



Architecture

Unified IC is a clustered, web-based reporting application that can extract and display data from external data sources using Structured Query language (SQL) statements. It supports external [security](#) integration and SNMP-based management tools, which make it very easy to administer and deploy.

This chapter contains the following topics:

- [Platform, page 1](#)
- [Browser clients, page 1](#)
- [Unified Intelligence Center administration, page 2](#)
- [Unified Intelligence Center Reporting node, page 3](#)
- [Cluster support, page 3](#)
- [Data source server, page 7](#)
- [Failover using Application Control Engine load balancing module, page 9](#)
- [Local and wide area networks, page 9](#)

Platform

Unified Intelligence Center is supported as an appliance and is packaged with its complete environment, including its base platform—the Cisco Unified Communications Manager Operating System (Unified CM OS). The Unified CM OS actively discourages user modifications or access to the underlying components except through supported administrative interfaces.

Unified Intelligence Center is also shipped co-resident with Unified Contact Center Express application. This is a seamless integration of the two applications with support for single sign-on across the two applications. Unified Intelligence Center supports only UCCX historical reports in the 9.0(1) release.

Browser clients

The Unified Intelligence Center application is designed to be accessed over Hypertext Transfer Protocol (HTTPS) from web browser clients having JavaScript support. After logging in, HTTPS is used for all other pages.

The Internet Explorer and Mozilla Firefox clients are supported in this release. For more information, see the [Cisco Unified Intelligence Center Bill of Materials](#).

Unified Intelligence Center administration

The Unified Intelligence Center Administration server is the administrative and management component for Unified Intelligence Center and provides operations, administration, maintenance, and provisioning (OAMP) functions. The Administration is the primary interface for configuring Unified Intelligence Center-specific application settings for a cluster, as well as provisioning devices in a cluster. The Administration works in conjunction with the other Unified CM OS web applications to manage a Unified CM OS device.

Administration is a mandatory component and is deployed and accessible only on the primary (controller) node in the cluster. If you install only a single (standalone) node, that node is the controller and has both the Administration and the Reporting applications.

The Administration is a web-based application and requires System Application User or Super User login through HTTPS.

Use the Administration to:

- Configure reporting nodes in the cluster
- Manage licensing
- Download the Real Time Monitoring Tool (RTMT)
- Configure Super User and policies
- Configure security-related settings such as LDAP configuration
- Manage serviceability settings such as trace levels and log file size, which are used for troubleshooting and debugging
- Configure SMTP
- Start or stop the Unified Intelligence Center web reporting application on other nodes in the cluster
- Synchronize (integrate) Unified CCE users



Note

Platform management applications for managing system-level functions such as network options, certificates, upgrades, and SNMP and Alert settings are not handled by the Unified Intelligence Center Administration application.



Note

Access three Cisco Unified Communications Manager (Unified CM) tools (Cisco Unified OS Administration, Cisco Unified Serviceability, and Disaster Recovery System) from any node. For information about these platform management applications, see the [Administration Console User Guide for Cisco Unified Intelligence Center](#).

Unified Intelligence Center Reporting node

The Reporting node contains the core of the Unified Intelligence Center and contains all the features for reporting. Deploy a standalone single Reporting node in a Controller node or in the cluster of a Controller node and up to seven Reporting Member nodes.

The following software components are deployed in the reporting node:

- Firewall
- Web Server (Tomcat 6.0) that runs the Unified Intelligence Center application
- JAVA services and JSP pages that translate the web requests into HTML
- Unified Intelligence Center Database (Informix) with replication support within the cluster
- The Administration (OAMP) application (when the reporting node is also the only node of the cluster).

Cluster support

When the number of concurrent logins and reporting needs exceed the capacity of a single (controller) server, deploy Unified Intelligence Center in a clustered mode. In this mode, multiple member (reporting) nodes present a unified application to the user and allow a consistent, shared configuration across the controller and all member nodes.

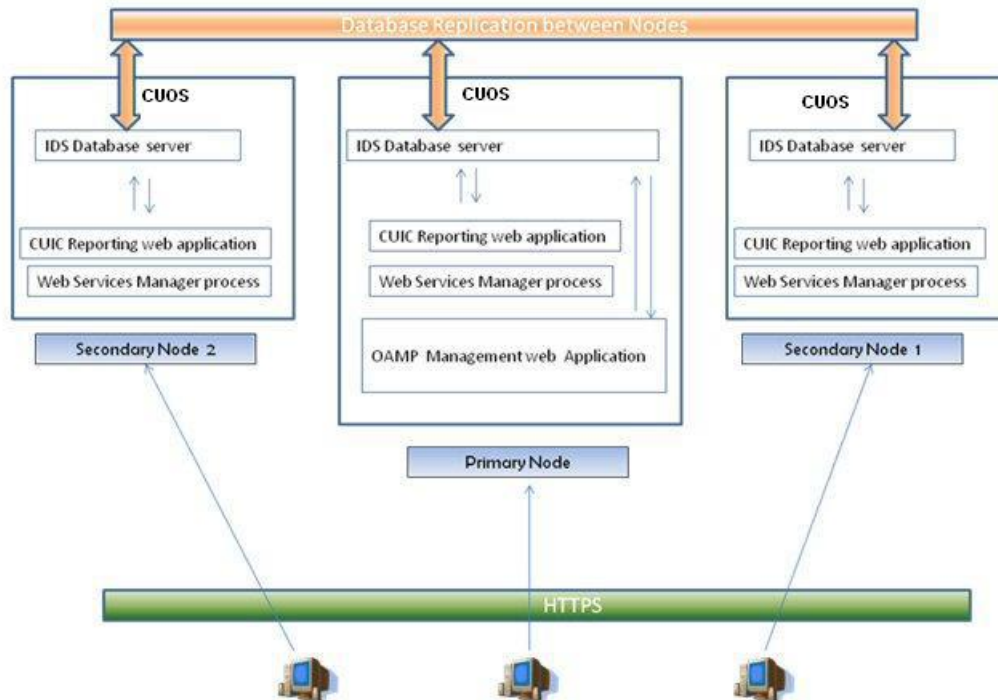
Multiple nodes bring better scalability to the reporting infrastructure than would be possible using a single, vertically-scaled reporting server. Cluster support also presents an opportunity to gradually scale the solution, depending on the needs of the deployment.

Clustering in Unified Intelligence Center is made possible by the IBM Informix Unified Intelligence Center database, which serves as the backend data store for the Unified Intelligence Center reporting application.

Enterprise Replication database connects and synchronizes the multiple nodes participating in a cluster. Administrative changes, configuration updates, and so forth are communicated throughout the cluster and all users are presented with a common interface object, irrespective of the node used to serve their requests and actions.

In addition, nodes communicate with each other directly by using IP multicast, which happens independently of database replication.

Figure 1: Database Replication



WAN overview

The Unified Intelligence Center cluster is a group of independent nodes, each of which has a replicated database that is kept in sync with the other nodes in the cluster. With Unified Intelligence Center Release 8.0.2 (and later), you can distribute the Unified Intelligence Center cluster over the WAN. When participating in a cluster, either over a LAN or WAN, the configuration objects created on one node automatically replicate to other nodes. This chapter describes the sizing considerations made when considering clustering over the WAN.

In general, size the WAN so that it supports the customer-expected performance characteristics. Consider the following network characteristics:

- **Bandwidth:** is a measure of the amount of data that can be sent per second.
- **Duplex:** is a measure of how data can be transmitted. Full duplex indicates traffic can be transmitted at full bandwidth speed in both directions. Unified Intelligence Center requires a full duplex WAN connection.
- **Latency:** is a measure of the time it takes for a data packet to propagate from one side of the WAN to the other. For Unified Intelligence Center, this must be 100 ms or less, which leads to a Round Trip Time (RTT) of 200 ms or less.
- **Jitter:** is a measure of how much variance there is with respect to latency. Cisco expects that network jitter is minimized and that the latency accounts for all the jitter.

- **Loss:** (sometimes referred to packet loss) is a measure of how much data can be lost or dropped. This should be minimized and it is expected that loss is at or near 0%. Anything else leads to retransmissions and effective loss of bandwidth (as the same data is being retransmitted).
- **Duplication:** is a measure of how much data can be duplicated. This should be minimized and it is expected that duplication is at or near 0%. Anything else leads to retransmissions and effective loss of bandwidth (as the same data is being retransmitted).
- **Corruption:** is a measure of how much data can be corrupted. This should be minimized and it is expected that corruption is at or near 0%. Anything else leads to retransmissions and effective loss of bandwidth (as the same data is being retransmitted).

As previously stated, each Unified Intelligence Center node is independent of every other node and all database updates (that is, configuration data) are replicated as to all other nodes. When the Unified Intelligence Center cluster is running over the WAN, it should be noted that this replication utilizes the WAN connection and therefore is a function of the network characteristics. As an object is created on one node, it will be instantly available users on that node, but may take a few seconds before the object can replicate to other nodes (see the following for time estimates based on object size). The only objects that are replicated are configuration objects.

Configuration objects include the following:

- **Data Sources:** encapsulates configuration data for each data source object
- **Dashboard:** encapsulates configuration data for dashboard object
- **Report Definition Filter:** encapsulates configuration data for each report definition filter (Note: Running a report with an initial or new filter causes replication of the filter object, but refreshing a report has no impact on replication.)
- **Report Definitions:** encapsulates report definition configuration information
- **Reports:** encapsulates a report object, which points to a report definition and corresponding views
- **Report Types:** encapsulates configuration information for the type of report (grid, chart, gauge)
- **Views:** encapsulates column names, groupings, hidden fields
- **Users:** encapsulates a user
- **Categories:** encapsulates configuration information for categorization of reports, report definitions, and dashboards
- **Value Lists:** encapsulates configuration information for the value list
- **Collections:** encapsulates configuration information for a collection that includes its values
- **DataSets (from Scheduled Reports):** encapsulates the data set for a scheduled report. All dataset information (report data not configuration data) is replicated to all subscriber nodes

In general, Cisco expects that most users configure the Unified Intelligence Center system with a set of commonly accessed reports and run these reports periodically. In addition, Cisco expects that users create new objects (for example, reports), but this should occur less frequently. Any configuration object that is modified or created needs to be replicated and will be replicated as quickly as possible, but since we are expecting users to be modifying or creating objects less frequently than executing or accessing an object, this should minimize the need for large amounts of WAN bandwidth.

Because every node requires a connection to every other node, replication is estimated by assuming that each node receives a portion or channel of the WAN bandwidth, which is a function of the number of nodes in Site

A and the number of nodes in Site B. Or more specifically, the WAN bandwidth requirements differ depending on how the cluster is setup.

For more information, see [Chapter 4: Bandwidth and Performance Recommendations](#).

Organizing Sites

Because the maximum size of a cluster supported by Unified Intelligence Center is eight nodes, it is impossible to have a fully redundant clustering solution unless each site is at most four nodes. As qualified, each Unified Intelligence Center supports up to 200 users given the Reporting Standard Profile defined in Chapter 4.

In addition, only the Unified Intelligence Center Primary node can provide the following services:

- OAMP
- Scheduler

The Primary node must be located at the customer's primary site (as these services will be unavailable during a failover situation). Furthermore, it is crucial that you back up the Primary node periodically in either a LAN or WAN environment.

Failures

Each Unified Intelligence Center node buffers replication data to send to other nodes in the cluster. When communication is lost with other nodes in the cluster (or a node fails) then the data is queued until contact with the other node or nodes is restored. Each node continues to work independently even during connectivity failure--it just does not have access to objects created or modified on other nodes. Even though the queue is quite large (1600 MB), it is not unlimited, therefore, fills up the queue during prolonged failure. As the buffer fills and starts to reach capacity, an alarm is sent (CiscoAlarm30) notifying administrators of the potential buffer exhaustion condition. If connectivity is restored before the buffer is filled, then it synchronizes at a rate proportional to the amount of data in the buffer and the connection bandwidth.

If connectivity is not restored before the buffer fills, then replication is reset. Resetting replication allows the node to continue running reports and working independently. If the node is a secondary node then it requires full synchronization with the primary node (primary database backup and restore on secondary node) when connectivity with the primary node is restored. If replication is reset, then anything created/modified on the secondary node is rolled-back to the state of the primary database. If the primary node fails, you have to re-install and revert to a saved backup. Make sure that you back up the primary node periodically so that no data is lost. It is very important in WAN environments to periodically back up the Primary node. Prolonged WAN outage, failure of primary node, and failure to perform recommended backups could result in a cluster that must be reinstalled or else there could be data lost.

If connectivity is restored before the buffer fills, then all data is automatically replicated. Depending upon the connection between the nodes and the amount of data that has accumulated, it may take some time for the nodes to fully synchronize. At any time, administrative users can view the status replication (including the number of bytes in the replication queue) using the CLI command: `utils dbreplication runtimestate`.

For more information, see the *CLI documentation*.

Configuring WAN

After you install and configure all systems, restart each system. This allows clustering to synchronize and to merge on network partitions. This only needs to be done the first time after configuring a new node.

Database replication in the cluster

The Unified CM OS installs Informix as part of the base platform installation. Installation-specific scripts create the Unified Intelligence Center application database on each reporting node and set up the enterprise replication among participating nodes.

The Unified Intelligence Center database is the main data store for the Unified Intelligence Center reporting web application. It holds configuration information relating to users, reports, and user access rights for each node in the cluster.

Unified Intelligence Center uses the update-anywhere enterprise replication model based on a topology of Fully Connected Database servers. Each database server has connection to all other database servers. Data can be transferred from any machine to all others. Data is replicated immediately, as soon as changes are made. Changes that are apparent in the User Interface (such as the addition or deletion of a dashboard or a change to a report definition) are evident across all devices when the page is refreshed, provided that all devices in the cluster have all the required network connections and CPU power to do the processing.

An automated daily purge that runs every day at midnight takes care of database maintenance activities. The purge schedule can be controlled or changed through the command line interface. Purge is the only local database maintenance apart from the backup that is required for local Unified Intelligence Center databases.

Data source server

Data source servers are used by the Unified Intelligence Center reporting component to extract data requests from the queries that are defined in the report definitions. Each data source server contains the database and the schema that store the reporting data that populate the reports. A Unified Intelligence Center data source server must be a valid JDBC compliant database. Unified Intelligence Center is currently validated with SQL server 2005 and Informix 11.5 database versions.

**Note**

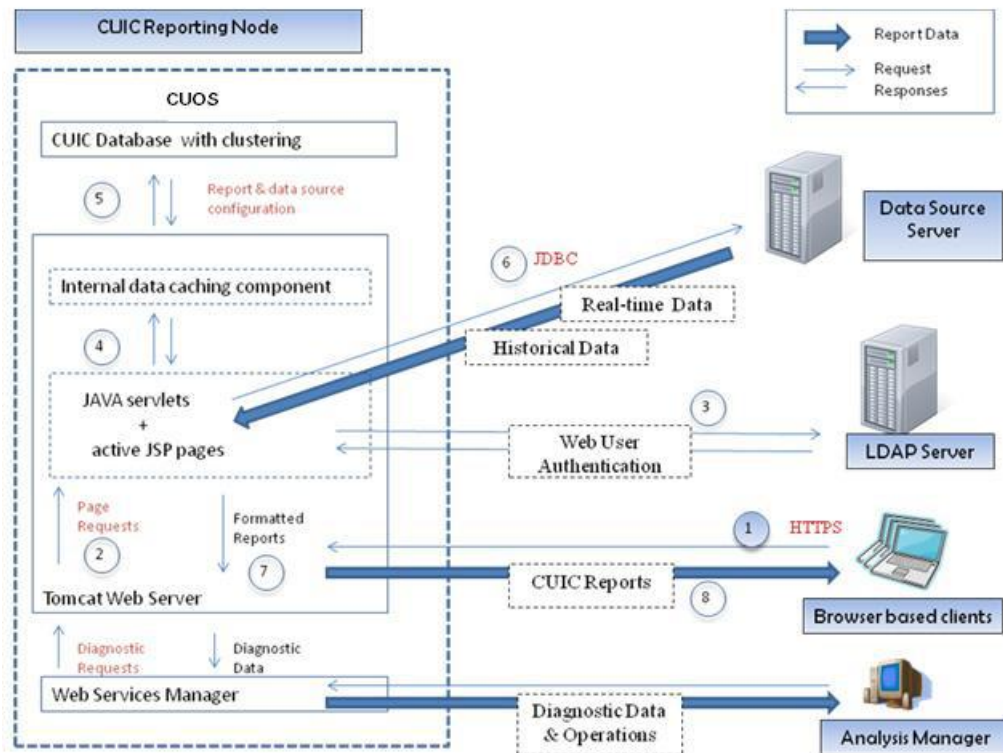
Set SQL Server Authentication to *SQL Server and Windows Authentication mode*. This release does not support Windows Authentication to a SQL Server database.

These data source servers are supported for the Unified CCE reports, importable Unified CVP, Finesse, Social Miner, EIM, and WIM reports, and UCCX historical reports in co-resident mode.

- Unified CCE AW-HDS, which is part of the [Cisco Unified Contact Center Enterprise deployment](#).

- Unified CVP reporting server, used in [Cisco Unified Customer Voice Portal](#) deployments.

Figure 2: Unified Intelligence Center Architecture

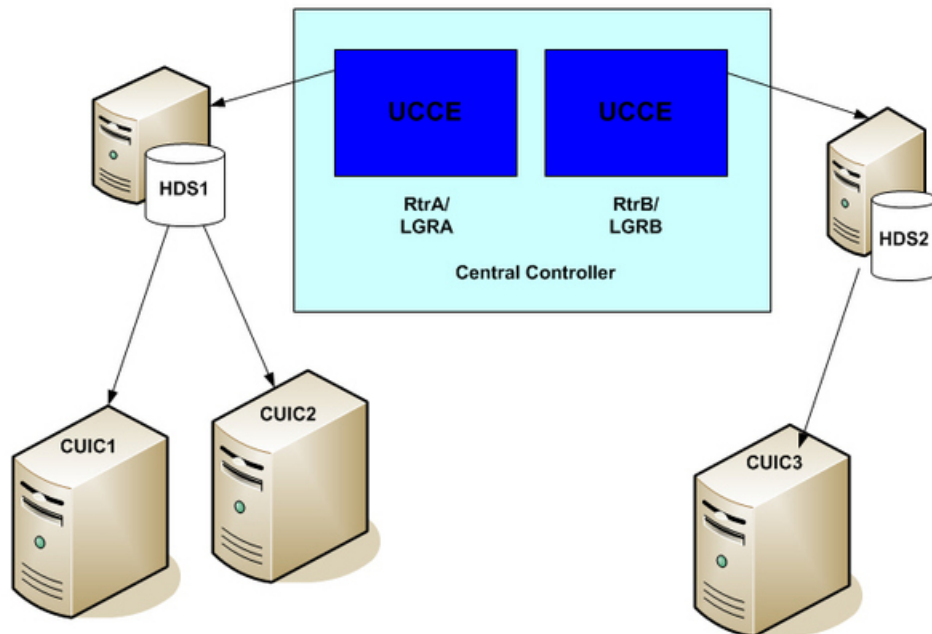


Nodes mapped to multiple databases

The number of data sources that you must deploy for Unified Intelligence Center depends on the number of concurrent reports to be supported and the number of concurrent logins. You can distribute the reporting load to several Unified CCE AW_HDS databases using the command line interface and conventional name resolution. If there is a need to direct a specific member node to a database host other than the one in configured on the data sources interface, you can use the `set cuic-properties host-to-ip` command to resolve the data source name differently on each node.

You can direct any given reporting member node to a specific Unified CCE database.

Figure 3: Nodes Mapped to Multiple Databases



Failover using Application Control Engine load balancing module

The Cisco Application Control Engine (ACE) is a multifunction module for Cisco Catalyst 6500 Series Switches and Cisco 7600 Series Routers that is used for maximizing availability, acceleration, and Server Load balancing. It has a native understanding of multiple protocols such as HTTPS, which is used for content-based traffic redirection, depending on various traffic predictors, source address / round robin / server load, and so forth.

Unified Intelligence Center uses ACE to present a single URL to access the application and to perform server load balancing in the cluster, by redirecting the HTTPS requests across the multiple nodes as required. When deployed along with ACE, the failure of a single node does not affect the application, and new requests are automatically routed to the remaining servers.

Coupled with database replication, which replicates the application state across the nodes, ACE presents an elegant failover solution for the overall application.

Local and wide area networks

Data Sources have been qualified over local area networks (LAN) and wide area networks (WAN).

**Note**

Distributing data source servers over a WAN can increase the report execution times, depending on the latencies and traffic conditions over the WAN.

Related Topics

[Firewall integration in a wide area network](#)