Internet-Draft
draft-ietf-ldapext-ldap-taxonomy-05
Expires in six months
Track: Informational

Ryan Moats
Lemur Networks
Roland Hedberg
Catalogix

July 2001

# A Taxonomy of Methods for LDAP Clients Finding Servers

**Filename: draft-ietf-ldapext-ldap-taxonomy-05.txt**

Status of this Memo



Abstract

There are several different methods for a LDAP client to find a LDAP server. This draft discusses these methods and provides pointers for interested parties to learn more about implementing a particular method.

## 1. Introduction

The Lightweight Directory Access Protocol (LDAP) [1] can be used to build "islands" of servers that are not a priori tied into a single Directory Information Tree (DIT.) Here, it is necessary to determine how a client can discover LDAP servers. This documents discusses the currently available methods and provides pointers for interested parties to learn more about implementing a particular method.

This draft documents only those methods that are currently being pursued in the IETF. Other methods have been considered for this

problem and the history of these other methods are presented in the Appendix.

## 2. Methods

### 2.1 Client Configuration

The simplest method of enabling a LDAP client to discover LDAP servers is for the client administrator to configure the client with a list of known LDAP servers (and associated base objects) to send queries to. While this method has the advantage of being correct (initially), it adds the requirement that the list of initial servers be kept small and constant. Otherwise, the required client update process won't scale.

### 2.2 Well known DNS aliases

If the DIT uses a naming scheme similar to that in RFC 2377 [2], then it is possible to build the DNS names of potential servers using well known DNS aliases, like those documented in RFC 2219 [3]. When a different naming scheme is used, it is also possible to build potential server names based on the client's fully qualified domain name or local (within the organization or country) environment.

One shortcoming of this method are that it is not exact. Multiple DNS lookups and LDAP protocol operations may be necessary to find the proper LDAP server to serve the client requests. To support client roaming, it is necessary that either the RFC 2377 (or similar) naming scheme be used or that roaming be implemented through tunnels.

Because this method uses DNS, it inherits all the security considerations of using DNS to discover LDAP servers: see the security consideration in [3] for more details.

### 2.3 Service Location Protocol

If a client supports the service location protocol [4], it could use a SLP query for LDAP servers. The SLP template that is used to describe LDAP servers is presented in [5], and requires that the servers announce themselves using SLP and this template.

Using this method inherits the scaling and security considerations for the service location protocol, which are documented further in [4].

### 2.4 Referrals

In LDAPv3, servers can return referrals to the client if the server has knowledge of where a query might be satisfiable. Two ways of deploying referral information are deploying a LDAP knowledge server or exchanging CIP index objects [6] between servers.

A LDAP knowledge server would hold cross references to possibly hundreds of other LDAP servers, so that a client would only need to know about its local LDAP server and the knowledge server. As an optimization, the local LDAP server could also act as a knowledge server.

If CIP index objects are exchanged between LDAP servers, then those objects can also carry URL information for providing referrals to clients. Here, the client would only need to know about the local server. Using CIP index objects inherits the security considerations of CIP: see [6, 7, 8] for more details.

In either of these cases, the local LDAP server could be determined using another of the methods discussed.

### 2.5 Using SRV records

RFC 2052 [12] defined SRV records for DNS, which bound a host name and port to a label in the DNS. This makes it possible for a client to look up information about a supported protocol for a domain and get back a weighted list of fully qualified domain names and ports for where that protocol is supported. For more information, see [13].

## 3. Implementation

The Norwegian Directory Forum plans to start a service based on a central LDAP service containing contact information for every organization within Norway [10]. If an organization has more information about its sub-units, employees or functions that it wants to publish it can do so by placing this information in a publicly available LDAP server and providing the management of the central service with a pointer (URL) to this server.

The TISDAG project is running a test service based on the TISDAG specification [11]. This service gathers indices from connected White Pages Service Providers using CIP Tagged Index Objects [9]. The rationale for this service is that by supplying the name of a person or a function/role to the service it will return pointers to where more information can be found about persons/functions with that name.

The European cofunded project DESIRE (www.desire.org) is designing a system to use a LDAP server that communicates with a referral index that in turn, uses CIP Tagged Index Objects [9] and is fed by LDAP crawlers. DANTE plans to set up a European infrastructure of such referral index servers.

## 4. References

Request For Comments (RFC) and Internet Draft documents are available from numerous mirror sites.

[1] M. Wahl, T. Howes, S. Kille, Lightweight Directory Access Protocol (v3), RFC 2251, December 1997.

[2] A. Grimstad, R. Huber, S. Sataluri, M. Wahl, Naming Plan for Internet Directory-Enabled Applications, RFC 2377, September 1998.

[3] M. Hamilton, R. Wright, "Use of DNS Aliases for Network Services," RFC 2219 (Also BCP 17), October 1997.

[4] E. Guttman, C. Perkins, J. Veizades, M. Day, "Service Location Protocol, Version 2," RFC 2608, June 1999.

[5] J. Wood, R. Tam, "The LDAP Service Type," SVRLOC Template http://www.isi.edu/in-notes/iana/assignments/svrloc-templates/naming-directory_ldap.1.0.en

[6] J. Allen, M. Mealling, "The Architecture of the Common Indexing Protocol (CIP)," RFC 2651, August 1999.

[7] J. Allen, M. Mealling, "MIME Object Definitions for the Common Indexing Protocol (CIP)," RFC 2652, August 1999.

[8] J. Allen, P. Leach, R. Hedberg, "CIP Transport Protocols," RFC 2653, August 1999.

[9] R. Hedberg, B. Greenblatt, R. Moats, M. Wahl, "A Tagged Index Object for use in the Common Indexing Protocol," RFC 2654, August 1999.

[10] R. Hedberg, H. Alverstrand, "Technical Specification, The Norwegian Directory of Directories (NDD)," http://www.catalogix.se/doc/draft-hedberg-alvestrad-ndd-01.txt

[11] R. Hedberg, L. Daigle, "Technical Infrastructure for Swedish Directory Access Gateways (TISDAG)," Internet Draft (work in progress), February 2000.

[12] A. Gulbrandsen, P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)," RFC 2052, October 1996.

[13] M. Armijo, L. Esibov, P. Leach, "Discovering LDAP Services with DNS," Internet Draft (work in progress), July 1999.

## 5. Author's Addresses

Ryan Moats
Lemur Networks
15621 Drexel Circle
Omaha, NE 68135
USA
Email: rmoats@lemurnetworks.net

Roland Hedberg
Catalogix
Dalsveien 53
0775 Oslo
Norway
Email: roland@catalogix.se

## Appendix A. Historical Methods

### A.1 Discovery

The discovery approach was to use a combination of other methods presented in this taxonomy along with storing either the search DN or a related URL in the DNS in some way. Using both TXT or NAPTR records in the DNS were considered. This approach requires an administrator to configure the DNS with necessary information. Further, the idea of storing standards based information (either a DN or an URL) in a DNS RR has been an extremely controversial one in the IETF.

### A.2 DHCP extensions

Another proposed method was to use DHCP to deliver information about LDAP server to a DHCP client. This would require that such information be configured into the DHCP server and that the client use DHCP to load host configuration information. While there has been some nascent interest in this method, there has been no interest in implementation of this approach.