

2.0 DIRECTORY SERVICE ARCHITECTURE

2.1 Directory System Agent Topology

2.1.1 Overview

The X.500 Directory is a set of complex, distributed processes for managing an on-line directory and providing directory services. The first edition of the Directory specifications was published in 1988 as the International Telegraph and Telephone Consultative Committee (CCITT) X.500 Series of Recommendations and in 1990 as the International Standards Organization (ISO)/International Electrotechnical Commission (IEC) 9594. The second edition of the specifications was published in 1993 as the ITU-T X.500 Series of Recommendations and in 1995 as ISO/IEC 9594. The two groups are now collaborating on a third edition of the Directory specifications which will include system management of the Directory and enhanced security.

The primary purpose of the X.500 directory service is to provide an enterprise-wide translation of user or network resource names to electronic-mail addresses, thereby facilitating transfer of information in various mediums between end-users. Given the size and complexity of the Federal Government, it is envisioned that each agency or department will have its own directory.

X.500's object-oriented design, scalability, and distributed nature make it well suited for its intended use as a global directory. The architecture of X.500 defines a single "tree" for the entire world, referred to as the global DIT, with various countries holding different parts of the tree and subordinate entities holding parts of each national tree. Naturally, each part of the tree must be able to share information with other parts, thereby necessitating a synchronization of directories belonging to different agencies and departments.

2.1.2 Principle Components

The X.500 directory model consists of three principle components:

- Directory Information Base (DIB) -- The DIB is the database component containing the collective information held in the directory. The DIB is composed of directory entries, consisting of a collection of information about one object, e.g., person, computer, organization. For each object there is precisely one directory entry. Each directory entry in the DIB consists of a set of attributes that provide information about the entry to which they correspond, e.g., name, address, telephone number, e-mail address. The DIB is logically structured and represented as a hierarchical information tree known as a DIT. This logical hierarchy provides the naming context by which distinguished names are constructed to uniquely identify directory entries.
- Directory System Agent (DSA) -- The DSA is the directory component that stores and maintains the DIB. The DSA is an application process that provides access to information about the directory entries contained in the DIB. Multiple DSAs may be associated with a DIB. Therefore, if multiple DSAs are employed then each DSA will maintain a portion of the DIB. However, these DSAs work in concert to share information in a fashion that will enable the user to view and browse the collection of DSAs as a single directory. The DSA also provides capabilities for redirecting requests for information it does not contain through the use of knowledge information. Chaining and referral functions enable DSAs to process and route directory information requests to the appropriate distributed DSA containing the requested data.
- Directory User Agent (DUA) -- Directories are typically accessed through a DUA (essentially a user interface). The DUA is the component that communicates with DSAs and provides the means to access information contained in the DIB. The structure of the underlying directory information is hidden from the user. The DUA enables the directory user to perform various functions such as browsing directory listings, performing keyword searches, and other functions such as viewing, adding, modifying, and deleting directory entries. A DUA may be used by an individual or an application process. The DUA interface may also take on many different forms, e.g., Graphical User Interface, touch screen, voice activated, and/or system-generated calls. All services are provided by the directory in response to requests from DUAs. Security mechanisms may be

employed to limit the ability of directory users to view, modify, or delete information based on access control information stored in the directory.

2.1.3 Protocols

The DUAs and DSAs require protocols to communicate with each other. There are four primary protocols associated with X.500.

- **Directory Access Protocol (DAP)** -- DUAs interface and communicate with DSAs using DAP which define the exchange of requests and outcomes between a DUA and DSA to identify and retrieve information stored in the DSA. The directory can be accessed through a series of service ports, as defined by the X.500 series of recommendations. Each port provides a series of services defined as a unique access point. Services define the directory operations, e.g., bind to the directory, read, search, and modify directory entries. In addition, the DAP returns result and error codes.
- **Lightweight Directory Access Protocol (LDAP)** -- LDAP is a protocol used for X.500 directory access over TCP/IP networks. Similar in function to a DUA, it provides access to the X.500 directory for the end-user, or client, in a way that is greatly simplified. In its current form, it does not support the full functionality of DAP, but is useful in applications where less data is needed, and where anonymous access, which may be used for public citizen access, is needed.
- **Directory System Protocol (DSP)** -- DSP is used between DSAs to service user queries that require information that might be distributed over multiple DSAs. The DSP consists of service ports supporting chained operations. Chaining enables a directory information request to be forwarded (through multiple DSAs) to the appropriate DSA containing the requested information. Once the appropriate DSA is identified, the specified operation can then be performed on the target DSA through a DAP operation with attached chaining arguments and chaining results. Chaining is transparent to the user and is performed by progressively forwarding a query through a number of DSAs, each collecting and evaluating results until the data is retrieved and passed back through the "chained" DSAs.

- **Directory Operational Binding Protocol (DOP)** -- DOP is used between any two DSAs that are entering into an association agreement, either to shadow information or to keep knowledge reference pointers up to date. The DOP enables DSAs to negotiate the nature of the binding agreement and define the parameters that will be used to govern their association.
- **Directory Information Shadowing Protocol (DISP)** -- DISP is used between shadowing DSAs to transfer information from one DSA to another, as well as to transmit shadow updates. DSAs must enter into a binding agreement using the DOP before the DISP can be used to replicate and update information.

The 1988 X.500 Series of Recommendations incorporated DAP and DSP while the 1993 extensions expanded the protocol suite to include DISP and DOP to handle replication. Exhibit 2-1 depicts the above mentioned protocols and their interaction.




Exhibit 2-1
X.500 Protocol Suite

2.1.3.1 LDAP

Recently, LDAP has become more popular as an X.500 directory access protocol, although it has been supported in X.500 directory products for several years. The recent popularity of LDAP is due to its recent adoption by Netscape and by many other vendors who have announced support for LDAP in their products. LDAP provides a simplified mechanism for TCP/IP-based clients to query and manages a database of hierarchical attribute/value pairs over a TCP/IP connection. Current versions of LDAP support mainly directory entry Read operations with very limited modify capabilities.

In the next version of LDAP, version 3, additional functionality will be added to the protocol, such as more complete entry modify, referral and replication functions. LDAP is clearly rooted in X.500 in that it uses a hierarchical structure for organizing the data, and uses very similar attribute names to X.500.

LDAP clients may access X.500 DSAs through a collocated LDAP process running on a DSA server, or it may directly access LDAP servers that may contain information. In the near term, LDAP access will be limited to anonymous access to directory data, due to its lack of support of access controls, referrals, and management interfaces. When a client connects to a DSA running an LDAP process, the LDAP client essentially has complete access to directory data stored in other DSAs. This complete access is accomplished via the X.500 directory DSP protocol which allows DSAs to obtain information from other DSAs and pass it on to the original client that requested it. The scenario of LDAP access in the Government X.500 Directory is shown in Exhibit 2-2.

The Government X.500 Directory will support LDAP in its current and future releases. In its current state, LDAP would best be utilized in smaller, intranet-scale directories initially, because many design and development issues remain to be seen by the X.500 marketplace.

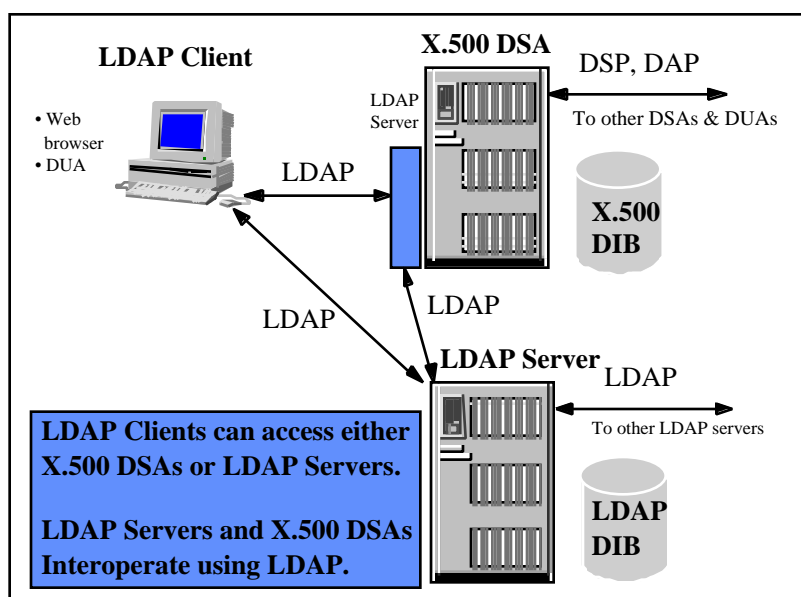


Exhibit 2-2
LDAP Access

2.1.4 Directory Information

The Government DIB is the repository of information to which the Government Electronic Directory provides access. The DIB may have its information distributed across a large number of DSAs; each holding a fragment. Additionally, each DSA may hold copies of other portions of the DIB. The distribution of the DIB will be transparent to the end-user, giving the effect that the entire DIB is held within a single DSA, as depicted in Exhibit 2-3. In support of this appearance, it is necessary for each DSA to be able to gain access to the information associated with a name (either distinguished or aliased) that is held in the DIB. If the DSA does not contain an entry associated with a distinguished name (or shadowed copy), then the DSA should be able to interact with other DSAs to identify the appropriate DSA containing the data for retrieval.




Exhibit 2-3
Accessing the DIB

2.1.5 Directory Information Tree

The DIB is logically structured as a hierarchical information tree, called the Directory Information Tree (DIT). Exhibit 2-4 presents the top level DIT for the Government Electronic Directory. Each entry in the DIB corresponds to a vertex of the DIT. The DIB refers to the physical realization of the X.500 directory, whereas the DIT defines the logical representation of the name space.

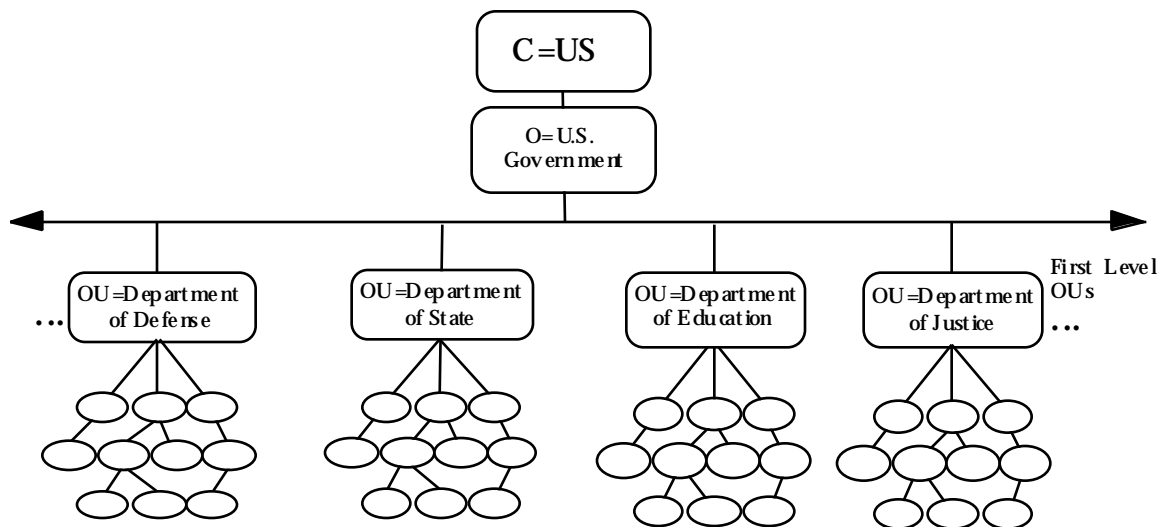


Exhibit 2-4

Top Level DIT for the Government Electronic Directory

Each directory entry in the DIT is uniquely identifiable through the ordered sequence of attributes called relative distinguished names (RDNs) that form the distinguished name (DN). The sequence of RDNs from the root of the DIT to the object being named forms the DN. Exhibit 2-5 provides an example of how DNs are formed.

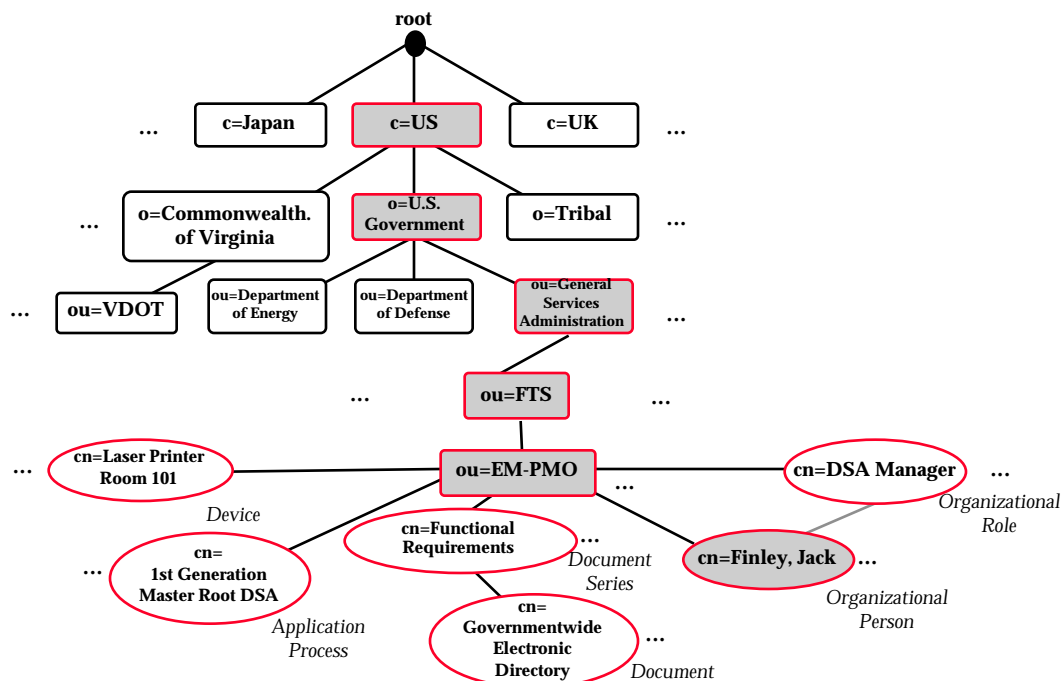


Exhibit 2-5

Distinguished Name Resolution

There is no uniqueness requirements associated with individual RDNs that comprise a DN. However, the sequence of RDNs when combined to form the DN must be unique. Section 3 provides guidance in developing naming constructs. The unique DN for the directory entry of Jack Finley shown above consists of the following sequence of RDNs:

```
C=US; O=U.S. Government; OU=General Services Administration;
OU=FTS; OU=EM-PMO; CN=Finley, Jack
```

2.1.6 Naming Contexts

A naming context is a subtree of the DIT, all entries of which have a common administrative authority and are held in the same master DSA. A naming context starts at the vertex of the DIT (other than the root) and extends downward to leaf and/or non-leaf vertices. Support vertices constitute the border of the naming context. The DIT is therefore partitioned into disjoint naming contexts, each under the administrative authority of a single master DSA. Exhibit 2-6 depicts an example of naming contexts distributed over multiple DSAs within the Government Electronic Directory.

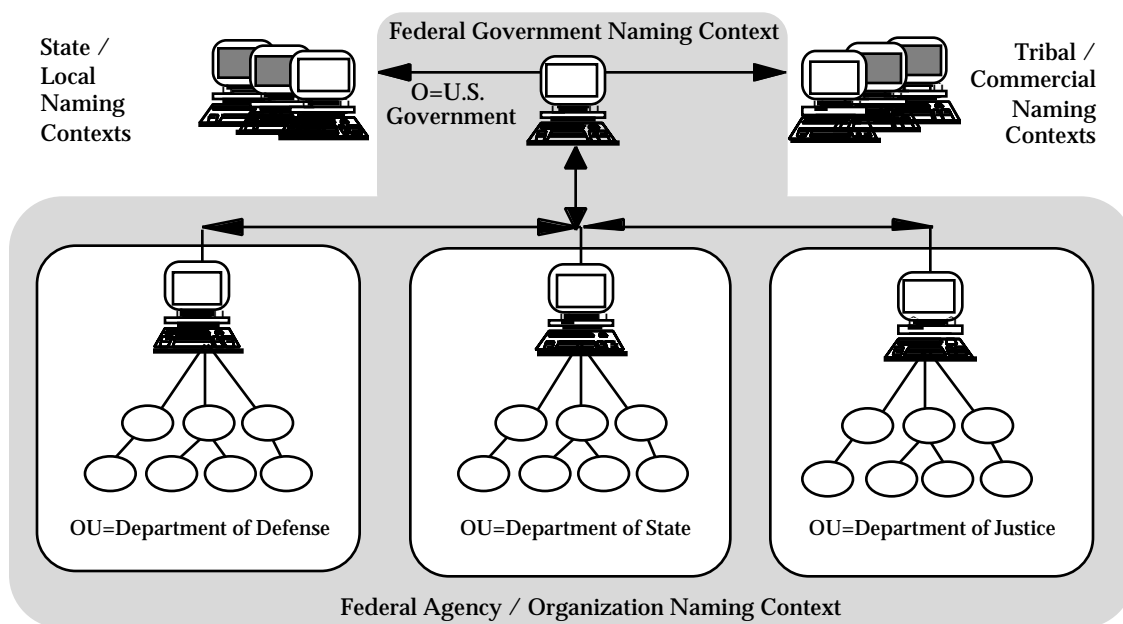


Exhibit 2-6
Example Naming Contexts

2.1.7 Replication Between DSAs

An X.500 directory allows information contained in a DSA to be replicated among numerous DSAs. This distributed architecture improves information availability and performance. An X.500 directory uses a replication model which describes two types of replication: *caching* and *shadowing*. Caching is the simplest form of replication in that a static copy of the information is held by the DSA. However, there is no intrinsic update mechanism to ensure the cached information is kept current and most importantly, access controls mechanisms may not apply to cached data. Shadowing is a form of replication which has intrinsic update mechanisms that are discussed below.

2.1.7.1 Master and Shadow DSAs

A DSA will hold a portion of the DIT to which changes can be made directly and for which it will be designated as the master DSA. Other DSAs may hold copies of this same subtree information and are known as shadow DSAs.

2.1.7.2 Primary and Secondary Shadowing

Primary shadowing is an arrangement where the shadow consumers directly obtain their information from the master DSA. Secondary shadowing is an extension of primary shadowing in that an initial consumer of information from the master DSA acts as a shadow supplier to other consumers. The main advantage of secondary shadowing is that the master DSA does not have to supply shadow copies to all the consumers. The Government Electronic Directory will implement primary and secondary shadowing whose concepts are presented in Exhibit 2-7.

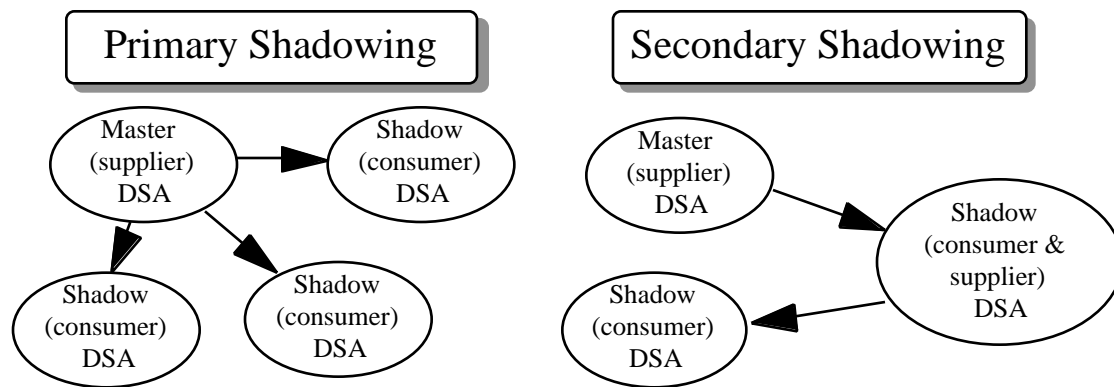


Exhibit 2-7

Primary and Secondary Shadowing

2.1.7.3 Shadowing Agreements

A shadowing agreement is a formal written specification of what types of information or knowledge will be replicated in the DIT and how frequently updates will occur. The agreement specifies the units of replication and the area, attributes, and knowledge that are going to be replicated. The agreement also designates the master DSA for each particular item to be replicated. The utilization of formal specifications allows review and analysis of the redundancy of the DIT structure, knowledge information, and data entries to prevent single points of failure. Since shadowing agreements exist between multiple DSAs, modification to shadowing agreements will require coordination with other DSA administrators.

2.1.8 Government Directory Knowledge

It is a requirement of the Government Electronic Directory that the distributed architecture of the directory be rendered transparent to its users thereby giving the effect that the entire DIB appears within each DSA. In order to support this requirement, it is necessary that each DSA be able to gain access to the information held in the DIB associated with any name (i.e., any object's distinguished or alias name). If the DSA does not itself hold an object entry or object entry copy matching the name, then it must be able to retrieve it from the DIB through direct and/or indirect interactions with other DSAs.

Knowledge information is the DSA operational information held by a DSA that represents a partial description of the distribution of DIB entries and copies of entries held in other DSAs. Knowledge is used by DSAs to determine an appropriate DSA to contact when requests from DUAs or DSAs cannot be satisfied with locally held information. Knowledge consists of knowledge references associating, either directly or indirectly, the name of a directory entry with a DSA holding the entry or a copy of it.

2.1.9 Knowledge Reference Types

The distributed meta-information providing the knowledge reference relationships for the Government Electronic Directory will be stored and maintained in each of the DSAs comprising the distributed directory. A general discussion of knowledge reference types follows:

- **Superior Reference** - A knowledge reference containing information about the DSA considered capable or resolving (i.e. finding) any entry within the whole of the DIT. This reference type consists of the access point of a DSA.
- **Immediate Superior Reference** - A knowledge reference containing information about a DSA that holds a specific superior entry or entry copy. This reference type will consist of the context prefix of a naming context that is immediately superior to one held by the DSA holding the reference and the access point of the DSA holding that naming context.
- **Subordinate Reference** - A knowledge reference containing information about the DSA that holds a specific subordinate entry or entry copy. This reference type consists of the context prefix of a naming context that is immediately subordinate to one held by the DSA holding the reference and the access point of the DSA holding that naming context.
- **Non-specific Subordinate Reference** - A knowledge reference that holds information about the DSA that holds one or more unspecified subordinate entries or entry copies. This reference type consists of the

access point of the DSA that holds the entries (or copies) of one or more immediately subordinate naming contexts.

- **Cross Reference** - A knowledge reference containing information about the DSA that holds an entry or entry copy. This is used for optimization. An entry does not require a superior or subordinate relationship to any entry in the DIB.

2.1.10 DUA / DSA Interaction Modes

The Government Directory will support three modes of interactions between DUAs and DSAs when executing operations. These operations may be inhibited based on agency specific policies and applications. These modes are chaining, multicasting and referrals.

During chaining operations DSAs interact directly with one another using DSP. Requests for information are forwarded from one DSA to another. This process is transparent to the end user or DUA that initiated the request. Queries are progressively forwarded to DSAs until one is reached that contains the information being requested. Exhibit 2-8 presents this concept.




Exhibit 2-8
Chaining

Multicasting is a special case of chaining where queries are forwarded to multiple DSAs in parallel. This parallel operation does not imply simultaneous “broadcasting” of a request. This type of DUA/DSA interaction is analogous to “transmitting” a query to multiple recipient DSAs with the expectation that one or more will be able to satisfy the request.

Referrals operate by having the DUA or DSA progressively contact each DSA to return the requested information. Each DSA contacted returns any portion of the query result as well as a pointer to another DSA that holds requested information or at least a portion of it. After receiving this pointer information, the DUA or DSA must then communicate with the next DSA to continue the operation. Exhibit 2-9 presents this concept.




Exhibit 2-9
Referrals

2.1.11 Government DSA Topologies

2.1.11.1 Background

Because of the size and complexity of the Federal Government, the corresponding X.500 DSA topology will also be highly complex. There are numerous Federal agencies and departments, each with a need for one or more DSAs. The unique requirements of each agency will dictate the number of DSAs that will be implemented. The interrelationships of these DSAs in the context of the Government Root DSA will be examined and discussed below.

2.1.11.2 Design Factors

The design of the Federal Government's X.500 DSA topology is dependent on several factors, including:

- Size of the agency
- Size of departments within the agency
- Reliability of the telecommunications network over which the traffic will pass

- Distance between directory data and the requester
- Degree of traffic between Federal agencies/departments
- Performance requirements
- Type of network connectivity between agencies/departments
- Security concerns of the individual agencies/departments.

The degree to which the personnel of two agencies communicate directly impacts the methods selected to transfer directory information among the DSAs. For example, if two agencies' missions are closely interwoven and there is considerable communication between their respective staffs, it is most likely that the DSAs would exchange directory information via replication or shadowing. If, on the other hand, the two agencies staffs rarely interact, the method for sharing directory information would more likely be via chaining.

One area of concern is the use of X.500 in ascertaining LAN electronic mail addresses. Given the proliferation of communication networks interconnecting the Federal Government, a common desire is to know the electronic-mail address of a particular Government organization or employee. The end-user's ability to acquire the e-mail address on-line depends on whether the DSA topology supports directory synchronization across the affected agency LANs.

2.1.11.3 Basic Topology

The Government DIT is the hierarchical representation of the DIB represented as a tree, whose vertices are the Directory entries. In other words, it is a logical definition of the hierarchical relationships between entries of different object classes. At the top of the DIT is the "root" naming context. This root context plays a significant role in defining the naming resolution rules at the top level of the DIT. The root is not an actual directory, it is a logical construct consisting of a set of entries that make up the immediate subordinates of the top level root naming context.



The Government Domain is comprised of all organizations which are represented within the Government Electronic Directory. This domain includes all Federal, State, Tribal, Commonwealth and Local Governments, and is also known as the Government Naming Context.

The U.S. Government domain is comprised of all US Federal Government agencies and organizations and their personnel, publications, and services. This domain, or portion of the DIT, is a component of the Government Electronic Directory. Each State, Tribal and Commonwealth Government will have their own domain, or portion of the DIT, each of which is also a component of the Government Electronic Directory.

In order to interoperate effectively and efficiently the individual domains which comprise the Government Electronic Directory must cooperatively develop and adapt standard guidelines, policies, procedures and methods to establish, manage and efficiently use the Government Electronic Directory.

Exhibit 2-10 depicts the Government DSA topology on a global scale. Single DSAs in each domain depict Well Known Entry Point (WEP) points in the DIT. It is envisioned that each agency will operate at least one WEP DSA.

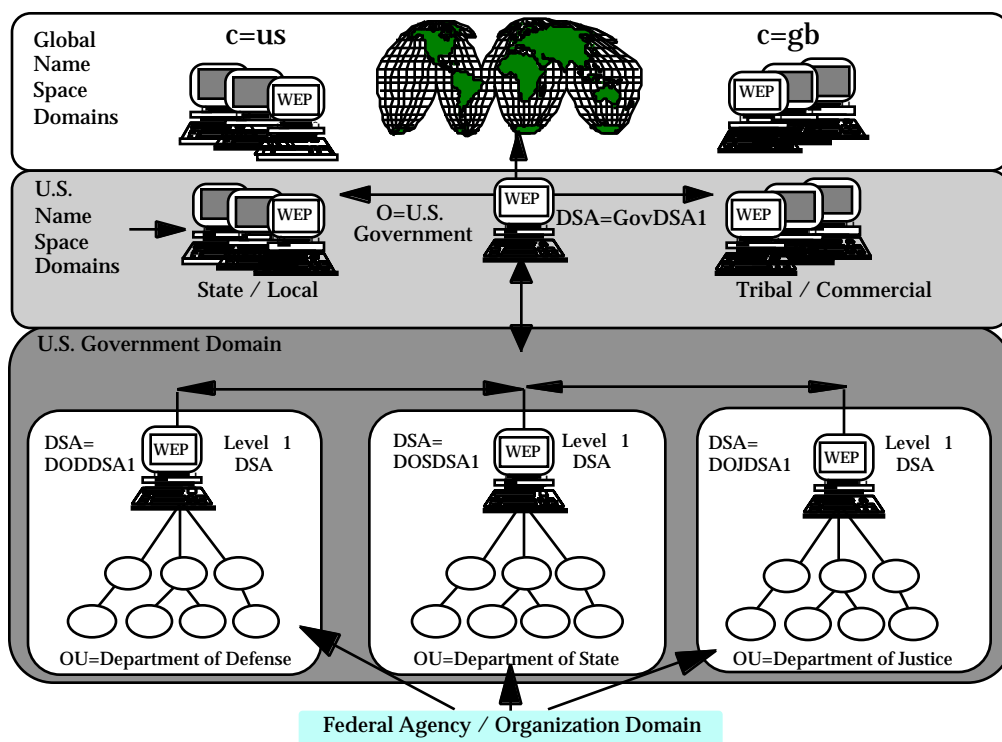


Exhibit 2-10
Global DSA Topology

The WEP DSAs contain sufficient knowledge information for superior, subordinate and cross reference navigation around the DIT, as depicted in Exhibit 2-11. The Top Level DIT will be distributed among the WEP DSAs. This is a topology of information, not of components, therefore depending on network configurations, usage patterns, and security considerations, there may be more than one WEP DSA per domain. Shadowing and replication agreements also affect component topology.

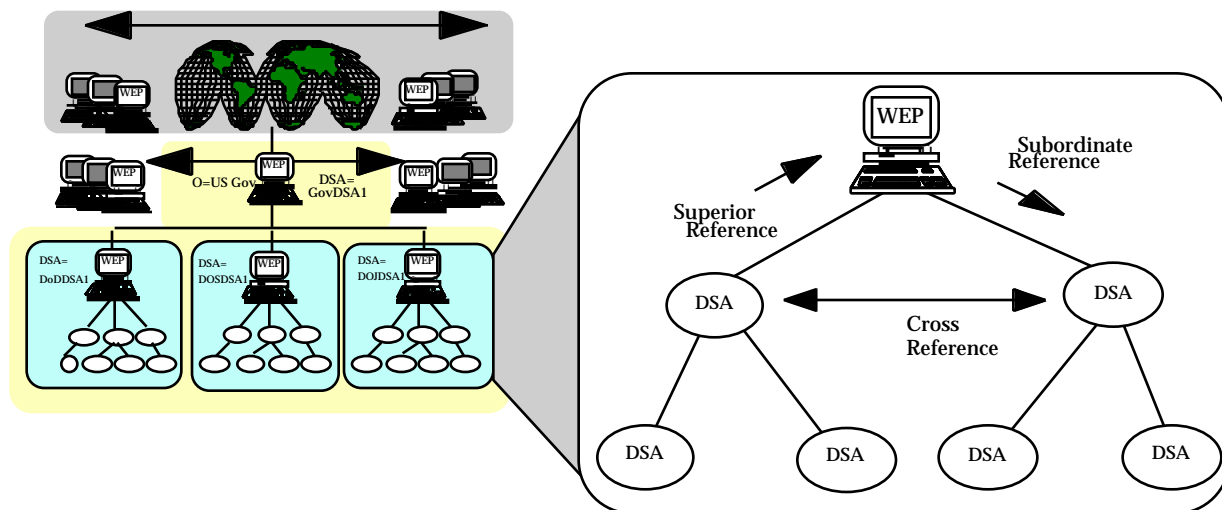


Exhibit 2-11

Navigating the DIT

Exhibit 2-12 depicts a DSA information topology on an international (global) scale. The Government root WEP DSA(s) will hold sufficient knowledge to act as a chaining/referral agent for collateral US domestic domains and as an access point for other global DSAs. Federal Agency DSAs will be distributed with the top level of the DIT. Agency DSA topology knowledge references and DSA DIT distribution characteristics will be determined internally. Additionally, chaining, referral, and replication will be determined by the agency based on their local operational considerations.

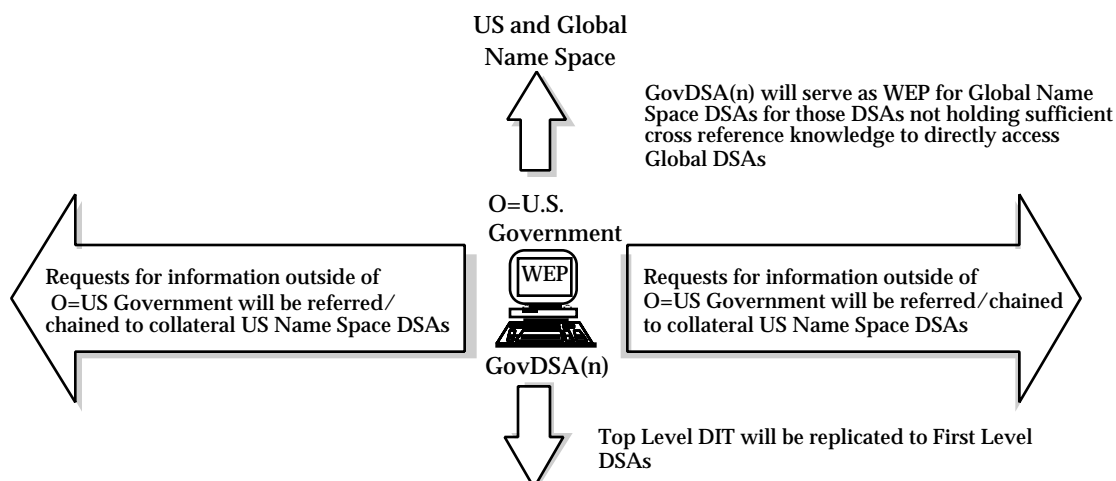


Exhibit 2-12

Global DSA Information Topology

2.1.11.4 Directory Synchronization

The Government Electronic Directory will support synchronization of directories (DIRSYNC) as a transition path to complete X.500 interoperability; primarily supporting the exchange of electronic mail addresses among the installed base of first generation, file-based electronic mail systems. The recommended topology for DIRSYNC is one that uses the X.500 DIB as the repository for e-mail addresses, along with other white pages information about Organizational Persons in the Government.

The Government Directory will support the exchange of X.400 and Request for Comments (RFC) 822 Internet Simple Mail Transport Protocol (SMTP) electronic mail addresses with Local Area Network (LAN) and Host-based electronic mail systems through gateways between Agency WEP DSAs and electronic mail systems. This gateway will be defined by the electronic mail systems in use by the Agency, and will typically involve numerous gateways, one for each electronic mail system that will exchange electronic mail address information. This will allow agencies to communicate and exchange e-mail addresses through the Government Electronic Directory to present an X.500 “face” to the Root DSA, and other Agency DSAs. The Root DSA will provide the “glue” unifying the DIB across multiple DIRSYNC implementations. Additionally, WEP DSAs may exchange RFC 822 and X.400 address information with other WEP DSAs using X.500 protocols (1993). This information may then be used in the DIRSYNC topology to populate LAN-based electronic mail system with addresses from other agencies. Agencies are responsible for synchronizing legacy system directories with the Government Electronic Directory, as is shown in Exhibit 2-13.

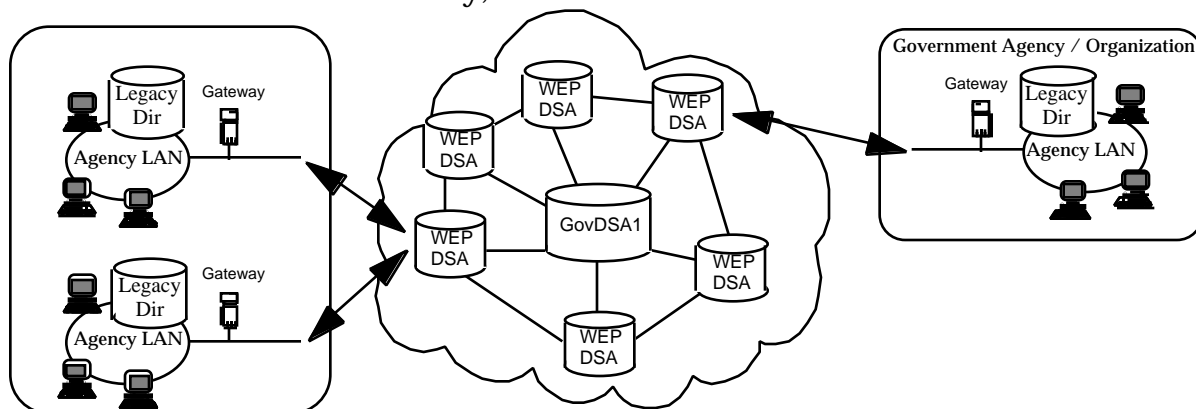


Exhibit 2-13
Directory Synchronization

Access by users to directory information from electronic mail systems may be performed in a number of ways, dependent upon the DIRSYNC implementation selected for installation. All of the following will be available:

- Access to LAN-based electronic mail directories (addresses appear in the format of the LAN-based electronic mail system)
- DUAs using LDAP
- Web Clients, such as Mosaic and Netscape
- Proprietary address books (provides a proprietary view of X.500 information to LAN-based electronic mail clients using LAN protocols).

2.2 Information Interface Topology

2.2.1. Introduction

In order to provide a universally accessible Government Electronic Directory there will be two interface topologies implemented:

- End-user computer applications -- Provide graphical user interfaces (GUI) to humans so they can read or modify information stored in the directory service and related services,
- Client applications -- Use the directory service to locate the network addresses and other similar types of information on the server applications they wish to connect with.

Exhibit 2-14 illustrates an end-user interface interacting with the Government Electronic Directory.




Exhibit 2-14
End-User Interface

Exhibit 2-15 depicts a client/application program accessing the Government Electronic Directory to retrieve information necessary for connecting with a server.

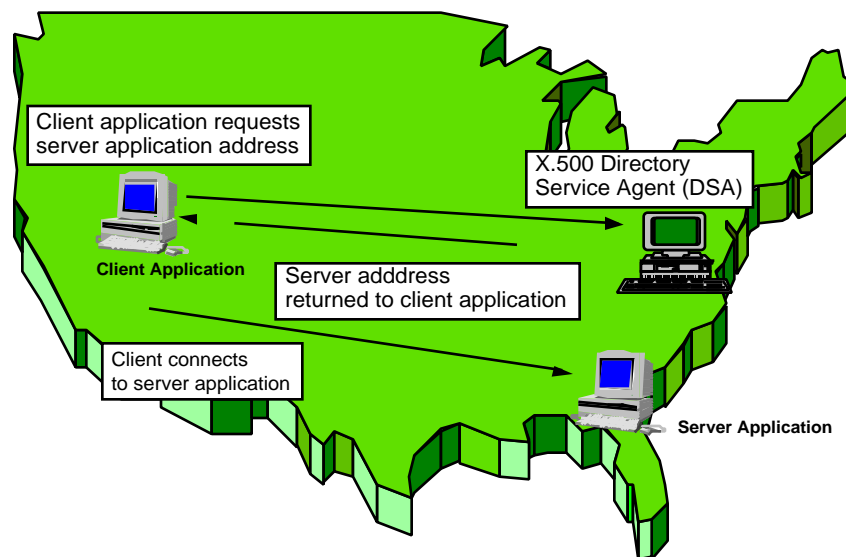


Exhibit 2-15
Application Interface in a Client/Server Environment

End-user interfaces will be of four primary types:

- Kiosk Interface -- Publicly placed end-user interfaces provided by the government which allow the general public to access information about the government such as organization contacts, services and forms.
- Public Domain Software -- End user interface software made available to the public (e.g., non-government organizations, businesses or citizens). For example, this software may be installed and run on business' or citizens' computers in order to allow access to the Government Electronic Directory and related services. The level of directory information access granted to the business or private user will be dependent on the access controls defined for the user (see Section 4.0, Security Architecture).
- Government Employee Interfaces -- This interface may be customized to support the unique requirements specified by government organizations. It will provide more advanced search and retrieval mechanisms for the types of information which assist government employees in performing daily work assignments. This information will include government-wide employee directory look-ups for telephone numbers, fax numbers, e-mail addresses and other similar information necessary to enable government-wide communications.
- Government Administrative Directory Interfaces -- This directory interface is specialized so that Directory Administrators can fully manage the directory service. Directory management consists of adding/deleting/modifying government employee or organization entries, replication configurations and similar types of administrative operations.

End-user directory interfaces will be multi-purpose client applications supporting a range of protocols necessary to access the Government Electronic Directory and related information bases. These protocols can include any or all of the following:

- Directory Access Protocol (DAP) - ITU-T standard for accessing X.500 directories over OSI networks,
- Lightweight Directory Access Protocol (LDAP) - De facto standard for accessing X.500 directories over TCP/IP networks,
- Hyper-Text Transfer Protocol (HTTP) - De facto Internet standard for communicating with World Wide Web servers,
- Hyper-Text Markup Language (HTML) - De facto Internet standard for linking objects over HTTP and the Web,
- Z39.50 Wide Area Information Service (WAIS) - ANSI/NISO-defined, Internet-based distributed document storage and retrieval service.

The above protocols are supported by the Government “Information Interface” so that a wide range of related information can be presented to the user from a single interface. The information interface topology describes the necessary types of interfaces, their approximate physical location and required resources to enable access to the X.500 Directory Service and related services. This document will describe the recommended information interface topology in order to support end-user and client/server application access to the Government Electronic Directory and related services.

2.2.2 Public Information Interface Topology

Public information interfaces will be of two types:

- Kiosk Interfaces -- A hardware configuration located in a publicly accessible area
- Public Domain Software -- Non-government organization, business and citizens' home computer use of publicly available software.

2.2.2.1 Kiosk Interfaces

The kiosk interface is a general purpose, publicly reachable information interface. This interface will be hypertext enabled so that it can access information stored in on-line government Web servers via HTML/HTTP. In addition, the kiosk interface will support Z39.50 for integrated access to on-line government document repositories. Exhibit 2-16 illustrates a potential kiosk interface.

| | | |
|---------------|----------------------|---|
| Topic: | <input type="text"/> | Person |
| | <input type="text"/> | Agency (Enter Text or F1 for List) |
| | <input type="text"/> | Document |
| | <input type="text"/> | Form |
| | <input type="text"/> | Project |
| | <input type="text"/> | Map |
| | <input type="text"/> | Keyword(s) (Enter Text or F1 for List) |

Figure 2-16
Kiosk Interface

Kiosks will be deployed in public places such as libraries, U.S. Post Offices and potentially other similar locations. Kiosk interfaces will be generalized enough to enable the public to obtain white, blue, yellow and green pages information regarding government persons, organizations, services and documents. HTML and Z39.50 support will allow users to select information for retrieval using a mouse point-and-click mechanism such as a kiosk-mounted rollerball. Kiosk interfaces may also provide the capability to print information retrieved, such as tax forms, government services listings and similar types of information. Therefore, the kiosk may need to be fitted with a laser printing device.

2.2.2.2 Kiosk Interface Connectivity

Kiosk interfaces will be configured to use LDAP for connection to the closest DSA over TCP/IP asynchronous connections. Local DSAs will be configured to retain a copy of most remote government organizations' directory subtrees through replication,

or the DSA may chain the directory query to the appropriate DSA for resolution. Exhibit 2-17 illustrates this concept.

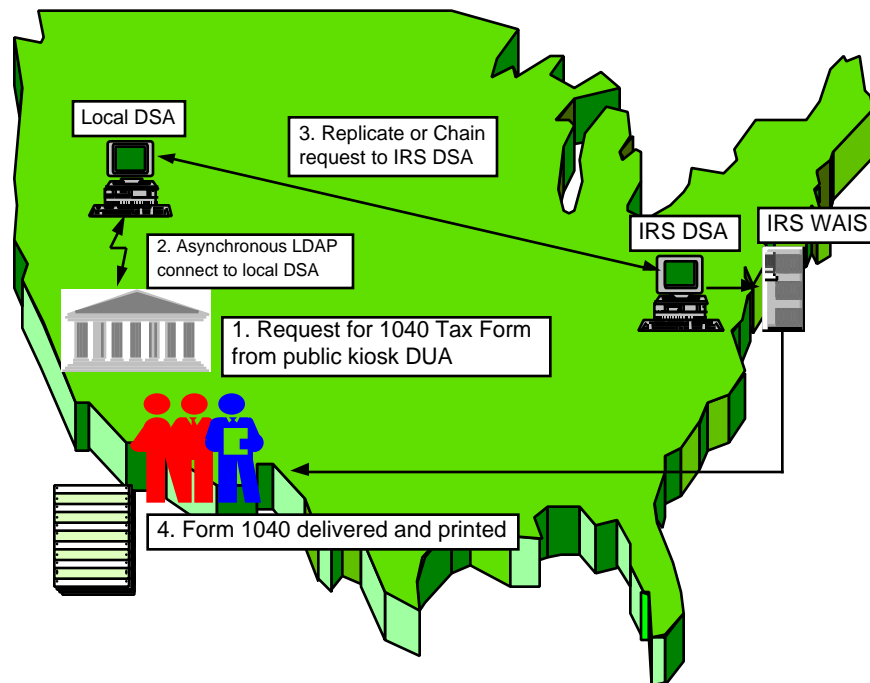


Exhibit 2-17

Kiosk Interface Connectivity

In the above example, a citizen uses a public kiosk in California to request a 1040 Tax Form from the Internal Revenue Service. The kiosk interface accesses the local DSA to search for the actual location of the tax form. The local DSA chains the query to the IRS DSA which retrieves the document from the IRS on-line archives, perhaps maintained on the Z39.50/WAIS server, in Washington, D.C. The form is then printed at the kiosk upon the user's request.

2.2.2.3 Public Domain Software

In addition to making the Government Electronic Directory available in public kiosks, the Federal Government may elect to make a version of the information interface software available for use by non-government organizations, businesses or the public. This software is similar to the kiosk interface software in the following respects:

- Employs LDAP and asynchronous network connections for local DSAs

- Supports Z39.50 for document search and retrieval services provided by the government
- Supports HTML and HTTP for linking to on-line government Web servers
- Provides a general purpose interface to allow white, blue, yellow and green pages look-ups
- Allows retrieved information to be printed on an attached print device.

2.2.3 U.S. Government Directory Interface Topology

2.2.3.1 Government Employee Information Interface

Government employees and DSA Administrators will access the Government Electronic Directory via more specialized directory interfaces than those provide to the public. Government employees (non-X.500 Administrators) will require access to a broader range of employee entry information, organizational information, government services listings and documents (white, blue, yellow and green pages). Though the government employee interface may be similar to the one described for public/kiosk use (e.g., supports LDAP, HTML, HTTP, Z39.50), access to this broader range of information will be directly related to the Government Schema Definition regarding authentication and access controls. As a result, the government employee information interface may prompt the employee for a password or private encryption key in order to authenticate the user and grant appropriate access to the information stored in the Government Electronic Directory.

2.2.3.2 Location and Connectivity of Government Employee Interfaces

It is recommended that each government employee with a need to regularly access the directory and related services have a government employee information interface on his/her desktop computer with the following properties:

- Uses LDAP to connect to the local DSA over TCP/IP LAN connection

- Versions available for Microsoft Windows, UNIX Motif, UNIX command line and Macintosh platforms
- Access to directory via dial-up asynchronous connection supported
- Support for HTML, HTTP, and Z39.50.

Exhibit 2-18 depicts the Government Employee information interface topology.

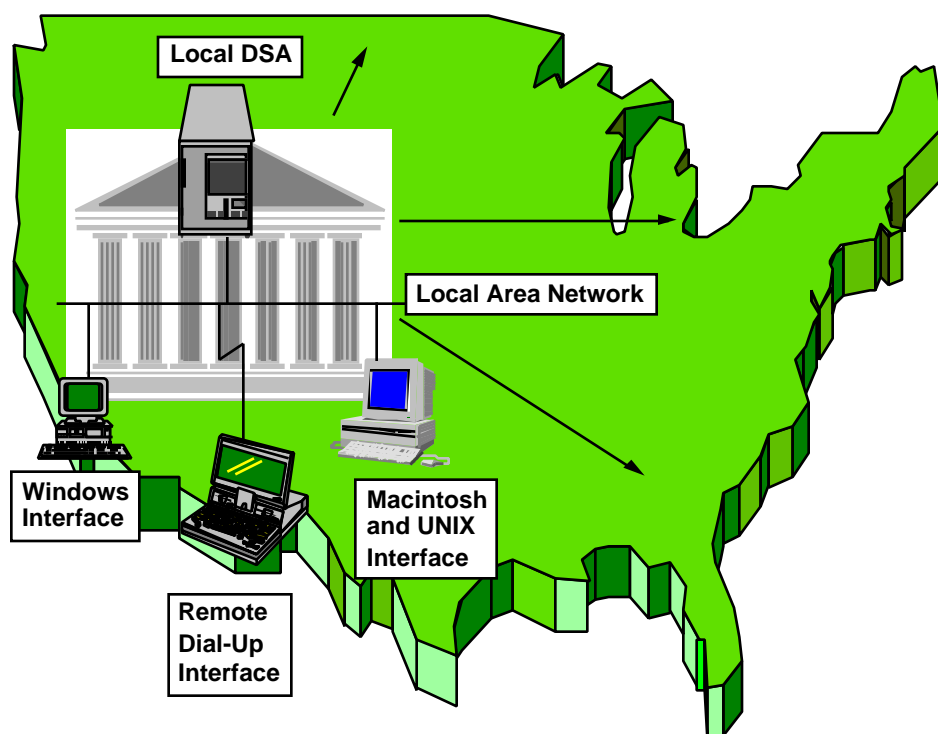


Figure 2-18

Government Employee DUA Topology

This DUA topology would be similarly deployed at all federal agencies in order to allow government employee access to the Government Electronic Directory.

2.2.3.3 Government Administrative Directory User Agent

The Government Administrative DUA (ADUA) will enable X.500 Administrators to manage the directory service. The ADUA will allow administrators to perform the following functions:

- Add, delete, modify entries in the X.500 DSA
- Configure knowledge information regarding superior, subordinate and cross references
- Configure authentication passwords and procedures
- Configure access control information
- Configure shadow update procedures
- Monitor the activities and logs of the DSA.

2.2.3.4 Location and Connectivity of ADUA

Each government location which is maintaining a DSA must have an ADUA locally or remotely connected. A DSA cannot be administered without an attached ADUA. ADUAs should use the full DAP protocol (rather than LDAP) in order to support full functionality. Because full DAP is required, the ADUA uses a partial OSI stack (application layer through session layer) and supports RFC-1006 for TCP/IP transport/network layer connectivity to the DSA. The ADUA platform will be an object oriented, icon-driven interface such as Microsoft Windows or UNIX Motif. The overall Public and Government Information Interface topology to the Government X.500 Directory is depicted in Exhibit 2-19.

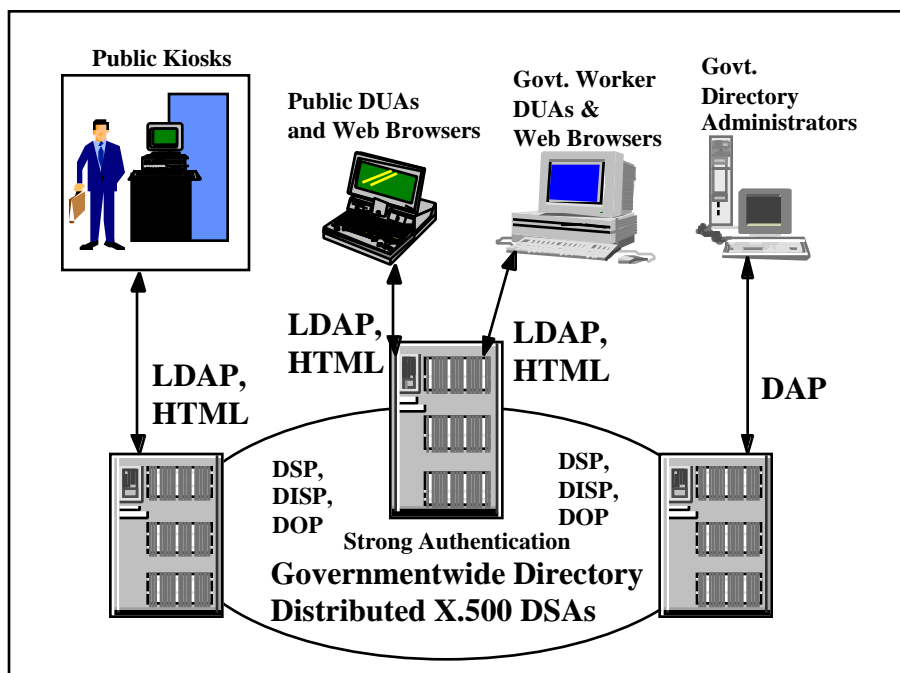


Exhibit 2-19

Government X.500 Directory Interface Suite

2.2.4 Client/Server Application Interfaces

In the distributed computing, client/server environment being deployed within the Federal government, client applications will access the Government Electronic Directory. Client applications require the network addresses of the server applications they wish to connect to. Using X.500 in this manner allows server applications to be deployed on network servers without the server address being 'hard-coded' into the client application. This fundamental principle of distributed computing lets server applications move from machine to machine, with no impact on the large number of client applications already deployed.

Applications will be registered as entries in the Government Electronic Directory as application process objects. One of the mandatory attributes for each application process is its network address, including TCP/IP port and socket addresses. Client applications will need to be developed or re-engineered to employ LDAP, DAP or X/Open Directory Services (XDS) API calls in order to interact with the directory. In addition, the applications may need to support HTML, HTTP and/or Z39.50 to access related services.