




System Architecture and Reporting

- [Central Controller, on page 1](#)
- [Peripherals and Peripheral Gateways, on page 3](#)
- [Administration & Data Server and Administration Client, on page 5](#)
- [Historical Data Server, on page 6](#)
- [Cisco Unified Communications Manager, on page 8](#)
- [Reporting Servers, on page 8](#)
- [Agent and Supervisor Desktops, on page 8](#)
- [Voice Response Units, on page 8](#)

Central Controller

Figure 1: Central Controller



The Central Controller serves as the clearinghouse for routing and reporting data. It does this by receiving route requests, making routing decisions, and monitoring data messages about what is happening in the system.

The Central Controller is installed on one or more servers and comprises three major components: the CallRouter (Router), the Logger, and the Central Database.



Note The CallRouter, the Logger, and the Central Database can be installed on the same computer—or—the CallRouter can be installed on one computer, and the Logger/Central Database can be installed on another computer. The Logger and the Central Database are always co-located on the same computer.

• CallRouter (Router)

The CallRouter receives notification from a routing client (such as a Network Interface Controller or a Peripheral Gateway) that a call is in need of some form of routing. It then runs a user-defined routing script that specifies how the routing client is to handle the call.

These routing scripts are created on the Administration & Data Server, replicated and stored in the Central Database, and loaded into CallRouter program memory.

In addition to receiving routing requests, the CallRouter receives messages from all Peripheral Gateways that monitor real-time status events in the network.

These messages update the system's current representation of agents and system resources. Awareness of the current status of these resources is essential to the routing scripts.

The CallRouter serves as a real-time server by immediately forwarding this data directly to the Administration & Data Server so that it is available to appear in reports. The CallRouter also writes records to the Central Database on the Logger.

• Logger and Central Database

The Logger receives data from the CallRouter (such as detail messages about calls and summary messages that have been computed by the Peripheral Gateways) and serves as the interface between the CallRouter and the SQL Server database manager.

The following data-management processes occur at the Logger:

- Data is written first to temporary tables.
- Data is then written to actual tables on the Central Database.
- Historical records on the Central Database are replicated to the Historical Data Servers on one or more Administration & Data Servers.

The Central Database serves as a buffer where data is committed to quickly support the performance of the CallRouter. The Central Database stores the following data:

- Configuration data, as entered and changed on the Administration & Data Server
- Routing scripts, as entered and changed on the Administration & Data Server
- Summary historical data passed from the CallRouter
- Termination and CallRouter call detail data



Note The Central Controller Database stores no real-time data.

Related Topics

[Call detail data](#)

[Historical Data Server](#), on page 6

[Historical and interval data](#)

[Peripherals and Peripheral Gateways](#), on page 3

Peripherals and Peripheral Gateways

The Central Controller obtains the routing and reporting data that it processes by communicating with each network peripheral.

A peripheral is a device (such as an ACD, a Private Branch Exchange (PBX), or an Interactive Voice Response (IVR)) that receives calls that are routed by Unified CCE software. A peripheral can also distribute calls or nonvoice media contacts.

The Central Controller communicates with each peripheral through a monitoring node called the *peripheral gateway* (PG).

Each peripheral requires a connection to a PG, and Unified CCE software has unique PGs for each device it supports. PGs connect to Voice Response Units (VRUs). Media Routing (MR) PGs send routing requests from multichannel options that are integrated into the system.


An example of specific PGs that connect to ACDs is the Definity Enterprise Communications Server (ECS) PG connects to an Aspect and Avaya ACD, respectively, for enterprise routing and reporting.

If multichannel options or Outbound Option are integrated into the system, the configuration also includes Media Routing Peripheral Gateways (MR PGs) used to send routing requests from the multichannel applications to Unified CCE software. A single MR PG can support multiple applications; you configure a separate Peripheral Interface Manager (PIM) for each application.

It is important to understand the type of peripheral gateway used in your deployment. In real time, the CallRouter receives performance and monitoring information from each PG every few seconds. The CallRouter holds this data in memory and uses it to make routing decisions. New data constantly overwrites this real-time information in the CallRouter memory.

The following illustration shows a Unified CCE, and the peripheral is an ACD.

Figure 2: Peripherals and Peripheral Gateways



Processes on the PG interpret messages on the peripheral and provide data to Unified CCE as follows:

1. By extracting status information from the peripheral through the peripheral's proprietary CTI interface,
2. By normalizing that information and converting it into the format that Unified CCE uses,
3. By forming database objects (Call object, Agent objects, Routing objects, and so forth) from the information, and
4. By passing the normalized data to the CallRouter.

Related Topics

[Different Unified ICM/CCE deployments](#)

Peripherals in Unified CCE Deployments

In Unified CCE deployments that use a Generic PG (which allows multiple peripherals of different types to reside inside a single PG), or separate PGs for Unified CM and the VRU, Unified CM and the VRU appear as separate peripherals to the software. Each time a task switches between Unified CM and VRU peripherals, it appears as a new task to the system.

From a reporting perspective, this separation has an impact on how and when data is collected:

- A call that comes into Unified CM, is then transferred to the VRU, and then back to an agent looks like three separate tasks (calls). A Termination_Call_Detail is written for each call.
- A call that is queued to a skill group or Precision Queue and later answered by an agent is not considered as offered to a skill group or Precision Queue until the call is answered.

In this deployment, data is collected as follows:

- A call that comes into Unified CM, is then transferred to the VRU, and then transferred back to an agent, looks like a single call to the Unified CCE software. A single Termination_Call_Detail is written.
- A call is considered as offered to a skill group or Precision Queue when the call is queued to a skill group or Precision Queue.

Administration & Data Server and Administration Client

Administration & Data Server

The role selected at setup for the Administration & Data Server determines the data that is available for reports. The Administration & Data Server was formerly named the Distributor Admin Workstation (AW).

Depending on the setup selection, the Administration & Data Server can capture some or all the following:

- Real time data
- Historical data
- Configuration data
- Detail Data Store (DDS)

A DDS comprises:

- Call Detail data (Termination Call Detail and Route Call Detail for custom reporting)
- Call Variable data


Unified CCE can support multiple Administration & Data Servers.

Administration Client

The Administration Client (formerly named the Client AW) allows you to access the Configuration Manager tools without installing a full Administration & Data Server with databases. The Administration Client must point to an Administration & Data Server.

The Administration Client is typically installed on a laptop or personal desktop where installation of a full Administration & Data Server is not desirable. There is a separate, small installer for the Administration Client, which provides a configuration utility specific to the Administration Client.

Figure 3: Administration & Data Server



Historical Data Server

A Historical Data Server (HDS) is required if you plan to use historical reports. The Historical Data Server (HDS) must reside on an Administration & Data Server.

The HDS is enabled at setup and created using the Intelligent Contact Management database administration (ICMDBA) tool.



Note As a fault-tolerant strategy, two Administration & Data Servers are typically set up at a site as HDS machines, each with its own HDS database.

Depending on your selection at setup, the HDS can contain only historical data or both historical data and call detail data forwarded from the Logger.

This historical data is not accessed directly, but rather through views that exist in the local Administration & Data Server database. To retrieve information for historical reports, the reporting application connects to the Administration & Data Server where the HDS resides.

Follow these guidelines to ensure that your Historical Data Server is configured to meet reporting needs:

1. **Decide how many Historical Data Servers you require.**

The number of Historical Data Servers that you configure depends on how long the HDS takes to back up and on your reporting demands. If you store large amounts of data, backup can take several hours. If you want to run reports while the HDS is backing up, configure at least one additional HDS to use to run reports.

2. Determine the types of HDS, based on the deployment size and the role of the Administration & Data Server.

3. Determine the size of the HDS.

The size of the database depends on the size of your configuration and on how long you want to retain data.

Configuration that affects the size of the HDS includes the number of call types, skill groups, agents, skills per agent, routing clients, trunk groups, services, peripherals, scripts, calls routed daily, and calls terminated daily.

The larger the configuration, the bigger the HDS must be to store data. For example, the historical call type database tables store data for each call type for each 5-minute and 15- or 30-minute interval.

The amount of time that you want to retain data on the HDS also affects database size. Decide how long you want to retain reporting data before it is automatically purged from the databases. Data beyond the configured retention time is purged each day at 12:30 PM.

You can use the ICMDBA tool to estimate the sizes of your databases. The tool prompts you for your configuration information and the amount of time that data is retained in the databases.

4. Determine how you want to back up the HDS.

You can back up the HDS either while the HDS is running or while it is offline (generally when the contact center is closed or during a time with low call volume).

Performing a backup during peak hours while the HDS is running can impact performance, especially if you are backing up a large amount of data. While the HDS database is being backed up, new data from the Logger is stored in the transaction log. If the transaction log reaches maximum capacity before the HDS has completed the backup, updates to the database stop until an administrator manually empties the log.

It is appropriate to back up at a regularly scheduled time when the contact center is not busy. You can also take the HDS offline and perform a backup. However, the HDS is not available for reporting when it is offline. If you plan to back up the HDS database while it is offline, you might want to configure a secondary HDS to use for reporting during the backup interval.

5. Determine the HDS backup schedule and the number of days for which data is retained on the Logger.

You can configure the number of days for which data is stored in the Logger Central Database and the HDS.

The Logger stores data for less time than the HDS. For example, you might store two weeks of data on the Logger and a year of data on the HDS. Configure the amount of time that data is stored on the Logger in relation to the schedule for HDS backups to ensure that you do not lose data in the event that the HDS goes offline.



Note When the HDS database is newly created, the replication from the Logger to HDS does not copy the first record for historical tables. The impact to historical reporting is negligible, as the records are queried based on the interval datetime.

Related Topics

[Administration & Data Server and Administration Client](#), on page 5

[Call detail data](#)

[Data flow from Logger to Historical Data Server](#)

[Historical and interval data](#)

Cisco Unified Communications Manager

For the Unified CCE environment, Cisco Unified Communications Manager (Unified CM) provides features comparable to those of a traditional PBX system to Voice over IP telephony devices such as Cisco IP phones and VoIP gateways.

Unified CM handles the switching requirements of the Unified Contact Center system and allows deployment of voice applications and the integration of telephony systems with Intranet applications.

Reporting Servers

Depending on the deployment model, there might be from one to eight Unified Intelligence Center servers.

The Unified Intelligence Center web server application on the Unified Intelligence Center server(s) is configured to connect to a Unified CCE Administration & Data Server and to populate reports with the databases on that data source.

To select and generate reports, reporting users log in to the Unified Intelligence Center web application from their browser.

Agent and Supervisor Desktops

Cisco Unified CCE supports Cisco Finesse Agent and Supervisor desktops for contact center agents and supervisors.

Voice Response Units

Voice Response Units (VRUs) are computers that run Interactive Voice Response telephony applications.

Your enterprise might implement one or more types of IVR applications on a VRU platform to serve several purposes:

- **Information gathering.** The IVR prompts for certain information through dual-tone multi-frequency (DTMF) digit or ASR (Automatic Speech Recognition) collection. The information is used in the routing decision and is passed to the agent desktop.

- **Self service.** The IVR prompts for and provides certain information to the caller, such as account balance. The entire call transaction might take place within the VRU.
- **Queuing.** The VRU serves as the queue point by playing announcements or music to the caller until an agent is available.

VRUs can be integrated into Unified CCE software in several ways: at an enterprise (network) level, as a premise-based VRU for an ACD, or as a Virtual VRU on an ACD/PBX. The way in which a VRU is integrated into Unified CCE systems affects the flow of call processing and determines the type of data Unified CCE can collect from the IVR.

For example, a Network VRU provides data used in call routing, monitoring, and reporting. Only Service Control VRUs can be used as Network VRUs. A Service Control VRU is a VRU that implements the Service Control Interface protocol. The Service control protocol allows the VRU to utilize ICM to control call treatment and queuing; for example, by providing the capability of running VRU scripts as commanded by the ICM. A VRU that has an interface only to the ACD has more limited capabilities.

Because VRUs support different features and behave differently, reporting data is affected by the type of VRU you have deployed in your system.

Related Topics

[VRU application reporting](#)

