Database Management

Cisco Unified Customer Voice Portal (Unified CVP) provides access to database maintenance tasks such as database backups and data purges through the Operations Console.

Although the Reporting Service does not directly perform database administrative and maintenance activities such as backups or purges, familiarize yourself with the database management concepts discussed in this chapter.

Note Operations such as purge and changing user passwords are executed through OAMP by the Administrator cvp_dbadmin.

Passwords

Passwords on the reporting server *must* be created and updated as part of the Unified CVP installation or by means of the Operations Console. Do not use any other means to create or update passwords.

Passwords for the Unified CVP Reporting Server are kept encrypted locally on the OAMP server and are set on the reporting server. If these passwords are not in synch (are not identical), the Operations Console will be unable to communicate with the reporting server, and restricted operations such as the purge will fail.

If you implement a password expiration policy, remember to use the Operations Console to change the Database Administrator, and Database Users passwords before the passwords expire to avoid the possibility of data loss or downtime.

Note Using the Operations Console to change passwords and to renew passwords before they expire ensures that all dependencies are synchronized.

Caution Changing passwords outside of the Operations Console can result in a failed connection between the reporting server and the database.

Reporting passwords are subject to both the Unified CVP password policy *and* the password policy enforced by the operating system of the computer on which the reporting server resides and must meet the requirements of the more restrictive policy.

Database Users

Unified CVP defines three categories of database users: database administrator, application user, and reporting user.

Database Administrator

The `cvp_dbadmin` creates, updates, and owns the database.

This user can create and delete reporting users and perform database administrative activities, such as purge and backup.

This account should not be used to run the database or to run reports against the system.

Note If this administrator's password expires, then data insertion and purge will fail, which could result in data loss.

Application User

The Unified CVP JDBC uses `cvp_dbuser` to access the Informix database. This user has the rights to connect, insert, update, and delete records in the Unified CVP database. If this user's password expires, then data insertion and purge will fail, which could result in data loss.

The User ID and password for the application user is required to access the Cisco Unified Intelligence Center data source.

Reporting User

The Unified CVP OAMP has a UI page for creating Reporting users who have read-only database access to the Unified CVP Informix reporting database.

After the Active Directory configuration for these users is enabled in the Unified IC Administration Console, they can log in to Unified IC reporting with their AD credentials.

They have the basic "Login User" user role only, until the Unified IC Security Administrator assigns additional roles and privileges to them.

Data Categories and Data Retention

Using the Operations Console, users are able to select the time of day to run database purge and to set the number of days of data to be retained by data category. During schema creation, default data retention values are specified for each data category. Note that a high-level category, such as Call, cannot have a lower retention time than a dependent category, such as Call Event.

Increased database space availability, as documented in the *Solution Design Guide for Cisco Unified Contact Center Enterprise*.

Default Data Retention for Data Categories

The following data categories exist for Unified CVP. Note that a high level category, such as Call, cannot have a lower retention time than a dependent category, such as CallEvent. For each category, the default data retention times, in days, is given within parentheses.

level 1: `Call` (30)

level 2: `-Call Event` (30)

level 2: `-VoiceXML Session` (30)

level 3: `--VoiceXML Element` (15)

level 4: `---VoiceXML ECC Variable` (15)

level 4: `---VoiceXML Interact Detail` (15)

level 4: `---VoiceXML Session Variable` (15)

level 4: `---VoiceXML Element Detail` (15)

Database Purge

This section explains how to schedule purges, the difference between midday and nightly purges, and how to run emergency purges.

To allow for rapid space management, all Unified CVP Reporting Server data is kept in date-specific fragments. On a daily basis, new day fragments are created for incoming data. This allows the Unified CVP Reporting Server to quickly drop old data by dropping the disk fragment in which that data resides.

This means that new fragments must be created on a regular basis to ensure that they can be rapidly disposed of when their retention period expires. All of this is handled by the purge.

In a quiet environment, the purge can execute in less than one second. In a busy environment, the purge might take considerably longer (15-20 minutes) while the purge navigates around running processes.

To allow for situations where the purge may be unable to execute, space is allocated two days before it is needed. This allocation is triggered by the first purge to run after a date boundary has been crossed.

If the nightly purge is scheduled to run at 1:00 a.m., it will typically perform this task. If the nightly purge is scheduled to run at 11:00 p.m., then the next day will not occur until 11:00 a.m. the next morning, which may not be an optimal time to execute the process.

For this reason, schedule the nightly purge to occur after midnight.

Schedule Purges

To run database purge from the Operations Console:

Step 1 Select **Device Management** > **CVP Reporting Server**.

Step 2 Select a reporting server by clicking on the link in its name field or by clicking the radio button preceding it and then clicking **Edit**.

Step 3 At the Edit Reporting Server Configuration window, select the Database Administration menu in the toolbar, then select **Data Delete**.

Step 4 On the Reporting Server - Data Delete page, change the data retention time for each category of data.

Step 5 Select the hours and minutes to run the purge each day. This defines the time for the primary (nightly) purge and sets the Midday purge to run 12 hours later.

Step 6 Click **Save & Deploy**.

What to do next

See the *Administration Guide for Cisco Unified Customer Voice Portal* at https://www.cisco.com/en/US/products/sw/custcosw/ps1006/products_user_guide_list.html for information on categories of reporting data that can be purged and the default number of days to retain data before purging.

Nightly and Midday Purges

When you schedule a purge from the Operations Console, two jobs are scheduled on the reporting server - the nightly purge and the midday purge:

- The **nightly purge** job runs at the time (hour and minute) that you define in the Operations Console. The nightly purge performs a purge if necessary (as required by a data retention value, or for an emergency purge—see the following section on emergency purges), in addition to other tasks updating the database statistics.

  If a purge *is* required and performed, the statistics are updated after the purge.

  In addition, on Sundays, the nightly purge also copies the Informix log file to a backup folder, creates a new log file and deletes the old.

  The nightly purge should be scheduled after midnight.

- The **midday purge** is automatically scheduled to run 12 hours after the nightly purge. So, for example, if you schedule a purge at 2 a.m., then the nightly purge is run at 2 a.m. and the midday purge at 2 p.m.

  Midday purge also serves as a backup for the nightly purge. If the nightly purge fails to allocate new fragments for new data, this will be taken care of by the midday purge. A midday purge is not system intensive in the same way that the nightly purge is.

In the event that data volume spikes during the day and an emergency purge is required, it will be handled at midday.

Emergency Purge

If the number of days of data that you chose to retain cannot be contained within the database, then the database initiates emergency purge. It purges the old data to create space for new data. Emergency purge is a critical safety mechanism for Unified CVP.

If used space exceeds the systems threshold, a Simple Network Management Protocol (SNMP) trap message notifies the user after the emergency purge is complete. The SNMP notification alerts the user of the loss of data and request that they reduce their retention days data settings.

Reduce the number of days of data retained so that emergency purge is not required. Also, you can reduce the data generated by using data filters (for VXML Server application detail data filtering).

Emergency Purge is triggered based on the database size and the free space value of the database. The **partitionparameters** table in the ciscoadmin database has the fields **dbsize** and **pctfree** which indicates the database size and percentage of the database size which is free respectively. For example, If the database size is 100 GB and free space is 17%, then the emergency purge starts when the database is 83% filled. It purges the oldest data until the required free space is achieved.

Guidelines for Purge

- **Data Granularity** - The CVP Reporting database houses records of calls handled by Cisco Unified Voice call servers. The amount of data captured for each call is managed by filters specified in the Operations Console. The granularity of data captured depends on these settings.
- **Data Retention** - As limited space is available to capture this data, the purge mechanism uses retention settings to govern how long data is retained.

  If there is insufficient space to retain data for the desired time frame, the oldest data is purged in one-day increments until there is sufficient space for the reporting server to remain operational.

  If more data is captured on a daily basis than can be stored, the purge mechanism will be unable to remove data because it operates only on a daily basis. If this is the case, consider installing a larger reporting server.
- **Database backup and purge** cannot run at the same time. Purge should be scheduled at least 30 minutes before a backup. These jobs, as well as on-demand backup, should be run at low call and reporting volume times. From the perspective of Unified CVP, database backups are optional, data purges are mandatory. However, from the perspective of the user, database backups should *not* be considered optional.
- **Reporting Server** - During a database purge operation, the reporting server disconnects from the database (though for no more than 10 minutes) and starts buffering messages in memory until the purge is finished. The same memory limitations apply as described in the section Reporting User.
- **Reporting users** may be disconnected from the database if they are holding locks that contend with purge. Notify reporting users not to run reports at this time.
- **Upgrades** - Turn off scheduled purge before performing an upgrade.
- **Windows Scheduled Tasks**- The database backup and purge maintenance tasks are created as Windows Scheduled Tasks, and can be viewed in the Scheduled Tasks window. (**Start** > **Programs** > **Accessories** >

System Tools > **Scheduled Tasks**.) Periodically, you should check the Scheduled Tasks to ensure the Last Run Time was as expected and no status messages exist.

Database Backup

Unified CVP lets users turn the scheduling of data backups on or off, and to run backups on demand. Backups are made to the reporting servers local file system. By default, scheduled backups are turned off.

Caution Unified CVP backup scheduling is an optional feature. Backup is the user's responsibility. Data loss may occur if the backing up of files is not managed properly by the user.

If Unified CVP backup scheduling is turned on, the backup occurs once per day. Backups must be scheduled to run no sooner than 30 minutes after the scheduled purge job.

Users can run a backup on demand—as long as another backup, or a purge, is not already running. Database backups are performed and stored on the local machine. Due to space limitations, a maximum of two backups and a minimum of one backup are available on the local machine. Retaining two backup files is critical. If the system fails while writing a backup, and a restore is necessary, the older backup file is required for restore.

Follow these guidelines:

- Keep a given backup for at least two weeks.
- Check the integrity of the backup periodically.
- Run a backup before an upgrade.

Unified CVP uses the Informix backup utility ontape (for both backup and restore).

When a new backup launches—either scheduled, or on demand from the Operations Console—the new file is named `cvp_backup_data.gz`. The Unified CVP backup script copies the previous `cvp_backup_data.gz` backup file and renames it to `cvp_backup_data.old.gz`. This always leaves two backup files on the local system and makes it easy for Unified CVP administrators to script copy jobs to move the files. The backup script ensures that two backups cannot be launched at the same time.

Note The backup script also ensures that a backup cannot be launched if a purge is underway, and vice versa.

Important You must manually or automatically (by creating an automated job) copy the `cvp_backup_data` files to a separate machine at a separate location. This will prevent accidental deletion of the files when the CVP machine fails or when CVP needs to be installed / uninstalled.

Warning Only the `cvp_backup_data.old.gz` file can be copied. The `cvp_backup_data.gz` file cannot be copied. Attempting to copy the `cvp_backup_data.gz` file locks the file and prevent another backup from running.

Database backup updates the log file when the backup finishes. If the database server goes down during backup, the backup file gets corrupted.

To check if Informix is up and running and to validate the backup, execute the `cvpverifybackup.bat` file located at `%CVP_Home%\bin\cvpverifybackup.bat.`

While executing the script, you are prompted with following message:

```
Please put in Phys Tape 1.
Type <return> or 0 to end:
Press 0 and press Enter
```

Note Based on the size of the database, the prompt keeps changing as Tape 2, Tape 3 and so on.

This process execution takes a long time (based on the database size) to validate and the results are displayed on the console.

In Cisco Unified CVP, there is a supported script to perform a database restore.

Restoring a backup image is required when older data on a backup image needs to be recovered. It is also required when a machine is rebuilt after a hardware failure and you need to recover data.

Note Although it is possible to restore a backup image from one reporting server to another, such a restoration is not supported with the CVP restore process.

The restore process in Unified CVP is as follows:

- Stop the CallServer process (Reporting Server).
- Execute the script: `%CVP_Home%\bin\cvprestore.bat.`
- Restart the CallServer process.

Caution Using a third-party backup utility to back up the Informix database is ineffective and may be dangerous to the integrity of the reporting database. The only effective way to perform a reporting database backup is with the backup process provided by the OAMP interface.

For information on configuring backups, see the *Administration Guide for Cisco Unified Customer Voice Portal*.

Backup and Purge Retries

Occasionally, a backup or purge cannot run when scheduled. For example, if an on-demand backup is running when a purge is scheduled to run, the purge will be prevented from running.

Retries of scheduled backups or purges are performed according to the following rules.

Note There are no retries for an on-demand backup.

- A scheduled backup retries every 10 minutes, for up to 4 hours.
- A purge retries every 10 minutes, for up to 6 hours.
- At the end of 4 hours (for a backup) or 6 hours (for a purge), if the operation has not succeeded, retries stop and an SNMP alert is sent.
- If both a backup and a purge are retrying simultaneously, there is no guarantee as to which operation will run first.
- If a lock (the mechanism preventing a backup or purge from running) is more than 12 hours old, the system clears it.

Database Recovery

Unified CVP database recovery returns the database to the state of the most recent complete backup. For example, if the user schedules a backup at 01:00 and restores the database at 23:00, the same day, the restored database is in the state it was in at 01:00.

During a database restore, the database will go offline for the duration of the restore operation.

Note Data loss occurs if the reporting server is turned off and the message bus exceeds its temporary persistence capabilities.

Caution Before following a database restore, the following steps must be performed:

1. Before the restore, disable scheduled tasks (backup, purge).
2. After the restore, re-enable scheduled tasks.

Failure and Restoration

If the reporting server fails, messages destined for the reporting server are buffered by the Call Server, in memory, up to 200,000 messages. After that limit is reached, all new message detail information is dropped.

If the database connection fails, the reporting server sends out an SNMP alert and starts persisting messages to a file, up to a user-specified limit. During this time the reporting server stays *In Service*. When 75% of the specified limit is reached, a warning is written to the log file. Once 100% of the limit is reached, an SNMP alert is sent out and the reporting server goes into *Partial Service*—any new messages may be dropped.

When the database connection comes back up, the reporting server goes into recovery mode and changes its state to Partial Service if it is not in that state already. It then starts reading messages from the file and committing them to the database. Depending on the size of the file, it may take a long time (sometimes hours) to commit all of the data to the database. Any new messages that come in during recovery will be buffered in memory. There is, however, a limit to the number of messages that the reporting server can buffer. This is true regardless of the mode or state it is in. When the number of buffered messages reaches 100,000, an SNMP alert is sent out to warn the user. At 200,000 another SNMP alert is sent out and all new messages detail information is dropped—keeping only basic data like call, call event, and session information. Also at 200,000, the reporting server changes its state to *Partial Service*, if it is not already in that state. After the total number of buffered messages reaches 300,000, another SNMP alert is sent out and all new messages are dropped from that point forward.

When the number of messages in memory drops back below 50,000, an SNMP alert is sent out stating that the queue size is back to normal, and the reporting server's state goes back to *In Service*.

If, on startup, a persistent file exists, the reporting server stays in *Partial Service* and goes into recovery mode as previously described.

During a database purge operation, the reporting server disconnects from the database and starts buffering messages in memory until the purge is done. The same memory limitations as previously described apply in this case as well.

⚠ **Caution** **When the reporting server is in *Partial Service*, there are no guarantees that new messages will be kept and committed to the database. As many as possible will be buffered in memory, but at some point they may be dropped either partially or fully.**