

Directory Database Service

This chapter provides an overview of use, functionality, and architecture for the Directory Service optional feature for the Controller.

Overview

The Directory Service is based on the industry standard, Lightweight Directory Access Protocol (LDAP), which is an infrastructure that supports the messaging requirements of certain Controller applications. The Directory Service acts as a link to a *directory* or database to provide quick mailbox information retrieval services. The directory clients Unified Message Portal (UMP), Mobile Message Portal (MMP), Admin Web (AW), and Voice Profile Internet Messaging (VPIM), use the Directory Service.

The Directory Service has a hybrid database architecture; a centralized database architecture combined with a distributed database architecture. This hybrid architecture accommodates large numbers of subscribers and fast mailbox information retrieval.

The Directory Service network is divided into one or more regions based on the number of subscribers per geographical location and the geographical location of the Controllers.

Functionality

The Directory Service function is based on a client-server model. The two main parts of the Directory Service are Directory Clients and Directory Servers.

The Directory Client software is part of the Data Access Manager (DAM) package. A Directory Client serves client(s)' requests by establishing connections to a pre-configured Directory Server where the requested subscriber mailbox information is stored and managed.

A Directory Server is a software module residing in a separate application processor card or a separate Controller that executes the requests made by the Directory Client(s). The Directory Server can also be housed in an external platform.

Controller directory client applications make a request to the Directory Service for mailbox information that is stored in a Directory Server database. The Directory Service in turn spawns a Directory Client that makes a request to the Directory Server. The Directory Server either completes the request or passes back a search reference to the Directory Client which contains the name of the Directory Server that can complete the request. The Directory Client then forwards its original request to the Directory Server named in the search reference.

Two types of requests are sent by the Directory Client:

- Query requests are used to retrieve subscriber's information, such as the subscriber's E-mail address, phone number, user id, domain name, and so forth.
- Provisioning requests are used to add, modify, or delete the subscriber's mailbox and personal information.

Search for an Entry in the Centralized Database

Figure 1-1 shows how a mailbox entry in the New York's Directory Server is searched from San Jose.




Figure 1-1 Searching for a Mailbox Entry

- A Directory Client in San Jose requests information on mailbox 212xxxxxxx from San Jose's Directory Server.
- San Jose's Directory Server answers with a search reference to New York's Directory Server.
- The Directory Client in San Jose then contacts New York's Directory Server.
- New York's Directory Server provides the requested data to the Directory Client in San Jose.

The updating of Directory Server database entries is integrated with the updating of the VoiceRecord database. Any related requests to add, delete, and modify mailbox information in the VoiceRecord database are forwarded to Directory Servers.

Search for an Entry in Distributed Databases

A combination of the mailbox numbers' area codes and prefix are used as branch points (BP) in the Directory Information Tree (DIT) under which entries are stored. In addition to the local area code (local BP), each Directory Server has every other Directory Server in the network's prefix in a prefix branch point. However, the San Jose Directory does not have the New York mailbox entries under the 212 prefix branch point. Instead, the BP contains the IP address of the New York Directory Server. If the San Jose Directory Client requests mailbox information about a 212 mailbox, the San Jose Directory Server passes a search reference to it, where the search reference contains the IP address and the Branch Point on the New York Directory Server. Then the San Jose Directory Server will reroute its request to the New York Directory Server.

Figure 1-2 shows the diagram of referral that is described below. In this example, two Directory Servers are located in San Jose and in New York. The local mailbox entries of each Directory Server starts with the prefix. In the example below, the NPA or area code is used for the prefix branch point.

- San Jose – Area code 408
- New York – Area code 212

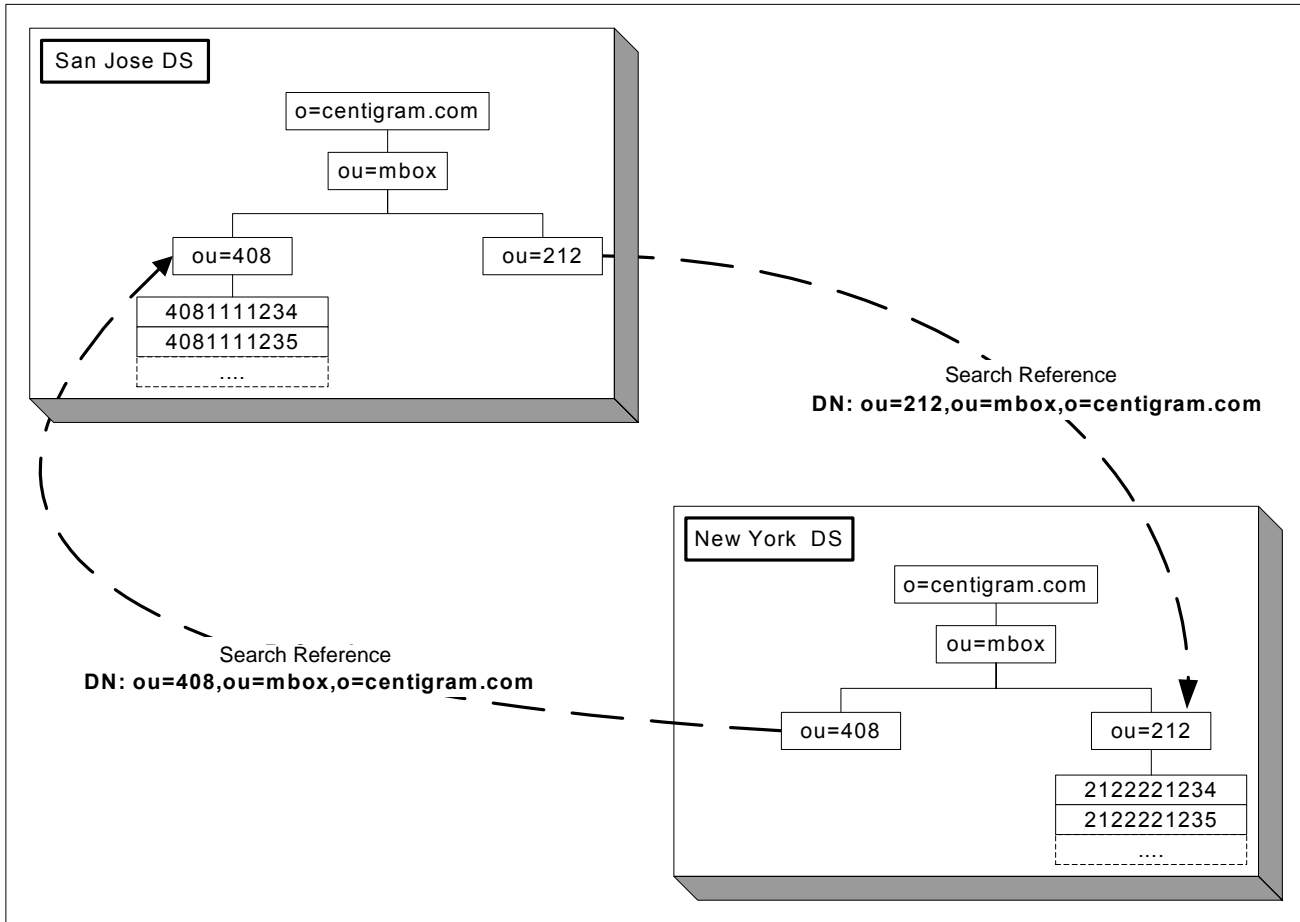


Figure 1-2 Prefix Branch Points

Directory Service Components

The basic components of the Directory Service are:

- Directory Server

- Directory Client
- Directory

Directory Server

The Directory Server manages the relationships between all Directory Servers in the network and also manages its subscriber database. It serves the query and provisioning requests that are generated by Directory Clients.

The Directory Server has its own set of menus as well as a command-line interface. The menus are used for installation, configuration, and administration, and the command-line interface can be used for batch mode processing.

An important part of a Directory Server is its Replication Server; this is a daemon running on a Directory Server that copies or *replicates* database entry modifications from the Master Directory Server (or a Slave Directory Server) to Slave Directory Servers. This replication task is called *synchronization*. The Master Directory Server records the replication activity in the *replication log*.

To implement the distributed network architecture, you will create different *host types* among the Directory Servers by configuring a *replication agreement*. The replication agreement between two Directory Servers allows one Directory Server to synchronize its database entries with another. A replication agreement must be configured for each Directory Server.

The three different Directory Server host types are

- Global Master Directory Server
- Regional Master Directory Server
- Slave Directory Server

Global Master Directory Server

Master Directory Servers are the Directory Servers that accept provisioning requests directly from a Directory Client and replicate modifications to other Directory Servers in the network. One Global Master Directory Server exists for the entire network. One Regional Master Directory Server per region.

Within its region, the Global Master functions as a specialized Regional Master and can have many Slave Directory Servers. Outside its region, the Global Master updates and replicates the global (non-mailbox) database. Since the global database is infrequently updated after the initial setup of the network, this function of the Global Master does not affect its performance as a Regional Master.

The Global Master must have the IP addresses or host names of all Regional Masters in order to replicate the global database to all Directory Servers in the network.

Regional Master Directory Server

A Regional Master Directory Server is the only Directory Server in a region that accepts provisioning requests directly from a Directory Client. A Regional Master performs mailbox provisioning requests in its region and then replicates the modifications to its Slave Directory Servers. A Slave Directory Server may also be configured to replicate modifications to other Slaves.

Each Regional Master must be registered with the Global Master in order to replicate the global (non-mailbox) database successfully. All Slave Directory Servers must be registered with their Regional Master in order to successfully replicate the mailbox database.

Slave Directory Servers

A Slave Directory Server is a Directory Server that does not accept provisioning requests directly from a Directory Client. When a Slave receives a provisioning request, it sends an update referral to the Directory Client. The update referral contains the IP address of the Regional Master Directory Server. The provisioning request for a mailbox entry is handled by the Regional Master Directory Server. A request for global data provisioning is forwarded to the Global Master via the Regional Master.

A Slave Directory Server can be configured to replicate database modifications to other Slave Directory Servers. Though the Slave has the same ability to replicate as a Master does, it does not accept provisioning requests directly from a Directory Client but only from the Master Directory Server. A Slave Directory Server can only accept query requests directly from a Directory Client. The number of Slave Directory Servers is unlimited in a network of Directory Servers, but at least one Slave per region is strongly recommended.

Directory Client

The Directory Client software is the interface between Controller applications and the Directory Service. It serves the information requests from various directory client applications that need information held in a Directory Server database. It establishes connections to a Directory Server and makes a query or provisioning request. At least five Directory Clients are usually connected to a Directory Server at any given time. The Directory Client uses a Distinguished Name (DN) to specify which entry it wants to look up or update.

The Directory Client must be configured to make requests to the Directory Server that is closest geographically, and has the most reliable connection: the *preferred server*. The Directory Client is also configured to send requests to a second Directory Server if it receives no response from the preferred server: the *secondary server*.

A Directory Server Slave can satisfy query requests, and can, therefore, be used as a Directory Client's preferred server.

The Directory Service software, which includes the Directory Client software, is installed on the Controller *after* the Directory Server software has been installed, configured and verified.

Directory

A directory is a specialized database, which stores and manages subscriber's mailbox information. This type of database is used mostly for information queries.

The Directory Server database contains information that is organized in a DIT. The blocks of information are called entries. Each entry has a globally unique DN that is made up of the entry's attributes linked into a single string and used to refer to the entry. A DN presented to one Directory Server references the same entry at another Directory Server. Typically, entries include attributes such as the subscriber's name, E-mail, and namespeech.

The hybrid architecture of the Directory Service has two types of databases. One database is centralized, and the other is distributed. These databases contain different entries; the centralized database contains non-mailbox information that doesn't change often, and the distributed database contains mailbox information that changes much more often.

Centralized Database

The centralized database in the Directory Service architecture contains non-mailbox information, such as product, users, and cluster. This non-mailbox information is kept centralized to reduce processing time of query requests. Updates to this database are replicated throughout the network.

Note: The term database is used in the rest of this manual to describe the directory.

Distributed Databases

The distributed portion of the Directory Service architecture is a mailbox database. This database contains mailbox numbers of subscribers that belong in a defined region of a network; therefore, each distributed database in the network is unique.

Only mailbox database information is distributed to accommodate the increasing number of mailboxes in a region.

As shown in the figures earlier in this chapter, San Jose has a Directory Server and New York also has a Directory Server. Each Directory Server has its own, unique database of mailboxes.

The Directory Client is referred by the Directory Server not containing the requested data to a Directory Server that does. Search Reference and Referral works only within SS8 Networks' Directory Service network because the information on the DIT cannot be referred to a third-party Directory Server.

A portion of a mailbox number is used in the prefix branch point of the mailbox DIT. A prefix is a selected group of numbers that precede the last four digits of a mailbox number. The prefix branch point may include not just the NPA (area code), but also part of the NXX (known as the prefix in the U.S.) of a mailbox number. Distribution and referral are handled by using the prefix of mailbox entries. This has two advantages:

- The number of entries is unlimited if more Directory Servers are added to the network of Directory Servers.
- Mailbox entries can be searched without having to replicate all entries to all Directory Servers.

Directory Schema

The schema defines the structure of the database. It defines the server's name, the ports it operates on, and the root domain of the object tree that the server will reference. The schema defines the type of information that can be stored in the database and gives each type of information a name. The schema contains the attributes of each entry and their name.

SS8 Networks' schema is automatically installed when the Directory Server software is installed. No additional steps to install the directory schema file are required after installing the Directory Server software.

DIT

A Directory Information Tree or DIT is the structure of mailbox information in a database. In a Directory Server, there are two DITs, one for regional mailbox information, and one global or non-mailbox information. The mailbox DIT contains one prefix branch point. The prefix branch point contains all the prefix entries in the network and one more level of entries that contain the regional Directory Server's mailbox entries. The prefix branch point entries for other regions contain references (IP addresses) of all the Regional Master Directory Servers.

Network Architecture Guidelines

Below is a list of guidelines for your network architecture.

Design your network of Controllers to accommodate different groups of subscribers. The network should be divided into regions based on the concentration and geographical locations of subscribers. Large concentrations of subscribers in one geographical area can be split into more than one Directory Service region.

The network design can include several Controllers, some with a Directory Server installed, and some without. The Controllers without a Directory Server installed can access a Controller that contains a Directory Server. The Directory Service software must be installed in all the Controllers that use it. Directory Service does not support the SSL and SASL features of LDAP. In designing the Directory Service network architecture, you must determine:

- The regional divisions of the Directory Service
- The Directory Server from each region that will be the Regional Master
- The Directory Server that will be the Global Master of the entire network
- The preferred Directory Server of each Directory Client
- The Directory Servers in each region that will be Slaves

The designation of the Regional Master and the Global Master depends on the configuration. As a general rule, the most reliable network/location should be the master, since it must be accessible for update requests from the other slaves.

Redundancy

A single Directory Server is the master Directory Server for an entire network. The Master Directory Server can replicate its data to any number of Slave Directory Servers. All updates to data occur on the Master Directory Server.

For redundancy, allocate at least two Directory Servers per region, one as a Regional Master and the remaining as Slaves. When two Directory Servers are used (one as the Master and the other as the Slave), the two databases are identical, with changes to the Master's database replicated to the Slave. However, if the Master Directory Server fails, only query requests are successful on the Slave Directory Server. Requests for entry modifications to the Slave Directory Server, such as add, delete, or modify, fail.

You can configure both Applications Servers and the UMP/VPIM Gateways with a primary and a standby Directory Server to user for queries. For large networks, a

single Master Directory Server is configured, along with many Slave Servers. The Master Server is used only for updates, and all queries are performed against the Slave Servers.

Directory Service Regions

Note: If there is only one region, then the Global Master and the Regional Master will still be there, except that the Global Master will also be a Regional Master but without any Slaves.

Each region can serve 30,000 mailbox entries with name speech of 16KB. The size of each database entry may vary based on the input of attributes. Since up to 16 regions can be defined in the system, the maximum number of subscribers is 4.8 million. Plan ahead for future expansion of the entries and consider the geographical and network proximity of the Directory Servers and the Directory Clients of each region.

A reliable and fast wide-area network is necessary to handle cross-regional searches with referral. Network failure in one region results in entries inaccessible from other regions. The impact is minimal on normal operations of other regions except that cross-regional messages, such as directory client, cannot be delivered during network failure. For this reason, SS8 Networks recommends having redundant network paths among regions.

Figure 1-3 shows the relationships between the servers of two regions. The following lists describe the information held by each component and its origin:

Directory Servers	6
Directory service network	2 regions: A and B
Directory Servers per region	3
Region A	1 server is the Global Master (A0) 2 servers are Slaves (A1 and A2)
Region B	1 server is the Regional Master (B0) 2 servers are Slaves (B1 and B2)

Global Master A0	Regional Master B0 replicates the global database Slave A1 replicates both the global and mailbox databases
Regional Master B0	Global Master A0 receives replication update of the global database from the Global Master Slave B1 propagates both the mailbox and global databases
Slave A1	Global Master A0 provides the clients requesting an update Slave A2 replicates the database

Slave B1	Regional Master B0 provides the clients requesting an update Slave B2 replicates the database
Slaves A2 and B2	A1 and B1 provide a referral to the clients requesting an update

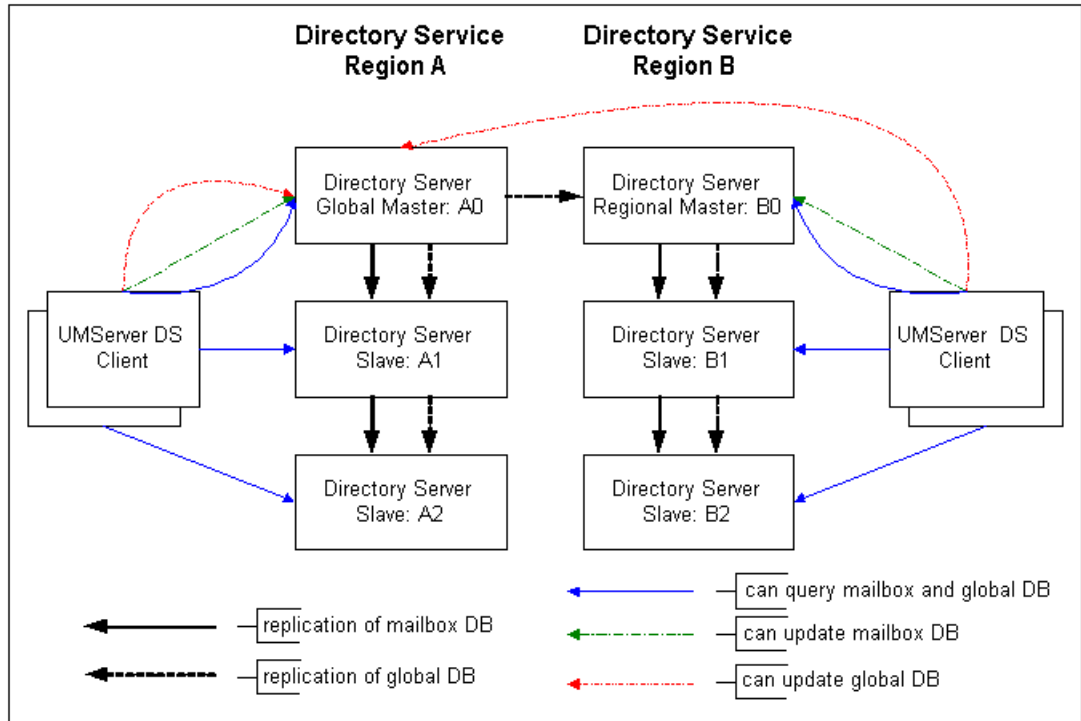


Figure 1-3 Replication Process

The Directory Server(s) has only one Global Master and one Regional Master per region. Only the Global Master provides service to update the global database from directory-enabled clients. The Regional Master provides service to update the mailbox database from directory-enabled clients. With the replication agreement setup, only the Global Master can update the Regional Masters. The Slave provides service only for query requests.

Applications that Use the Directory Service

Several optional features require the Directory Service. [Table 1-1](#) shows what applications require the Directory Service and those that perform better with the Directory Service. One Directory Server can serve multiple optional feature applications: only one pair of Directory Servers is required per Directory Service region for an architecture that runs directory client(s).

Table 1-1 Controller Applications Using the Directory Service

Directory Service Required	Directory Service Recommended
Unified Message Portal (UMP)	EasyAdmin Web (EAW)
Voice Profile for Internet Mail (VPIM)	Mobile Message Portal (MMP)

Note: At this time, MMP is not supported by the Directory Service, and EAW is supported via AC30SP/AC30LP connection, not by direct query from EAW to a Directory Server.

Feature Bit

When FCOS bit 320 is added to the FCOS group set of the mailboxes that use name speech, an attribute for name speech can be created in the Directory Server's database and used by directory client applications.