.)
  - If oraclescdb is the top running process, using a lot of CPU resources:
    Run: **sqlplus portal/portal@//localhost:1521/scdb @/opt/aptare/database/tools/list_quiery_by_pid**
    This task will prompt for the process id (pid) and show the query that is running. If you run this several times over the course of 5-10 minutes and it returns the same query, you have a long-running query. Note that this will only work if you use process IDs associated with oraclescb processes.
- Query the report database for long-running jobs.
  - Select from **apt_query_execution** table to see what reports have been running and for how long. This will help identify a report that someone may have inadvertently scheduled to run every five minutes.
  - In **/opt/aptare/database/tools**, use other sql queries:
    list_long_running_queries.sql
    list_running.sql
    list_running_queries.sql

long_running.sql
Run the above queries and capture the output to send to the Cohesity Support.

Portal upgrade performance issues

When the entropy of the system is very low, cryptographic functions takes considerable amount of time.

Following are the examples of low entropy:

- while adding new data collector on Portal, takes longer time to generate the key file.
- Upgrade of Portal hangs when upgrading internal objects.
- Aptare agent service takes longer time when started to get the `collectorconfig.xml` from data receiver side.
- `checkinstall.sh` file execution takes longer time than expecting.

These issues are observed on **Linux Platform**.

The following solution is recommended:

Note: Download and install **rng-tools rpm** on the Portal.

For **RHEL/OEL**, execute the following steps to install the **rng-tools** and start the services:

1. Access command prompt.
2. Type `yum install rng-tools` to install the rng-tools.
3. Type `systemctl start rngd` to start the services.
4. Type `systemctl enable rngd` to enable the services.

For **Suse**, execute the following steps to install the **rng-tools** and start the services:

1. Access command prompt.
2. Type `zypper install rng-tools` to install the rng-tools.
3. Type `systemctl start rng-tools` to start the services.
4. Type `systemctl enable rng-tools` to enable the services.

Appendix A

Kerberos based proxy user's authentication in Oracle

This appendix includes the following topics:

Overview

Kerberos is a computer network security protocol that authenticates the communication of nodes over a non-secure network to prove their identity to one another in a secure manner. It uses secret-key cryptography and a trusted third party for authenticating client-server applications and verifying users' identities.

This section helps you to configure Kerberos-based authentication for proxy user/s in Oracle. Kerberos authentication allows to connect to Oracle without specifying the username / password credentials. The authentication is done externally. Proxy authentication in Oracle, allows connection to a target database username via another database user (the proxy user).

For example you can authorize a user with a development account to connect to the application owner account using his/her credentials (thus no need to expose the application user's password). This section helps to configure for Kerberos and proxy authentication: you can provide a mean to connect to any given DB user via

a Kerberos-authenticated user, i.e. without specifying the password at connection time.

The section is categorized into Oracle related modifications and Portal related modifications.

Note: k1portal is the Kerberos username is referred. It can vary from environment to environment.

Note: scdb referred to as ORACLE_SID in this documentation

Pre-requisite

The following packages must be installed on the IT Analytics portal and pre-stashed Kerberos ticket:

- krb5-libs
- krb5-workstation

Exporting service and user principal's to keytab file on KDC

The Authentication Server needs a keytab file and validate a user when refreshing session credentials.

This section describes how to export and copy a keytab file.

1 Create oracle service principal.

```
# kadmin.local
  addprinc -randkey <oracle SID>/<oracle server host name>@<domain
 realm name>
```

For example: `addprinc -randkey scdb/abc.example.com@EXAMPLE.COM`

2 Extract oracle service principal to keytab file

```
# kadmin.local
ktadd -k <keytab file path> <oracle SID>/<oracle server host
name>@<domain realm name>
```

For example: `ktadd -k /tmp/keytab scdb/abc.example.com@EXAMPLE.COM`

3 Create Kerberos user principal

```
#  kadmin.local
addprinc <kerberos user name>
```

For example: `addprinc k1portal`

4 Extract Kerberos user principal to keytab file.

```
# kadmin.local
ktadd -k <keytab file path> <kerberos user name>
```

For example:`ktadd -k /tmp/keytab k1portal`

Modifications for Oracle

Following are the steps to perform the Oracle related modifications:

Oracle modifications

1 Copy `/etc/krb5.conf` from KDC to Oracle server `/etc/krb5.conf` path.

2 Copy the `keytab` file from KDC to Oracle as `/etc/v5srvtab`

Note: The exported keytab file can be removed from KDC once it has been copied to oracle server.

Note: For more information, see See “Exporting service and user principal’s to keytab file on KDC” on page 306.

---

3 Grant appropriate permissions to files and directories so that Oracle service can read those files/directories mentioned in `sqlnet.ora`

```
# chown <oracle user>:<oracle group> /etc/krb5.conf /etc/v5srvtab
# chmod 444 /etc/krb5.conf /etc/v5srvtab
```

For example: `# chown aptare:dba /etc/krb5.conf /etc/v5srvtab`

4 Add the following entries to /opt/aptare/oracle/network/admin/sqlnet.ora file.

- SQLNET.AUTHENTICATION_SERVICES=(BEQ,KERBEROS5)
- SQLNET.AUTHENTICATION_KERBEROS5_SERVICE=scdb
- SQLNET.KERBEROS5_CONF=/etc/krb5.conf
- SQLNET.KERBEROS5_CONF_MIT=TRUE

- SQLNET.KERBEROS5_REALMS=/etc/krb5.conf
- SQLNET.KERBEROS5_KEYTAB=/etc/v5srvtab
- SQLNET.FALLBACK_AUTHENTICATION=TRUE
- SQLNET.KERBEROS5_CC_NAME=/tmp/kcache
- SQLNET.KERBEROS5_CLOCKSKEW=300

Note: The Oracle server and KDC should have the same time and Timezone settings. If there is slight time mismatch, add the below entry to sqlnet.ora to cover the time mismatch. For example within 20mins. The default value is 300. `SQLNET.KERBEROS5_CLOCKSKEW=1200`. Cohesity recommends to configure both the servers to sync time from time servers.

---

5 Create a user in Oracle to be used with Kerberos. We created *k1portal* user principal in KDC.

```
# su - aptare (login as Oracle service user)
# source <INSTALL_PATH>/aptare/bin/aptare_env.sh
# sqlplus / as sysdba
SQL> alter session set container=scdb;
CREATE USER <kerberos user name> PROFILE DEFAULT IDENTIFIED
EXTERNALLY AS
'<Kerberos user name>@<domain realm name>' DEFAULT TABLESPACE
APTARE_TBS_DATA_1M TEMPORARY TABLESPACE APTARE_TEMP_TBS;
```

For example:

```
SQL> CREATE USER k1portal PROFILE DEFAULT IDENTIFIED EXTERNALLY AS
'k1portal@EXAMPLE.COM' DEFAULT TABLESPACE
APTARE_TBS_DATA_1M TEMPORARY TABLESPACE APTARE_TEMP_TBS;
SQL> EXIT;
```

6 Verify the parameter `OS_AUTHENT_PREFIX=""` is added in the parameter file ("`pfile`" or "`spfile`") of the database.

If not added, the parameter needs to be added to the parameter file.

For example: In the `/opt/aptare/oracle/dbs/initscdb.ora` file (initialization file for the service) add the following entry `OS_AUTHENT_PREFIX=""`

7 In the /opt/aptare/oracle/dbs/initscdb.ora file (initialization file for the service) add the following entry

```
OS_AUTHENT_PREFIX=""
```

8 Restart Oracle service.

9 **Note:** This step is optional if using pre-stashed Kerberos ticket.

Execute the following commands to verify and to fetch initial TGT for k1portal user; login as Oracle service user

```
# su - aptare (login as Oracle service user)
# source <INSTALL_PATH>/aptare/bin/aptare_env.sh
# oklist -k -t /etc/v5srvtab
# okinit -k -t /etc/v5srvtab k1portal
```

10 Connect to Oracle DB using the following command.

```
# sqlplus /@scdb
```

11 Create a trigger for the Kerberos users corresponding to portal to alter the session which will set current schema as PORTAL

```
# su - aptare ( login as Oracle service user)
# source <INSTALL_PATH>/aptare/bin/aptare_env.sh
# okinit -k -t /etc/v5srvtab k1portal;
```

The above command is optional for pre-stashed Kerberos ticket.

```
# sqlplus /@scdb
SQL> show user; # This must return k1portal
SQL> CREATE OR REPLACE TRIGGER set_portal_schema
after logon on k1portal.schema begin execute immediate
'alter session set current_schema=PORTAL';
end;
/
```

12 Grant RO user to connect through PORTAL user

```
sqlplus / as sysdba
SQL> alter session set container = scdb;
SQL> alter user aptare_ro grant connect through k1portal;
SQL> GRANT CREATE SESSION to aptare_ro;
Verify with RO user login
SQL> sqlplus [aptare_ro]/@scdb;
SQL> show user;
This must return aptare_ro
SQL> exit;
```

Modifications for Portal

Following are the steps to perform the portal related modifications:

Portal Modifications

1. Create a copy of `/etc/krb5.conf` from KDC to Portal server `/etc/krb5.conf` path.
2. Copy the `keytab` file from KDC to Portal at `/etc/v5srvtab`.

   Note: The exported keytab file can be removed from KDC once it has been copied to portal server.

   For more information, see See "Exporting service and user principal's to keytab file on KDC" on page 306.
3. Modify the owner and permission of above copied two files using the following commands:

   ```
   chown <oracle user>:<oracle group> /etc/krb5.conf /etc/v5srvtab
   # chmod 444 /etc/krb5.conf /etc/v5srvtab
   ```

   For example: `# chown aptare:dba /etc/krb5.conf /etc/v5srvtab`
4. Add the following entries to `/opt/aptare/oracle/network/admin/sqlnet.ora` file
   - SQLNET.AUTHENTICATION_SERVICES=(BEQ,KERBEROS5)
   - SQLNET.AUTHENTICATION_KERBEROS5_SERVICE=scdb
   - SQLNET.KERBEROS5_CONF=/etc/krb5.conf
   - SQLNET.KERBEROS5_CONF_MIT=TRUE
   - SQLNET.KERBEROS5_REALMS=/etc/krb5.conf
   - SQLNET.KERBEROS5_KEYTAB=/etc/v5srvtab
   - SQLNET.FALLBACK_AUTHENTICATION=TRUE
   - SQLNET.KERBEROS5_CC_NAME=/tmp/kcache
   - SQLNET.KERBEROS5_CLOCKSKEW=300

5 Modifications in the property file is required because when JDBC try to make multiple connections to Oracle DB, Oracle application treats this as a replay attack and errors out.

To avoid the error, ensure that the *[libdefaults]* section in the Kerberos configuration file */etc/krb5.conf* on KDC and client machine is configured `forwardable = false`.

To update, restart kdc and admin service on KDC server using the following commands:

```
systemctl restart krb5kdc.service
```

```
systemctl restart kadmin.service
```

6 Create cache file for portal user.

For example, the command to generate cache file: `kinit -k -t <Key Tab File> <kerberos user@domain realm name> -c <cache file name>`

```
kinit -k -t <Key Tab File> <kerberos user name>@<domain realm name>
 -c <cache file name>
# su - aptare (login as oracle user)
#source <INSTALL_PATH>/aptare/bin/aptare_env.sh
# kinit -k -t /etc/v5srvtab k1portal@EXAMPLE.COM
-c /tmp/portal_kcache;
```

7 Tomcat user must have read privileges to the cache file. To ensure that the Tomcat OS user is able to make a JDBC connection to Oracle DB, use the following commands:

```
.
# chown <portal user>:<portal group> /tmp/portal_kcache;
# chmod 444 /tmp/portal_kcache;
```

For example: `chown aptare:aptare /tmp/portal_kcache;`

8 The following properties must be added or updated in `/opt/aptare/portalconf/portal.properties`

- **db.url**=jdbc:oracle:thin:@(DESCRIPTION= (ADDRESS=(PROTOCOL=tcp)(HOST=localhost) (PORT=1521))(CONNECT_DATA=(SERVICE_NAME=scdb))) Host and Service name could be different here.
- **db.user**=<kerberos user name>@<domain realm name>

For example: db.user=k1portal@EXAMPLE.COM Combination of kerberos portal user name and domain realm name

- **db.auth.scheme**=kerberos
  This property must be defined to enable kerberos authentication and is case-insensitive
- **db.kerberos.keytab.path**=/etc/v5srvtab
  This is absolute path of keytab file
- **db.driver**=oracle.jdbc.OracleDriver
- **db.kerberos.portal_kcache.path**=/tmp/portal_kcache
  This is absolute path of portal user cache file
- db.connection.max=25
- db.connection.min=25

- db.connection.expiration=5

9 Similar changes are required in the Data-receiver property file `/opt/aptare/datarcvrconf/datrarcvrproperties.xml`.

Add or updated the bold perperties.

```
<dataSource>
<Driver> oracle.jdbc.driver.Oracle</Driver>
<URL>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)
(HOST=localhost)(PORT=1521))(CONNECT_DATA=(SERVICE_NAME=scdb)))</URL>
<UserId><kerberos user name>@<>
<domain realm name></UserId>

For example: <UserId>k1portal@EXAMPLE.COM</UserId>

<Password>Z0Q5W+lQD2jreQaLBoYsviYO21WGOq5iTEo0Ad2uUj/e0GtqPkOtXFblKxCse
KXO4VhpIQwwfrSfe59nGy156DV8lYoa7HWmL0hF+kAZXOoXfIN5YRAGfqDbCwrKQdtPY7pQh
uTkZMPLl0d9Kzy6sLGMb/33L4hKuEl0ZZN2FG5US26JZ/uSOBF7T69ppqxGqXMleZ19QBcv
UElLwJTn52SurL+K3RjCY7Xi0VJb4wLkax07xCkpSK9dJ6NMFJS3ybWP4jNs3rC3roudZak8
wGqLNhAacyXgW4pMpgigVjGwNr0N8rJIgcGmXgAxSNs0qmQItuXPIyqGf+nWWEfScQ==
</Password>

<oracle_service_name>scdb</oracle_service_name>

<ro_user>aptare_ro</ro_user>

<ro_password>U9a7a+af94q0CUaIfzaVmYl1P1DhdQW96CQiYWgxUGSV5sfVVsxoWF5Riy
V85MD8V0Ogy7UJo1sFmAL36KjDy8LA61pKeO4X39hRK/g8vvl/xNnG5bBYIF04/1LwD2FTz
0lJERWopKVZ6pd6TkT0mGeKrnu2oYi97GtlW4J73tPGTFRhHyVw7yZKMmaxbs/FBwrz5aIf
je3rT0w85m7Obtrjf2nJ2HjsaHnmToh0Ua96xlshjrE75UbaLMu0QEcF3PYF3qufYVIegn
4VGSHcpsU/AFzurKpr0JTsU/6VqvdE4veBLv4FH5D05bRetaOA0SGKCazWA50
xiirwocvgyw==
</ro_password>

<MaxConnections>125</MaxConnections>

<MinConnections>5</MinConnections>

<ConnExpirationTime>5</ConnExpirationTime>
<authScheme>kerberos</authScheme>
<portalKcacheFile>/tmp/portal_kcache</portalKcacheFile>
<kKeyTabFile>/etc/v5srvtab</kKeyTabFile>
</dataSource>
```

Before the upgrade

The following are the steps to be performed before the upgrade.

1. Provide dba grant to Kerberos portal user.

```
# su - aptare ( login as Oracle service user)
# source <INSTALL_PATH>/aptare/bin/aptare_env.sh


# sqlplus / as sysdba

SQL> alter session set container=scdb;
SQL> GRANT DBA TO k1portal;
SQL> GRANT EXECUTE ON SYS.DBMS_JOB TO k1portal;
SQL> GRANT EXECUTE ON SYS.DBMS_LOB TO k1portal;
SQL> GRANT EXECUTE ON SYS.DBMS_SCHEDULER TO k1portal;
SQL> GRANT SELECT ON DBA_OBJECTS TO k1portal;
```

2. Ensure portal cache file is valid and Tomcat user must have read permission.

```
# chmod 444 /tmp/portal_kcache;
chown <portal user>:<portal group> /tmp/portal_kcache
```

   For example:`# chown aptare:aptare /tmp/portal_kcache`

Post upgrade

The following are the steps to be performed after the upgrade.

1. Revoke DBA role and grant a specific list of privileges to Kerberos users after a successful upgrade. k1portal is the Kerberos username here. It can be varied from environment to environment.

   Under sys user performs below revoke tasks:

```
# su - aptare (login as oracle user)
# source <INSTALL_PATH>/aptare/bin/aptare_env.sh


# sqlplus "/ as sysdba"

SQL> alter session set container=scdb;
Session altered.
```

```
SQL> REVOKE DBA FROM k1portal;
Revoke succeeded.
```

2. Again under sys user runs individual PLSQL scripts to grant a list of required privileges to Kerberos-enabled users for the normal functioning of ITA application.
3. Ensure that the correct Kerberos username is given as arguments to the script.

```
# su - aptare
# source <INSTALL_PATH>/aptare/bin/aptare_env.sh


sqlplus "/ as sysdba"
SQL> alter session set container=scdb;
SQL> @/opt/aptare/database/ora_scripts/kerberos_grants_portal.plb;
Enter value for db_object_schema: portal
Enter value for kerberos_schema: k1portal
SQL> @/opt/aptare/database/ora_scripts/
metadata_grants_to_kerberos_user.plb
Enter value for kerberos_user_name: k1portal
SQL> exit;
```

4. Restart tomcat-portal and tomcat-agent and verify NBU ITA portal.

```
/opt/aptare/bin/tomcat-portal restart

/opt/aptare/bin/tomcat-agent restart
```

Note: Kerberos cache file should not be expired, Tomcat and Aptare users must have access to the cache file, for this add a script in crontab to re-generate cache file as below :

```
# cat krb_cache_refresh.sh
su - aptare (login as oracle user)
source <INSTALL_PATH>/aptare/bin/aptare_env.sh
okinit -k -t /etc/v5srvtab k1portal
kinit -k -t /etc/v5srvtab k1portal@EXAMPLE.COM
 -c /tmp/portal_kcache
chmod 444 /tmp/portal_kcache;
chown <portal user>:<portal group> /tmp/portal_kcache
```

For example: `chown aptare:aptare /tmp/portal_kcache`

Appendix B

Configure TLS-enabled Oracle database on IT Analytics Portal and data receiver

This appendix includes the following topics:

About Transport Layer Security (TLS)

Transport Layer Security (TLS) is a security protocol designed to enable privacy and data security over the internet. It provides privacy (confidentiality), integrity, and authenticity through cryptography, such as using certificates to authenticate the communication between two or more communicating computer applications. TLS runs in the presentation layer and is made up of two layers:

- TLS record

- TLS handshake protocols

TLS in Oracle environment

A network connection that is initiated over TLS requires a TLS handshake between the client and server before performing the authentication. The handshake process includes the following:

- The client and server choose the cipher suites and the encryption algorithms to use for data transfers.
- To verify the identity of the server, the server sends its certificate to the client and the client verifies whether it was signed by a trusted CA.
- If a client authentication is required, the client sends its certificate to the server, which the server verifies whether it was signed by a trusted CA.
- The client and server use public key cryptography to exchange key information and each generates a session key. Both the keys are generally used for data encryption during a single communication session. Session keys are typically used to encrypt network traffic. A client and a server can negotiate a session key at the beginning of a session and that key is used to encrypt all the network traffic between the parties of that session. A new session key is generated if the client and server communicate over a new session. The subsequent communications between the client and server are encrypted and decrypted with the use of the session key and the negotiated cipher suite.

Authentication process

Authentication process includes the following:

- The user initiates a Oracle Net connection from the client to the server using TLS.
- TLS performs the handshake between the client and the server.
- After a successful handshake, the server verifies whether the user has appropriate authorization to access the database.

Configure TLS in Oracle with IT Analytics on Linux in split architecture

In a split architecture, database and IT Analytics Portal are on different systems.

Step 1: Configure Oracle wallet on the server side.

1. Login as Oracle user.

```
su -aptare
```

2. Create a directory `server_wallet` on the server system to store the server wallet.

```
mkdir /opt/aptare/oracle/network/server_wallet
```

3. Create an empty wallet for the Oracle server with auto login enabled.

```
orapki wallet create -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password>
-auto_login
```

4. Add a self-signed certificate in the wallet. A new pair of private/public keys is created at this stage.

```
orapki wallet add -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password> -dn
"CN=<server_machine_name>" -keysize 2048 -self_signed -validity
<# of days>
```

5. Check the contents of the wallet. Verity whether the self-signed certificate is a trusted certificate.

```
orapki wallet display -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
```

6. Export the certificate so that it can be loaded into the client wallet later.

```
orapki wallet export -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password> -dn
"CN=<server_machine_name>" -cert
<SERVER_WALLET>\<server-certificate-name>.crt
```

7. Check whether the certificate is exported to the above directory.

Step 2: Configure Oracle wallet for client application

1 Login as Oracle user.

```
su - aptare
```

2 Create a directory `client_wallet` on the client system to store the client wallet.

```
mkdir /opt/aptare/oracle/network/client_wallet
```

3 Create a wallet for the Oracle client. Create an empty wallet with auto login enabled.

```
orapki wallet create -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
-auto_login
```

4 Add a self-signed certificate in the wallet. A new pair of private/public keys are created at this stage.

```
orapki wallet add -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password> -dn
"CN=<client_machine_name>" -keysize 2048 -self_signed -validity
<# of Days>
```

5 Check the contents of the wallet. Verify that the self-signed certificate is both a user and a trusted certificate.

```
orapki wallet display -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
```

6 Export the certificate so that it can be loaded into the server wallet later.

```
orapki wallet export -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password> -dn
"CN=<client_machine_name>" -cert
<CLIENT_WALLET>\<client-certificate-name>.crt
```

7 Check whether the certificate is exported to the above directory.

8 Make sure the Oracle service user can access the wallet file `cwallet.sso`.

Step 3: Perform client-server exchange certificate process.

1 Repeat these steps on each of the database client systems:

- Copy `<server-certificate-name>.crt` from the server system to the client system `/opt/aptare/oracle/network/client_wallet` folder.
- Copy `<client-certificate-name>.crt` from the client system to the server system `/opt/aptare/oracle/network/server_wallet` folder.
- After copying certificate file, change the owner to Oracle service user on the client system.
  On the client system:

  ```
  chown aptare:aptare <server-certificate-name>.crt
  ```

  On the server system:

  ```
  chown aptare:aptare <client-certificate-name>.crt
  ```

2 Load the server certificate into the client wallet.

```
orapki wallet add -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
-trusted_cert -cert
/opt/aptare/oracle/network/client_wallet/<server-certificate-name>.crt
```

3 Check the contents of the client wallet. Note that the server certificate is now included in the list of trusted certificates.

```
orapki wallet display -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
```

4 Load the client certificate into the server wallet.

```
orapki wallet add -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password>
-trusted_cert -cert
/opt/aptare/oracle/network/server_wallet/<client-certificate-name>.crt
```

5 Check the contents of the server wallet. Note that the client certificate is now included in the list of trusted certificates.

```
orapki wallet display -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password>
```

Step 4: Configure the Oracle database to listen for TCPS connection (Server/Oracle system). In the steps below, host is Oracle server IP address and `/opt/aptare/oracle/network/server_wallet` **is the server wallet location.**

1 Stop Oracle listener.

```
lsnrctl stop
```

2 Modify the `listener.ora` (`/opt/aptare/oracle/network/admin/listener.ora`).

```
LISTENER =
  (DESCRIPTION_LIST =
     (DESCRIPTION =
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
      (ADDRESS = (PROTOCOL = TCPS)(HOST = xx.xx.xx.xx)(PORT =
2484))
     )
  )
```

Append the below line at the end of the file.

```
SSL_CLIENT_AUTHENTICATION = FALSE
SECURE_PROTOCOL_LISTENER=(IPC)
WALLET_LOCATION =
    (SOURCE =
      (METHOD = FILE)
      (METHOD_DATA =
         (DIRECTORY = /opt/aptare/oracle/network/server_wallet)

      )
    )
/opt/aptare/oracle/network/server_wallet
```

3 Modify the `sqlnet.ora` file (`/opt/aptare/oracle/network/admin/sqlnet.ora`).

```
SSL_CLIENT_AUTHENTICATION = FALSE
  WALLET_LOCATION =
    (SOURCE =
       (METHOD = FILE)
       (METHOD_DATA =
          (DIRECTORY = /opt/aptare/oracle/network/server_wallet)

       )
    )

SSL_CIPHER_SUITES = (SSL_RSA_WITH_AES_256_CBC_SHA,
SSL_RSA_WITH_3DES_EDE_CBC_SHA) SQLNET.WALLET_OVERRIDE = TRUE
```

4 Modify the `tnsnames.ora` file (`/opt/aptare/oracle/network/admin/tnsnames.ora`)

```
SCDB =
     (DESCRIPTION =
        (ADDRESS=
            (PROTOCOL=TCPS)
            (HOST=xx.xx.xx.xx)
            (PORT=2484)
        )
        (CONNECT_DATA=(SERVICE_NAME=scdb)(SID=SCDB))
    )
```

5 Start the Oracle service.

```
lsnrctl start
```

6 Check the listener status.

```
lsnrctl status
```

Step 5: Configure the Oracle database to listen for TCPS connection on the client system. Configure the `listener.ora` and `sqlnet.ora` files on the database server using the following steps. In the procedure below, host is

Oracle server IP address and `/opt/aptare/oracle/network/server_wallet` **is the server wallet location.**

1 Modify the `listener.ora` (`/opt/aptare/oracle/network/admin/listener.ora`) and add the below contents.

```
LISTENER =
   (DESCRIPTION_LIST =
      (DESCRIPTION =
         (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
         (ADDRESS = (PROTOCOL = TCPS)(HOST = xx.xx.xx.xx)(PORT =
 2484))
       )
   )
```

Add below line at the end of file:

```
 SSL_CLIENT_AUTHENTICATION = FALSE
    SECURE_PROTOCOL_LISTENER=(IPC)
    WALLET_LOCATION =
     (SOURCE =
       (METHOD = FILE)
       (METHOD_DATA =
         (DIRECTORY = /opt/aptare/oracle/network/client_wallet)

     )
  )
/opt/aptare/oracle/network/client_wallet
```

2 Modify the `sqlnet.ora` file (`/opt/aptare/oracle/network/admin/sqlnet.ora`).

```
SSL_CLIENT_AUTHENTICATION = FALSE
    WALLET_LOCATION =
       (SOURCE =
         (METHOD = FILE)
         (METHOD_DATA =
          (DIRECTORY = /opt/aptare/oracle/network/client_wallet)

         )
       )

SSL_CIPHER_SUITES = (SSL_RSA_WITH_AES_256_CBC_SHA,
SSL_RSA_WITH_3DES_EDE_CBC_SHA) SQLNET.WALLET_OVERRIDE = TRUE
```

3 Modify the `tnsnames.ora` (`/opt/aptare/oracle/network/admin/tnsnames.ora`)

```
SCDB =
   (DESCRIPTION =
       (ADDRESS=
           (PROTOCOL=TCPS)
           (HOST=xx.xx.xx.xx)
           (PORT=2484)
       )
       (CONNECT_DATA=(SERVICE_NAME=scdb)(SID=SCDB))
    )
```

4 Test Oracle connection using sqlplus.

```
sqlplus username/password@dbService
```

Step 6: Load Oracle server wallet certificate to the portal and upgrader Java KeyStore.

1 Login as a root user.

2 Add server certificate in portal Java.

```
cd /usr/java/bin
keytool -import -trustcacerts -alias ora_server_cert -file
/opt/aptare/oracle/network/client_wallet/server-cert-db.crt
-keystore /usr/java/lib/security/cacerts
password: changeit
```

3 Add server certificate in upgrader Java.

```
cd /opt/aptare/upgrade/jre/bin
keytool -import -trustcacerts -alias ora_server_cert -file
/opt/aptare/oracle/network/client_wallet/server-cert-db.crt
-keystore /opt/aptare/upgrade/jre/lib/security/cacerts
password: changeit
```

Step 7: Modify connection URL in the portal and receiver property file.

1 Stop portal and agent services.

```
/opt/aptare/bin/tomcat-portal stop
/opt/aptare/bin/tomcat-agent stop
```

2 Modify database URL in `/opt/aptare/portalconf/portal.properties`.

```
db.url=jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))
```

3 Modify database URL in `/opt/aptare/datarcvrconf/datrarcvrproperties.xml`.

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))
```

4 Start portal and agent services.

```
/opt/aptare/bin/tomcat-portal start
/opt/aptare/bin/tomcat-agent start
```

Configure TLS in Oracle with IT Analytics on Linux in non-split architecture

In a non-split architecture, Database and portal are on the same system.

Step 1: Configure Oracle wallet on the server side.

1 Login as Oracle user.

```
su - aptare
```

2 Create a directory on the server machine to store the server wallet as `server_wallet`. Run the make directory command below at `/opt/aptare/oracle/network` folder.

```
mkdir /opt/aptare/oracle/network/server_wallet
```

3 Create an empty wallet for the Oracle server with auto login enabled.

```
orapki wallet create -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password>
-auto_login
```

4 Add a self-signed certificate in the wallet (a new pair of private/public keys is created).

```
orapki wallet add -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password> -dn
"CN=<server_machine_name>" -keysize 2048 -self_signed -validity
<# of Days>
```

5 Check the contents of the wallet. Notice the self-signed certificate is both a user and trusted certificate.

```
orapki wallet display -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password>
```

6 Export the certificate so it can be loaded into the client wallet later.

```
orapki wallet export -wallet
"/opt/aptare/oracle/network/server_wallet" -pwd <password> -dn
"CN=<server_machine_name>" -cert
/opt/aptare/oracle/network/server_wallet/<server-certificate-name>.crt
```

7 Check whether the certificate has been exported to the above directory.

Step 2: Configure Oracle wallet for client application.

1 Login as oracle user.

```
su - aptare
```

2 Create a directory on the client system to store the client wallet. Call it `client_wallet`. Create it under the `/opt/aptare/oracle/network` folder.

```
mkdir /opt/aptare/oracle/network/client_wallet
```

3 Create a wallet of the oracle client. Create an empty wallet with auto login enabled.

```
orapki wallet create -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
-auto_login
```

Step 3: Load server certificate into the client wallet.

1 Load the server certificate into the client wallet

```
orapki wallet add -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
-trusted_cert -cert
/opt/aptare/oracle/network/client_wallet/<server-certificate-name>.crt
```

2 Check the contents of the client wallet. Note that the server certificate is now included in the list of trusted certificates.

```
orapki wallet display -wallet
"/opt/aptare/oracle/network/client_wallet" -pwd <password>
```

Step 4: Configure the Oracle database to listen for TCPS connection: Configure the `listener.ora`**,** `tnsnames.ora`**, and** `sqlnet.ora` **files on the database server**

using the following steps. In these steps, host is the oracle server IP address and the server wallet location is `/opt/aptare/oracle/network/server_wallet`**.**

1 Stop the Oracle listener before updating the files.

```
lsnrctl stop
```

2 Modify the `listener.ora` (`/opt/aptare/oracle/network/admin/listener.ora`)

```
LISTENER =
    (DESCRIPTION_LIST =
       (DESCRIPTION =
         (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
         (ADDRESS = (PROTOCOL = TCPS)(HOST = xx.xx.xx.xx)(PORT =
 2484))
       )
     )
```

Add below line at the end of file:

```
SSL_CLIENT_AUTHENTICATION = FALSE
  SECURE_PROTOCOL_LISTENER=(IPC)
WALLET_LOCATION =
(SOURCE =
(METHOD = FILE)
(METHOD_DATA =
(DIRECTORY = /opt/aptare/oracle/network/server_wallet)
)
)

/opt/aptare/oracle/network/server_wallet
```

3 Modify the `tnsnames.ora` (/opt/aptare/oracle/network/admin/tnsnames.ora).

```
SCDB =
   (DESCRIPTION =
      (ADDRESS=
        (PROTOCOL=TCPS)
        (HOST=xx.xx.xx.xx)
        (PORT=2484)
      )
      (CONNECT_DATA=(SERVICE_NAME=scdb)(SID=SCDB))
   )
```

4 Start the Oracle listener.

```
lsnrctl start
```

5 Check the listener status

```
lsnrctl status
```

6 Test the Oracle connection using sqlplus.

```
sqlplus username/password@dbService
```

Step 5: Load oracle server wallet certificate to the portal and upgrader Java KeyStore.

1 Add server certificate in portal Java.

```
cd /usr/java/bin
keytool -import -trustcacerts -alias ora_server_cert -file
/opt/aptare/oracle/network/server_wallet/server-cert-db.crt
-keystore /usr/java/lib/security/cacerts
password: changeit
```

2 Add server certificate in upgrader Java

```
cd /opt/aptare/upgrade/jre/bin
keytool -import -trustcacerts -alias ora_server_cert -file
/opt/aptare/oracle/network/server_wallet/server-cert-db.crt
-keystore /opt/aptare/upgrade/jre/lib/security/cacerts
password: changeit
```

Step 6: Modify connection URL in the portal and receiver property file.

1 Login as a root user.

2 Stop portal and agent services.

```
/opt/aptare/bin/tomcat-portal stop
/opt/aptare/bin/tomcat-agent stop
```

3 Modify database URL in `/opt/aptare/portalconf/portal.properties`.

```
db.url=jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))
```

4 Modify database URL in
`/opt/aptare/datarcvrconf/datrarcvrproperties.xml`.

```
<URL>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))</URL>
```

5 Start portal and agent services.

```
/opt/aptare/bin/tomcat-portal start
/opt/aptare/bin/tomcat-agent start
```

Configure TLS in Oracle with IT Analytics on Windows in split architecture

In a split architecture, database and portal are on different systems.

Step 1: Configure Oracle wallet for server side.

1 Create a directory on the server machine to store the server wallet in the `C:\opt\oracle\network\` folder.

```
mkdir C:\opt\oracle\network\server_wallet
```

2 Create an empty wallet for the Oracle server with auto login enabled.

```
orapki wallet create -wallet "C:\opt\oracle\network\server_wallet"
 -pwd <password> -auto_login
```

3 Add a self-signed certificate in the wallet (a new pair of private/public keys is created).

```
orapki wallet add -wallet "C:\opt\oracle\network\server_wallet"
-pwd <password> -dn "CN=<server_machine_name>" -keysize 2048
-self_signed -validity <# of Days>
```

4 Check the contents of the wallet. Notice the self-signed certificate is both a user and trusted certificate.

```
orapki wallet display -wallet
"C:\opt\oracle\network\server_wallet" -pwd <password>
```

5 Check whether the certificate has been exported to the above directory.

6 Make sure the oracle service user can access the wallet file `cwallet.sso` (READ permission).

Step 2: Configure Oracle wallet for client application.

1 Create a directory on the client machine to store the client wallet. Call it `client_wallet`. Create it under the `C:\opt\oracle\network` folder.

```
mkdir C:\opt\oracle\network\client_wallet
```

2 Create a wallet for the Oracle client. Create an empty wallet with auto login enabled.

```
orapki wallet create -wallet "C:\opt\oracle\network\client_wallet"
 -pwd <password> -auto_login
```

3 Add a self-signed certificate in the wallet (a new pair of private/public keys is created).

```
orapki wallet add -wallet "C:\opt\oracle\network\client_wallet"
-pwd <password> -dn "CN=<client_machine_name>" -keysize 2048
-self_signed -validity <# of Days>
```

4 Check the contents of the wallet. Note that the self-signed certificate is both a user and a trusted certificate.

```
orapki wallet display -wallet
"C:\opt\oracle\network\client_wallet" -pwd <password>
```

5 Export the certificate, so it can be loaded into the server wallet later.

```
orapki wallet export -wallet "C:\opt\oracle\network\client_wallet"
 -pwd <password> -dn "CN=<client_machine_name>" -cert
C:\opt\oracle\network\client_wallet\<client-certificate-name>.crt
```

6 Check whether the certificate is exported to the above directory.

Step 3: Perform client-server exchange certificate process. These instructions are for the exchange server and client public keys.

1 Repeat these steps on each of the database client systems.

- Copy `<server-certificate-name>.crt` from the server system to the client system `/opt/aptare/oracle/network/client_wallet` folder.

- Copy `<client-certificate-name>.crt` from the client system to the server system `/opt/aptare/oracle/network/server_wallet` folder.

2 Load the server certificate into the client wallet.

```
orapki wallet add -wallet "C:\opt\oracle\network\client_wallet"
-pwd <password> -trusted_cert -cert
C:\opt\oracle\network\client_wallet\<server-certificate-name>.crt
```

3 Check the contents of the client wallet. Note that the server certificate is now included in the list of trusted certificates.

```
orapki wallet display -wallet
"C:\opt\oracle\network\client_wallet" -pwd <password>
```

4 Load the client certificate into the server wallet.

```
orapki wallet add -wallet "C:\opt\oracle\network\server_wallet"
-pwd <password> -trusted_cert -cert
C:\opt\oracle\network\server_wallet\<client-certificate-name>.crt
```

5 Check the contents of the server wallet. Note that the client certificate is now included in the list of trusted certificates.

```
orapki wallet display -wallet
"C:\opt\oracle\network\server_wallet" -pwd <password>
```

Step 4: Configure the Oracle database to listen for TCPS connection (Server/Oracle system). In the steps below, host is Oracle server IP address and `C:\opt\oracle\network\server_wallet` **is the server wallet location.**

1 Stop the Oracle listener.

```
lsnrctl stop
```

2 Modify the `listener.ora` (`C:\opt\oracle\network\admin\listener.ora`)

```
LISTENER =
  (DESCRIPTION_LIST =
     (DESCRIPTION =
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
      (ADDRESS = (PROTOCOL = TCPS)(HOST = xx.xx.xx.xx)(PORT =
2484))
     )
  )
```

Append the below line at the end of the file.

```
SSL_CLIENT_AUTHENTICATION = FALSE
SECURE_PROTOCOL_LISTENER=(IPC)
WALLET_LOCATION =
    (SOURCE =
      (METHOD = FILE)
      (METHOD_DATA =
         (DIRECTORY = C:\opt\oracle\network\server_wallet)
       )
    )
C:\opt\oracle\network\server_wallet
```

3 Modify the `sqlnet.ora` file (`C:\opt\oracle\network\admin\sqlnet.ora`)

```
SSL_CLIENT_AUTHENTICATION = FALSE
  WALLET_LOCATION =
    (SOURCE =
       (METHOD = FILE)
       (METHOD_DATA =
          (DIRECTORY = C:\opt\oracle\network\server_wallet)
       )
    )

SSL_CIPHER_SUITES = (SSL_RSA_WITH_AES_256_CBC_SHA,
SSL_RSA_WITH_3DES_EDE_CBC_SHA) SQLNET.WALLET_OVERRIDE = TRUE
```

4 Modify the `tnsnames.ora` file (`C:\opt\oracle\network\admin\tnsnames.ora`)

```
SCDB =
     (DESCRIPTION =
        (ADDRESS=
            (PROTOCOL=TCPS)
            (HOST=xx.xx.xx.xx)
            (PORT=2484)
        )
        (CONNECT_DATA=(SERVICE_NAME=scdb)(SID=SCDB))
    )
```

5 Start the Oracle service.

```
lsnrctl start
```

6 Check the listener status.

```
lsnrctl status
```

7 Test Oracle connection using sqlplus.

```
sqlplus username/password@service_name
```

Step 5: Configure the Oracle database to listen for TCPS connection on the client system. Configure the `listener.ora` and `sqlnet.ora` files on the database server using the following steps. In the procedure below, host is

Oracle server IP address and `C:\opt\oracle\network\server_wallet` is the server wallet location.

1 Modify the `listener.ora` (`C:\opt\oracle\network\admin\listener.ora`) and add the below contents.

```
LISTENER =
   (DESCRIPTION_LIST =
      (DESCRIPTION =
         (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
         (ADDRESS = (PROTOCOL = TCPS)(HOST = xx.xx.xx.xx)(PORT =
 2484))
      )
   )
```

Add below line at the end of file:

```
 SSL_CLIENT_AUTHENTICATION = FALSE
    SECURE_PROTOCOL_LISTENER=(IPC)
    WALLET_LOCATION =
     (SOURCE =
       (METHOD = FILE)
       (METHOD_DATA =
         (DIRECTORY = C:\opt\oracle\network\client_wallet)
     )
  )
C:\opt\oracle\network\client_wallet
```

2 Modify the `sqlnet.ora` file (`C:\opt\oracle\network\admin\sqlnet.ora`).

```
SSL_CLIENT_AUTHENTICATION = FALSE
    WALLET_LOCATION =
       (SOURCE =
         (METHOD = FILE)
         (METHOD_DATA =
           (DIRECTORY = C:\opt\oracle\network\client_wallet)
         )
       )

SSL_CIPHER_SUITES = (SSL_RSA_WITH_AES_256_CBC_SHA,
SSL_RSA_WITH_3DES_EDE_CBC_SHA) SQLNET.WALLET_OVERRIDE = TRUE
```

3 Modify the `tnsnames.ora` (`C:\opt\oracle\network\admin\tnsnames.ora`)

```
SCDB =
   (DESCRIPTION =
       (ADDRESS=
           (PROTOCOL=TCPS)
           (HOST=xx.xx.xx.xx)
           (PORT=2484)
       )
       (CONNECT_DATA=(SERVICE_NAME=scdb)(SID=SCDB))
    )
```

Step 6: Load Oracle server wallet certificate to the portal and upgrader Java KeyStore.

1 Login as a root user.

2 Add server certificate in portal java.

```
cd C:\opt\jre\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\jre\lib\security\cacerts
cd C:\opt\jdk\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\jdk\lib\security\cacerts
password: changeit
```

3 Add server certificate in upgrader Java.

```
cd C:\opt\aptare\upgrade\jre\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\aptare\upgrade\jre\lib\security\cacerts
password: changeit
```

Step 7: Modify connection URL in the portal and receiver property file.

1. Stop portal and agent services.
2. Modify database URL in `/opt/aptare/portalconf/portal.properties`.

   ```
   db.url=jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
   (HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))
   ```

3. Modify database URL in `/opt/aptare/datarcvrconf/datrarcvrproperties.xml`.

   ```
   <URL>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
   (HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB))</URL>
   ```

4. Start portal and agent services.

Configure TLS in Oracle with IT Analytics on Windows in non-split architecture

In a non-split architecture, the database and portal are on the same system.

Step 1: Configure Oracle wallet on the server side.

1. Create a directory on the server system to store the server wallet as `server_wallet`. Run the make directory command below at `C:\opt\oracle\network\` folder.

   ```
   mkdir C:\opt\oracle\network\server_wallet
   ```

2. Create an empty wallet for the Oracle server with auto login enabled.

   ```
   orapki wallet create -wallet "C:\opt\oracle\network\server_wallet"
    -pwd <password> -auto_login
   ```

3. Add a self-signed certificate in the wallet (a new pair of private/public keys is created).

   ```
   orapki wallet add -wallet "C:\opt\oracle\network\server_wallet"
   -pwd <password> -dn "CN=<server_machine_name>" -keysize 2048
   -self_signed -validity <# of Days>
   ```

4 Check the contents of the wallet. Notice the self-signed certificate is both a user and trusted certificate.

```
orapki wallet display -wallet
"C:\opt\oracle\network\server_wallet" -pwd <password>
```

5 Export the certificate so it can be loaded into the client wallet later.

```
orapki wallet export -wallet "C:\opt\oracle\network\server_wallet"
 -pwd <password> -dn "CN=<server_machine_name>" -cert
C:\opt\oracle\network\client_wallet\<server-certificate-name>.crt
```

6 Check whether the certificate has been exported to the above directory.

7 Make sure the Oracle service user can access the wallet file `cwallet.sso` (READ permission).

Step 2: Configure Oracle wallet for client application.

1 Login as oracle user.

```
su - aptare
```

2 Create a directory on the client system to store the client wallet. Call it `client_wallet`. Create it under the `/opt/aptare/oracle/network` folder.

```
mkdir C:\opt\oracle\network\client_wallet
```

3 Create a wallet of the oracle client. Create an empty wallet with auto login enabled.

```
orapki wallet create -wallet "C:\opt\oracle\network\client_wallet"
 -pwd <password> -auto_login
```

Step 3: Load server certificate into the client wallet.

1 Load the server certificate into the client wallet

```
orapki wallet add -wallet "C:\opt\oracle\network\client_wallet"
-pwd <password> -trusted_cert -cert
C:\opt\oracle\network\client_wallet\<server-certificate-name>.crt
```

2 Check the contents of the client wallet. Note that the server certificate is now included in the list of trusted certificates.

```
orapki wallet display -wallet
"C:\opt\oracle\network\client_wallet" -pwd <password>
```

Step 4: Configure the Oracle database to listen for TCPS connection: Configure the `listener.ora`, `tnsnames.ora`, and `sqlnet.ora` files on the database server using the following steps. In these steps, host is the oracle server IP address and the server wallet location is `/opt/aptare/oracle/network/server_wallet`.

1 Stop the Oracle listener before updating the files.

```
lsnrctl stop
```

2 Modify the `listener.ora` (`C:\opt\oracle\network\admin\listener.ora`)

```
LISTENER =
    (DESCRIPTION_LIST =
       (DESCRIPTION =
         (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
         (ADDRESS = (PROTOCOL = TCPS)(HOST = xx.xx.xx.xx)(PORT =
 2484))
        )
   )
```

Add below line at the end of file:

```
SSL_CLIENT_AUTHENTICATION = FALSE
  SECURE_PROTOCOL_LISTENER=(IPC)
     WALLET_LOCATION =
       (SOURCE =
       (METHOD = FILE)
       (METHOD_DATA =
         (DIRECTORY = C:\opt\oracle\network\server_wallet)
       )
     )

C:\opt\oracle\network\server_wallet
```

3 Modify the `tnsnames.ora` (C:\opt\oracle\network\admin\tnsnames.ora).

```
SCDB =
   (DESCRIPTION =
      (ADDRESS=
        (PROTOCOL=TCPS)
        (HOST=xx.xx.xx.xx)
        (PORT=2484)
      )
      (CONNECT_DATA=(SERVICE_NAME=scdb)(SID=SCDB))
   )
```

4 Start the Oracle listener.

```
lsnrctl start
```

5 Check the listener status

```
lsnrctl status
```

6 Test the Oracle connection using sqlplus.

```
sqlplus username/password@service_name
```

Step 5: Load oracle server wallet certificate to the portal and upgrader Java KeyStore.

1 Add server certificate in portal Java.

```
cd C:\opt\jre\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\server_wallet\server-cert-db.crt -keystore

C:\opt\jre\lib\security\cacerts
cd C:\opt\jdk\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\server_wallet\server-cert-db.crt -keystore

C:\opt\jdk\lib\security\cacerts
password: changeit
```

2 Add server certificate in upgrader Java

```
cd C:\opt\aptare\upgrade\jre\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\server_wallet\server-cert-db.crt -keystore

C:\opt\aptare\upgrade\jre\lib\security\cacerts
password: changeit
```

Step 6: Modify connection URL in the portal and receiver property file.

1 Stop portal and agent services.

2 Modify database URL in `C:\opt\aptare\portalconf\portal.properties`.

```
db.url=jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))
```

3 Modify database URL in
`C:\opt\aptare\datarcvrconf\datrarcvrproperties.xml`.

```
<URL>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))</URL>
```

4 Start portal and agent services.

Configure TLS in user environment

On Windows

Step 1: Load oracle server wallet certificate to portal and upgrader Java KeyStore. This step is required only if the wallet certificate is self-signed.

1 Add server certificate in portal Java.

```
cd C:\opt\jre\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\jre\lib\security\cacerts
cd C:\opt\jdk\bin
keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\jdk\lib\security\cacerts
password: changeit
```

2 Add server certificate in upgrader Java.

```
cd C:\opt\aptare\upgrade\jre\bin

keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\aptare\upgrade\jre\lib\security\cacerts
password: changeit
```

Step 2: Ensure the Oracle service user has READ access to `cwallet.sso` file of the server wallet. To provide the permission:

1 Right-click on the `cwallet.sso` file of the server wallet and select **Properties**.

2 Go to the **Security** tab and click **Edit** of the group or user names.

3 Click **Add**, search for Oracle service user, and click **OK**.

4 Select **READ** permission and click **OK**.

Step 3: Modify connection URL in portal and receiver property file.

Note: Host, port, and SERVICE_NAME can be different.

---

1. Stop the portal and agent services.
2. Modify database URL in `C:\opt\aptare\portalconf\portal.properties`.

   `db.url=jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)(HOST=x.x.x.x)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SID)))`

3. Modify database URL in `C:\opt\aptare\datarcvrconf\datrarcvrproperties.xml`.

   `<URL>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)(HOST=x.x.x.x)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SID)))</URL>`

4. Start portal and agent service.

On Linux

Step 1: Load oracle server wallet certificate to portal and upgrader Java KeyStore. This step is required only if the wallet certificate is self-signed.

1. Login as a root user.
2. Add server certificate in portal Java.

```
cd /user/java/bin

keytool -import -trustcacerts -alias ora_server_cert -file
/opt/aptare/oracle/network/server_wallet/server-cert-db.crt
-keystore
/usr/java/lib/security/cacerts
password: changeit
```

3. Add server certificate in upgrader Java.

```
cd C:\opt\aptare\upgrade\jre\bin

keytool -import -trustcacerts -alias ora_server_cert -file
C:\opt\oracle\network\client_wallet\server-cert-db.crt -keystore

C:\opt\aptare\upgrade\jre\lib\security\cacerts
password: changeit
```

Step 2: Modify connection URL in portal and receiver property file.

Note: Host, port, and SERVICE_NAME can be different.

---

1 Stop the portal and agent services.

```
/opt/aptare/bin/tomcat-portal stop
/opt/aptare/bin/tomcat-agent stop
```

2 Modify database URL in `/opt/aptare/portalconf/portal.properties`.

```
db.url=jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB)))
```

3 Modify database URL in `/opt/aptare/datarcvrconf/datrarcvrproperties.xml`.

```
<URL>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)
(HOST=xx.xx.xx.xx)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=SCDB))</URL>
```

4 Start portal and agent services.

```
/opt/aptare/bin/tomcat-portal start
/opt/aptare/bin/tomcat-agent start
```

Appendix C

IT Analytics for NetBackup on Kubernetes and appliances

This appendix includes the following topics:

Configure embedded IT Analytics Data collector for NetBackup deployment on appliances (including Flex appliances)

You can configure a Data Collector on the primary server pod. Following are the steps for respective configurations.

Note: From NetBackup version 10.3 Cloud Scale release, Data Collector on primary server pod is supported.

Configure IT Analytics for NetBackup deployment by configuring Data Collector on a primary server

The below steps to configure the IT Analytics Data Collector on a primary server must be performed as a root user. On a Flex appliance, connect to the primary server pod first and then switch to the root user using `sudo`. On a NetBackup Appliance, access shell by creating NetBackup CLI user.

To configure IT Analytics for NetBackup deployment

1. Create a DNS server entry in such a way that IP of the IT Analytics Portal must be resolvable to a single FQDN. IP of the IT Analytics Portal must be resolved to:

   ```
   itanalyticsagent.<yourdomain>
   ```

   Note the following:

   - If the Portal URL is `itanalyticsportal.<yourdomain>`, then ensure to add the DNS entries for the following hostnames:

     ```
     itanalyticsagent.<yourdomain>
     ```

   - If the Portal URL is `itanalyticsportal.<yourdomain>`, then ensure to add the DNS entries for the following hostnames:

     ```
     aptareagent.<yourdomain>
     ```

2. Collect the `<your-collector-name>.key` file for the new Data Collector by accessing the Portal link and creating a collector and copy it to the host machine from where NetBackup primary is deployed.

   For more information, refer to the *Data Collector Encryption section in IT Analytics User Guide*.

3. Collect the `<your-collector-name>.json` file for the new Data Collector by accessing the Portal link and creating a collector and copy it to the host machine from where NetBackup primary is deployed.

   For more information, refer to the *Data Collector Encryption section in IT Analytics User Guide*.

4. Create a new folder **analyticscollector** at persisted location (for example, `/mnt/nbdata/`) using the following commands:

   ```
   cd "/mnt/nbdata/"
   mkdir analyticscollector
   ```

5. Copy the `<your-collector-name>.key` file to `/mnt/nbdata/analyticscollector` inside the NetBackup primary host or container as the case may be.

6. Exit from the container and copy the `<your-collector-name>.json` file to `/mnt/nbdata/analyticscollector` inside the NetBackup primary host or container as the case may be.
7. In case if the data-receiver is configured with self-signed certificate (https), you must add the certificate in the data collector.

   See "Configure the Data Collector to trust the certificate" on page 261.
8. Connect to the NetBackup primary host or the container.
9. Navigate to `/usr/openv/analyticscollector/installer/` location and perform the following:

   - Open the `responsefile.sample` and add the following parameters:
     If the Portal version is lower than 11.3, create the response file with the following contents.

```
COLLECTOR_NAME=<your-collector-name>
COLLECTOR_PASSCODE=<your-password>
DR_URL=<http>/<https>://itanalyticsagent.<yourdomain>
COLLECTOR_KEY_PATH=<path to your-collector-name.key>
HTTP_PROXY_CONF=N
HTTP_PROXY_ADDRESS=
HTTP_PROXY_PORT=
HTTPS_PROXY_ADDRESS=
HTTPS_PROXY_PORT=
PROXY_USERNAME=
PROXY_PASSWORD=
PROXY_EXCLUDE=
```

     If the Portal version is 11.3 or later, create the response file with the following contents.

```
COLLECTOR_REGISTRATION_PATH=<keyfile path>
HTTP_PROXY_CONF=N
HTTP_PROXY_ADDRESS=
HTTP_PROXY_PORT=
HTTPS_PROXY_ADDRESS=
HTTPS_PROXY_PORT=
PROXY_USERNAME=
PROXY_PASSWORD=
PROXY_EXCLUDE=
```

10. Configure the Data Collector with the IT Analytics Portal as follows.

Note: If the Data Collector installed is of a lower version than the IT Analytics Portal, wait for the Data Collector auto-upgrade to finish before you proceed.

For NetBackup Appliance version 5.3 or later:

- Run the following command as a NetBackup CLI user:

  ```
  /usr/openv/analyticscollector/installer/dc_installer.sh -c
  /usr/openv/analyticscollector/installer/responsefile.sample
  ```

- To verify the Data Collector integration with IT Analytics Portal, run:

  ```
  /usr/openv/analyticscollector/mbs/bin/checkinstall.sh
  ```

For NetBackup Appliance version 5.1.1:

- Run the following command as a NetBackup CLI user:

  ```
  sudo /usr/openv/analyticscollector/installer/dc_installer.sh
  -c /usr/openv/analyticscollector/installer/responsefile.sample
  ```

- To verify the Data Collector integration with IT Analytics Portal, run:

  ```
  sudo /usr/openv/analyticscollector/mbs/bin/checkinstall.sh
  ```

If you are on Flex Appliance:

- Connect to the primary server container and then switch to roo user using `sudo` and run:

  ```
  /usr/openv/analyticscollector/installer/dc_installer.sh -c
  /usr/openv/analyticscollector/installer/responsefile.sample
  ```

- To verify the Data Collector integration with IT Analytics Portal, run:

  ```
  /usr/openv/analyticscollector/mbs/bin/checkinstall.sh
  ```

  If the Data Collector is configured with the Portal, it will display **SUCCESSFUL**.

Note: If there is a version mismatch of `aptare.jar` between Data Collector and Portal, execution of `checkinstall.sh` command will trigger an auto-update of the Data Collector.

11 Check the Data Collector services status by running the following command and ensure that the following Data Collector services are up and running:

```
/usr/openv/analyticscollector/mbs/bin/aptare_agent status
```

Output of the above command:

```
IT Analytics WatchDog is running (pid: 13312).
IT Analytics MetaDataCollector is stopped.
IT Analytics EventDataCollector is stopped.
IT Analytics DataCollector process is running (pid: 13461).
IT Analytics On-demand process is running (pid: 13463).
IT Analytics Message Relay Server process is running (pid: 13471)
```

For more information about IT Analytics Data Collector policy, see *IT Analytics User Guide*.

Configure IT Analytics for NetBackup deployment by configuring Data Collector on a separate host machine

IT Analytics Data Collector can be configured to use with NetBackup primary server in this environment. IT Analytics can be configured at the time of primary server deployment or user can update the primary server custom resource to configure NetBackup IT Analytics.

Note: Remote data collection from Netbackup primary server cannot be configured when multi-factor authentication is configured on Netbackup Appliance.

To configure IT Analytics for NetBackup deployment:

1 Using the `ssh-keygen` command, generates public key and private key on IT Analytics data collector.

IT Analytics Data Collector uses passwordless ssh login.

2 Create and copy NetBackup API key from NetBackup web UI.

3 On IT Analytics Portal:

- Navigate to **Admin > Collector Administration > Select respective Data Collector > Add policy > Veritas NetBackup > Add**.

- Add required options, specify the NetBackup API in the **API Key** field, and then click **OK**.
- Select newly added primary server from **NetBackup Primary Servers** and provide **nbitanalyticsadmin** as **Primary Server User ID**.
- Provide **privateKey=<path-of-private-key>|password=<passphrase>** as **Primary Server Password** and **Repeat Password** whereas **<path-of-private-key>** is the private key created using ssh-keygen in earlier steps and **<passphrase>** is the passphrase used while creating private key via ssh-keygen.
- Provide appropriate data to Data Collector policy fields and select collection method as **SSH or WMI protocol to NetBackup Primary Server**.

Configuring the primary server with IT Analytics tools is supported only once from primary server custom resource.

For more information about IT Analytics Data Collector policy, see *Add a Veritas NetBackup Data Collector policy* section.

For more information about adding NetBackup Primary Servers within the Data Collector policy, see *Add/Edit NetBackup Primary Servers within the Data Collector policy* section in *IT Analytics Data Collector Installation Guide for Backup Manager*.

To change the already configured public key

1. Connect to the NetBackup Primary host or container.
2. Copy the new public keys in the `/home/nbitanalyticsadmin/.ssh/authorized_keys` and `/mnt/nbdata/.ssh/nbitanalyticsadmin_keys` files.
3. Restart the **sshd** service using the `systemctl restart sshd` command.

Configure IT Analytics for NetBackup deployment on Kubernetes

You can configure Data Collector on the primary server. Following are the steps for respective configurations.

Note: From NetBackup version 10.3 Cloud Scale release, Data Collector on primary server is supported.

Configure IT Analyticsfor NetBackup deployment by configuring Data Collector on a primary server

The following steps to configure the IT Analytics Data Collector need to be performed as root user. Connect to the primary server pod and then switch to root user using `sudo`.

To configureIT Analytics for NetBackup deployment

1. Create a DNS server entry in such a way that IP of the Portal must be resolvable to a single FQDN. IP of the IT Analytics Portal must resolve to:

   ```
   itanalyticsagent.<yourdomain>
   ```

   Note the following:

   - If the Portal URL is `itanalyticsportal.<yourdomain>`, then ensure to add the DNS entries for the following hostnames:

     ```
     itanalyticsagent.<yourdomain>
     ```

   - If the \ URL is `aptareportal.<yourdomain>`, then ensure to add the DNS entries for the following hostnames:

     ```
     aptareagent.<yourdomain>
     ```

2. Collect `<your-collector-name>.key` file for the new Data Collector by accessing the Portal link and creating a collector and copy it to the host machine from where Kubernetes cluster is accessed.

   For more information, see *Data Collector Encryption* section in *IT Analytics User Guide*.

3. Copy the `<your-collector-name>.key` file to `/mnt/nbdata/analyticscollector` inside hte NBU Primary pod using the command `kubectl cp <keyfile-name> <namespace>/<primary-pod-name>:/mnt/nbdata/analyticscollector`.

4. In case if the data-receiver is configured with self-signed certificate (https). User must add the certificate in the Data Collector.

   See “Configure the Data Collector to trust the certificate” on page 261.

5. Execute the following command on the primary server pod:

   ```
   kubectl exec -it -n <namespace> <primaryServer-pod-name> -- bash
   ```

6 Navigate to `/usr/openv/analyticscollector/installer/` location and perform the following:

- Open the `responsefile.sample` and add the following parameters:

```
COLLECTOR_NAME=name_of_the_data_collector
COLLECTOR_PASSCODE=passcode_for_the_data_collector
DR_URL=data_receiver_URL
COLLECTOR_KEY_PATH=path_to_the_key_file
HTTP_PROXY_CONF=N
HTTP_PROXY_ADDRESS=
HTTP_PROXY_PORT=
HTTPS_PROXY_ADDRESS=
HTTPS_PROXY_PORT=
PROXY_USERNAME=
PROXY_PASSWORD=
PROXY_EXCLUDE=
```

- Run `/usr/openv/analyticscollector/installer/dc_installer.sh -c /usr/openv/analyticscollector/installer/responsefile.sample` command to configure Data Collector with IT Analytics portal.

7 Validate Data Collector integration with IT Analytics by performing the following:

- Navigate to `/usr/openv/analyticscollector/mbs/bin/` location.
- Run the following command:
  `./checkinstall.sh`
  If Data Collector is configured with portal, it will display as **SUCCESSFUL**.

Note: If there is a version mismatch of `aptare.jar` between Data Collector and Portal, execution of `checkinstall.sh` command will trigger an auto-update of the Data Collector.

8 Check the Data Collector services status by running the following command and ensure that the following Data Collector services are up and running:

```
/usr/openv/analyticscollector/mbs/bin/aptare_agent status
```

Output of the above command:

```
IT Analytics WatchDog is running (pid: 13312).
IT Analytics MetaDataCollector is stopped.
IT Analytics EventDataCollector is stopped.
IT Analytics DataCollector process is running (pid: 13461).
IT Analytics On-demand process is running (pid: 13463).
IT Analytics Message Relay Server process is running (pid: 13471)
```

For more information about IT Analytics Data Collector policy, see *IT Analytics User Guide*.

Configuring IT Analytics for NetBackup deployment by configuring Data Collector on a separate host machine

IT Analytics can be configured to use with NetBackup primary server in this Kubernetes environment.IT Analyticscan be configured at the time of primary server deployment or user can update the primary server custom resource to configure IT Analytics.

To configure IT Analytics for NetBackup deployment

1 Using the `ssh-keygen` command, generates public key and private key on IT Analytics Data Collector.

IT Analytics Data Collector uses passwordless ssh login.

2 Update the primary server custom resource, copy public key generated in previous steps to “itAnalyticsPublicKey” section.

- Apply the primary server custom resource changes using `kubectl apply -f environment.yaml -n <namespace>`.

  On successfully deployment of primary server, describe the primary server custom resource using `kubectl describe PrimaryServer <primary-server-name> -n <namespace>`
- In the **Status** section, verify **IT Analytics Configured** is set to **true**.

For more information, refer to the *NetBackup™ Web UI Administrator's Guide*.

3. Create and copy NetBackup API key from NetBackup web UI.

4. On IT Analytics Portal:
   - Navigate to **Admin > Collector Administration > Select respective Data Collector > Add policy > Veritas NetBackup > Add**.
   - Add required options, specify the NetBackup API in the **API Key** field, and then click **OK**.
   - Select newly added primary server from **NetBackup Primary Servers** and provide **nbitanalyticsadmin** as **Primary Server User ID**.
   - Provide **privateKey=<path-of-private-key>|password=<passphrase>** as **Primary Server Password** and **Repeat Password** whereas **<path-of-private-key>** is the private key created using ssh-keygen in earlier steps and **<passphrase>** is the passphrase used while creating private key via ssh-keygen.
   - Provide appropriate data to Data Collector policy fields and select collection method as **SSH or WMI protocol to NetBackup Primary Server**.

Configuring the primary server with IT Analytics tools is supported only once from primary server custom resource.

For more information about IT Analytics Data Collector policy, see *Add a Veritas NetBackup Data Collector policy* section.

For more information about adding NetBackup Primary Servers within the Data Collector policy, see *Add/Edit NetBackup Primary Servers within the Data Collector policy* section in *IT Analytics Data Collector Installation Guide for Backup Manager*.

To change the already configured public key

1. Execute the following command in the primary server pod:

   ```
   kubectl exec -it -n <namespace> <primaryServer-pod-name> --
   /bin/bash
   ```

2. Copy the new public keys in the `/home/nbitanalyticsadmin/.ssh/authorized_keys` and `/mnt/nbdata/.ssh/nbitanalyticsadmin_keys` files.

3. Restart the **sshd** service using the `systemctl restart sshd` command.