

AE White Paper: Alternative Distributed Directory Models

APS Engineering, Inc.

<http://www.aeinc.com>

Overview

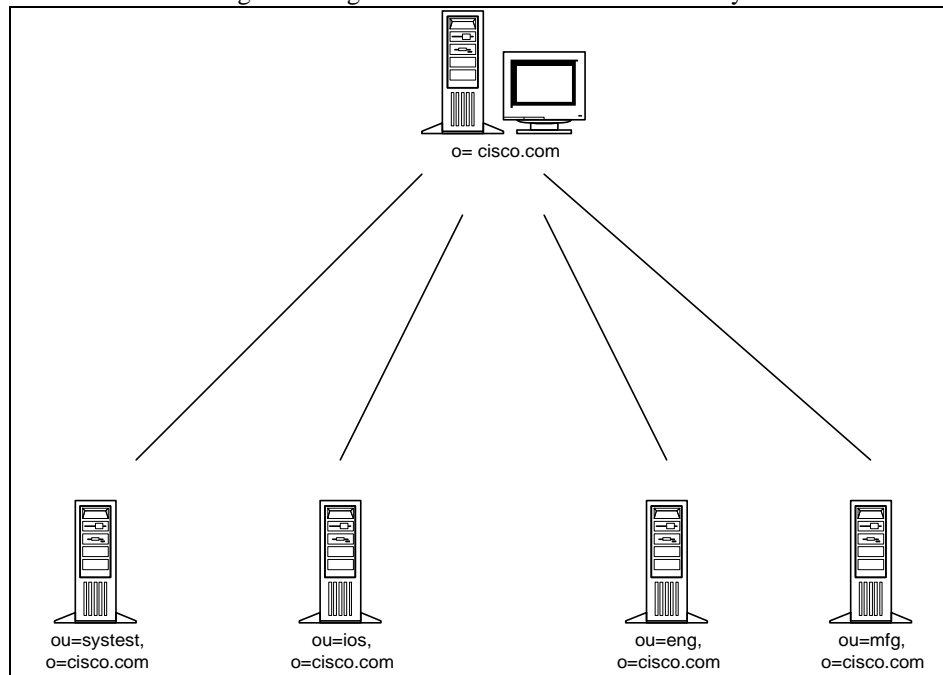
Large enterprise directory implementations often require distributed directories involving multiple servers. This paper identifies mechanisms to accomplish this using AE SLAPD (slapd.com). The purpose for selecting a distributed topology instead of a single monolithic database (i.e., everything on one server) is twofold:

1. Improve performance
2. Improve fault tolerance - reduce points of failure

In both cases, single points of administration should be represented in the architecture. Either all administration is done centrally, or administrative responsibilities are local to the directories. The rule is that administration should *never* be duplicated or left to manual synchronization.

Given a large corporate network, there may be several divisions or dozens (even hundreds) of departments, commonly represented as *organizational units*. In each case, the physical database distribution commonly reflects resources along those organizational boundaries (see figure 1).

figure 1: Organizational Unit view of a Directory



Distributed Tree Model with Referrals

A common method of distributing the workload for a busy server is to implement **referrals**. There are two elements of referrals: *parent* and *node* referrals. If figure 1 represented a referral based distribution, each

computer would have its own AE SLAPD server controlling a unique database. The top level in the tree (o=cisco.com) would be the parent, hence it may have few local entries, but would contain special embedded referral entries in the database to point to each node. A sample LDAP interchange format (LDIF) file excerpt is shown in example 1 below.

example 1: node referral database entry

```
dn: ref="ldap://host1.cisco.com/ou=mfg,o=cisco.com",o=cisco.com
objectclass: referral
ref: ldap://host1.cisco.com/ou=mfg,o=cisco.com
```

Similar entries would have to exist for each organization unit. The main advantage of this structure is that **queries are initiated via the central server**. Each of the remote servers act as the workhorses to actually *resolve* queries by searching their individual databases.

The above database link is *one-way*. It assumes that only the main server is accessed by a client to query the directory. If, for example, each department had access to their servers, each would have to be able to resolve searches starting with the top-level of the tree. These backward reference would be a parent referral. Unlike node references, these are specified once in the slapd.conf configuration file. If the main host name for o=cisco.com (following convention) was *ldap.cisco.com*, the slapd.conf for each of the departmental hosts would contain a line as in example 2.

example 2: parent referral in slapd.conf

```
referral ldap://ldap.cisco.com
```

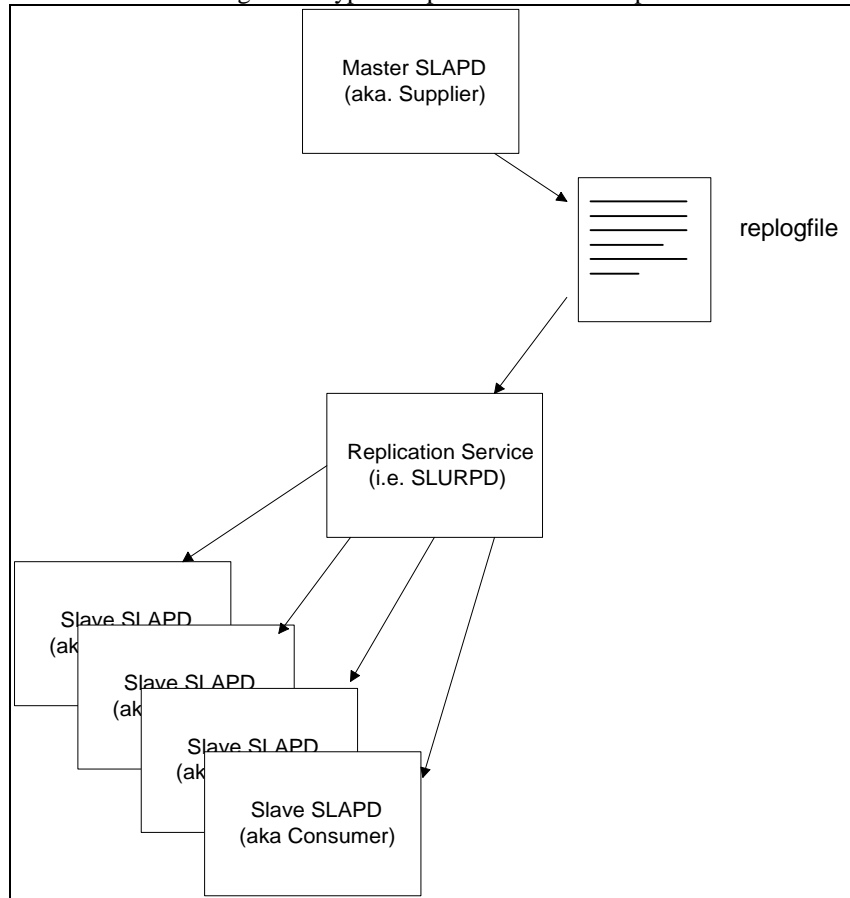
Two-way referrals (parent and node referrals) allow for N-way searches to occur. This is optimal because both the network traffic and the database search workload is distributed.

Replicated Tree Model

Typically, replication is used to off-load a centrally administered server so multiple servers can improve the overall performance or availability of the directory. In this mode, the top level server is master (a.k.a. supplier), and contains a **repllogfile** directive to identify the file used for logging, and a **replica** directive instructing the replication server (i.e., slurpd) which AE SLAPD slaves (a.k.a. consumers) will receive the updates as the replication log is written.

A conventional setup for replication is shown in figure 2.

figure 2 - typical replication mode setup



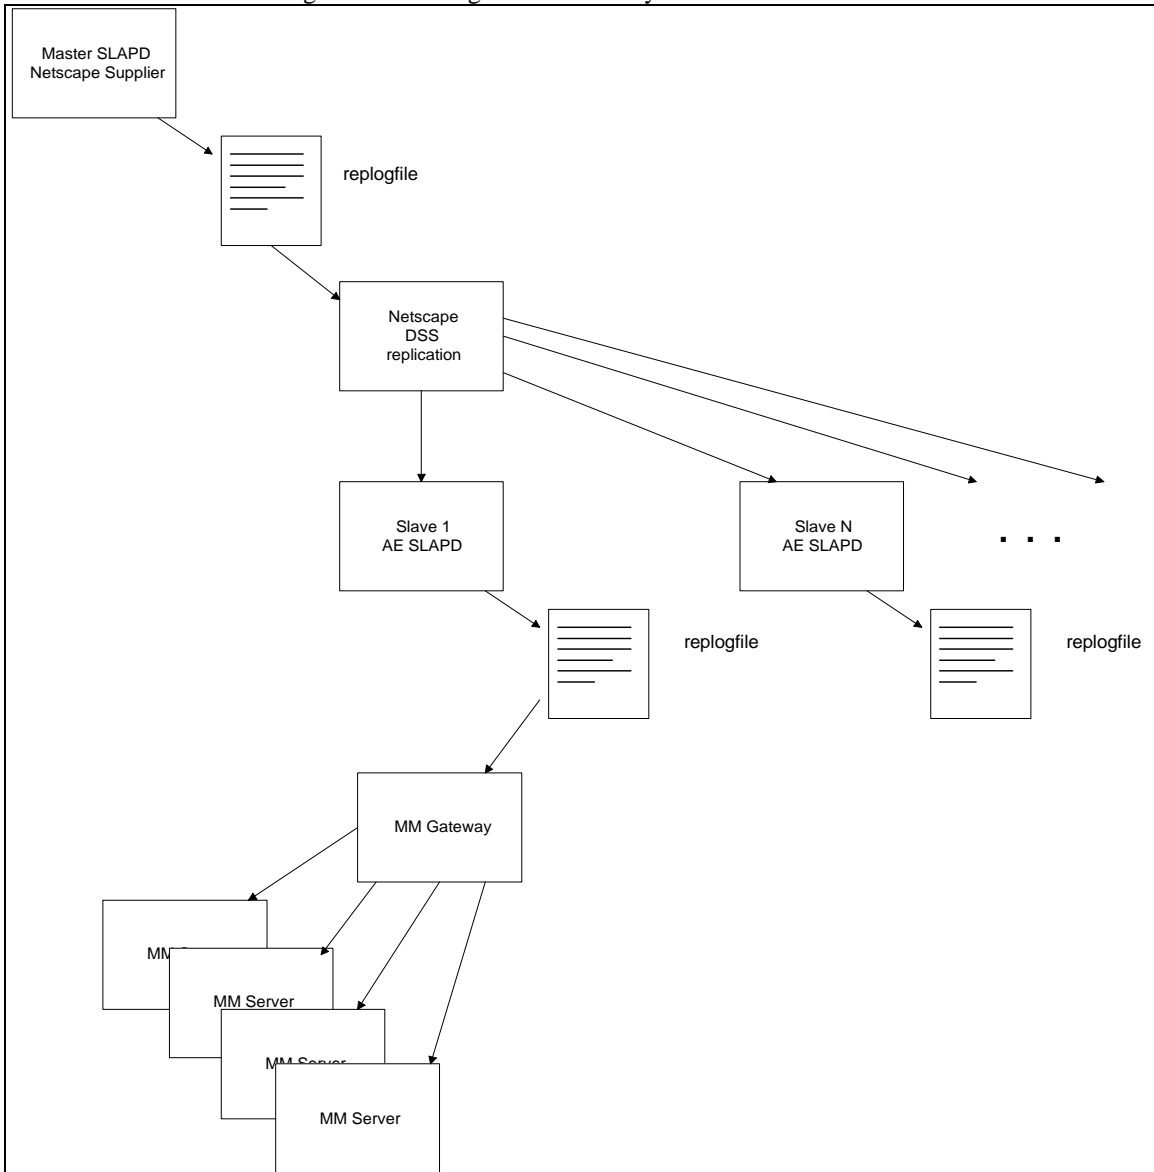
Single master single database replication is an atypical case. It is better to configure the master server for **subtree replication** so that each slave copies just part of the database.

Meeting Maker Configuration

The configuration prescribed by Meeting Maker , Inc. for gateway involves use of replication (via the Netscape Database Synchronization Service) from the Netscape Directory as Supplier (master), off-loaded by a number of remote AE SLAPD Consumer (slaves). The Meeting Maker Gateway acts as a second level replication service, to propagate slave replication logs to one or more Meeting Maker Servers (see figure 3).

This architecture requires that each slave act as replica master to the MM Gateway (repllogfile defined in slapd.conf). This is known as **chaining**. For *performance* and *latency* considerations, chaining is usually not recommended. In the event of failure by one of the slave servers, re-synchronization of the databases and repllogfile purging may be complicated. (NOTE: MM Gateway truncates the slave replication log files, so if multiple AE SLAPD / MM Gateways are used in this configuration, it may be necessary to shutdown and re-sync all databases if one server fails.

figure 3 - Meeting Maker Gateway LDAP Architecture



Novell Critique of Netscape Architecture

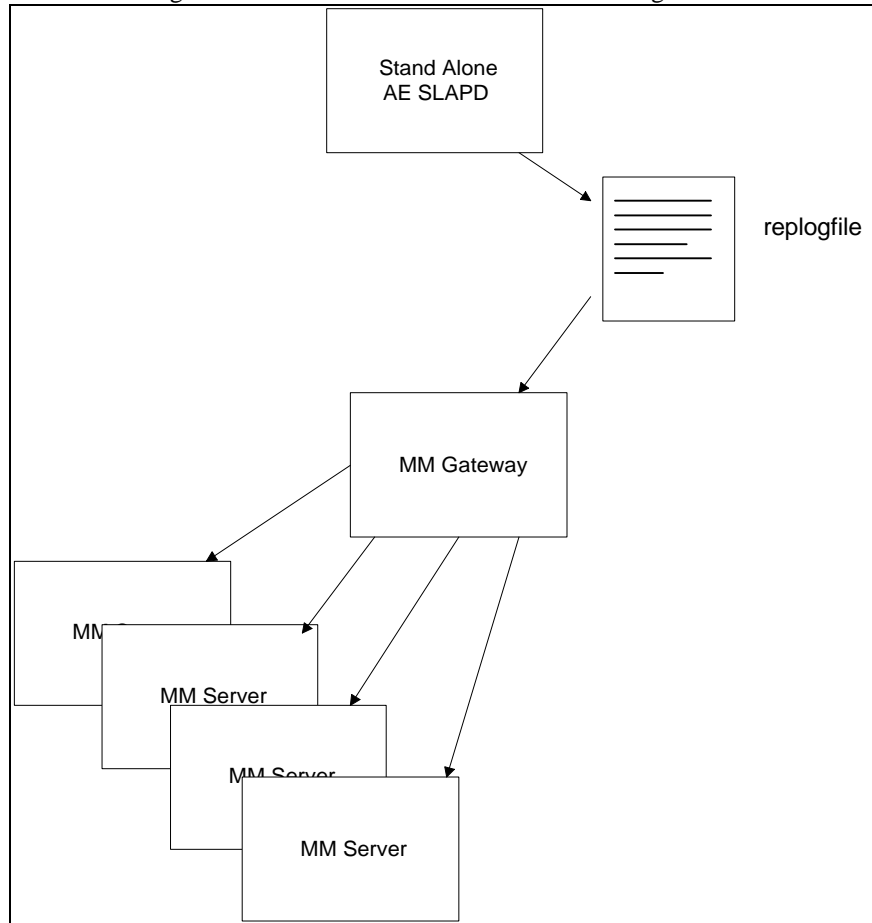
(ref <http://www.novell.com/products/nds/netscape2.html>)

“Netscape: Although promoted as a distributed enterprise directory, Netscape Directory Server has serious scalability limitations. Netscape’s directory uses a master-slave model, which limits scalability because it forces all data to come from a single server, creating a single point of failure. Since all changes happen at the master server before they can be replicated to the others, your administration must be completely centralized--that is, unless you don’t mind crossing a WAN link to administer or access your network. The master-slave model is not a very robust and dynamic method of changing information for the Internet or any other distributed “net.” Analysts agree that Netscape Directory Server is best suited for highly centralized environments or departmental applications. “

AE Recommended Configuration

The AE prescribed configuration would actually use both replication log (for the MM Gateway) and referrals. Instead of a single database, the organization tree would be distributed (per the previous section) and each individual (departmental) server would contain the additional relogfile definition to cause the stand-alone AE SLAPD server to dump its replication log as each change is made to the database. This is a more practical directory layout in organizations that require multiple servers.

figure 4 - Stand-alone AE SLAPD server configuration



One more implication of this strategy is that the MM Gateway would most likely **deploy 1:1 with the Meeting Maker Servers**. Although a single gateway could facilitate an entire organization, it is driven by the administrative topology of the enterprise. In an organization that may store thousands of entries per server, the LDAP servers, and Meeting Maker servers may be distributed along organizational boundaries. Where it makes sense *administratively* to have distinct MM servers, it probably will make sense to have an LDAP server/MM gateway.

Finally, if LDAP/MM Gateways are co-resident with the Meeting Maker server host, the complication of added equipment acquisition can be greatly reduced.