

# Volume I, Chapter 4 – Table of Contents

<b>4. Network Services .....</b>	<b>4-1</b>
4.1 Network Time Protocol.....	4-2
4.1.1 Service Description.....	4-2
4.1.2 Applicable Standards, Policy, and Guidance .....	4-3
4.1.3 Requirements .....	4-3
4.1.4 Assumptions .....	4-3
4.1.5 Service Architecture .....	4-3
4.1.6 Roles and Responsibilities .....	4-5
4.2 Domain Name Service.....	4-5
4.2.1 Service Description.....	4-5
4.2.2 Applicable Standards, Policy, and Guidance .....	4-6
4.2.3 Requirements .....	4-6
4.2.4 Assumptions .....	4-6
4.2.5 Service Architecture .....	4-6
4.2.6 Roles and Responsibilities .....	4-8
4.3 Enterprise Directory Service .....	4-8
4.3.1 Service Description.....	4-8
4.3.2 Applicable Standards, Policy, and Guidance .....	4-9
4.3.3 Requirements .....	4-10
4.3.4 Assumptions and Observations.....	4-10
4.3.5 Service Architecture .....	4-11
4.3.6 Roles and Responsibilities .....	4-25
4.4 Electronic Mail.....	4-25
4.4.1 Service Description.....	4-25
4.4.2 Applicable Standards, Policy, and Guidance .....	4-26
4.4.3 Requirements .....	4-27
4.4.4 Assumptions .....	4-27
4.4.5 Service Architecture .....	4-28
4.4.6 Addressing Conventions .....	4-29
4.4.7 Routing Architecture.....	4-30
4.4.8 User Interface.....	4-31
4.4.9 Roles and Responsibilities .....	4-32
4.5 Network News Service using NNTP.....	4-32
4.5.1 Service Description.....	4-32
4.5.2 Applicable Standards, Policy, and Guidance .....	4-33
4.5.3 Requirements .....	4-33

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

4.5.4	Assumptions .....	4-33
4.5.5	Service Architecture .....	4-33
4.5.6	Roles and Responsibilities .....	4-35
4.6	Web Hosting.....	4-35
4.6.1	Service Description.....	4-35
4.6.2	Applicable Standards, Policy, and Guidance .....	4-36
4.6.3	Requirements .....	4-36
4.6.4	Assumptions .....	4-38
4.6.5	Service Architecture .....	4-39
4.6.6	Roles and Responsibilities .....	4-41
4.6.7	Security Guidelines.....	4-42
4.7	File Transfer Protocol.....	4-45
4.7.1	Service Description.....	4-45
4.7.2	Applicable Standards, Policy, and Guidance .....	4-45
4.7.3	Requirements .....	4-46
4.7.4	Assumptions .....	4-46
4.7.5	Service Architecture .....	4-46
4.7.6	Roles and Responsibilities .....	4-48
4.8	Public Key Infrastructure (PKI) .....	4-48
4.8.1	Service Description.....	4-48
4.8.2	Applicable Standards, Policy, and Guidance .....	4-48
4.8.3	Requirements .....	4-49
4.8.4	Assumptions .....	4-49
4.8.5	Service Architecture .....	4-49
4.8.6	Roles and Responsibilities .....	4-52
4.9	Remote Access .....	4-52
4.9.1	Service Description.....	4-52
4.9.2	Applicable Standards, Policy, and Guidance .....	4-53
4.9.3	Requirements .....	4-53
4.9.4	Assumptions .....	4-53
4.9.5	Service Architecture .....	4-53
4.9.6	Roles and Responsibilities .....	4-57
4.10	General Voice.....	4-57
4.10.1	Service Description.....	4-57
4.10.2	Applicable Standards .....	4-57
4.10.3	Requirements .....	4-57
4.10.4	Assumptions .....	4-59
4.10.5	Service Architecture .....	4-60
4.11	Shipboard Voice.....	4-61
4.11.1	Service Description.....	4-61
4.11.2	Applicable Standards, Policy, and Guidance .....	4-61

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

4.11.3	Requirements .....	4-61
4.11.4	Assumptions .....	4-62
4.11.5	Service Architecture .....	4-62
4.11.6	Roles and Responsibilities .....	4-63
4.12	Secure Voice .....	4-63
4.12.1	Service Description.....	4-63
4.12.2	Applicable Standards .....	4-64
4.12.3	Requirements .....	4-64
4.12.4	Assumptions .....	4-66
4.12.5	Service Architecture .....	4-66
4.13	Multimedia .....	4-67
4.13.1	Service Description.....	4-68
4.13.2	Applicable Standards and References .....	4-69
4.13.3	Requirements .....	4-69
4.13.4	Assumptions .....	4-70
4.13.5	Service Architecture .....	4-70
4.13.6	Roles and Responsibilities .....	4-72
4.14	Common Operating Environment Applications .....	4-73
4.14.1	Service Description.....	4-73
4.14.2	Applicable Standards, Policy, and Guidance .....	4-74
4.14.3	Requirements .....	4-74
4.14.4	Assumptions .....	4-74
4.14.5	Service Architecture .....	4-74
4.14.6	Roles and Responsibilities .....	4-77

This page intentionally left blank.

## **4. Network Services**

---

The previous chapter provides design guidance for the network connectivity of the DON enterprise network. It includes the physical connections, the protocols that establish connections and maintain communication sessions between systems, and a description of how applications access and inter-operate with the lower-level network communication functions. The Wide Area Connectivity Plan and the Metropolitan Area Network and Campus Area Network templates provide design guidance for network connectivity across the entire Naval enterprise.

The organizations within the DON must receive basic Information Technology Infrastructure (ITI) services in order to support the information requirements pertaining to their basic missions. This chapter describes those basic ITI services that users in all functional areas require that must be accessible from the network on an enterprise basis. These services closely correspond to the Basic Network and Information Distribution Services (BNIDS) described in the DON Information Technology Support Guidance (ITSG). The ITSG is the companion document to this DON ITI architecture and should be used as the DON authoritative source for ITI standards.

All network services must have a common planning framework and consistent implementation strategy. Some services, such as Domain Name Services, must be implemented under an enterprise hierarchical plan. These basic services are described in this chapter with sufficient detail to allow the appropriate ITSCs to consistently plan and implement integrated services. This chapter should be used in concert with the ITSC template in determining ITSC implementation of services.

The network services defined within this chapter are as follows:

- Network Time Protocol (NTP) (Section 4.1)
- Domain Name Service (DNS) (Section 4.2)
- Enterprise Directory (Section 4.3)
- Electronic Mail (Section 4.4)
- Network News / Network News Transfer Protocol (Section 4.5)
- Web Hosting (Section 4.6)
- File Transfer Protocol (FTP) (Section 4.7)
- Public Key Infrastructure (PKI) (Section 4.8)
- Remote Access (Section 4.9)
- General Voice (Section 4.10)
- Shipboard Voice (Section 4.11)
- Secure Voice (Section 4.12)
- Multimedia (Section 4.13)
- Common Operating Environment (COE) (Section 4.14)

## **4.1 Network Time Protocol**

### **4.1.1 Service Description**

Network Time Protocol (NTP) is a protocol that rides on the Internet Protocol (IP) and assures accurate local time-keeping with reference to radio or atomic clocks. This protocol is capable of synchronizing distributed clocks within milliseconds over long time periods. Several DON services require that system clocks are set accurately so that all servers and services have a consistent time. This is especially important for logging servers, file servers, and some security services.

Hosts are assumed to have no other means to verify time other than NTP itself. Although each host typically contains an internal battery-backed clock, a number of factors, including environmental (e.g., temperature), hardware imperfections (e.g., imperfect oscillator), and time setting/resetting inaccuracy cause errors in the reported time. Because there is no synchronization among hosts within the network, local hosts should use internal battery-backed clocks only to confirm the sanity of the NTP time-keeping system, not as the source of the system time.

While some local clocks maintain time-keeping accuracy to a published and trusted standard (“truechimers”), others are consistently slow, consistently fast, or unpredictable (“false tickers”). The accuracy achievable by NTP depends upon the precision of the local clock hardware and stringent control of device and process latencies. NTP provides the means to set and adjust local clocks, thus correcting for offset-slewing, frequency compensation, and other errors in the local clock function. The local clock is then usable by resident applications for the current time/calendar functions, and periodically updates itself based upon previously measured differences between the true time and the local clock time. In this way, a local host continuously “slews” its internal clock to the correct time using calculations from recent NTP updates. Clocks that are relatively stable in frequency need less frequent updates from NTP servers because they can slew to the correct time based upon predictable “offsets” based on error rates. Unreliable clocks “drift” and require more frequent NTP updates to achieve desired clock accuracy.

Clock synchronization over a network requires long periods and multiple comparisons in order to maintain accurate time-keeping. While only a few measurements are usually adequate to reliably determine local time to within a tolerance of 1-2 seconds, periods of many hours and dozens of measurements are required to resolve oscillator skew and maintain local time on the order of a millisecond. Thus the accuracy achieved is directly dependent on the time taken to achieve it. Fortunately, the frequency of measurements can be quite low and almost always non-intrusive to normal network operations. Correctly implemented, local clocks can use NTP to maintain time to within 15 ns and frequency to within 0.3 ms per day.

The DON enterprise time service delivers accurate time via NTP to all servers that wish to subscribe. Even client machines can use this service, if desired. This architecture offers stratum 1 and stratum 2 service at each of the ITSCs. Campus networks are encouraged to install their own stratum 2 or stratum 3 servers for distribution of time service locally.

## **4.1.2 Applicable Standards, Policy, and Guidance**

RFC 1305, Network Time Protocol (Version 3) Specifications, Implementation, and Analysis

A number of applicable references are at <http://tycho.usno.navy.mil>.

ITSG Chapter 6.6.1

RFC1035 – Network Time Protocol (version 3)

## **4.1.3 Requirements**

Must provide a low stratum-level service to any hosts that wish to subscribe to this service. This implies a high degree of scalability, since hopefully all DON enterprise servers will subscribe to this.

Network Time Service must have no single point of failure.

## **4.1.4 Assumptions**

GPS antennas can be installed on the roof at the ITSC site.

## **4.1.5 Service Architecture**

The DON enterprise time service will establish a stratum 1 time source at each ITSC to obtain time from the GPS network. For redundancy, two such time servers will be installed at each ITSC. Because these stratum 1 servers do not necessarily have sufficient performance and scalability to deliver time directly to all subscribers in the region, a hierarchy is established in this guidance.

At an ITSC, stratum 2 servers (a minimum of two) will slave directly to the stratum 1 servers, and to each other, and to NTP sources at one or more ITSCs in other regions. These stratum 2 servers will provide network time to hosts at the ITSC. They can also provide time to any stratum 3 servers in the region.

At the campus level, time servers can be installed if the subscriber population on that campus justifies it. The campus time servers should operate at stratum 2 and slave to the servers in the ITSC. Very small campuses can have their hosts obtain time directly from the servers at the ITSC.

An option currently being explored is to get terrestrial NTP service directly from the U.S. Naval Observatory (USNO). Currently, the network path to USNO is congested and dispersion values over the link are high, thus limiting its usefulness.

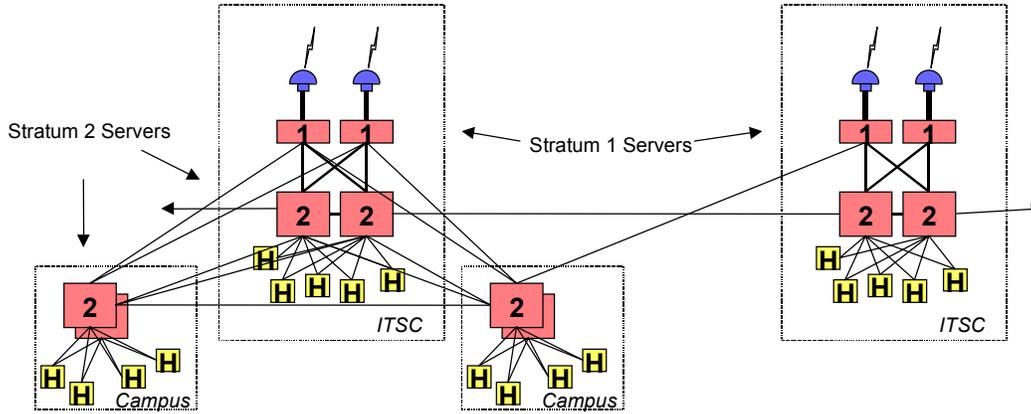


Figure 4-1. Time Service Architecture for ITSC

The architecture diagram in Figure 4-1 illustrates the time service architecture within an ITSC, how campus level stratum 2 servers can peer with the ITSC, and how those servers can deliver time to hosts on the campus.

The DON NTP design is such that accidental or malicious data modification (tampering) or destruction (jamming) of a time server should not result in time keeping errors elsewhere in the DON synchronization subnet. In addition to obtaining time information from multiple, topologically-distributed stratum 1 time servers, all regionally-operated time servers implement access controls within the stratum 2 peer group using the IP security protocol (IPSEC). Redundant timeservers and diverse network paths also increase the quality of the time service within each region. The protocol itself is self-healing because the synchronization hierarchy reconfigures itself after it is determined that a timeserver is no longer in service. Vulnerability is minimized by allowing only designated regional time servers to become synchronization sources for subordinate time servers and denying synchronization requests to/from unknown and untrusted time servers. The NTP standard identifies the available security mechanisms.

The two stratum 1 servers in a region should have DNS names (or aliases) of “tick” and “tock”, i.e., “tick.pacsw.navy.mil”.

## Regional Issues and Considerations

Each region will implement redundant stratum 1 and stratum 2 NTP servers as described above. This server will be located inside the DON enterprise firewall.

Regional scalability needs to be considered. Any one server should not be oversubscribed.

## Campus and Operational Node Issues and Considerations

Each campus may implement one or more stratum 2 servers using the regional stratum 1 and stratum 2 servers as its primary time source and a neighboring region as the alternate. These time servers should be located on or near the individual subnetworks or LANs containing end systems requiring accurate clocks. For large bases, alternative architectures may be appropriate, including stratum 3 servers on individual LANs within large organizations. Implementation agreements for establishing NTP peer groups for time synchronization are coordinated through the base IT service center. For example, all stratum 3 servers synchronize with each other and lower stratum clocks on the campus and with the regional clocks.

## **Deployed Forces Issues and Considerations**

While in port, deployed forces synchronize with the regional stratum 2 servers just as other operational nodes do. However, because ships are not a full-time participant in the campus time synchronization, they receive synchronization information from other base operational nodes but typically do not send synchronization information to them.

While at sea, deployed forces operational nodes continue to synchronize internal clocks, but terrestrial time sources are unavailable. Instead, ships obtain accurate clock information from GPS or other satellite sources for extended deployments when there is concern that internal synchronization may have unacceptably drifted from actual time.

### **4.1.6 Roles and Responsibilities**

The ITSC will be responsible for providing accurate time within a region. Regional network engineers will be cognizant of the number of time subscribers and ensure that there is adequate hierarchy and redundancy to support the demand.

The ITSC will publish the DNS names of the NTP servers and will provide the instructions for configuring higher stratum servers and for hosts that wish to subscribe to time.

## **4.2 Domain Name Service**

### **4.2.1 Service Description**

DNS is the service that translates domain names to IP addresses and vice versa. A domain name is a mechanism that gives unique names to network devices to eliminate the need for users to remember numerical IP addresses. The DNS service is implemented as a hierarchical distributed database and is accessed using a client/server model. The server component of DNS is the subject of this discussion.

The concept of domain names was introduced at a time when Internet hosts used a flat name space. For example, if one person chose the host name of “eagle,” then no one else on the entire Internet could use that same host name. Clearly, this was not scaleable. Alternatively, the solution was to establish name space domains at the local level where name uniqueness and control only have to be controlled at that level. Domains were assigned to organizations (for example, somecommand.navy.mil). Name uniqueness was guaranteed within that domain because it was controlled within that organization, and uniqueness was guaranteed across the Internet because the name was accompanied by that domain’s name qualifier (for example, eagle.somecommand.navy.mil).

The domain naming scheme is hierarchical. Each level in the hierarchy can assign sub-domains within an assigned name space. For example, the people who control the “navy.mil” domain can allocate subdomains to those Naval commands that have a need for their own domain. This process of establishing sub-domains and assigning control and responsibility to subordinate organizations is called “delegation of authority” and is performed using “NS” resource records in the DNS.

For a given domain, there is always more than one server. One of these domain servers is known as the “primary”, and all others are referred to as “secondary” servers. The primary server contains the master files for the domain (a.k.a. “zone”). The process by which the primary server updates the secondary servers is called “zone transfer”.

## 4.2.2 Applicable Standards, Policy, and Guidance

See section 6.3 of the ITSG.

See the DISA DNS policy memo for 2<sup>nd</sup> level domains. (DISA WASHINGTON DC//D//, DTG 162151Z MAR 98)

Applicable RFCs:

- ◆ RFC974 – Mail Routing and the Domain System
- ◆ RFC1034 – Domain Names – Concepts and Facilities
- ◆ RFC1035 – Domain Names – Implementation and Specification
- ◆ RFC1996 – A Mechanism for Prompt Notification of Zone changes (DNS NOTIFY)
- ◆ RFC2065 – Domain Name System Security Extensions

## 4.2.3 Requirements

In the context of the Naval enterprise, the DNS service must provide the following capabilities:

- Translation of Naval domain names to IP addresses (e.g., hq.navy.mil → 164.224.250.80)
- Translation of Naval IP address space to their respective domain names (e.g., 138.147.50.5 → spider.ncts.navy.mil)
- The above services must be delivered to the entire Internet, not just the Naval enterprise. Therefore, the registered primary servers must have good access to the full Internet.
- Every computer in the Naval enterprise is a DNS client and must associate with one or more DNS servers for general domain lookup services. Therefore, a scaleable architecture is required in order to support 1 million clients.

## 4.2.4 Assumptions

The assumptions for Domain Name Service will be further refined in future IPT iterations.

## 4.2.5 Service Architecture

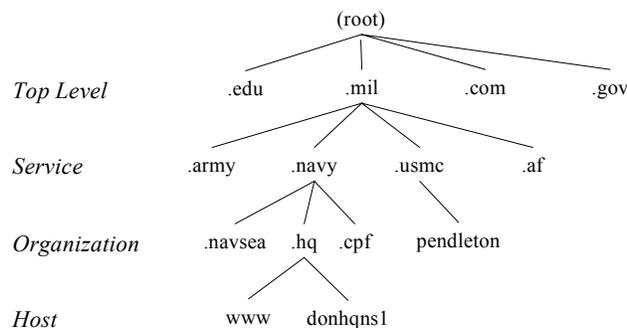


Figure 4-1. Naval Domain Name Hierarchy

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

The diagram in Figure 4-1 shows the domain naming hierarchy. Organization by service level has been the DoD standard since the mid-1980s. This is strictly a naming hierarchy; it has nothing to do with physical topologies (a common misconception) or distribution of servers.

For each *zone* in the domain name hierarchy (a vertex in the above graph, e.g. `hq.navy.mil`, `pendleton.usmc.mil`, etc.), multiple servers must provide the translation services for that zone. This is a fundamental rule, for reliability purposes, and is enforced by the Network Information Center (NIC). The higher in the naming hierarchy, the more important this is. Therefore, for most organizations, 2 or 3 servers for a zone may be sufficient. But for the second level zones (i.e., `navy.mil` and `usmc.mil`), many more should be established. Note that a single DNS server can support hundreds or thousands of zones. Dedicated servers per zone are not required.

The architecture for deployment of DNS servers will be as follows. Each ITSC will have at least two DNS servers, for both the unclassified and classified (SECRET) levels. These servers will do nothing but DNS. So that these servers can be considered “hardened” for security purposes, all nonessential services (such as send mail) will be disabled and all security patches will be installed. These servers will receive their zone information through zone transfers from the master data source identified for their zone (more about that later). The zone transfers will be protected by DNS Sec. Of the total number of Naval servers (4 at each ITSC), at least 12 should be authoritative (i.e. delegated authority through use of Name Server (NS) resource records in the `.mil` zone). Those 12 (or more) should be chosen for good connectivity and geographic separation. Specific sites will be determined at a future time.

The two primary DNS servers at each ITSC serve as the master servers for the ITSC and the region. All clients in the ITSC will configure their resolvers to point to the master servers. All subordinate DNS servers in a region will be configured to use these primary DNS as forwarders. The primary DNS servers must provide caching services (for performance), as well as recursive queries.

The master data source (per zone) is a machine on which additions and changes are made to the zone information. This is not a critical machine – no clients use it for DNS services. It can be rebooted or restarted at will. It does not need to be a high performance server. The master data source is where the zone administrator makes and tests all zone updates. The primary DNS servers that serve this zone can then obtain the information using DNS Sec-authenticated zone transfers, either in response to an automatic update (based on time-out values in the Start of Authority (SOA) resource record for the zone) or in response to a zone update message from the master data source. This decouples the critical DNS service from the database editing function, and thereby increases overall stability of the DNS.

All of the unclassified primary DNS servers will serve the `navy.mil` and `usmc.mil` zones. Similarly, all of the classified primary DNS servers will serve the `navy.smil.mil` and `usmc.smil.mil` zones. A given ITSC will also service the zones that correspond to commands within their respective region. The authoritative servers for a given organizational zone should include one of the DNS servers at the ITSC, but should also include one or two DNS servers at other ITSCs. In other words, `spawar.navy.mil` might be served by DNS servers in San Diego, Norfolk, and Hawaii. Good geographic separation is achieved in this example.

Naval DNSs outside the DON firewalls will point all queries to the firewalls. This helps to hide the details of the Naval network structure.

Adherence to one rule of naming consistency is required. For every domain name that is returned by the DNS in response to a query to translate an IP address to a name, that name must be a valid domain that can be queried in the DNS, and that query must produce the same IP address as was queried. For example,

if a lookup of the name that goes with IP address 138.147.50.5 returns spider.ncts.navy.mil, then a query of spider.ncts.navy.mil must return the address 138.147.50.5.

## **Regional Issues and Considerations**

Organizations in a region can operate a master data source (primary server). This server must provide accurate data via zone transfers to the secondary server function at the ITSC and must allow local editing of zone information.

## **Campus and Operational Node Issues and Considerations**

A campus may want to install local DNS caching servers to serve on-campus clients.

## **Deployed Forces Issues and Considerations**

The implementation must allow on-board updates, even while disconnected from the network. In other words, the on-board DNS server is the “truth” for that zone. When connected, it provides updates to the NOC or ITSC via zone transfers.

### **4.2.6 Roles and Responsibilities**

The Navy NIC assigns administration of the navy.mil zone. It is currently delegated to UARNOC. For the Marine Corps, it is delegated to the Marine Corps NOC at Quantico.

Administration of organizational zones is assigned to responsible individuals within that organization (if that is the desire of the organization); otherwise it will be provided by the ITSC.

The ITSC is responsible for maintaining the DNS servers (care and feeding of the servers).

The ITSC may be responsible for performing the updates to given organizational zones if that responsibility has been assigned by the respective organization. For example, smallcommand.navy.mil may not have the necessary IT personnel to understand how to run a master DNS, therefore it will obtain this service from the ITSC.

It is the responsibility of anyone connecting devices to the network to properly register IP addresses and names. (ITSCs will registers these devices with the ITCC on behalf of the organizations they service.)

## **4.3 Enterprise Directory Service**

### **4.3.1 Service Description**

A Directory Service provides a function similar to that of a phone book, but generally offers more than just a list of people and their phone numbers. It serves as a repository for “people information” and can be searched to find phone numbers (office, fax, pager, mobile, STU, etc.), e-mail addresses, and mailing addresses. Once the service is in place, it can be used to contain other information attributes about individuals such as passwords, digital certificates, and emergency contact information.

Technologies and standards exist today in support of directory architectures and implementations. Over the years, well-defined protocols and schemas have been established for organizing information in a directory database and for retrieving that information. Modern client applications are becoming “directory aware” and use the standard protocols to locate information. The growing dependence of these client applications on directories is increasing the importance of directories in the technology infrastructure.

Directories can exist in many forms and serve a variety of purposes, but the focus here is primarily on information about people and organizations. The specified DON enterprise directory service will provide a function similar to the “white pages”-- providing the ability to find people in the Navy and Marine Corps by searching on their name. It will also provide the components of a “blue pages” service by providing the ability to “drill down” through an organizational hierarchy and locate groups or individuals based on their organization or billet. A third basic directory function that can be provided is something similar to a “yellow pages” service-- providing information organized by attributes, such as skill, position, mission, or responsibility. The DON implementation of directory service will be called the Naval Enterprise Directory Service (NEDS).

Not addressed here are directory functions such as a directory of file and print services and system configuration information for desktop management that can be found in various Network Operating Systems (NOSs). At present, these directory functions are more of a regional or site issue. When appropriate, these enterprise management functions may be addressed in future versions of this DON architecture document.

### **4.3.2 Applicable Standards, Policy, and Guidance**

- X.500 and LDAP standards – See ITSG section 6.3.2.
- Defense Message System (DMS) standards for Directory Information Tree (DIT) and attributes
- Public Key Infrastructure (PKI) standards for DIT and attributes
- Applicable RFCs:
  - ◆ RFC1777 – Lightweight Directory Access Protocol (v2)
  - ◆ RFC1778 – The String Representation of Standard Attribute Syntaxes
  - ◆ RFC1779 – A String Representation of Distinguished Names
  - ◆ RFC1960 – A String Representation of LDAP Search Filters
  - ◆ RFC2247 – Using Domains in LDAP/X.500 Distinguished Names
  - ◆ RFC2251 – Lightweight Directory Access Protocol (v3)
  - ◆ RFC2252 – Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions
  - ◆ RFC2253 – Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names
  - ◆ RFC2254 – The String Representation of LDAP Search Filters
  - ◆ RFC2255 – The LDAP URL Format
  - ◆ RFC2256 – A Summary of the X.500(96) User Schema for use with LDAPv3

### **4.3.3 Requirements**

This architecture solution must satisfactorily address a mix of enterprise and local directory requirements:

- Must be accessible throughout the entire DON.
- Must be scaleable to support access from hundreds of thousands of clients.
- Should contain entries for all Naval civilian and military personnel and include the ability to add contractors and other associated individuals as appropriate. The resultant number of directory objects is estimated to be approximately 1 million. For each object, the attributes maintained include each individual's full or "common" name, SMTP e-mail address, office phone number, and location (typically city and organization).
- Information in the directory must have a high degree of accuracy based on authoritative sources, and it must be timely (update latency of no more than 24 hours).
- Information in the directory should be searchable based on a person's name. The ability should exist to limit the scope of searches to a major claimant or extend them to the entire DON. The ability must also exist to search by billet, such as "who is the N6 at CINCPACFLT?"
- There should be a natural integration with e-mail. The addressing function of e-mail clients should be able to directly query the directory using standard protocols without the user having to cut-and-paste the results through a separate application.
- The directory solution must support high availability, reliability, survivability, and performance.
- Sensitive information must have controlled access.
- Linkage with the DoD and DON directory hierarchy to the extent possible should support required synchronization of attributes. This requires a directory synchronization function and not a duplicate maintenance of overlapping databases.
- The architecture should enable organizations within the DON to maintain their own directories to meet local requirements, thus accommodating local attributes and control of information.
- Must support a Public Key Infrastructure (PKI) by storing X.509(v3) certificates as attributes to individual directory entries. These certificates will most likely come from the DISA PKI. (The model associated with PKI will most likely grow to hold multiple certificates, such as military identification card and electronic commerce.)
- The directory should be extendable to include new attributes as required by various applications making use of the directory.

### **4.3.4 Assumptions and Observations**

The following important observations and assumptions are relevant to current directory implementations:

- Once an enterprise directory is populated to the extent that it reaches a critical mass, new applications may arise that must be supported by the directory. This dictates that the directory be flexible and have an extendable architecture and implementation. For clarity, "extendable" as it relates to directories is defined to be one that can meet the peculiar directory-related requirements of the Naval or local organization. It also pertains to meeting the special needs of particular mission applications.

- DMS will have a directory solution, but that solution will not meet the DON requirements identified for directories. The DMS directory will be populated only with high-level organizational information in the near term because its primary purpose is to support organizational messaging. In the DON, e-mail services will support the individual sailor and marine, and e-mail must be supported by the directory. There is no requirement for the DON to exactly align with the DMS DIT, although alignment at OU=Navy and above will support synchronization between the DMS directory and the DON enterprise directory.
- DISA is constructing a directory as part of its PKI initiative. While DISA envisions that it will eventually contain all DoD personnel, it will not be able to meet many of the DON-specific requirements listed above. The DISA directory exists primarily to support PKI and organizes directory information under OU=PKI in the DISA DIT.
- There are initiatives in the DON to link selected MS Exchange address books via replication, and this is viewed by some as a potential enterprise solution. However, such a solution is not scalable to the DON, nor does everyone in the DON use MS Exchange. Also, such replication occurs over proprietary protocols that will be blocked at many firewalls. This replication idea is not a candidate DON solution, however, individual organizations within the DON can continue to perform such replication among themselves to meet local requirements.
- A Navy directory currently exists that is populated with all Naval personnel and receives its data input from the personnel system. Additional populating of directory e-mail addresses and other attributes is often performed by individuals. The result is inaccurate, out of date, or missing information on many individuals. As such, it is not a practical DON solution.

### **4.3.5 Service Architecture**

Because this section is very complex and introduces many new concepts, a number of less-familiar terms are defined below.

- ***Naval Enterprise Directory.*** The logical architecture and the information used to populate the directory described in this architecture.
- ***Naval Enterprise Directory Service (NEDS).*** Sometimes shortened to just “enterprise directory.” This is the implementation of the Naval Enterprise Directory. It includes the physical components necessary to provide this directory information as a service to all users. It will be implemented in FY99.
- ***Authoritative Source.*** This is an accurate source of information for populating the enterprise directory. Many such sources will be employed for implementation of the enterprise directory because any individual source has a limited scope and range of attributes. From the view of the enterprise directory, the claimancy directories (see below) will be used as authoritative sources.
- ***Claimancy Directory.*** Some of the major claimants in the DON have initiatives underway to provide a claimancy-wide directory or address book. This is often a consolidation of many sources of information from within the claimancy, including various e-mail systems, personnel systems, and NOS directories. These claimancy directories contain most, if not all, of the personnel employed by or associated with a claimancy. Each claimancy should have its own directory. A claimancy directory implementation may exist as a central server or as a set of distributed servers that are replicated from a master, depending on the particular needs of the claimancy. It is strongly recommended that such claimancy directories be “LDAP-enabled” so that information in the claimancy directories can be pulled into the enterprise directory.

**Department of the Navy Chief Information Officer**  
**Information Technology Infrastructure Architecture, Version 99-1.0**  
**16 March 1999**

- **Command Directory.** Like claimancy directories, some lower echelon commands may have initiatives to synchronize disparate e-mail address books into a consolidated directory. These commands should be feeding their results “up” to their own claimancy directory if it exists, and if not, to the enterprise directory.
- **Regional replica.** A “clone” of the enterprise directory that supports the needs of a given region. These are deployed to the various regions, or Naval concentration areas, as required for performance and scalability.
- **Campus replica.** A “clone” of the enterprise directory that supports the needs of a given campus. These exist mainly for performance reasons, if access to the regional replica does not meet minimum performance criteria. The campus replica may contain only a subset of the objects in the enterprise directory as required by that particular campus.
- **NOS directory.** This is a directory that comes with a network operating system (NOS) for administering things such as user workstations and devices from a network rather than from a host perspective. Examples include NT domain controllers, Novell Directory Service (NDS) in a NetWare environment, and Sun NIS. They often contain people information, in particular, the users of the systems supported by the particular NOS. This can serve as an authoritative source of information for other directories such as command, claimancy, and enterprise.
- **E-mail address book.** Most e-mail systems include an address book function. Often these address books are server-based and contain many names and e-mail addresses of the e-mail system users within a given organization or beyond. This can serve as an authoritative source for e-mail address information for some sets of users.
- **LDAP.** The lightweight directory access protocol. This is the standards-based network protocol used to communicate with the directory. In the client/server model, this is the protocol that the client uses to talk to the directory service.
- **Replication.** Describes the process by which a directory, or a portion of a directory, is “cloned” elsewhere, typically on another server. This allows multiple directories containing copies of some master directory to remain up-to-date. Replicating directories allows scaling to support more users and allows the directory service distribution point to be closer to the user for performance reasons.
- **Synchronization.** Describes the process by which multiple directories containing overlapping information can exchange directory information and remain synchronized.
- **The Directory Information Tree (DIT).** This is the logical hierarchy in which the directory information (objects) is organized.
- **Schema.** Defines the rules, naming conventions, and structure of information in a given directory entry.
- **O=, OU=.** These are X.500 abbreviations used to give names to the levels in a DIT structure. “O” stands for organization. “OU” stands for organizational unit.
- **DISA PKI directory.** This is an instance of an enterprise directory within DoD. It exists for the purpose of supporting the DISA PKI initiative and is where the PKI certificates are stored and from where they should be imported into the Naval Enterprise Directory.
- **DMS directory.** The Defense Message System is a system for performing organization-based message traffic within DoD. DMS has its own directory and is based on X.500 standards.

The directory architecture will be based on a client/server model. Users will operate client applications that obtain directory information by accessing the enterprise directory (server). Two forms of access will

be provided. The “native” directory protocol LDAP (version 3) will be the primary and preferred means for accessing the directory. For situations in which a function cannot be performed using LDAP or in which an LDAP client is not available, a web browser-style interface will be used.

The enterprise directory will be centrally managed with a distributed implementation. The distributed implementation allows increased scalability, reliability, and performance by locating the service close to the end user or client application. “Close” is a relative term and may imply “same hemisphere”, “same region”, or “same campus” as dictated by network paths and other considerations. It must be centrally managed to provide a consistent view across the enterprise.

The distributed components of the enterprise directory will be replicas of a central master directory.

The enterprise directory will leverage efforts of the major claimants by using claimant directories as authoritative sources. Claimants will be encouraged to begin or continue internal directory efforts as appropriate.

### 4.3.5.1 Logical Architecture

A challenge for the architecture is the poor alignment between the X.500 data structures (schema, Directory Information Tree (DIT)) and the real world requirements. The DIT dictates the boundaries for replication, control, logical structure, search root, and others. In any Naval implementation, the boundaries must be established based on the organizational entities that control the attributes. So, the DIT structure is a compromise that best meets requirements.

The DON directory will use the DIT hierarchy depicted in Figure 4-1.

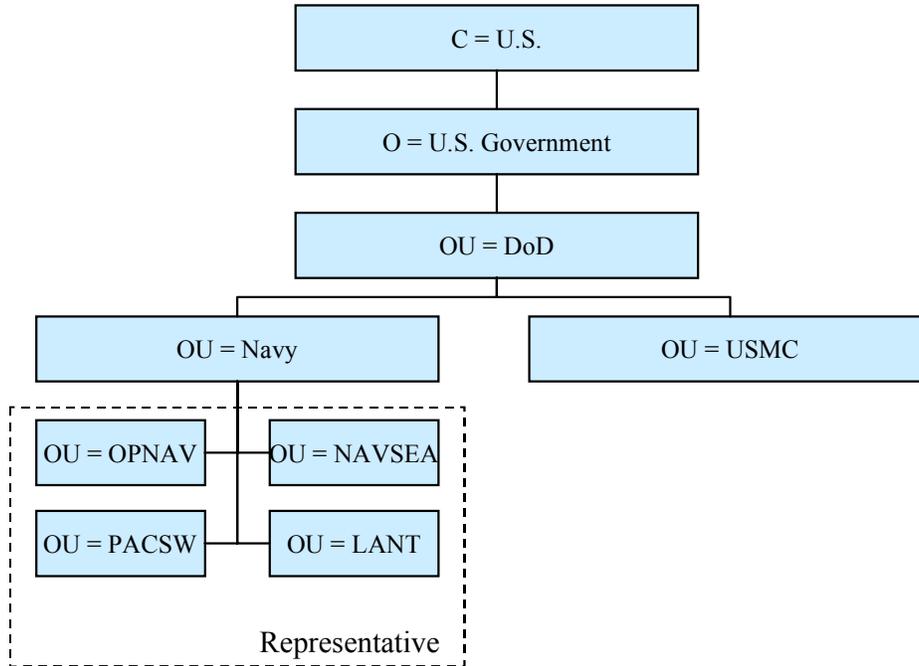


Figure 4-1. DIT Hierarchy

The DIT hierarchy is similar to the DMS structure except that there is no “OU=Organizations” or “OU=Locations” level. The level under “OU=Navy” applies to echelon 1 and 2 commands (major claimants). Below this level, the DIT structure is implemented at the discretion of the particular organization. It is recommended at the lower levels (below major claimant level) that the structure categorize people under “OU=People”, groups of people categorized as “OU=Groups”, and the internal structure of the organization under “OU=Organization”. Other conventions will be established in the future.

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

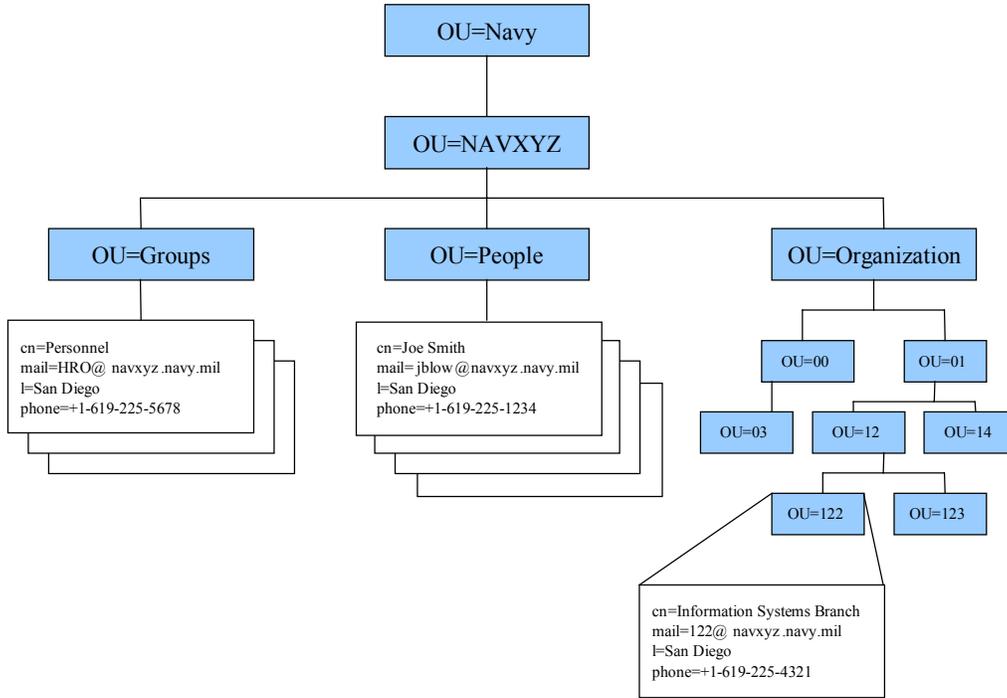


Figure 4-2. Example of a DIT structure below claimancy level

At the same level of the major claimants, there will be OU entries for the various Naval regions such as “OU=PACSW”. This will accommodate regional specific information that does not necessarily need to be replicated globally.

Every object stored in the directory has a name that uniquely identifies it within the directory. The naming scheme is analogous to fully-qualified filenames in a file system in which the name of the file and the whole path must be identified. In a directory this name is called the “Distinguished Name”, or simply “DN”. The DN for any object will include a valid attribute for the object itself plus the entire path within the DIT to this object.

With the structure outlined in Figure 4-2, all “person” objects in this directory will have the same “path” from the root of the hierarchy down to the “People” level. Accordingly, there must be some attribute in each “person” object that identifies that object uniquely at that level. The answer is the common name (CN) attribute, which is created from a combination of the person’s name plus his/her e-mail address. This CN can be used to uniquely identify an entry when combined with the DIT path and thus fully qualify the object name. The format of this special CN attribute is as follows:

```
CN=John Q. Public <jqp@somecommand.navy.mil>
```

The information inside the “<” “>” characters is the user’s official e-mail address. (Note that this is a valid RFC-822 format, and is also the PGP key rings format used for publishing PGP directory keys.)

The fully-qualified DN includes the CN and other attributes that fully qualify the CN relationship in the hierarchy. The fully-qualified DN is illustrated in the following:

```
CN=John Q. Public <jqp@somecommand.navy.mil>, OU=SomeCommand,  
OU=Navy, OU=DoD, OU=U.S. Government, C=US
```

It should be noted that this DN scheme is different from the DMS scheme. The DMS scheme generates a unique identifier for each user and includes that in the CN. The DON scheme proposed does not require a unique identifier.

Within the DON directory naming scheme, the attributes that must be populated for each directory entry are described in Figure 4-3.

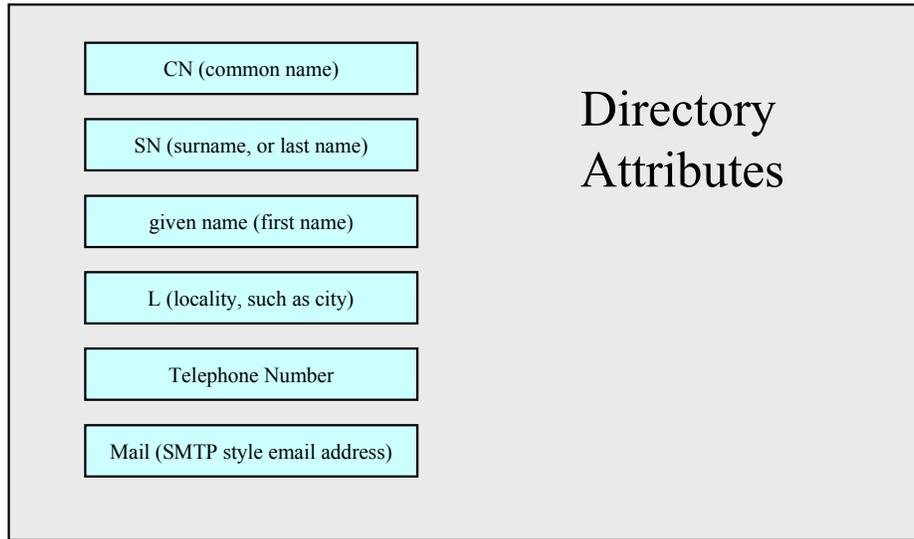


Figure 4-3. Directory Attributes

The CN attribute is multi-valued. Each form of a person's name should be included in the CN attribute to accept queries by all potential searches for that person's name:

CN: William A. Jones

CN: William Jones

CN: Bill Jones

It is envisioned that the major claimants will maintain directories and these will have significant additional attributes to meet local requirements. The ones listed above represent the subset of attributes that are of enterprise-wide interest. The list of enterprise attributes may grow over time as additional enterprise requirements are identified. As additional attributes are incorporated into the enterprise directory, ownership of those attributes needs to be delegated to the authoritative source(s) for those attributes.

If the enterprise directory evolves to the point where unique data is stored in the directory (that is, there is no separate authoritative source), then those attributes will need the appropriate access controls to

facilitate delegation of control to an individual(s) or group that can manage the information in a timely and accurate manner.

#### 4.3.5.2 Physical architecture

The following describes the specific physical implementation of the directory architecture.

- Directory users should be able to point to multiple local redundant directories (local may be on-base or within the region) consistent with user demand and available regional bandwidth. A large campus may want its own directory (or multiple directories) for performance and reliability reasons (similar to multiple DNS servers to provide redundancy).
- A campus replica might include only a sub-tree of the enterprise directory (for example, only the campus's claimancy) because that is the primary information the campus users need to access.
- Each region (fleet concentration area) should have a fully-populated replica of the enterprise directory.
- Replication masters which, in aggregate, fully define a region, should input to the regional directories. There should be multiple masters, such as east coast, west coast, PAC, EUR, and CENT, that replicate down to the 15 regional directories.
- This directory implementation will include a global service that is replicated to the regions. Users can query the regional directory to perform searches and other functions. Major commands that maintain their own directories feed attributes "up" to this global (meta-) directory. The regional directories can replicate only a user-required sub-tree or the whole directory down to a campus, or the regional directory can provide content to a directory administered by a local command or major claimant.

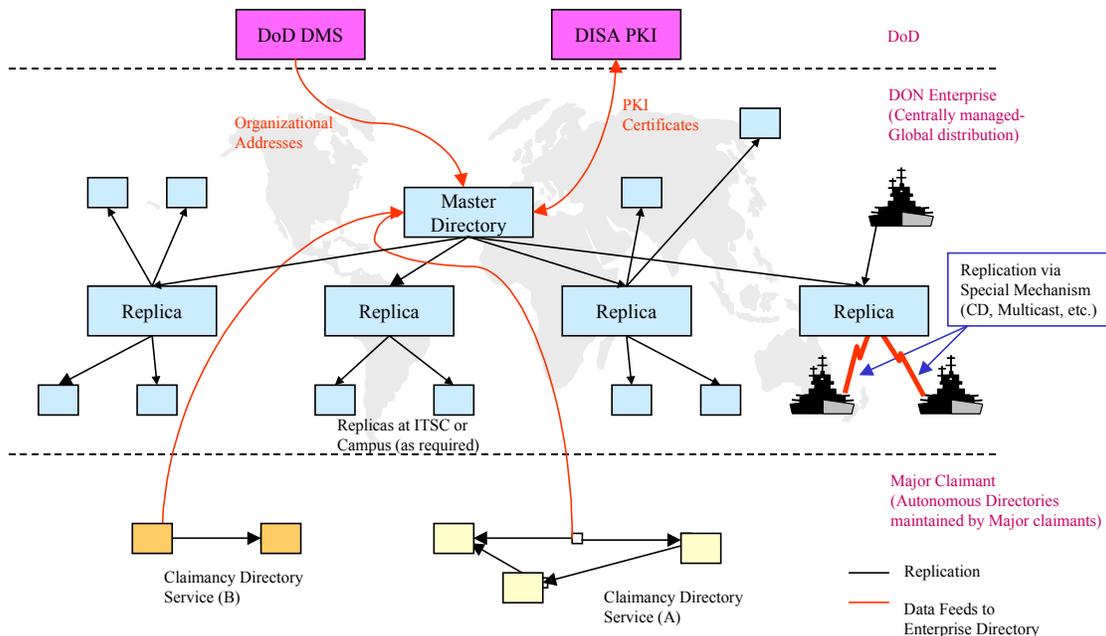


Figure 4-1. Naval Enterprise Directory Physical Architecture

The Naval enterprise directory consists of a three-layer hierarchy. The middle layer as depicted in Figure 4-1 represents the Naval-wide global directory and contains all Naval personnel and associated

individuals. Above the Naval-wide global layer is the rest of DoD. At the layer below are the directories maintained by the major claimants in the DON.

The replication process to the ships will be a modified process which prioritizes the updates by segment and by attribute and makes use of broadcast technologies besides the normal IP connections.

Each major claimant will maintain its own directory (if not now, then certainly in the future). The claimant directories will typically include more information (attributes) than the enterprise level. The loose coupling between the claimant directories and the enterprise directory allows for flexibility and claimancy control. There is a normalization or translation process that feeds the enterprise portions of the claimant directory information up to the enterprise directory. In like manner, the Naval enterprise directory will be coupled to the DoD directory and, with required transformations, provide service input. The DoD directory will import some subset of information from the enterprise directory, and the enterprise directory will import PKI certificates from the DoD directory.

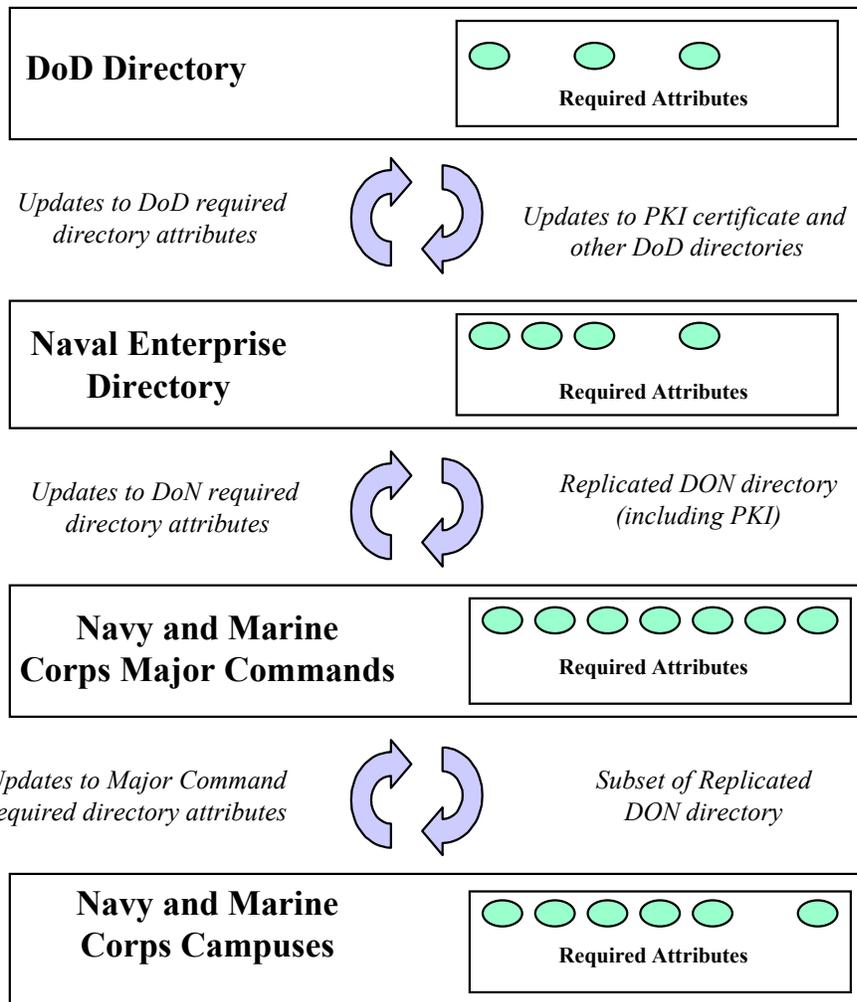


Figure 4-2. Hierarchical Exchange of Directory Attributes

Figure 4-2 shows the replication and normalization of the directory attributes across the multiple information hierarchies. The figure emphasizes the differences in each hierarchy's attribute requirements.

There may be an additional requirement to feed attributes down from the enterprise directory to the claimant directories, but this can be determined on a case-by-case basis because the claimant directory requirements are expected to vary widely.

#### 4.3.5.2.1 Alternative Strategies for Campus Directories

A number of strategies exist for maintaining a campus directory, and the appropriateness of these depends on the individual requirements of the tenant organizations resident on the campus. Figure 4-1 depicts some principal sources of directory information and the following three cases illustrate them.

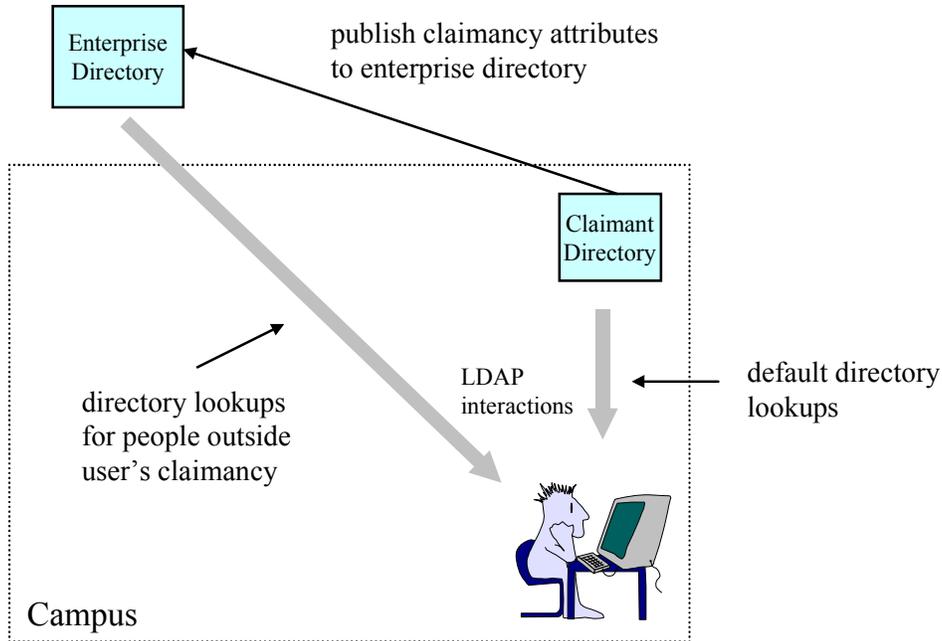


Figure 4-1. Potential Sources of Directory Information

**Case 1.** The first in Figure 4-1 shows a user that normally gets directory information from his own claimancy directory, which is located on the user's campus (although it could be located almost anywhere). In this case, most directory queries are for information that is within that user's own claimancy.

When the user needs to look for information in other claimancies, he or she goes to the enterprise directory for the information by manually selecting a part of the directory hierarchy that references the enterprise directory, or through a referral process that is implemented in the claimancy directory.

If performance of the enterprise directory is not acceptable, it is possible to obtain an on-campus replica of the enterprise directory (see next example).

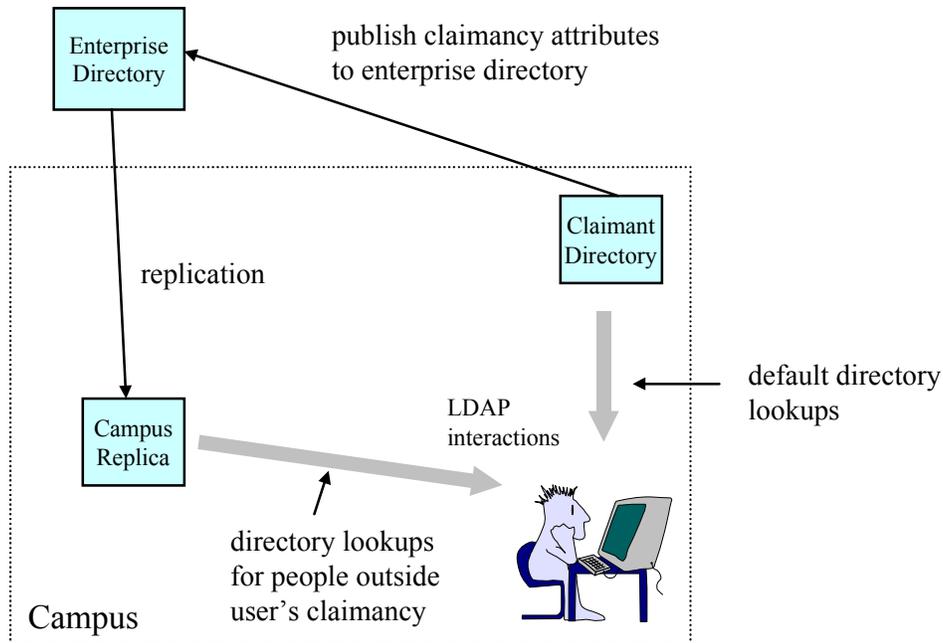


Figure 4-2. Potential Sources for Directory Information

**Case 2.** This situation is like the previous one except that here there is an on-campus replica of the enterprise directory. It contains the same information as the enterprise directory, but it is physically much closer, and will likely provide better performance.

Because the directory is local, lookups outside the user's claimancy will be performed against this directory.

The idea of merging the two on-campus directories is a possibility if they have consistent DIT structures.

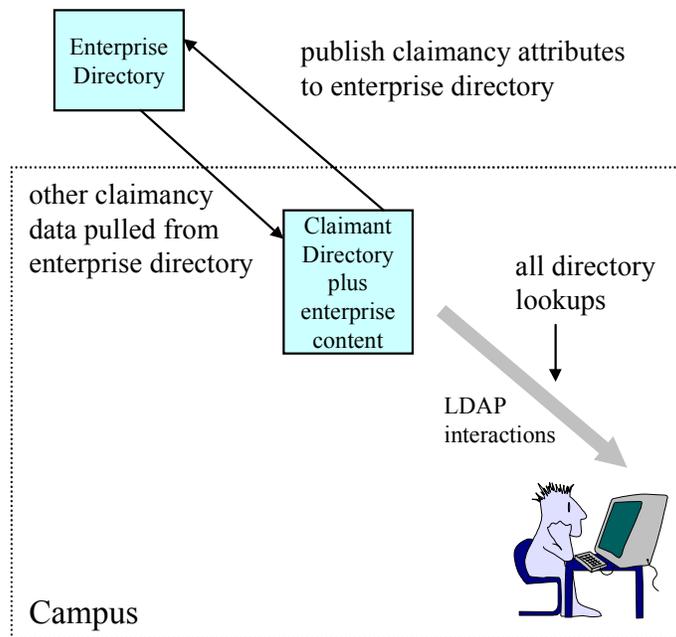


Figure 4-3. Potential Sources for Directory Information

**Case 3.** This situation in Figure 4-3 is similar to the previous case except the enterprise directory information is now merged into the claimancy directory. The information is “pulled” from the enterprise directory as required. There is no directory “join” process necessary because additional claimancies are added to a directory tree that is already established, assuming the claimancy directory conforms to the enterprise directory structure.

The claimancy still “owns” and manages the claimancy directory. From the enterprise perspective, it is not a true replica because the enterprise is not controlling the replication process, but that poses no problem for the claimancy. The claimancy has achieved a locally-controlled directory that is richly populated and still under claimancy control.

From the user perspective, they can find all DON information in a single directory. This may not be immediately apparent because they still have to navigate the DIT to select the scope of searches.

#### 4.3.5.2.2 Alternatives for Enterprise Directory Replication

The means by which the campus receives replication in Case 3 above depends on the required information and the performance of the network connecting the campus.

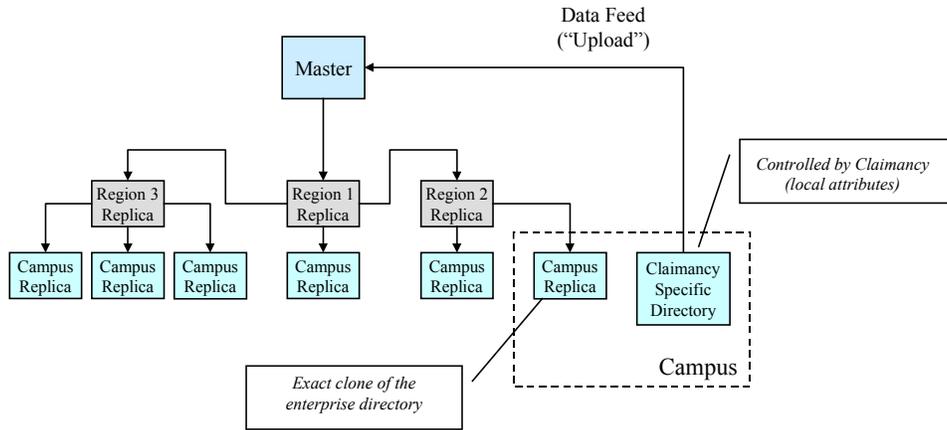


Figure 4-1. Directory Replication

Figure 4-1 shows the data feed from the campus to the enterprise directory. The corresponding feed from the enterprise to the campus should be tailored to the situation. If the updates are strictly controlled, such as a directory data pull situation from shipboard, the specific updates are obtained as required by the campus. Where frequent update is required and bandwidth supports it, a data push from the enterprise directory down to the claimancy or campus directory may be more appropriate. There are accompanying security implications here and these are supported by certificates to ensure integrity of the data.

#### 4.3.5.2.3 Naval Enterprise Directory Service Scalable Architecture

This architecture is very scalable and can grow in phases as demand increases. Initially it could be a single server and a backup for redundancy. Later, replicas could be distributed to other major locations and these replicas could feed other regional or campus replicas. The emphasis must be on making a fully populated Naval Enterprise Directory Service available to all Navy and Marine Corps organizations.

The rationale for locating replicas is based on a number of factors. The first and foremost is to provide reliable access for all users with response time (performance) measured in only a few seconds. Factors that could potentially contribute to poor performance are overloaded servers (support of too many simultaneous clients), congested networks, network outages, low bandwidth links, and large round-trip delays in the network. Factors that will contribute to good server performance are selecting a combination of software and hardware that can support 1 million entries, few-second response times, and a large population (tens of thousands) of users. To ensure that faulty networks do not degrade performance, replicas should be located near the users.

The replication strategy includes providing replication to all 15 or so fleet concentration areas, plus to the many campus networks or remote sites. Considerations for location of replicas include cost, management issues, and possibly political factors.

#### 4.3.5.2.4 Other Areas for Development

For reliability, the DNS name referring to the directory will have multiple "A" records. The DNS will return the names in round-robin fashion if load-balancing is desired or will return "primary first" in the case of a primary and backup server.

As regional ITSCs are established and IT support becomes less command-centric and more region-centric, new directories will be implemented in support of various functions and applications that are region-specific. Considerable thought is required to determine how best to evolve into an integrated directory solution that meets both the enterprise needs and the regional needs. For the near term, this will be done with multiple physical directories, with some loose coupling to the enterprise directory for attributes that should be shared. Designers and implementers are strongly encouraged to follow the enterprise DIT and schema to the extent possible – to facilitate long term migration to an integrated solution. Especially important is the need to organize people information by claimancy rather than by region, because the capability to search by claimancy should be preserved. It is acceptable to put people information under the regional part of the tree, but it must be associated with the proper claimancy.

When it is necessary to look up a name in the enterprise directory, there must be a means to limit the scope of the search. If the user searches for a “Smith” in the total Navy population, there will be too many hits, and it will be difficult to search through the results to find the right person. Alternatively, the major claimant that person belongs to will be known, for example, NAVSEA, and the search can be conducted for all the Smiths located in NAVSEA. In fact, for any search in the directory, it is appropriate to have some initial selector that specifies the claimancy in which the person is assigned. Also, there will normally be an LDAP client default to the user’s own claimancy. In the same way that you have multiple phone books sitting on your shelf, each one for a different organization, you will want a selector in your client interface that allows you to choose which claimancy to search in, or to search all of Navy. This is possible with modern e-mail/LDAP clients as long as the directory is organized in a way that supports this structure.

#### **4.3.5.2.4.1 Blue Pages**

The directory architecture includes a description of the DON “blue pages.” The directory should have a representation of the organizational hierarchy of each major claimant. Each node in that hierarchy includes the following:

- Name of that part of the organization
- Organizational code
- Office phone number and fax number
- Head (or acting head) of that organization
- E-mail address of the organization
- List of subordinate organizations

Each node will contain a list of the employees at that level in the organization. A web-based tool will be used to “drill down” into the organization in order to find the office of interest within the organization. (The challenge is to get accurate organizational structures and attributes from the major claimants.) The appropriate DIT and schema for this type of information are contained in Figure 4-1 above. These DITs provide a repository for this information as well as provide an excellent tool for the major claimants to maintain their organizational structure information. (This requires that good tools be developed so that the claimants can easily maintain their portion of the enterprise directory.)

#### **4.3.5.2.4.2 Security Issues**

There are security implications that accompany decisions about the information that should be displayed in the directory. A directory database that is rich enough to be useful also contains the same kind of

information that has been the subject of contention on DoD WWW pages. The appropriate directory security policy is not to remove the information but to restrict its access to authorized users and to make portions of the database available only to certain classes of users. For example, individual e-mail addresses should be a widely-distributed application, but conversely, emergency contact information contains non-public items such as home addresses and phone numbers, and should be screened from all external users.

One of the more important directory functions is to make public keys (x.509 certificates) available to everyone. Equally important is providing traceable authenticity -- if a means for spoofing the public keys is exposed, then authenticity is in question. An authenticity trail must go up the directory tree from the point where the public key is manufactured to the core database and back to anyone who needs to use it.

## **Regional Issues and Considerations**

Regional ITSCs will operate a replica of the enterprise directory.

Regional ITSCs will manage the region-specific portions of the directory.

## **Campus and Operational Node Issues and Considerations**

Local commands must input accurate information to their appropriate claimancy directory. This includes an authentic means for communicating x.509 certificate data from units to the ITSC.

## **Deployed Forces Issues and Considerations**

Deployed forces are constrained by limited and sometimes unavailable bandwidth. Directory synchronization across bandwidth-constrained boundaries must be carefully considered and avoided where possible. However, many such units need access to directory information. The following questions emerge: How much of the global directory needs to be immediately available to deployed units? How accurate and timely does the information need to be? To limit loading over the slow communications links, directory updates should be restricted to those objects and attributes that are most important (e-mail address, certificate revocation lists) and should avoid updating those that are less important (phone number, location, etc.). The full directories of regions where a deployed unit expects to operate should be synchronized prior to deployment. Subsequently, only the changes to the directory should be communicated until return to home port.

There are a number of options for distributing updates to the deployed forces. Initial loading of the directory or major updates can be distributed on a CD-ROM delivered to the ship or via a direct connection when the ship is connected at pier side. Another option is to distribute updates via reliable multi-cast over direct broadcast satellite. Yet another option would be for the NCTAMS to prepare updates of limited scope for delivery over RF links at times when bandwidth is available. Any further design discussion is beyond the scope of this document.

### **4.3.6 Roles and Responsibilities**

The ITSC should operate the regional replicas.

Major commands will need to implement and maintain their own directories and feed that information up to the global directory. Note that SPAWAR, NAVSEA, and NAVAIR already operate their own corporate directories.

DON CIO needs to encourage claimants to establish claimancy-wide LDAP capable directories that are highly accurate and provide the necessary “feeds” to the enterprise directory.

A DON level person or group (i.e. “Directory Architect”) should be established to be the arbiter of attribute and DIT conventions. These should be documented at the Navy NIC.

An interface needs to be established with the DoD level directory initiatives so that we can leverage off each other’s efforts. We will be able to offer content to the DoD level directories (they leverage our efforts) and in return will want to import attributes from the DoD level directories such as PKI certificates. This needs to be well thought-out and coordinated.

An overall “Directory Administrator” needs to oversee and administer the operational aspects of the directory.

## **4.4 Electronic Mail**

### **4.4.1 Service Description**

Electronic mail (e-mail) is the basic service for interpersonal and organizational messaging for use throughout the DON. It employs store-and-forward technology that does not provide real-time information exchange but does provide the capability for high-speed communication of small messages or file transfer. It is not a substitute for bulk transmission of large data files to data processing centers or between application servers requiring data replication.

The components of interpersonal e-mail service include the following:

- a user agent for submitting and retrieving messages
- a message store for temporary storage of messages pending delivery or receipt
- a message transfer agent for the reliable transmission through the store-and-forward messaging application network (which uses the DON enterprise telecommunication network)
- an addressing and routing scheme

The architecture that defines the directory components for search and retrieval of recipient e-mail addresses is provided in a separate services section.

All of these components interact to meet the messaging needs of the DON. The service architecture subsection describes the e-mail components, their configuration and interaction, and provides guidance for planning and implementation.

This e-mail service is distinct from the messaging service provided by Defense Messaging System (DMS), whose primary purpose is to process record message traffic. While the e-mail service described here applies to all Naval users, DMS is intended for use by a small segment of the DON with record message traffic requirements. Interoperability between e-mail service and DMS will be addressed in this architecture document so that the systems will be interoperable to the extent required by users. Architecture guidance for DMS is addressed in separate documentation.

#### **4.4.2 Applicable Standards, Policy, and Guidance**

- Section 6.2 of the ITSG document
- Defense Message System (DMS) Recommended System Design Architecture (SDA) Document Release 1.1 (Initial)
- Applicable RFCs:
  - ◆ RFC821 – Simple Mail Transfer Protocol
  - ◆ RFC822 – Standard for the format of ARPA Internet text messages
  - ◆ RFC974 – Mail routing and the domain system
  - ◆ RFC1651 – SMTP Service Extensions
  - ◆ RFC1652 – SMTP Service Extension for 8bit-MIME transport
  - ◆ RFC1653 – SMTP Service Extension for Message Size Declaration
  - ◆ RFC1731 – IMAP Authentication Mechanisms
  - ◆ RFC1891 – SMTP Service Extension for Delivery Status Notifications
  - ◆ RFC1892 – The Multi-part/Report Content Type for the Reporting of Mail System Administrative Messages
  - ◆ RFC1893 – Enhanced Mail System Status Codes
  - ◆ RFC1894 – An Extensible Message Format for Delivery Status Notifications
  - ◆ RFC2045 – MIME Part One: Format of Internet Message Bodies
  - ◆ RFC2046 – MIME Part Two: Media Types
  - ◆ RFC2047 – MIME Part Three: Message Header Extensions for Non-ASCII Text
  - ◆ RFC2048 – MIME Part Four: Registration Procedures
  - ◆ RFC2049 – MIME Part Five: MIME Part Five: Conformance Criteria and Examples
  - ◆ RFC2060 – Internet Message Access Protocol (IMAP) – Version 4 rev 1
  - ◆ RFC2110 – MIME E-mail encapsulation of Aggregate Documents, such as HTML (MHTML)
  - ◆ RFC2298 – An Extensible Message Format for Message Disposition Notifications (MDN)
  - ◆ RFC2311 – S/MIME Version 2 Message Specification
  - ◆ RFC2312 – S/MIME Version 2 Certificate Handling
  - ◆ RFC2342 – IMAP4 Name space
  - ◆ RFC2425 – A MIME Content-Type for Directory Information
  - ◆ RFC2426 – vCard MIME Directory Profile

### **4.4.3 Requirements**

The following list of requirements is by no means exhaustive, in that the full list of capabilities and functionality available in today's leading e-mail products is assumed and not repeated here. However, this list highlights specific additional requirements of the Naval implementation which this architecture must support.

The e-mail system must comply with industry standards and remain consistent with those standards as they evolve. Important examples of these include Simple Mail Transfer Protocol (SMTP), Multipurpose Internet Mail Extension (MIME), Secure Multipurpose Internet Mail Extension (S/MIME), Post Office Protocol Version 3 (POP3), Internet Message Access Protocol Version 4 (IMAP4), and Delivery Status Notification (DSN). This is to ensure a high degree of interoperability with e-mail systems external to the DON.

The e-mail system must be compatible with and use the Naval Public Key Infrastructure (PKI) (section 3.8) for S/MINE and other certificate-based security mechanisms.

All user agents must be able to read and write messages in MIME format. Once a message is in MIME format, no component of this enterprise e-mail infrastructure should modify the message body during any part of the transport of that message or as it is stored on any mail store. The goal is to minimize content loss (data or format) through gateways into proprietary environments.

The e-mail infrastructure must support attachments of up to 10 Mbytes in size and be able to reject messages that exceed this size. Bandwidth-constrained environments are exempted from this requirement and will almost certainly need to restrict message sizes to some smaller value. User agents should be able to reject downloading of messages that exceed the user-specified size.

The e-mail system must provide adequate storage for user mailboxes consistent with user requirements. It should include a per-user quota mechanism to automatically manage the available storage space. It is recommended that the quota be set to 50 MB per user.

The e-mail infrastructure should be web-enabled so that users have the capability to get their e-mail with nothing more than a web browser in circumstances where no other client tool is available.

The e-mail infrastructure must filter out unsolicited commercial e-mail ("SPAM" e-mail) as much as possible.

Users should be able to choose an address style that is organization-independent so that the address can be maintained as users undergo reassignment.

### **4.4.4 Assumptions**

- Between DON users in separate organizations, SMTP will be the primary message transport system for interpersonal e-mail.
- Between DON users within the same organization, this architecture is indifferent to the e-mail system used.

- Defense Messaging System (DMS) will be used initially for organizational record message traffic only.
- Interpersonal e-mail does not require the enhanced security mechanisms (high assurance) offered by DMS.
- The e-mail infrastructure described here is a basic utility service and is highly standards-based, as is the case of other basic infrastructure components. It is not the only e-mail service provided within the DON. It is assumed that ITSCs will offer e-mail services that are integrated with groupware environments such as Microsoft Exchange or Lotus Notes for those users that require those environments. The basic e-mail utility described here is available to all DON users, while the integrated or proprietary e-mail environments will be offered only where required.

#### 4.4.5 Service Architecture

The primary architectural components of the e-mail system include a message store (mail server), an SMTP relay (mail transfer agent or MTA), user agents (UA), and a directory providing “white pages” service. It also includes the protocols required to allow interaction among the various components -- SMTP between the MTAs, IMAP4 and POP3 between UAs and the mail server, and Lightweight Directory Access Protocol (LDAP) for accessing the directory. The e-mail service architecture accommodates interpersonal as well as legacy DoD organizational e-mail, which is currently being migrated from AUTODIN to DMS throughout the DoD. Because DoD defined organizational messaging for the services, it is described here only to the extent necessary to clarify the relationship between the DON service architecture for interpersonal e-mail and the DoD architecture for organizational messaging.

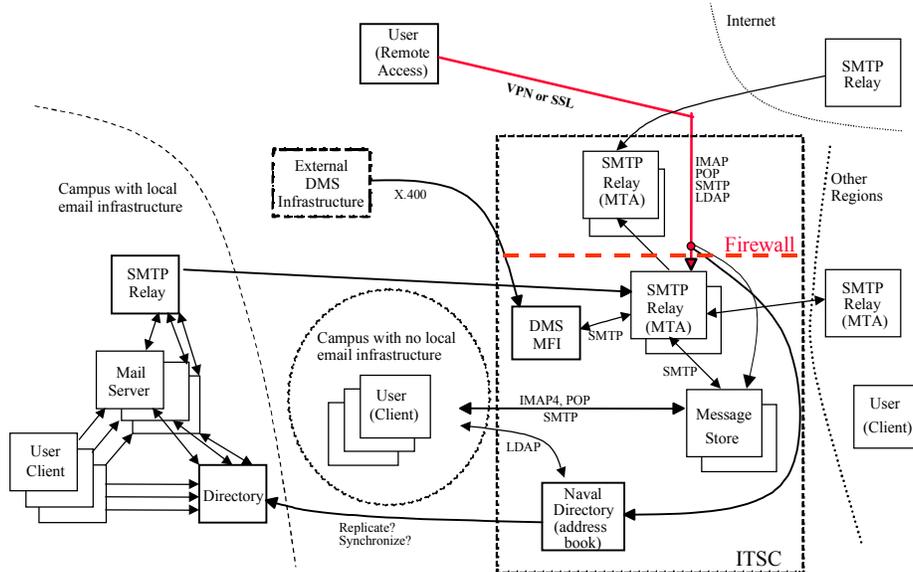


Figure 4-1. Enterprise E-mail Architecture

Figure 4-1 represents the top-level e-mail architecture that shows the interface between the enterprise, the ITSCs, and various user communities. Campuses with and without local e-mail infrastructures are both supported seamlessly. It also shows the interface between the enterprise and external networks. In this

regard, network security is provided by various firewalls in accordance with the Defense in Depth strategy.

#### **4.4.6 Addressing Conventions**

The addressing conventions used in the e-mail service must first be compatible with RFC-822 and SMTP addressing standards in which the machine-readable address uses the style user@domain. Within the DON, the right-hand side of the “@” sign must be one of the following: navy.mil, usmc.mil, or organization.navy.mil, where “organization” is the name of the organization within navy (e.g. cpf.navy.mil). Sub-domains below the level of organization.navy.mil are possible, but not encouraged. The goal is to keep things flat below the navy.mil level.

For users who require a permanent, unchanging e-mail address despite career/organizational changes (e.g., military personnel), he or she will have an e-mail address in the navy.mil domain. For users who generally stay with a single organization (e.g., civilian personnel), they should have an organization-based address. In both cases, it is imperative that “user” names on the left-hand side of the “@” sign be unique within a given domain.

In the navy.mil domain, the convention for names on the left hand side of the “@” sign will be Firstname.Lastname with the flexibility to allow for conflict resolution or other personal preferences. The rule will be “first come, first served”. This name should not include rank or other designations that are likely to change.

In the organization.navy.mil domains, the conventions for names on the left hand side of the “@” sign will be left to the organizations’ discretion. It is recommended that login names be used in this case because they already need to be unique within that domain. These should be assigned on a first come, first served basis.

All DON e-mail addresses must be registered in the Naval Enterprise Directory (see section 4.3). For each user in the directory, the common name attribute should reflect all names that he or she goes by (e.g. “William F. Smith” and “Bill Smith”) so that directory searches can successfully match all possibilities.

For display names on the “From:” header of actual e-mail messages, the RFC-822 specification allows a number of conventions. Within the DON, the following format will be used:

From: Bill Smith <smithb@clf.navy.mil>

The address inside the “<” “>” characters must be the officially registered e-mail address for the user, and should not contain the mail server name:

(wrong)      From: Bill Smith <smithb@mailserver3.clf.navy.mil>

Rank or other titles in the display name are acceptable:

From: Capt James T Kirk <Jim.Kirk@navy.mil>

Characters in the display name (such as commas and periods) that force one to quote the name are discouraged, but will be supported:

(discouraged)    From: “Smith, William F.” <smithb@clf.navy.mil>

## 4.4.7 Routing Architecture

The DON e-mail routing architecture is concerned with proper routing and delivery of messages. Issues for consideration are reliability and/or redundancy, security, relationship to DNS, and the processing of errors.

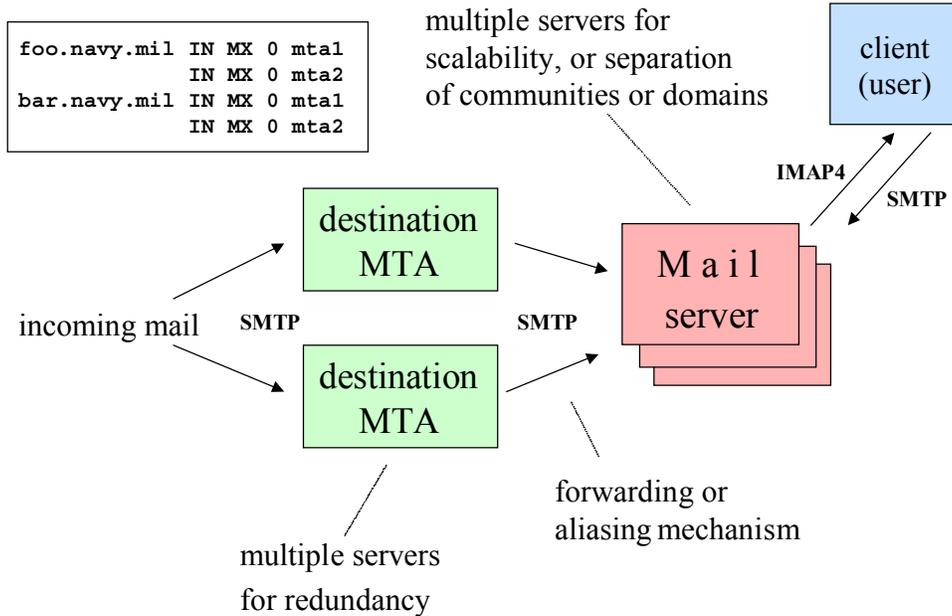


Figure 4-1. E-mail Routing Architecture

For purposes of redundancy and reliability, each e-mail domain will include a minimum of two SMTP servers for receiving e-mail as Figure 4-1 indicates. The DNS will be configured so that all such servers are equally capable of receiving and forwarding inbound e-mail to the destination mail server. Mail Exchanger (MX) records for the e-mail domain must be configured in the DNS and will point at the respective MTAs. The MTAs will be configured so that one can be taken off line with no disruption in service other than a possible degradation in performance.

The destination MTAs and mail servers must be able to support multiple domains simultaneously. The destination MTA forwards the e-mail to the appropriate mail server based on the SMTP destination address of the message. In the process, it hides the internal structure of the e-mail service (i.e. names of mail servers). This allows for a rich and potentially complex internal structure while supporting a simple addressing structure externally. This implies that within a domain, every mail server must know the correct mail server for every user in that domain; accordingly, the server can route e-mail directly to its proper destination without the originator having to specify that in the message headers. A typical solution of this concept is to distribute alias files (mapping to mail server for every user in the domain) on a daily basis to all the mail servers in a domain. A more leading-edge solution is for mail servers to base their routing information on a directory (see direction section) and learn from that central information source which mail server supports a given user. This latter approach is more instantaneous than the alias distribution approach previously described.

A firewall may separate the MTA function into multiple sub-functions. The destination MTA will be configured to disallow relaying between external domains in order to minimize its use in relaying spam.

A special device known as a "Secure Mail Guard", or SMG, is used to pass e-mail between unclassified and secret networks. They are very restrictive and only allow carefully formatted messages to pass and only between certain source/destination combinations. Attachments generally cannot pass through an SMG. An SMG will be installed at the ITSC in support of this limited capability.

A function that might be implemented in the firewall as part of the e-mail proxy services is a dirty-word search mechanism to allow early detection of classified e-mails leaking to the unclassified side due to human error or lapse of due care.

Users operating an e-mail client will use the IMAP4 protocol to retrieve their mail from the mail server and will use SMTP to send mail via the mail server.

#### **4.4.8 User Interface**

The e-mail system must accommodate multiple modes of access to e-mail by individual messaging users. This access can range from direct, on-line access in a LAN or WAN environment to off-line access used when users are unable to establish a network connection (e.g., in flight) and to remote access from the Internet using a web browser on public access workstations (similar to those found in public libraries). The Internet message access protocol (IMAP) defines a client-server standard by which vendors can develop products with the required functionality. Microsoft Exchange, Netscape Messenger, and Lotus Domino are examples of server products that can be configured as IMAP servers. Netscape Communicator is a popular example of an e-mail client that supports IMAP and LDAP.

The e-mail client must be able to work seamlessly with a wide variety of attachment types, including those that are officially registered with Internet Assigned Numbers Authority (IANA) and those that are de facto standards. The user agent must be able to launch the correct application to interpret or display the attachment based on file name or content type. This must include various image types (e.g., gif, jpeg, mpeg, tiff, bmp), document types (e.g., pdf, postscript, rtf), packaging types (e.g., zip, tar), and business application types (e.g., Word, Excel, Powerpoint). At a minimum, plain ASCII, HTML, and RTF must be fully supported in all DON e-mail systems. For other standards, refer to the ITSG.

### **Regional Issues and Considerations**

The mail service provided within a region must scale to accommodate hundreds of domains and hundreds of thousands of users. Domains that are administered within a region must be cognizant of the requirement for uniqueness of user IDs within a given domain.

### **Campus and Operational Node Issues and Considerations**

For any number of reasons, a campus could choose to operate its own mail server(s). It still may choose to use the regional MTAs. The campus network will need to coordinate the mail-forwarding function with the regional provider in this case.

## **Deployed Forces Issues and Considerations**

There are some unique problems with respect to e-mail delivery and routing to deployed forces. The major problem is limited bandwidth. Because of this, restrictions on size of messages need to be applied. Also, message prioritization is a necessary capability to make sure that a large low-priority message does not get in the way of delivering a small urgent message.

Afloat platforms are not always connected. Therefore, a destination MTA must be operated ashore to accept and queue mail for eventual delivery to the afloat platform. The afloat platform should then perform a selective “pull” of queued e-mail when communications are restored.

There are many other issues and details here (in relation to DNS, firewalls, etc.) that are beyond the scope of this document.

### **4.4.9 Roles and Responsibilities**

The DNS administrator is responsible for properly registering the domain and MX records for the MTA host.

The ITSC is responsible for operating a regional e-mail service as outlined above.

One of the ITSCs must be responsible for support of the “navy.mil” domain and should consider using the MTAs of another ITSC for reliability/redundancy of that domain.

E-mail addresses must all be registered in “the directory”.

The directory must provide a “white pages” service to the e-mail clients.

A registration desk must enable users to sign up for e-mail. There is one for each organization and one for the “navy.mil” level.

## **4.5 Network News Service using NNTP**

### **4.5.1 Service Description**

The Network News Service is an information distribution service with which users can selectively gain access to and watch “netnews.” It is different than an e-mail subscription in that the information content is not delivered directly to user mailboxes. That approach does not scale well. With NNTP, all content is stored on a “news server” and replicated to the degree necessary to provide reasonable local access and performance. Client tools “pull” this content from the news servers and make it available to the users for presentation on demand, based on the user’s particular selection of “news groups.” News clients typically can be configured to present only new messages to the user and to organize it by discussion threads within a news group. It also allows contributing to a discussion through “posting” of messages to a news group.

This service is strongly recommended for mass distribution of information or discussion content, especially when the subscriber base is very large or dynamic. NNTP meets the mass distribution requirement much better than the e-mail system because each news article is stored only on a news server and is not duplicated for each recipient.

Discussion groups are organized in a hierarchy. For example, a news group that discusses the NT operating system would be found under “Microsoft systems” under the “computer” category, and would be named comp.sys.microsoft.nt. News groups specific to Navy issues could also be created, e.g. navy.doncio.ipt.iti.

Network news communications between servers and between clients and servers uses the Network News Transport Protocol (NNTP).

## **4.5.2 Applicable Standards, Policy, and Guidance**

ITSG Section 6.8 Network News Transport Protocol

RFC 977 Network News Transfer Protocol

## **4.5.3 Requirements**

Must provide access to the public Internet news groups, of which there are many. May also be required to filter out newsgroups that are not work-related (rec.pictures.dirty).

Must provide the means to create new news groups as required. There must be a means to restrict this capability to a small set of news administrators.

Must provide a “moderator” capability for those news groups that require it.

Must provide the ability to restrict who can “post” articles to a given news group.

Must provide a means to provide authenticated access to certain news groups, if required. Must also provide the means to encrypt information between clients and servers using Secure Sockets Layer (SSL).

## **4.5.4 Assumptions**

It is assumed that DON users will take advantage of this service. Currently, it is severely underused, primarily because people are not in the habit of tuning in to the news service and until recently the performance of news clients for the PC world was poor. With the advent of news support in Microsoft Outlook, Netscape Communicator, and other such systems, it should be easier for users to take advantage of the news service and integrate it into their daily routine.

## **4.5.5 Service Architecture**

The NNTP architecture is quite simple. Each ITSC will have a news server, and these news servers will all be interconnected with NNTP connections for distribution of news. Clients configure their news reader to connect to the nearest news server. Note that mobile users must always connect to the same server, not to the “nearest” one, so that client and server article numbers remain synchronized.

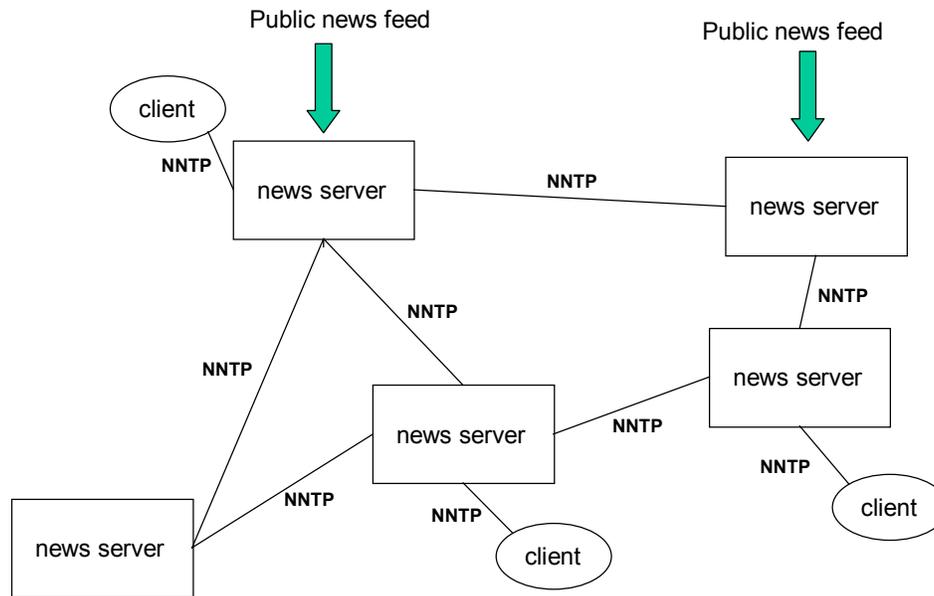


Figure 4-1. NNTP Server Architecture

Figure 4-1 illustrates that no NNTP server has fewer than two peers (to ensure reliability). The architecture includes duplicate public news feeds. One is probably sufficient, but a second is added for reliability or when competing news services offer different content. In each case, all the news will be distributed to all servers.

Most news server implementations provide control over what news groups are sent to peers. This can be useful for regional news groups.

It should be noted that a client can choose to connect to any of the news servers. Also, once a news server is selected, it should remain in use since the client caches information on articles that have been seen, and article numbers on different news servers are not the same.

## Regional Issues and Considerations

Each region will have a news server located at the ITSC. The machine should be named “news.domain” where domain is the regional domain name (news.pacsw.navy.mil).

## Campus and Operational Node Issues and Considerations

News servers will generally not be located at the campus. Users should have connectivity between their location and the nearest news server.

## Deployed Forces Issues and Considerations

The usefulness of the NNTP service to the deployed forces is unclear. NNTP is an interactive, on-demand, high bandwidth service. Certainly while ships are directly connected to the ashore infrastructure (pierside), they have full access to this service like their ashore counterparts. However, while afloat there are serious performance issues because the client/server interface is across the RF communications links. A workable strategy is to host the limited set of news groups of greatest interest on the ship’s news server.

In this way, the news articles are sent just once, and onboard subscribers can pull the information from the onboard news server as required. Where this is appropriate, consideration should be given to deployment of this service to the afloat community.

### **4.5.6 Roles and Responsibilities**

Administrators must be established for management of the news service. These administrators should have the authority to create or delete news groups, to administer access controls to those news groups, and to assign moderators for some news groups if necessary.

Moderators may be required to support some newsgroups.

The ITSC will provide operations support of at least one news server.

## **4.6 Web Hosting**

### **4.6.1 Service Description**

The World Wide Web (WWW) is the basic service for one-to-many information sharing throughout the DON and to the public Internet. It is a “pull” technology that allows individuals to retrieve shared information from servers. The information is stored in Hypertext Markup Language (HTML) or eXtended Markup Language (XML) documents that are moved via the Hypertext Transport Protocol (HTTP) for transport and are displayed across an IP network using a WWW browser.

Running a WWW site involves both the development of the content that will be hosted on the server and administering the web server host and software. Web content development is relatively easy given the COTS what-you-see-is-what-you-get HTML editors. Web site management (controlling the location and structure of the files that make up the web server’s content) is more difficult, but is also facilitated by COTS software. Web server administration (system administration) requires more specialized expertise to keep the web server and its contents secure and in top performance. Many organizations have sufficient hardware and personnel resources to properly maintain their own local web servers. However, for organizations that want a WWW presence without the burden of maintaining a WWW server, the DON Information Technology Service Centers (ITSC) will offer the solution of shared web servers. The ITSC hosted web servers may also be the solution of choice for organizations that want to reduce the risk and resource load of hosting their own public server, but do want to continue hosting their own internal and classified web services. Regardless of the location of the web server, there are basic architectural guidelines (as defined by this document) that should be followed.

Web servers can do more than just provide access to static HTML pages. Web servers can be extended through Common Gateway Interface (CGI) scripts, Java servlets, server side includes (SSI), and other technologies that link application software with the web server. For example, web sites can act as front-end user interfaces for database servers and give easy, forms-based access to data stored in databases to a widely distributed user population. Another common web server function is to serve up a “dynamic” web site. This is a site where most of the content is generated “on the fly” from templates and data is stored in databases based on the user’s specific data requirements. Commercial shopping sites are the most common example of this, but a DON logistics web application or a multilevel intelligence data server can work the same way. This advanced functionality will not initially be available on the ITSC web servers.

ITSCs may provide such specialized services at their discretion. Agreements for these services will be individually negotiated between the ITSCs and their client organizations.

This service is not intended, nor is it a substitute for, bulk transmission of large data files to data processing centers or between application servers requiring replicated data sets.

## **4.6.2 Applicable Standards, Policy, and Guidance**

- Section 6.7 of the ITSG for standards
- DON and local organization ADP security guidelines
- DON and local organization guidance for release of information
- RFC 2068 Hypertext Transfer Protocol 1.1
- RFC 2376 XML Media Types
- RFC 1866 Hypertext Markup Language 2.0
- Internet Engineering Task Force Transport Layer Security (TLS) draft standard (based on Netscape's Secure Sockets Layer protocol)
- Various user authentication, key management and other security standards discussed in other sections of this chapter (see e-mail, directory, and public key infrastructure)

## **4.6.3 Requirements**

No specific web server application or host operating system is specified. However, the web server application should comply with industry standards and remain consistent with those standards as they evolve. Traditionally, there has been little change in web server functionality. The basic web architecture has been flexible enough to support the great evolution in web client functionality without radical changes in the server implementations. Acceptable web server implementations range from "freeware" applications running on personal computers to expensive enterprise server software running on clusters of UNIX systems.

Any organization implementing a web site must comply with the following minimum requirements:

- Must provide a web server that supports the full suite of HTTP version 1.1 services.
- Must have a trained system administrator assigned to maintain the web server host. The system administrator maintains the hardware and operating system and application software on the server and ensures that the system is configured for optimum security.
- Must have a web master or web server administrator assigned to manage the web server software and the contents of the web server. The web server administrator is responsible for implementing proper security procedures as outlined in Section 3.6.7.

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

- Web servers providing access to the general public must be located outside the organization's firewall and must be on a separate machine from that hosting the organization's private (intranet) web site(s). The intranet web server(s) will be located inside the organization's firewall. Servers outside the firewall should be configured as web servers only to minimize vulnerabilities. The public web server and the firewall software should not be run on the same system. Organizations with high levels of access from other DoD sites may want to consider separating web servers containing that information from the intranet web servers.
- Web servers accessible to the general public should only have Distribution A material (Approved for Public Release) on them. There is also a need to consider the server access controls that apply to safeguarding information with other distribution codes.
- Must provide for access controls on the web site, web site sub-tree, or source directory level based on user name, user domain, or user IP address. Access controls are not required on all sites, but all sites need the ability to implement them. Restricted access controls will be implemented. (There is a problem in that there is not a single consistent mechanism available to all popular web servers to support access controls.) Applicable standards should be applied for implementing encrypted passwords, certificates to increase assurance of user identification, and access control lists (ACLs) for servers that tie to OS access control systems.
- Web server logging functions must be enabled and web server logs must be saved at daily or weekly intervals (depending on site traffic levels) and stored online for at least 30 days. Log files may be stored in compressed format to save space. Historical log files must be saved for at least 12 months. Off line (tape, CD-ROM, etc.) storage of historical logs is acceptable. Web server logs can be useful both in understanding what parts of a web site are being used (to focus future efforts) and to identify suspicious activity.
- Web server log files should be reviewed daily to identify suspicious activity (repeated login attempts, downloading large portions of the web site, ill-formatted URLs, etc.). For small sites with low traffic levels, manual review is possible. Larger sites may benefit from automated log file analysis tools to characterize web traffic patterns and suspicious events.
- The web server application software with all its configuration and support files must be backed up weekly. Web server data files must be backed up daily.
- A web server for public or other organization's consumption must have a robust (reliable and preferably redundant) network connection of sufficient bandwidth to provide reasonable response times to HTTP requests.

In addition, organizations are strongly encouraged to consider meeting the following additional requirements:

- Support certificates or other "strong" authentication means for both users and web servers.
- Support SSL authentication and session encryption.
- Organizations are strongly encouraged to implement change monitoring and/or management systems. These automatically check for unauthorized modifications to the content of the web server and facilitate "rolling back" the server's content to the correct files.

- Organizations should consider implementing integrated text search engines that work across the web site's contents. This functionality is included with the enterprise-level web server solutions and is considered to be essential by many web users.

#### **4.6.4 Assumptions**

This guidance is initially applicable to UNCLASSIFIED (NIPRNET) WWW servers. There are some specialized accreditation issues associated with SECRET (SIPRNET) client and server hardware, but the guidance provided here should still be applicable. The SCI level network (JWICS) and WWW system (INTELINK) have their own standards and procedures that are not covered by this document.

Individual organizations must determine the impact of the loss of web servers (internal and external) and the resulting reliability requirements on those servers. DON web servers that contain mission critical logistics or operational data have a high reliability requirement and must have alternate or backup servers in place. Organizational intranets also normally have a need for high availability. High reliability is not typically a requirement for public information servers. High security is a requirement for public servers to prevent embarrassing "hacks" of web sites to introduce false or misleading information. Publicly-accessible servers that are used in the conduct of DON business (contracts and electronic commerce, for example) must consistently be up.

Nothing in this document or the content management process implemented should bypass internal organization approval for publishing information. Access to web servers to manipulate web sites is typically limited to a few individuals who either have standing approval to update elements of a web site or who are the last step in the formal information "publishing" cycle.

A robust DNS exists to allow an organization to create "meaningful" web server names (for example, <http://public.cincpac.navy.mil> or <http://c4isr.spawar.navy.mil>). The DNS also needs to support distribution of an organization's web servers over more than one domain. In a case where an organization's public access server is hosted by an ITSC that is not part of the organization domain (such as the URL <http://public.cincpacfleet.navy.mil>), the DNS might actually take the user to a server located in the [disa.mil](http://disa.mil) domain).

A shared directory service will be implemented across the DON infrastructure to simplify the process of maintaining user level access control lists on web sites.

In the case where an organization is hosting web services remotely (at an ITSC), there should be a safe and convenient way to transfer those pages to the ITSC and to have those pages correctly loaded onto the organization's web site. When the author is local to the web server, root, or other privileged log in, this is feasible. This level of access is discouraged for remote hosts. Possible solutions include e-mail or FTP transfer of files to the ITSC for a local web master to load or the use of COTS tools that provide GUI interfaces for remote web site management.

Protecting systems and users from the impact of malicious code delivered via the web is beyond the scope of this initial ITI Architecture. Organizations must be aware of the potential damage that can be done by ActiveX, Java applets, JavaScript, and other executable software that can be delivered as part of web pages. DON policy on these "active" page elements must be developed and appropriate preventative action taken.

Also beyond the scope of this first document are considerations of intelligence and operations security impacts relating to web technology. Potential threats to national and personnel security have already been identified based on the value of information available on unrestricted web sites. Another impact to be

considered is the possible knowledge gained by an opponent through the analysis of individual or organizational web traffic patterns.

### 4.6.5 Service Architecture

The basic elements of a WWW (or simply web) service are the web client (or browser), the web server, and the local area network (LAN) or wide area network (WAN) that links them. As illustrated in Figure 4-1, the web browser issues a request for a web page from a web server by passing a Uniform Resource Locator (URL) statement onto the LAN or WAN. The routers on the network identify the domain name in the URL and route the URL to the proper web server. The web server parses the URL and places the requested HTML file onto the network addressed to return to the requesting client machine. Typically, web pages are made up of formatted text (the HTML page) and graphics, so the graphics files are also sent to the requesting client where the web browser assembles the web page and displays it. In more complex cases, the URL can instead result in the web server carrying out some action (like a database query) and then returning the results of that action as an HTML page.

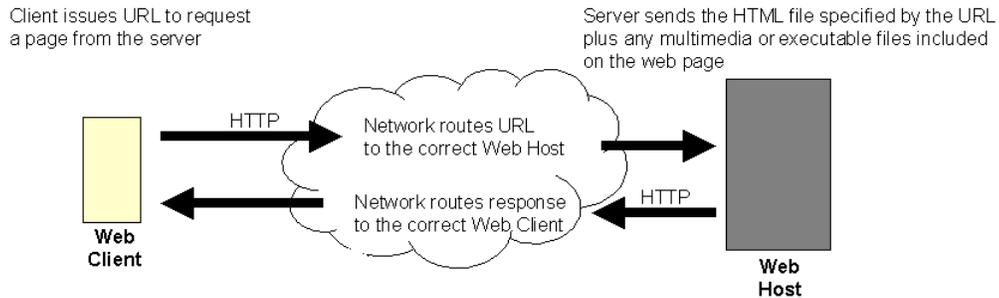


Figure 4-1. WWW Data Flow

For security reasons, most commercial and DoD activities now place web servers that are meant to be accessed by untrusted clients somewhere other than on their internal network. In the case illustrated in Figure 4-2, a firewall has been used to isolate a publicly accessible web server from the organization's internal network that includes a web server for the organization's internal use only. The firewall allows organization members to access web servers on the external wide area network but prevents untrusted users from accessing the organization's internal network and servers.

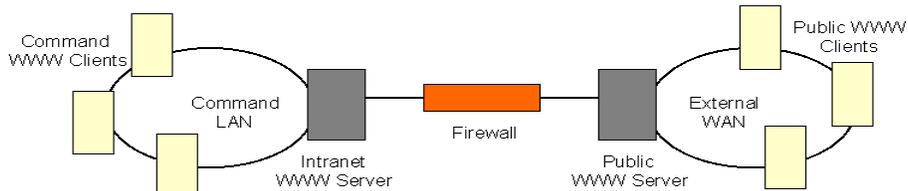


Figure 4-2. Basic WWW Architecture

Figure 4-3 describes a more complex web architecture where a Navy ship has a publicly accessible web site that is primarily used for morale and welfare information and an internal web site for ship's business. Because the morale and welfare site needs to be available when the ship is deployed and because it is designed for access by the general public, the organization has chosen to maintain this web site at a regional ITSC. The ITSC hosts multiple web sites for a number of organizations, so it has a "cluster" of web servers and supports multi-homing. Clustering allows the increased performance to clients on the network by having multiple computers appear as a single web server. Multi-homing allows the ITSC to host multiple web sites, complete with distinctive URLs, on its web server(s).

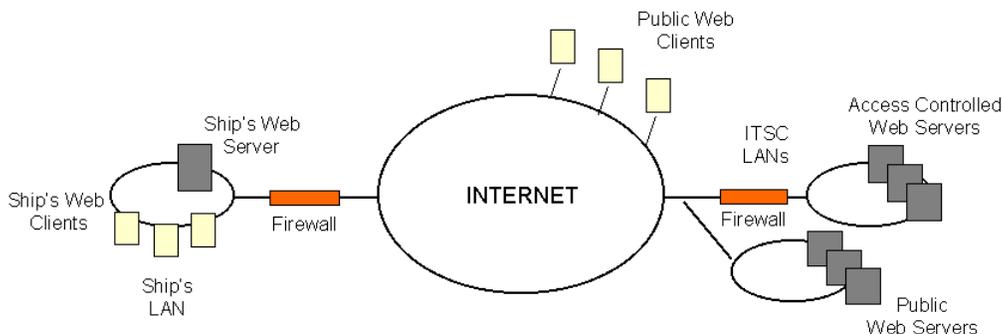


Figure 4-3. Distributed Web Servers

There are three possible web server combinations that an organization can employ:

- A web site hosted by another organization, typically for an organization that only needs a small public web site
- A mix of internal and external web servers - an external organization can host the public access server and the organization can host its own intranet and SIPRNET servers
- Hosting of all of its web services, including multiple servers for public and restricted access

To achieve optimum performance, a web server needs a high bandwidth connection to the network. Enterprise web servers located at ITSCs will all have high bandwidth connections over multiple communications paths. Local organizations may be able to provide sufficiently high bandwidth connections for both their intranet and public web servers.

Other factors that can influence the performance of a web server include the type of network connection, CPU speed, memory, disk space, and the number of other services provided by the server. If a single server is not able to meet client demands (volume of hits or traffic), then a virtual web server can be constructed from multiple server machines. This load-balancing across several web servers can be achieved through proper configuration of the DNS and web server applications. Conversely, several low-traffic web sites (each with their own distinctive domain names) can be combined on a single server by configuring the web server to support "multi-homing."

Web sites with a high reliability requirement should have redundant servers at the local site and a replicated server at a remote site.

Web server functionality is extended through the use of CGIs, servlets, SSIs, and applications that perform dynamic page generation. Each of these options provides increased functionality along with increased security vulnerabilities and complexity. A local organization may be able to host basic web services without any of these server extensions, but a regional facility is almost guaranteed to require them. Even a solution like Microsoft FrontPage, which solves some problems by making web content creation and web site management easier, creates new problems by requiring the FrontPage server extension CGI to be present to achieve all of its functionality.

Organizations should carefully consider implementing any web functionality (for example, Microsoft Active Server Pages) that makes the organization dependent upon one vendor's solution. If an organization insists on using non-standard web functionality and has web sites hosted at an ITSC, the additional costs will have to be borne by the organization.

## **Regional Issues and Considerations**

The web service provided within a region must accommodate multiple domains (multi-homing) and potentially thousands of access-controlled users and tens of thousands of anonymous users. Coordination among organizations to manage domain names and user IDs is essential. Web servers at regional sites must be scalable to provide reasonable response to user requests as the number of domains served increases. The addition of server add-ons (CGIs, servlets, etc.) will require special coordination among the organizations involved and will likely result in additional costs being passed on to hosted organizations.

## **Campus and Operational Node Issues and Considerations**

An organization can choose to operate its own campus web servers, host all its web sites on ITSC (regional) servers, or operate its web sites off a mix of local and remote servers. The variety of security, performance, and quality of service issues addressed elsewhere in this section need to be considered when determining an organization's web strategy.

## **Deployed Forces Issues and Considerations**

It is normally impractical for a deployed force to host a web server due to the constraints of low bandwidth and possibly discontinuous network connections. Deployed forces wishing to make information publicly available should set up a web site at an ITSC. For deployed forces that have a need to access remote DON web servers, caching servers and web subscription services should be employed to make the required information locally available without a need for a real-time connection.

### **4.6.6 Roles and Responsibilities**

#### **Regional**

The service provider establishes rigorous technical and administrative controls to ensure that only authorized persons may update published information. Additionally, the service provider makes reasonable attempts to limit access as described above, based upon IP network, domain, and other filtering techniques but cannot guarantee 100 percent success in restricting access. Similarly, it is the information producer's responsibility to abide by security, public affairs, and organization policies for the release of information and to consider the consequences of unwanted release of privacy act or other

sensitive information. The service provider operates and maintains WWW systems, but is not responsible for the information content.

### **Campus and Operational Node**

Campus service providers have the same responsibilities as regional providers.

### **Deployed Forces Issues and Considerations**

Deployed forces will typically not be providing web services. It is the responsibility of the organization to ensure an appropriate web presence is maintained during their deployment through off-board assets.

### **ITSC**

The ITSC is assumed to be the regional provider for web services.

## **4.6.7 Security Guidelines**

Web servers are security vulnerabilities because they are shared resources for “public” access. The security configuration of the web server application must be coordinated with that of the underlying operating system and the network security environment. Weakness in any one of these three areas renders ineffective the precautions taken in either of the other two. The consequences of lax security include loss of data, loss of service, corruption of data, or unauthorized entry into the balance of the organization’s ADP infrastructure through a web site vulnerability.

System and web server administrators must be aware of operating system and web server software vulnerabilities identified by the software vendors and must apply security patches as they are made available.

The following is from the WWW Organization’s Security FAQ available at <http://www.w3c.org> and compiled by Lincoln D. Stein (mailto: lstein@cshl.org).

If you are a webmaster, system administrator, or are otherwise involved with the administration of a network, the single most important step you can take to increase your site's security is to create a written security policy. This security policy should succinctly describe your organization's policies with regard to:

- who is allowed to use the system,
- when they are allowed to use it,
- what they are allowed to do (different groups may be granted different levels of access),
- procedures for granting access to the system,
- procedures for revoking access (e.g. when an employee leaves),
- what constitutes acceptable use of the system,
- remote and local login methods,
- system monitoring procedures, and

- protocols for responding to suspected security breaches.

For Web servers running on UNIX and NT systems, here are some general security precautions to take:

1. Limit the number of login accounts available on the machine. Delete inactive users.
2. Make sure that people with login privileges choose good passwords. The Crack program will help you detect poorly-chosen passwords:

<ftp://ftp.cert.org/pub/tools/crack/>

3. Turn off unused services. For example, if you do not need to run FTP on the web server host, get rid of the FTP software. Likewise for tftp, sendmail, gopher, NIS (network information services) clients, NFS (networked file system), finger, systat, and anything else that might be present. Check the file/etc/inetd.conf (UNIX) or service manager for a list of servers. Deactivate any that you do not use.

4. Remove shells and interpreters that you do not absolutely need. For example, if you do not run any Perl-based CGI scripts, remove the Perl interpreter.

5. Check both the system and web logs regularly for suspicious activity. The programs Tripwire (UNIX) and Internet Security Scanner (UNIX & NT) are helpful for detecting this type of activity:

Tripwire: <ftp://coast.cs.purdue.edu/pub/COAST/Tripwire>

Internet Security Scanner: <http://ciac.llnl.gov/ciac/ToolsUnixSysMon.html>

6. Make sure that permissions are set correctly on system files, to discourage tampering. On UNIX systems, the program COPS is useful for this:

<ftp://ftp.cert.org/pub/tools/cops/>

On Windows NT, consider Midwestern Commerce's Administrator Assistant Toolkit:

<http://www.ntsecurity.com>

7. Consider turning off the automatic directory listings feature of most web servers.
8. Consider turning off the symbolic link following feature of some web servers.
9. Consider turning off the "exec" form of server side includes.
10. Consider not supporting user-maintained directories.

## **4.7 File Transfer Protocol**

### **4.7.1 Service Description**

File transfer protocol (FTP) is used for bulk file upload and download between computers. It is a very simple and efficient protocol that has been in use even longer than e-mail has. From a client perspective, the user can connect to a remote computer and either “get” or “put” one or more files, as well as perform simple file manipulation commands.

E-mail is strictly a “push” technology and has many limitations. The user cannot “get” files at will from a server using e-mail. Also, many e-mail servers limit the size of e-mail messages to a range of 1 MB up to 10 MB. While e-mail is convenient from a sender’s perspective, recipients have little control over receipt of very large attachments that clog up personal mailboxes and degrade the performance of bandwidth-challenged networks.

FTP provides a solution for many of the e-mail shortcomings. Large files can be distributed by placing it onto an FTP server and then announcing a pointer to that location so recipients can download needed files at their convenience.

The FTP service can be URL-enabled. That is, contents of FTP repositories can be referenced with a URL so that easy access is provided through the web browser or similar interfaces.

The FTP service is a streaming protocol. This makes it fast and efficient when compared to other interactive protocols such as a network file service.

In summary, FTP service is useful in those cases in which files are too large to send via e-mail and when the user needs to create a repository of files for users to upload at will.

One important convention in the FTP world is the notion of “anonymous” access. By identifying one’s self to an FTP server as “anonymous” and using an arbitrary password, the user can get public access to FTP repositories. This is very useful when files to be distributed are truly public, and it obviates the need to use passwords and other access controls. Many client tools have this convention embedded as their default authentication, making these client FTP tools very simple to use.

For the DON, the service of FTP repositories is provided for use in distribution of files. It provides both unrestricted (anonymous) and restricted access (limited to authorized authenticated users) to the files in the repository. There is a mechanism for authorized personnel to place files in the repository. There is also a mechanism for unauthenticated (anonymous) users to “upload” to the repository and provide some instructions to the repository administrator for the disposition of such files. There is an administrator of the repository who maintains the overall repository by cleaning out old files, providing the required access to authorized users, monitoring disk usage, maintaining the file structure and indexes, and more.

### **4.7.2 Applicable Standards, Policy, and Guidance**

- ITSG Chapter 6.6.4
- RFC 959 File Transfer Protocol (also known as STD 9).

- [http://www.cert.org/ftp/tech\\_tips/anonymous\\_ftp\\_config](http://www.cert.org/ftp/tech_tips/anonymous_ftp_config) (How to configure an anonymous FTP server securely)

### **4.7.3 Requirements**

Must align with the service description outlined above.

Must provide both unrestricted (“anonymous”) and restricted access (limited to authorized authenticated users) controls.

Must be “well-connected” to provide high bandwidth access both to the DON network and externally.

Must have an administrator assigned to manage the contents of the repository.

Should follow the conventions used at other Internet FTP repositories in order to allow maximum alignment and utility of COTS products that support these conventions.

### **4.7.4 Assumptions**

High availability is not a requirement, however, a single point of failure probably is not sufficiently reliable.

### **4.7.5 Service Architecture**

The two major architecture considerations are performance and reliability.

To achieve high performance, the FTP server(s) should be well connected to the network. It should be located at an ITSC and be connected at a point that has a high bandwidth external path.

To achieve reliability on an enterprise scale, the FTP service should be replicated at one or more locations. Figure 4-1 depicts this replication process. For general Naval FTP service, there is a primary FTP server at one location, and critical files are replicated on a (very) few servers at other regional ITSCs. For files that are protected through access controls, the access control mechanisms must be synchronized as well between primary and backup servers. Additional work needs to be done to determine the most effective means of automating the dynamic selection of FTP servers, either through Uniform Resource Locators (URLs) listing multiple hosts or through DNS entries with multiple “A” records.

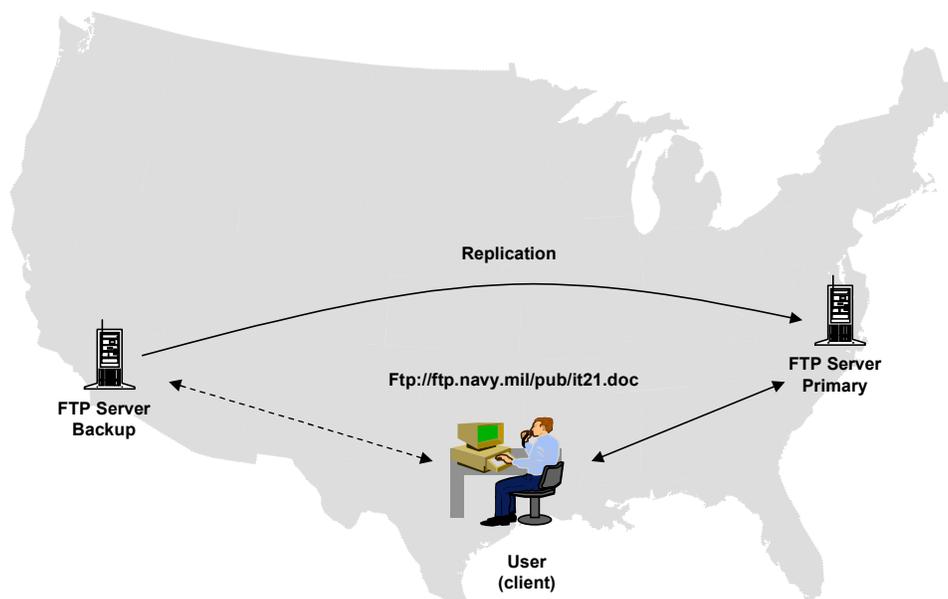


Figure 4-1. Interaction of FTP Components

## Regional Issues and Considerations

The FTP service should be provided locally by the regional ITSC. The ITSC operator will need to maintain the server and must conform to enterprise standards for providing this service. For files that need to be distributed strictly within a region, there should be local flexibility for providing this service to the regional customers.

## Campus and Operational Node Issues and Considerations

If an FTP service is established within a region, sites that operate existing FTP servers may be able to use the regional service instead. Because a site normally sets up an FTP server for the purpose of distributing files off-site, locating this service at the ITSC is an advantage because it places the service closer to the core of the network. This can off-load the network links between the campus and the rest of the world. Establishing instances of FTP service should be done once per region rather than be duplicated at each campus.

## Deployed Forces Issues and Considerations

How deployed forces use this service should take into consideration low bandwidth network links and potentially unreliable service to the deployed platforms. Deployed forces wishing to distribute files via FTP should upload them to one of the FTP servers at an ITSC. Deployed forces that wish to download files via FTP can do so under their own control, depending on available bandwidth. As a quality of service consideration, FTP priority should be lower than interactive traffic when running over IP. This is because FTP may consume all available bandwidth using large packets, which will result in very poor interactive performance (typically small packets). This can be avoided by giving precedence to smaller interactive packets over large streaming packets.

## **4.7.6 Roles and Responsibilities**

The ITSC operator will provide maintenance of the FTP server.

## **4.8 Public Key Infrastructure (PKI)**

### **4.8.1 Service Description**

Public Key Infrastructure (PKI) is the collection of technology, software, hardware, policy, procedures, authorities, and personnel that provides a set of cryptographic tools and the accompanying key management to support digital signature and encryption services to support required applications.

The term “certificate” in the PKI context refers to a data object that binds a public key to a person, a server, or other real-world entity and is cryptographically “signed” by a trusted third party. A certificate supports and extends public key cryptography and can be used to positively identify a person or device and can also be used to implement private channels of communications between individuals.

PKI service provides the mechanism both to generate certificates for individuals and servers and to publish them. There is a “certificate authority” (CA) which manages the life cycle of a certificate. It is the trusted third party that guarantees the authenticity of certificates.

A common directory is maintained to store the certificates for easy retrieval. Additional mechanisms are provided to revoke certificates when required. The CA will publish certificates into this directory.

Certificates are used to provide required services in multiple ways. Typically, they are used for encryption or for digital signatures. In the case of encryption, the PKI service includes “key recovery,” which allows recovery of data in the event of key loss. For digital signature, the service may include the characteristics of non-repudiation. This implies that an individual will have multiple key pairs, depending on the service required.

### **4.8.2 Applicable Standards, Policy, and Guidance**

ITSG Section 3.5 Public Key Infrastructure

PKIX standards:

- See <http://www.ietf.org/html.charters/pkix-charter.html> for a summary of PKIX standards efforts and a complete set of references.

PKCS standards:

- From the RSA FAQ: The Public-Key Cryptography Standards (PKCS) are a set of standards for public-key cryptography developed by RSA Laboratories in cooperation with an informal consortium originally including Apple, Microsoft, Digital Equipment Corporation, Lotus, Sun Microsystems, and the Massachusetts Institute of Technology.
- See <http://www.rsa.com/rsalabs/pubs/PKCS/> for a complete set of references to PKCS standards.

X.509v3 standard format for certificates:

- <http://www.ietf.org/internet-drafts/draft-ietf-pkix-ipki-part1-11.txt>

### **4.8.3 Requirements**

The DON implementation of PKI must meet the requirements of the DoD PKI initiative. The DoD requirements have already been established and are posted on the DoD web site. The basic requirements are summarized as follows:

- Generate certificates in support of Secure Sockets Layer (SSL) encryption for web pages.
- Provide digital signature capability in the e-mail system using the S/MIME protocol.
- Provide encryption capability in the e-mail system (for private or sensitive e-mail).
- Provide positive identification to web servers for access control.
- Provide object signing for Java applets.
- Provide digital signature for electronic forms in support of paperless office initiatives.
- Provide identification certificates in support of VPN access control.
- “Identity certificates” used for digital signature must have the feature of “non-repudiation.”

The DON PKI implementation must publish certificates in a DON directory and these certificates must be retrievable via Lightweight Directory Access Protocol (LDAP). The mechanisms must be established to register users, sign their certificates, and to revoke certificates when necessary.

### **4.8.4 Assumptions**

We assume that it is possible to meet Naval PKI requirements by using the DISA PKI implementation for the DoD. Until recently, this was not considered an option due to the limited scope of the DISA initiative. However, DISA is now committed to meeting Naval requirements and will be incorporated into DON’s plan accordingly.

A DON directory exists and that directory will include all users that intend to use the DON PKI. That directory will be sufficiently reliable, functional, and current to store and retrieve certificates.

### **4.8.5 Service Architecture**

The ITSG describes the components and processes for issuing and revoking certificates. Those details are not repeated here.

The components of the architecture are the Certificate Authority (CA), the Registration Authority (RA), Local Registration Authorities (LRAs), the enterprise directory, and client applications and servers that are PKI aware.

#### **4.8.5.1 Using the DISA PKI**

DISA operates the CA hierarchy for DoD. Each of the services has delegated “signing” authority for certificates. This is the “registration authority” function. In the Navy, this function has been delegated to DCMS. Each RA then delegates authority to local commands. This is called “local registration authority”

or LRA. That function will normally reside at the computer help desk, the personnel office, or the pass and decal office for each base.

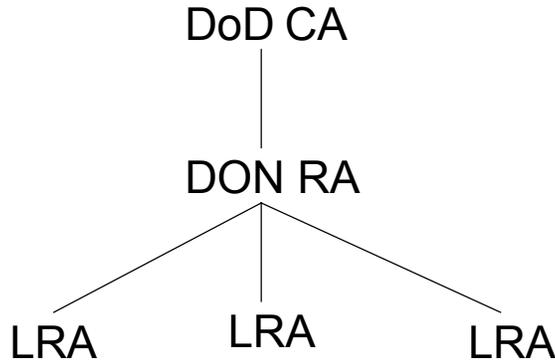


Figure 4-1. Certificate Authority Hierarchy

Figure 4-1 shows the DoD authority hierarchy described above. A similar hierarchy would be employed if the Navy and Marine Corps build their own PKI.

The LRA function will require a significant staffing effort within the DON. It may not require additional personnel, but will certainly require many personnel to take on additional duties. Each Naval location will need to perform the LRA function in support of the personnel at that location. At a minimum they will:

- positively identify each person with multiple picture IDs before issuing a certificate and
- “sign” the person’s public key and perform a process that gets the certificate published in the associated directory.

The exact process is dependent on the particular PKI implementation.

The DISA PKI uses its own directory for publishing certificates. At a minimum, the e-mail certificates generated by DISA must also be published in the Naval Enterprise Directory. This is necessary so that modern e-mail clients that are S/MIME- and LDAP-compatible can send encrypted e-mail messages. In a typical scenario, the e-mail client would retrieve the certificate of each recipient through LDAP access to the enterprise directory or replica. Then the e-mail client can encrypt the e-mail message in the recipient’s public key.

#### **4.8.5.2 Building a Naval PKI**

In the event that the DISA PKI does not meet Naval requirements, a Naval implementation will be used. Even though we use the term “implement”, this does not mean that the DON would write all the software from scratch. Rather, COTS products would be used and DON would serve as the integrator of a DON PKI implementation.

With the DON PKI, it is envisioned that multiple PKI implementations will exist. Each of these pilot or prototype implementations will have their own CAs. These are subordinate CAs and will be coordinated and authorized (by signature) by the DON CA. (A possible extension of this arrangement is for an even higher level CA (DoD or National Security Agency) to coordinate and authorize the DON CA.)

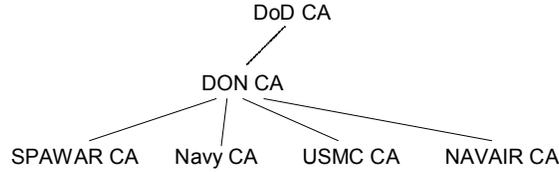


Figure 4-1. DON Certificate Authority Structure (Notional)

Figure 4-1 is a notional representation of a potential Naval CA hierarchy. The identification of the subordinate CAs requires additional study. What is clear is the need for a defined CA structure.

For each of the eventual subordinate CAs identified, the leaves of this hierarchy include the CA, an RA, and multiple local registration authorities (LRAs). These are depicted for the Navy in Figure 4-2 – those for the Marine Corps would be similar.

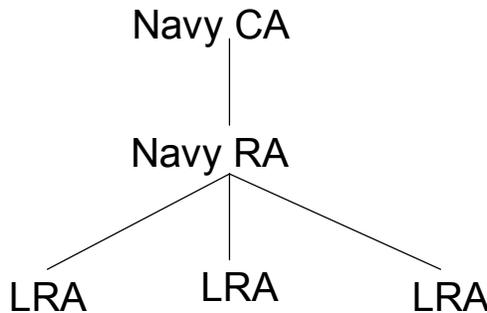


Figure 4-2. Navy CA Supporting Administration Structure

The RA delegates responsibility to the LRAs who perform the actual function of signing user and server certificates. The LRA function is performed at locations that normally provide user access, such as a computer help desk or a security office for a base or command.

The directory component is described in detail in section 4.3 of this chapter. The directory contains entries (objects and attributes) for each Naval user, and one of the stored attributes is each user's certificate. This close association of the PKI and directory heavily influences the PKI architecture.

PKI technology is still immature and rapidly evolving. In the DON, there are pilot efforts to gain experience and identify successful strategies for deployment of a PKI. The appropriate tack for this stage of the DON PKI strategy is to maintain a high level view and to appropriately update when required.

## **Regional Issues and Considerations**

The regional ITSC may provide an LRA function.

## **Campus and Operational Node Issues and Considerations**

Each campus will operate one or more LRA functions.

## **Deployed Forces Issues and Considerations**

The deployed forces present a unique challenge for PKI. If user agents obtain certificates from the directory, it implies that these certificates are supported by a directory that is accessible while afloat, or even while no external communications are available. This implies either that certificates must be cached locally or that the entire DON directory must be maintained on-board. Either could be executed – the challenge is to maintain directory synchronization with the master data.

An even greater challenge is the certificate revocation lists (CRLs) that are also stored in the directory. Keeping afloat directories fully synchronized may be difficult due to limited bandwidth, although some studies show that CRL bandwidth requirements may actually be very small and not very dynamic. Solutions must be engineered to ensure that systems which rely on digital signatures properly negotiate revoked certificates, even when there are long delays in distributing CRLs. One example is the case of signed e-mail using S/MIME. If the afloat directory is not getting updated in the correct sequence, the ashore e-mail infrastructure could request signature verification of the afloat unit before the directory updates are received. In that way, messages become invalidated by the ashore infrastructure whether or not the afloat systems stays synchronized. This is only one example. But before we get too concerned about the CRL issue, we must wait and see what the trends are under actual operating conditions. It is the view of some experts that CRLs may actually be very small and not very dynamic, so it may not be a problem.

### **4.8.6 Roles and Responsibilities**

A Naval PKI implementation would require that a DON organization must stand up the DON CA and RA.

PKI implementers must develop a Certificate Practice Statement (CPS) document and ensure that all DON organizations implement it.

Certification Authorities will need to perform policy coordination of CPSs in the event of cross-certification with other organizations.

The RA is responsible for delegating LRAs.

LRAs are responsible for signing certificates for local users.

The directory must allow the CA to store certificates in user objects.

## **4.9 Remote Access**

### **4.9.1 Service Description**

Each region provides a modem pool for telecommuters, travelers, and other users that require dial-up access to the DON enterprise. Regions cooperate to publish local access numbers for all metropolitan areas in all regions. This permits frequent travelers to dial local numbers for access to the enterprise. Secure remote access through the Internet is also provided.

The service provides a directory of local access numbers by region and metropolitan area. A 1-800 number is available for travelers who would otherwise need to call long distance to reach the DON

enterprise network. Virtual Private Network (VPN) service is provided for users connecting via external networks (e.g., those assigned temporarily at a defense contractor site with Internet access and commercial Internet service providers).

## **4.9.2 Applicable Standards, Policy, and Guidance**

RFC 1825, Security Architecture for the Internet Protocol

Interim Guidance for the DOD Public Key Infrastructure, OSD/C3I, 11 Aug. 1998

## **4.9.3 Requirements**

Travelers, telecommuters, and other remote users that lack “local” connectivity need access to the networks and other services described here. Dial-up access must be provided for access through the telephone network. Such dial-up capability must support modern high speed protocols (i.e. v.34, v.90). There must also be a means to access DON services via other Internet Service Providers (ISPs) regardless of whether they are based on dial-up, cable-modem, or wireless access.

There is a need to access “anything from anywhere at any time.” Any authorized individual should have the means to access any DON network or service from anywhere in the world at any time.

These requirements include both unclassified and secret access to DON networks and services.

Each region implements IPSEC-based access using a type-2 security association (see Sec. 4.5 of RFC 1825) between the remote workstation and the security gateway at the boundary between the Internet and the DON Enterprise Network. Regions may implement backup security gateways, but all such implementations shall use the same gateway products and configurations. Implementation agreements between regions ensure consistent configuration and security policies. Remote workstations are configured with IPSEC according to guidelines provided by the regional ITSC.

Authentication is used enterprise-wide so that every region is able to authenticate users throughout the DON. Once travelers obtain access to the DON Enterprise backbone they are able to access resources at their home operational node subject to the Zone 3 and Zone 2 protections, if there are any. Authentication to the DON Enterprise network is equivalent to authentication to their home region, even if the connection is made through another region.

## **4.9.4 Assumptions**

Most travelers will be located in concentrated areas. Therefore, local phone access will be the primary means of accessing the DON Enterprise Network.

All regions implement consistent Zone 4 security perimeter solutions.

## **4.9.5 Service Architecture**

The architecture for providing remote access service consists of the following components:

- Communication Server
- Virtual Private Network (VPN) Gateway

- Authentication Server
- Firewall
- Router

Remote users may access the network through the communication server for dial-up access or through the perimeter router for access from the Internet. In both cases, an authentication server behind the firewall verifies the identity of the remote user. Once the remote user is authenticated, their access to network resources is the same as any other local user except that the firewall may restrict certain functionality (typically based upon protocol, application, or network address) that is deemed unsafe to permit outside the enterprise perimeter. The VPN Gateway is not required in this scenario.

VPN capability is added to the model described to enhance functionality for remote users. The model configuration is shown in Figure 4-1 below. In this case, the remote user establishes a secure tunnel from his remote workstation to his home network as part of the authentication process.

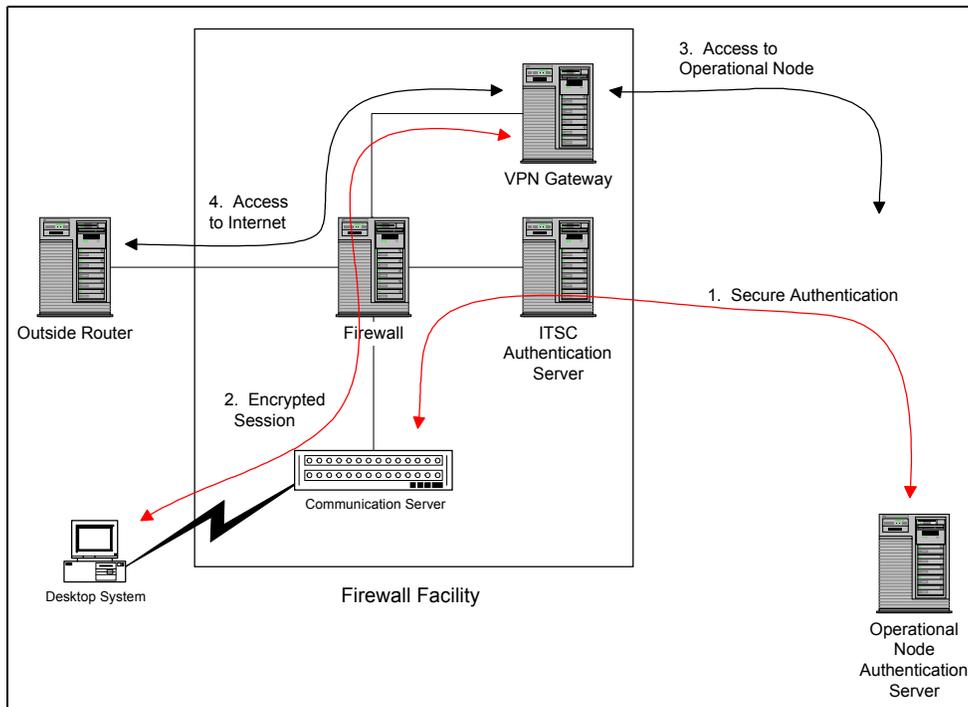


Figure 4-1. Remote Access Service Architecture

First, the user must be authenticated. As shown in Figure 4-1, the communication server establishes a secure channel (shown in red) to an ITSC authentication server. The user provides user ID and password using the remote authentication dial-in user service (RADIUS) protocol or stronger authentication such as SecurID or X.509 security certificate, if required, using a front-end security protocol (RADIUS or TACACS+). If necessary, the ITSC authentication server will employ a back-end security protocol (RADIUS proxy, TACACS+, Kerberos/DCE, Microsoft NT domain security, or Novell NetWare Bindery) to establish authentication within an operational node security context as shown in Figure 4-1. Once the communication server has verified the identity of the remote user, an IP address is issued to the remote user using the DHCP protocol.

Next, an IPSec security association (typically type 2) is established between the remote workstation and the VPN Gateway (shown in red) on his home network, which is located behind the firewall. All traffic (headers and data) between the VPN Gateway and the remote workstation is encrypted and permitted to pass freely through the firewall.

Finally, each end of the tunnel decrypts the received data and forwards it to the correct destination (in the case of the VPN gateway) or application (in the case of the workstation). If the correct destination is outside the firewall, it is again subject to firewall restrictions for outbound traffic (shown in black to indicate no encryption).

For load balancing and redundancy purposes, two (or more) VPN Gateways and Authentication Servers may be implemented.

Communication servers, VPN gateways, and ITSC authentication servers are always located within the regional firewall facility and managed by the regional ITSC staff. For security reasons, communication servers are located outside the firewall. The using community manages its own access control.

Travelers can determine the local phone number for the visited area by dialing a 1-800 number and responding to a menu-driven voice response system. Upon dialing the local phone number, the authentication request is forwarded by the communication server across the DON Enterprise Network to his home authentication server. Upon successful authentication, a secure tunnel is established between the remote user workstation and his home network.

Remote users with Internet service provider (ISP) access to the Internet are also able to establish a VPN connection to their home network through their home VPN Gateway.

To provide Secret access, a similar infrastructure must be put in place, with the addition of approved encryption devices at both ends of each dial-up connection, instead of the VPN solution described above. Secure Data Devices (SDDs) will be installed in-line between the modems and the secret terminal servers. They will need to be permanently keyed to allow for a hands-free auto answer capability. Users will need similar STU-III or SDD devices for secret access and will need these to be registered with the answering SDDs along with MOAs and accreditation paperwork as appropriate to insure proper secure operation and procedures.

## **Afloat Wireless Access**

A special case of remote access is required by ships underway with full term SATCOM or line-of-sight communications. Deployed units which dial in using SATCOM or cellular systems will be treated the same as other dial in remote users are. There are currently four classes of full-term connections that should be supported by fleet teleport facilities. These are distributed as follows:

### **4.9.5.1 Government SATCOM**

SHF - Carriers (CV/N), large deck amphibious ships (LHA/D), Cruisers (CG) (Future)

EHF - All ships except Frigates (FFG)

UHF - All ships

#### **4.9.5.2 Commercial SATCOM**

Challenge Athena - Carriers (CV/N) and large deck amphibious ships (LHA/D)

INMARSAT B HSD - All ships

#### **4.9.5.3 Line of sight (LOS) communications**

UHF LOS - Amphibious ships and Fast attack submarines (SSN)

DWTS - Amphibious ships

#### **4.9.5.4 Asymmetric communications**

GBS - All ships except Frigates (FFG)

Asymmetric submarine communications - Fast attack Submarines (SSN)

### **Regional Issues and Considerations**

Solution planners and designers should seek a single authentication service for dial-up access and Internet access to the DON Enterprise Network.

Regions must establish implementation agreements to ensure consistent (i.e., identical, except for address and other identification parameters) service offerings for remote access via the Internet.

The service architecture described above requires the ITSC to contract with the local telephone company provider to establish sufficient trunks to the ITSC in order to accommodate the region's remote access needs.

Regions maintain an interface to operational node access control databases using the back-end security protocols described above.

### **Campus and Operational Node Issues and Considerations**

Remote workstations accessing the DON Enterprise Network through the Internet implement IPSEC using a bump in the stack (BITS) technique. Regional ITSCs provide configuration guidance.

Access to operational node resources (e.g., file and print service resources) by remote users is facilitated by LAN designs that are IP-based and support access by authenticated external IP addresses.

Campus and operational nodes maintain their own access control databases and provide an interface to the authentication server at the ITSC using the back-end security protocols described above.

### **Deployed Forces Issues and Considerations**

The service architecture components associated with the firewall facility shown above are located in the fleet teleport facility or a shore-based ITSC serving the fleet. Because bandwidth is allocated first for

high-priority operational requirements, authentication and access to the operational node (the ship) may be at a reduced level of performance.

## **4.9.6 Roles and Responsibilities**

### **Regional**

The region maintains all equipment and system components in the firewall facility. It does not maintain individual accounts and access controls but is responsible for providing pass-through authentication to the operational nodes.

### **Campus and Operational Node**

Maintain accounts for secure remote access to local resources. Campus and Operational Nodes do not provide communication servers for remote access. All remote access to the Campus or Operational Node will be come through the regional ITSC firewall facility.

## **Deployed Forces Issues and Considerations**

Allocate bandwidth as needed to maintain high-priority operations.

### **ITSC**

Provide, operate, and maintain modem pools. The total number of modems and associated lines may be reduced significantly as existing capability for remote access is migrated to the ITSC. In order to accomplish this, a rotary system is needed so that all remote users dial a single phone number.

## **4.10 General Voice**

### **4.10.1 Service Description**

Voice is a technology that allows users to communicate interactively in real time through the transmission of sound between two or more users. Voice is a cost-effective tool for enhancing productivity and interpersonal communication between users separated geographically. Typical applications include person-to-person and multi-person real time collaboration, data transmission through the use of modulator-demodulators (modems), and facsimile transmission devices.

### **4.10.2 Applicable Standards**

Information Technology Standards Guidance (ITSG) 98-1.1.

### **4.10.3 Requirements**

The framework for voice will be an open system standard digital switch architecture supporting maximum global and/or regional centralization while ensuring maximum reliability, scalability, and flexibility. Centralization of and resource reductions in administrative functions such as switch management, billing, trouble desk, directory assistance, move-add-change (MAC) technicians, maintenance,

procurement/contract administration, and other voice networking related functions shall be a primary architecture requirement. When possible, the maximum use of existing network resources shall be adopted. However, the elimination of unnecessary or redundant switches, key systems, electronic equipment, or operational/management functions within the existing architecture is also a primary requirement. Sufficient redundancy shall be incorporated to maintain service reliability at or above current industry standards. Additionally, the architecture shall support but not be limited to the following features:

- Common Channel Signaling 7 and/or PRI Trunking
- Multi-Level Precedence Preemption (MLPP)
- Regional E911 Service
- Automated Attendant (If required)
- Centralized Trunking (Public Switch Telephone Network)
- Regional Ashore and World Wide Afloat Number Portability
- Capability of interface within Virtual Private Networks (VPN)

With other switches:

- Voice Mail available for identified users with storage times at current industry standards
- Grade of Service (GOS) at or above current industry standards
- Tail End Hop Off (TEHO)
- Supports Navy BLII initiatives
- Adheres to ITSG 98-1

Within the proposed architecture, regional switches shall be connected as a tandem network (TN) and interconnected via ISDN PRI facilities using switch-to-switch signaling in accordance with open system standards.

The city switch operating within this architecture can be used in a variety of switch exchange applications. It can serve as a local switch which provides end office services, a tandem switch, a toll inter-exchange carrier, an international gateway switch, and/or an Operator Service Position System (OSPS) for national and international calls (with the appropriate software and hardware installed).

The switch must inherently contain the ability, without any further modification, of capturing and displaying emergency on-base calls which shall include the caller's telephone number, the building from which the emergency call originated, as well as the floor and room number of the calling party at a central or consolidated emergency service location.

The city switch's TN feature permits PRI trunks to provide TN trunking between the tandem switch and other switch nodes within the TN. This does not change the implemented TN service for non-PRI trunks, but will partially expand its availability to ISDN PRI users for voice, data, and video traffic. In addition, use of switch-to-switch signaling does not impose any specific topology - the network can be a mesh, star, or main/satellite configuration.

Calls originating on the TN and arriving at the city switch through a PRI will have the following existing features available:

- Automatic Route Selection (ARS)
- Automatic Alternate Routing (AAR)
- Uniform Numbering Plan (UNP)
- Traveling Class Marks (TCM)
- Time-of-day for ARS

Switch-to-switch signaling will provide several identification services between calling and called parties, including:

- Calling Line Identification Presentation (CLIP)
- Connected Line Identification Presentation (COLP)
- Calling/Connected Line Identification Restriction (CLIR)
- Calling Name Identification Presentation (CNIP)
- Connected name identification Presentation (CONP)
- Calling/Connected Name identification Restriction (CNIR)

Other defined switch-to-switch Network Features include:

- Call Completion
- Call Forwarding and Diversion
- Call Interception
- Call Intrusion

The architecture shall also support centralized network management of telephone switches. This management system shall provide but is not limited to the following features:

- Based on COTS software
- Performance management through the collection of statistical data
- Configuration management
- Fault management through detection and isolation of problems
- Security management

The architecture shall also support and provide a mechanism for periodic technology updates based on industry standard timeframes. These periodic technological refreshments shall consider proven technological hardware and software advancements within the telecommunications industry such as voice over IP. However, technology refreshment recommendations must assure continued interoperability within the current architecture.

#### **4.10.4 Assumptions**

It is understood that no one business management scenario or overall technological topology will satisfy all the requirements. Due to the diverse mission requirements within the Navy, a case-by-case analysis must be done to determine which management scenario and topology best fits a given region or locality. Whether a regional central office switch, base switch, or Centrex service from the local exchange carrier provides the best solution should be evaluated during the business case analysis process. Additionally, whether a government owned/operated/maintained switch, government owned with contracted operation and maintenance, or completely outsourced service offers the best value to the Navy must be thoroughly examined.

#### **4.10.5 Service Architecture**

The proposed architecture features a combination of a city switch and remote campus switches connected together as a Tandem Network (TN). This network solution will provide the Navy with a highly reliable, scaleable, and flexible architecture with centralized network management.

The architecture shall consist of base area networks (BAN), regional metropolitan area networks (MAN), and global backbone networks connected in a hierarchical organization. Additionally, where BAN and MAN are not installed or available, point-to-point tie-line trunking will be used. Each level of the overall network shall support:

- Scalability
- Fault tolerance
- Multi-vendor Open System solutions
- Centralized Network Management

Centralized management shall support the concept of ITSCs as described in Chapter 10 of the ITSC. Each region shall maintain at least one ITSC for central management and administration of the regional network. If required, additional base level centers (ITOC) shall be established to assist the overall regional network management, operation, and maintenance of the regional network. Figure 4-1 provides a basic conceptual overview of the voice architecture.



Military unique functions include the following:

- Connectivity and throughput must be sufficient such that for mission critical functions there is no blocking of calls.
- Latency must be short enough that, when a function is keyed with a push-to-talk switch, the first syllable is never missed.
- Preemption is never permitted. Mission critical traffic can originate from anyone.
- Intelligibility must be high enough so that 95 percent of the time, 90 percent of what is spoken is understood. This can be mitigated by speaking in sentences, as is done during telephone conversations in the commercial world, where listeners can extrapolate missing information by using the context of the known conversation. However, on ships, communications are often in phrases or single words. For example, the average call-holding time on a portable communicator is 3 seconds.

#### **4.11.4 Assumptions**

Commercial systems and products will be the first choice for solutions and all evidence is that they meet most of the shipboard voice needs. Application of those systems in ships will be such that the mission critical requirement can be satisfied through the robustness of the network. There will be no single points of failure either physically or functionally.

Some requirements can only be met with unique military solutions because there is not a commercial equivalent. One example is that the shipboard voice has a high percentage of traffic in the form of meet-me conferences or nets. There are only a few products that meet that functionality, but solutions can be adapted from available products. The first choice will always be to seek a proven commercial solution that meets the shipboard requirements.

As other facets of communications tend to migrate from independent to systems to network to mission-centric solutions, the same is true of shipboard voice communications. At the highest level, it is the voice function that is needed and specified and not the infrastructure. Today, requirements are often described in terms of solutions, and they should not be.

#### **4.11.5 Service Architecture**

The order of priority for defining a satisfactory shipboard voice architecture is affordability, interoperability, and survivability. Within that framework, the reduction in workstation complexity will not only reduce the clutter, but it will also reduce the cost. To achieve that reduction, there needs to be a concurrent increase in the integration of shipboard voice with the other networks.

The goal is anywhere/anytime communications. There is a concurrent migration from wired to wireless communications so that the sailor is always connected to the network and is always available anywhere.

Tactical communications are characterized by communications at short range. Strategic communications tend to be long haul. Mission-critical communications can be tactical or strategic. The same is true of non-mission-critical communications. The need for the connectivity of both will be described in the remainder of this section when completed.

### **4.11.6 Roles and Responsibilities**

NAVSEA will host the Battle Group Engineer for the Navy. That role will be to assure that the focus is on integration of all systems from near term to long term and on requirements versus just solutions.

The top-level focus will be battle management, volume control, and sustained warfighting. These will be decomposed into other functional requirements, but all will have MOEs and metrics. The C4I infrastructure will operate in synergy with the other areas to produce affordable and interoperable battle group operations.

## **4.12 Secure Voice**

### **4.12.1 Service Description**

Secure voice communications for Naval forces ashore and afloat allows geographically-dispersed personnel and activities to securely communicate interactively in real time through the transmission of sound between two or more users. Typical applications include person-to-person and multi-person real time collaboration. Secure Voice equipment additionally supports ad hoc data transfer applications over circuit-switched connections and RF media.

## **4.12.2 Applicable Standards**

STE-210	Secure Terminal Equipment Signaling Plan- Interoperable Modes
SVS-210	Signaling Plan-Interoperable Modes
FSVS-211	Interface Control Document for STU-III Black Digital Interface
FNBDT	Future Narrowband Digital Voice Terminal
FED-STD-1015	Analog to Digital Conversion of Voice by 2400 Bits/Second Linear Predictive Coding (LPC-10E)
FED-STD-1016	Analog to Digital Conversion of Radio Voice by 4800 Bits/ Second Code Excited Linear Prediction (CELP)
CCITT G.721	32 kbits Adaptive Differential Pulse Code Modulation (ADPCM)
CCITT Q.931	Digital Subscriber Signal Subscriber No. 1 (DSS1) Network Layer, User-Network Management
ITU-T Q.921	Digital Subscriber Signal Subscriber No. 1 (DSS1) Data Link Layer SR-NWT-001937, Issue 1 National ISDN-1 SR-NWT-002120, Issue 1 National ISDN-2

## **4.12.3 Requirements**

Secure voice capability is required in all enclaves; the specific security mechanism employed is dependent upon both the communications capability and mission tasking of the customer. Interoperability among dissimilar networks is typically achieved via inter-working function devices and/or specifically engineered gateways. All systems employed by DON shall be supported in the Global Secure Voice System (GSVS) architecture. Employment by Allied/Coalition partners is subject to technology releasability considerations. Specific secure voice cryptographic equipment systems employed and/or planned for use by the DON include:

### **4.12.3.1 Secure Terminal Equipment (STE)**

Secure Terminal Equipment (STE) is the next generation Secure Telephony device for the U.S. Government. Transition from the current STU-III system is scheduled over the next five to seven years. The STE product line will incorporate four distinct secure telephony modes which each have unique connectivity requirements.

(1) **Secure Terminal Equipment Mode:** The STE mode is a voice service implemented on an ISDN data channel and requires an end-to-end Unrestricted Digital Interface (UDI) (which is 56-64 kbps). In order to use the STE mode, the telephone instrument currently requires an ISDN S/T interface. To support maximum functionality, the provisioning of ISDN service should conform to NI-1/2 industry standards. This mode can be extended to deployed forces via STE direct dial, which is in development.

Early testing of the STE has revealed that government telephone service contracts typically contain no provision for NI-1/2-compliant ISDN. Testing has also revealed that FTS-2000 voice trunking services do not consistently support end-to-end UDI.

(2) **Secure Telephone Unit (3<sup>rd</sup> Generation) (STU-III) Mode:** The STU-III mode is a voice service implemented on standard PSTN/DSN voice channels. Compression techniques employed on the voice channels must be engineered to support LPC-10 and CELP algorithms. It is fully interoperable with the current STU-III series of equipment and extended to deployed forces via various direct dial methods. It is important to note that the STU-III mode is inferior in quality to the STE mode and that if proper ISDN provisioning is not achieved, the unit will default to the STU-III mode of operation.

(3) **Future Narrowband Digital Voice Terminal (FNBDT) Mode:** FNBDT mode is an enhanced secure voice mode for narrowband connections (RF/Wireless). This mode will allow for interoperability between the STE and Condor (an emerging wireless secure product) secure product lines. Interworking functions required to support the FNBDT Mode have yet to be defined.

(4) **Allied/Coalition Mode:** Yet to be defined.

#### **4.12.3.2 Secure Telephone Unit – 3<sup>rd</sup> Generation (STU-III)**

STU-III equipment is the currently fielded secure telephony device in the U.S. Government. It utilizes LPC-10, CELP, and MRELP (Motorola) algorithms in the 2.4/4.8 and 9.6 kbps modes of operation, respectively. STU-III secure is a voice service implemented on standard PSTN/DSN voice channels. Compression techniques employed on the voice channels must be engineered to support operating algorithms.

#### **4.12.3.3 Advanced Narrowband Digital Voice Terminal (ANDVT)**

ANDVT equipment (KYV-5) is primarily used by deployed forces in support of tactical applications. Interface to the GSVS is accomplished through either Radio Wireline Interfaces installed at key communications facilities ashore (NCTAMS) or specifically configured Defense Red Switch Network (DSRN) facilities. ANDVT operations are supported on High Frequency (HF), Ultra-High Frequency (UHF) Satcom, Super-High Frequency Satcom (SHF), and Extremely High-Frequency Satcom. ANDVT utilizes the LPC-10e algorithm. Infrastructure requirements from RF facilities to either RWI or DRSN gateways require 2.4 kbps dedicated digital connectivity per communications channel.

#### **4.12.3.4 Digital Secure Voice Terminal (KY-68)**

DSVT (KY-68) is primarily used to provide secure telephony between commanders in the field and at sea. The Ground Mobile Force (GMF) voice architecture can support both 16 kbps and 32 kbps modes of operation. Bandwidth-constrained Naval Units typically operate at the 16 kbps mode. Primary interface to

the field is accomplished by relay via Standard Tactical Entry Points (STEP). Direct interface from the ship to the field is possible if direct satellite and/or Digital Wideband Transmission System (DWTS) is available. Shore interface requires 16 kbps NRZ signaling, while end terminal equipment requires Conditioned Di-Phase (CDI) interface.

#### 4.12.4 Assumptions

- National Security Agency STE and Condor initiatives will continue to drive secure telephony architecture.
- Joint Staff Global Secure Voice Network vision will continue to provide overarching architectural guidance to maximize infrastructure efficiency and minimize operating cost.

#### 4.12.5 Service Architecture

Figure 4-1 illustrates the topology of the DON ship-to-shore telephone connections. Three aspects of the figure are noteworthy.

- There are multiple RF circuits used to make such connections.
- The majority of telephone service is provided through specifically engineered gateway facilities (i.e., NCTAMS or STEP site).
- Connectivity to tactical telephone networks (i.e., TRI-TAC, MSE, etc.) is extended via these shore gateway facilities.

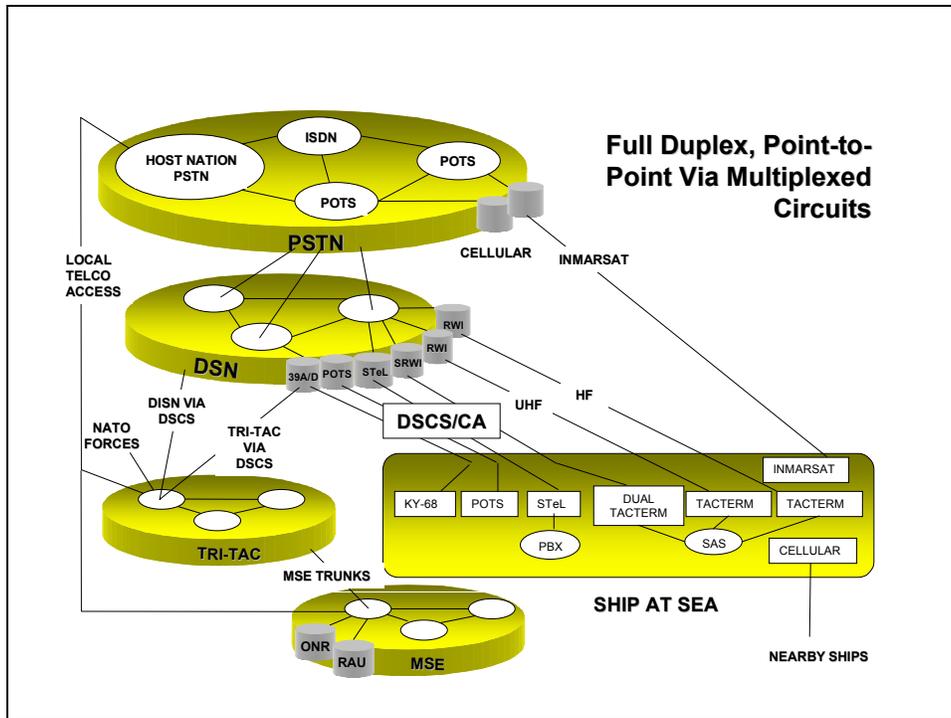


Figure 4-1. Secure Voice Notional Architecture

Figure 4-2 illustrates the components used to provide secure ship-to-shore telephone service. Direct dial secure telephone service is overlaid on basic ship-to-shore telephone connections.

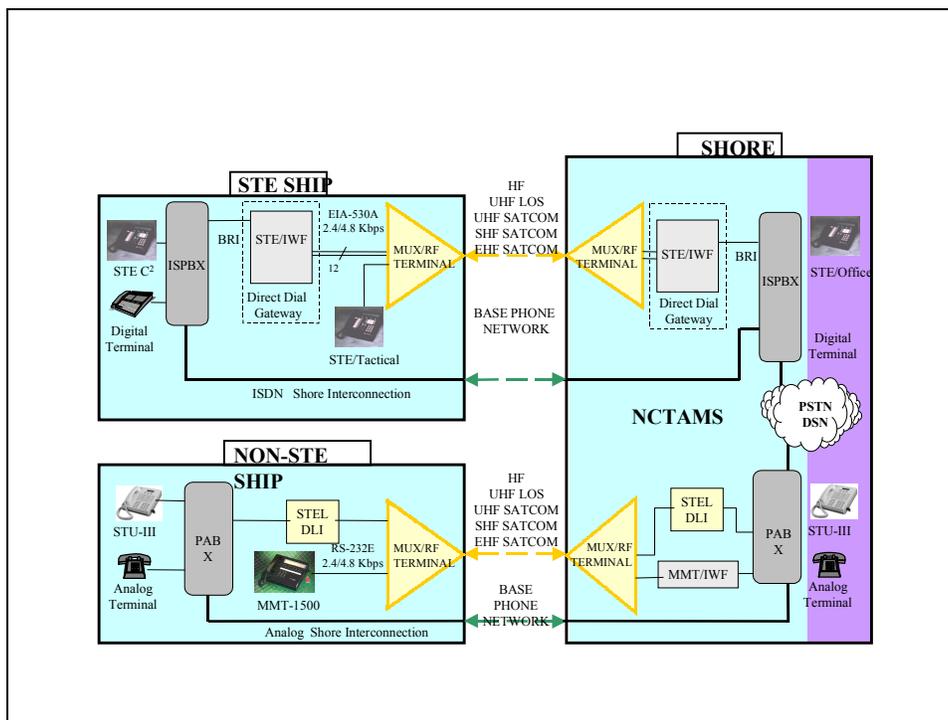


Figure 4-2. Direct Dial Telephony Architecture

## Regional Issues and Considerations

Regional planners and designers should ensure that NI-1-compliant ISDN telephone service is available in support of STE secure voice mode.

## Campus and Operational Node Issues and Considerations

STE fielding should be a prime consideration in sizing ISDN switches. Initial fielding estimates are being collected by SPAWAR PMW-161.

## Deployed Force Issues and Considerations

Existing Indirect DSN access is accomplished via FCC-100 FXS/FXO and Timeplex Voice Server Module (VSM) and FXS/FXO modules. Continued fielding of improved voice modules specifically designed to support STU-III algorithms (I.E. VSM.5) is required to support STU-III until STE Direct Dial is fielded.

## 4.13 Multimedia

Federal Standard 1037-C defines multimedia as “the processing and integrated presentation of information in more than one form, e.g., video, voice, music, and data.” In this services section,

multimedia services include video teleconferencing (VTC), video applications sharing, video teletraining, and video and image/graphics file servers. Also covered are VTC application-enhanced data services that allow users to share applications and documents and to participate in collaborative activities including video applications sharing, video document sharing, and “white boarding.” Figure 4-1 highlights the technologies supporting the various multimedia services.

	Video Conferencing	Video Application Sharing	Tele-training	Video/graphics file server
<b>Analog signal</b>				
<b>Digital signal</b>	X	X	X	X
<b>Real-time</b>	X	X	X	
<b>Stored image</b>			X	X
<b>Point-to-point</b>	X	X	X	X
<b>Multipoint</b>	X	X	X	
<b>Interactive</b>	X	X	X	X
<b>Non-Interactive</b>	X		X	X

Figure 4-1. Multimedia Services and Supporting Technologies

### 4.13.1 Service Description

**Video Teleconferencing:** Video teleconferencing for Naval forces ashore and afloat allows geographically-dispersed personnel and activities to conduct face-to-face meetings in real time through the transmission of images and sound. Current video teleconferencing systems range from reservation-based, dedicated boardroom systems to portable cart and desktop systems. Desktop video systems based on ATM are emerging. Two transmission models, typical of room-sized and desktop video conferencing, are provided in Figure 4-1.

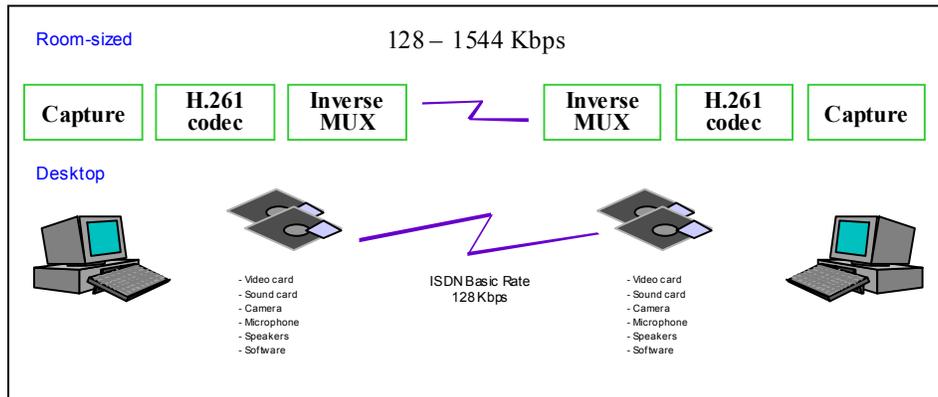


Figure 4-1. Room-sized and Desktop Video Conferencing transmission model

**Video Teletraining:** These systems are a special class of multimedia services and include interactive two-way video and audio systems, one-way video and two-way interactive audio, one-way video and one-way audio, and multimedia computer-based training applications.

**Video File Servers:** A video server is essentially a digital storage device (like a large digital video cassette recorder) designed to handle multimedia video content and deliver it to multiple simultaneous users on demand. The video server stores digital video, audio, and graphics in a compressed format that can be retrieved, sorted, and distributed over a communications network.

**Multimedia File Servers:** Multimedia servers are the more general systems designed to store and forward text, graphics, and images as well as some video and audio files.

### **4.13.2 Applicable Standards and References**

**Video Teleconferencing:** The video teleconferencing standards in use today were developed by public carriers to promote interoperability between desktop conferencing systems across public network transport service provided by Narrowband Integrated Services Digital Network (N-ISDN). N-ISDN does not specify the network, but does specify the interface to the network. Most existing standards are based on the concept of end user equipment connected to a public network that provides services and connectivity between users.

The DON Information Technology Standards Guidance (ITSG) addresses these end user interfaces and provides an overview of the ITU video teleconferencing standards in ITSG Table 9.5.

**Video Teletraining:** To be provided at a future date.

**Video File Servers:** To be provided at a future date.

**Multimedia File Servers:** To be provided at a future date.

### **4.13.3 Requirements**

**Video Teleconferencing:** Video teleconferencing must be provided as point-to-point and point-to-multi-point virtual connections for group and personal conferencing. The service must be easy to use and available to be invoked from anywhere. If a desktop video conference session is required, either scheduled or ad hoc, the user or VTC scheduling application simply “calls” another user (or users) to establish a session. The solution must provide for automatic session setup and tear down. The service must use the same transmission medium as voice and data. The service must be secure. The system latency must be low (less than 100 milliseconds) and jitter must be minimal (less than 150 milliseconds) as based upon quality of service guarantees.

**Video Teletraining:** Video teletraining must provide the following: group or personal interactive two-way video and audio systems, one-way video and two way interactive audio/data, one-way video and one-way audio (broadcast), and multimedia computer-based training.

**Video File Servers:** Video servers must provide services such as Video On Demand services where the end user has some control over the selection of the material to be displayed and the viewing time. Viewers must have control of the video stream similar to those found on a typical video cassette recorder for restart, rewind, pause, and fast forward. Most Video On Demand systems are point-to-point systems that require some signaling protocol between the end user and the video server. Video servers must also provide encoders to support real-time streaming video; e.g., broadcast video.

**Multimedia Servers:** Multimedia servers must store and forward text, graphics, and images, video clips, and audio files.

## Security Considerations

For persons or activities requiring NSA-approved type 1 link encryption, Figure 4-1 lists the required device(s) for video teleconferencing.

	<b>ISDN H.320</b>	<b>POTS H.324</b>	<b>IP/LANs H.323</b>	<b>ATM H.321</b>	<b>Hi-Res ATM H.310</b>
<b>External NSA Approved type 1 Encryption</b>	KG-194 KIV-7	TBD	KIV-7 at WAN Gateway from Classified LAN	Fastlane Between ATM Nodes;  KIV-7 at WAN Gateway	Fastlane Between ATM Nodes;  KIV-7 at WAN Gateway

Figure 4-1. Required Link Encryption Devices for Video Teleconferencing

### 4.13.4 Assumptions

- Naval Service voice, video, and data networks are currently separate, but will share the same transmission medium.
- Video will become a basic network service funded centrally for both capital investment and operations and maintenance.

### 4.13.5 Service Architecture

There are currently as many methods that provide efficient transport of multimedia conferencing information as there are that provide efficient transport of data traffic. Because few transport methods are capable of meeting the requirements of both, Naval organizations have built separate networks for voice, video, and data. Uncompressed video and audio information requires a high speed Constant Bit Rate (CBR) channel to avoid distorted speech and jerky motion. Data transfers tend to be intermittent or “bursty” in nature with periods of low activity followed by periods of high activity. Figure 4-1 depicts the target multimedia architecture for the Naval services.

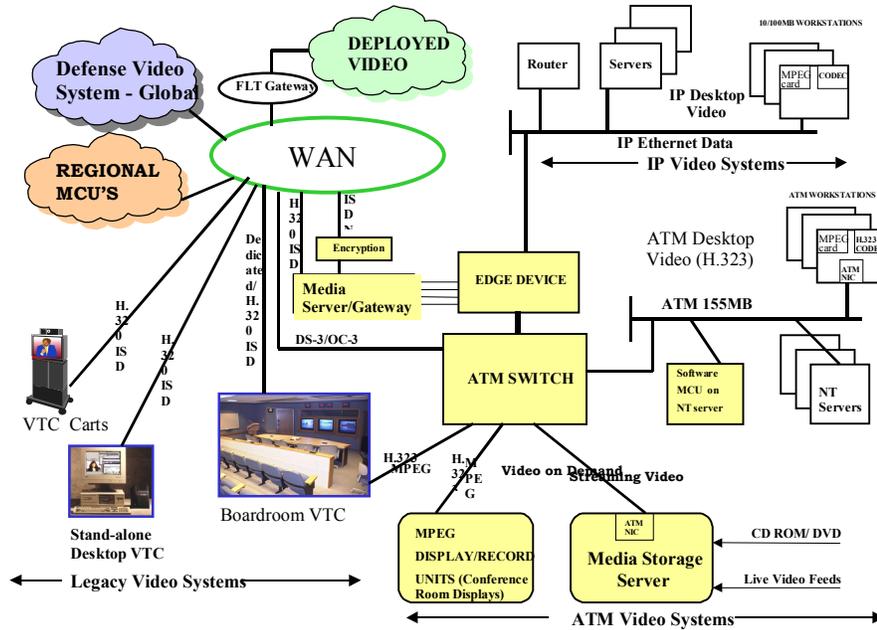


Figure 4-1. Naval Services Multimedia Architecture

## Enterprise Issues and Considerations

Interoperability must be ensured through the use of the multi-media/video teleconferencing standards consistent with those used for the Defense Video Service – Global (DVSG). The DON Information Technology Standards Guidance (ITSG) provides detailed information on the DVSG standards.

## Regional Issues and Considerations

- Video conferencing services (ad hoc and scheduled) among intra-regional entities or between intra-regional and inter-regional entities are best funded and managed by one using multi-point interpreters/translators.
- At first, all regional video services hubs must accommodate intra-region carrier diversity (e.g., ATM, Frame Relay, Switched Digital) but must transition to a single compatible service over a specified period of time. The target services are based on Asynchronous Transmission Mode (ATM) technology over Category 5 twisted pair or fiber optic cable.
- Multicasting, transmitting IP datagrams to intended recipients, is required for one-to-many or many-to-many applications such as video conferencing, applications sharing, and video teletraining. Existing Ethernet Network Interface Cards (NICs) may need to be replaced by NICs that filter multicasts. This approach prevents forwarding of multicasts to higher network layers (in the protocol stack) for filtering, which will thereby save CPU processing time. In addition, most switches and routers in use today forward all multicast transmissions to all ports, which places unnecessary demands on the network. Network components may need to be upgraded to support several multicast communications protocols that may not be supported in the current configuration. These include:
  - ♦ Internet Group Management Protocol (IGMP)
  - ♦ Distance Vector Multicast Routing Protocol (DVMRP)
  - ♦ Multicast Open Shortest Path First (MOSPF)

- ◆ Protocol-Independent Multicast (PIM) protocol

(While some combination of DVMRP, MOSPF, and PIM are being used today, the long-term goal is to migrate to native multicast using PIM.)

To support time-sensitive audio/video signals, the following additional protocols may be required:

- ◆ Resource Reservation Protocol (RSVP)
- ◆ Real-time Transport Protocol (RTP)
- ◆ Real-time Transport Control Protocol (RTCP)

Because multimedia services use large amounts of bandwidth and typically use the User Data Protocol (UDP) packets that are easier to spoof than standard TCP-based packets, multimedia services should be placed behind firewalls

## **Deployed Issues and Considerations**

Considering the bandwidth demands of multimedia services, stand-alone video teletraining, video file servers, and multimedia file servers is the recommended solution for a deployed shipboard environment.

## **Remote Shore-based Issues and Considerations**

For multi-point video teleconferencing, dial-in via ISDN to a Naval Service multimedia hub is recommended.

For point-to-point video teleconferencing, a dial-in connection can be established if the systems at each end are H.320-compatible or if they use H.323 through a H.320 gateway.

## **Campus and Operational Nodes Issues and Considerations**

- For scheduled board room, rollaway cart, and ad hoc desktop VTC, every system must be H.320-compatible or pass through a H.320 gateway if it uses H.323.
- Because multimedia services use large amounts of bandwidth and typically use the User Data Protocol (UDP) packets that are easier to spoof than standard TCP-based packets, multimedia services should be placed behind firewalls.
- For all new group video-conferencing systems, every system must support the H.261 and H.263 schemas for circuit switched systems to provide potential to use a single 2B+D line instead of three ISDN channels for many applications.

### **4.13.6 Roles and Responsibilities**

#### **Enterprise**

- DON multimedia policy, standards, and guidelines are developed, coordinated, and published by the DON Chief Information Officer (CIO).
- The DON CIO reviews waiver requests for the DON CIO multimedia policy, standards, and guidelines.

- Interoperability must be ensured through the use of the multi-media/video teleconferencing standards consistent with those used for the Defense Video Service – Global (DVSG).
- The SPAWAR System Center, Charleston is the DON multimedia/video teleconferencing lab.
- All Naval multimedia/video teleconferencing systems must be certified by the DON.

## Regional

- Multi-point Control Units (MCUs) must be placed as close as possible to the enterprise network backbone.
- For multi-point conferences, a Multi-point Control Unit (MCU) must be located somewhere within the network. Figure 4-1 depicts multi-point functionality through the use of multi-point control units.
- Multimedia service providers (Regional/Area/Campus) must establish Service Level Agreements (SLAs) with customers.

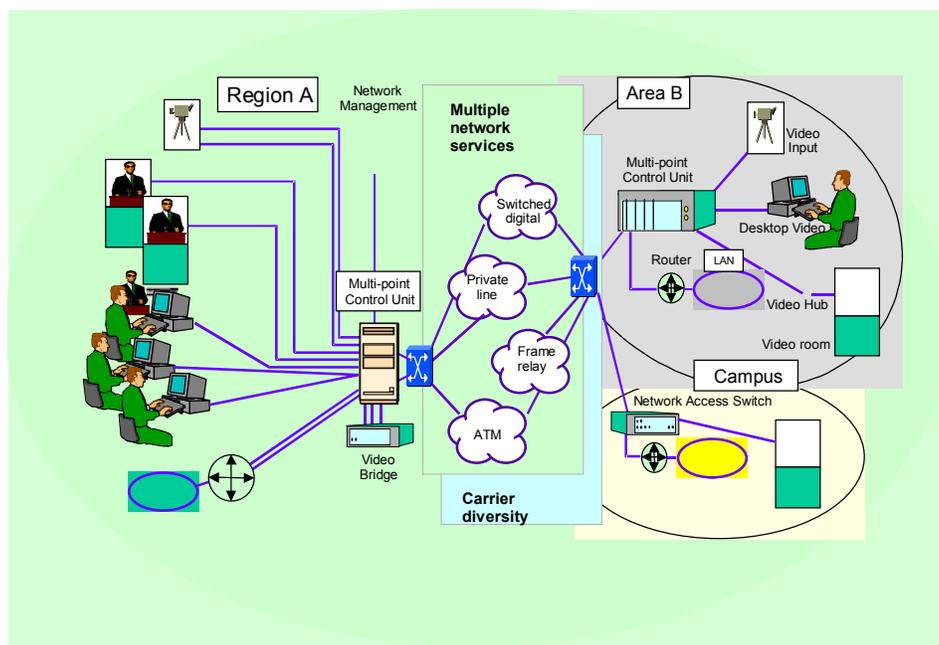


Figure 4-1. Multi-point Control Unit Functionality

## Shore-based Campus and Operational Nodes

The ITSC will provide operations support to at least one news server.

## 4.14 Common Operating Environment Applications

### 4.14.1 Service Description

The Defense Information Infrastructure (DII) Common Operating Environment (COE) is a set of DoD-wide guidelines, standards, and specifications for the development of software. The DII COE principally provides for software reuse, standardized “look and feel”, and improved interoperability within Joint and

Service software-based systems. The DII COE is compliant with DoD *Joint Technical Architecture* (JTA). The DII COE provides a modular open architecture for software development. It has been applied to software development in the functional areas of C4I, logistics, transportation, base support, health affairs, and finance.

The Office of the Secretary of Defense has issued a directive that all new C4I systems must be compliant with the JTA. The JTA, in turn, mandates the use of DII COE. Combat systems and weapons systems software development will be addressed within future versions of the JTA.

## **4.14.2 Applicable Standards, Policy, and Guidance**

See Joint Technical Architecture (Section 2.2)

See DII COE Integration & Runtime Specifications

See User Interface Specifications for the DII

DISA DII COE references are available at <http://spider.osfl.disa.mil/dii/>

## **4.14.3 Requirements**

Applications developed within the Naval enterprise must abide by the following COE requirements:

- Software development shall abide by the requirements set forth in the DII COE standards.
- Applications developed for operational use within the Naval enterprise must meet, at a minimum, DII COE Level 7 compliance requirements (DII COE levels are discussed later in this section).
- Applications developed for Advanced Technology Demonstrations (ATD), Advanced Concept Technology Demonstrations (ACTD), or Fleet Battle Experiments (FBE) must meet, at a minimum, DII COE Level 5 compliance requirements.

## **4.14.4 Assumptions**

DII COE will continue to be the guidance standards for DoD software development.

## **4.14.5 Service Architecture**

DII COE is structured to provide common services and promote shared database design within DoD-developed software as illustrated by Figure 4-1.

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

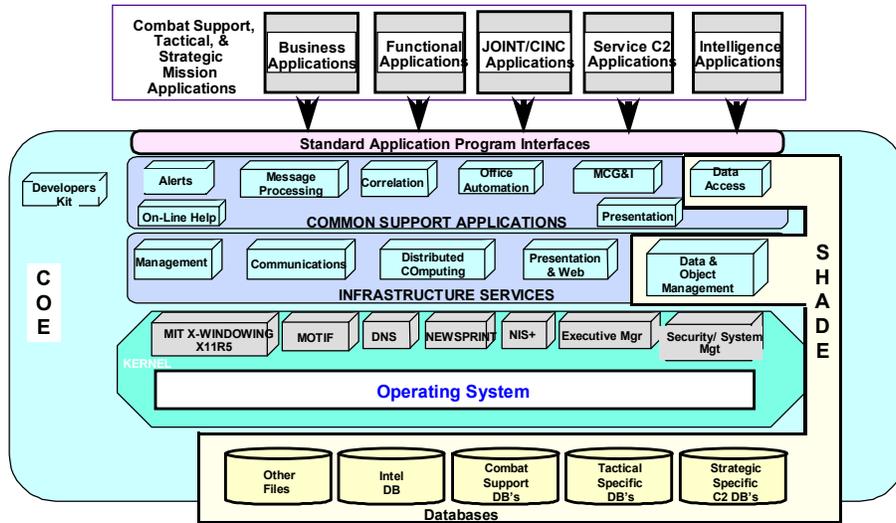


Figure 4-1. General Overview of DII COE Architecture

Adherence to the COE is designated by the depicted eight different levels of compliance. Each higher level of compliance includes all the requirements from the preceding level. Level 8 (Full Compliance) is the optimum level for software applications deployed within the Naval ITI. A minimum of Level 5 (Prototype Compliance) is acceptable for demonstration software which is deployed on a limited basis within the Naval Enterprise such as within an Advanced Concept Technology Demonstration (ACTD), Advanced Technology Demonstration (ATD), or Fleet Battle Experiment (FBE). A minimum of Level 7 (Interoperable Compliance) is the lowest level of compliance acceptable for software integration within the Naval ITI that serves operational (combat and combat support) units.

DII COE Level	COE Requirements	Naval ITI Requirement
Level 1 <b>Standards Compliance</b>	<ul style="list-style-type: none"> <li>- Software is based on either NT or POSIX 1.0 compliant OS</li> <li>- Supports DCE, SLIP, PPP, TCP/IP, UDP</li> <li>- GUI is either Motif (with X-Windows) or NT</li> <li>- Supports SQL</li> </ul>	Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure.
Level 2 <b>Network Compliance</b>	<ul style="list-style-type: none"> <li>- Is COE Level 1 compliant</li> <li>- Supports sockets</li> <li>- Supports DNS, NFS, NIS</li> <li>- Works with C2, BSM security modules</li> <li>- If NT, supports NTFS</li> <li>- Data base transactions implement strict two phase locking</li> </ul>	Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure.

**Department of the Navy Chief Information Officer  
Information Technology Infrastructure Architecture, Version 99-1.0  
16 March 1999**

	<ul style="list-style-type: none"> <li>- Requires no reserved IP addresses</li> <li>- Is not dependent on specific type of LAN</li> </ul>	
Level 3 <b>Workstation Compliance</b>	<ul style="list-style-type: none"> <li>- Is COE Level 2 compliant</li> <li>- Extensions to OS are clearly documented</li> <li>- Software does not modify associated COTS, GUI, or DBMS</li> <li>- Works with anonymous FTP</li> <li>- Limited "hard-coding" of ports</li> <li>- Software does not modify other software file structures</li> </ul>	Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure.
Level 4 <b>Bootstrap Compliance</b>	<ul style="list-style-type: none"> <li>- Is COE Level 3 compliant</li> <li>- Software is properly segmented</li> <li>- Software extensions do not conflict with other segments</li> <li>- Segment restrictions are clearly documented</li> <li>- Follows DII COE Directory structure</li> <li>- Software may be installed/de-installed without conflicting with other software segments</li> </ul>	Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure.
Level 5 <b>Minimal COE Compliance</b>	<ul style="list-style-type: none"> <li>- Is COE Level 4 compliant</li> <li>- Does not violate COE's UNIX root and login restrictions</li> <li>- Fully compliant with COE Style Guide</li> <li>- Only user interface is via the GUI</li> <li>- No developmental software tools are required</li> <li>- Minimal dependence/conflicts with other segments is present</li> <li>- Inter-segment communication is done via COE structures/processes</li> <li>- Software licensing requirements are satisfied</li> </ul>	Minimal level of COE compliance for prototype software deployed within the ITI for ACTD, ATD, and FBE. Deployment is not acceptable within operational portions of the ITI.
Level 6 <b>Intermediate Compliance</b>	<ul style="list-style-type: none"> <li>- Is COE Level 5 compliant</li> <li>- Data base access is via COE roles and groups</li> <li>- Segment uses COE web-server</li> <li>- Segment duplicates &lt; 50% of existing COE functions</li> <li>- Segment does not modify environment variables</li> </ul>	Segments meeting this level of COE compliance are acceptable for prototype software deployed within the ITI for ACTD, ATD, and FBE. Deployment is not acceptable within operational portions of the ITI.
Level 7 <b>Interoperable Compliance</b>	<ul style="list-style-type: none"> <li>- Is COE Level 6 compliant</li> <li>- NT segments use NT registry</li> <li>- Allows cut, copy, paste between segments</li> <li>- Does not replicate data present in Shared Data Environment (SHADE)</li> <li>- Eliminates file permission vulnerabilities</li> </ul>	This is the minimum level of COE compliance suitable for software deployment in operational portions of the ITI.

	<ul style="list-style-type: none"> <li>- Rules-based segments have rules within a rules data base</li> <li>- Supports frame-based web services</li> <li>- Does not duplicate COE functions</li> <li>- Less than 25% of COE functions are accessed via private API</li> </ul>	
<p>Level 8 <b>Full Compliance</b></p>	<ul style="list-style-type: none"> <li>- Is COE Level 7 compliant</li> <li>- Fully compliant with COE style guide</li> <li>- Uses Joint data elements</li> <li>- No private APIs are used</li> <li>- Uses COE DBMS</li> <li>- Does not duplicate functionality of other segments</li> </ul>	<p>This is the objective level of COE compliance for all software applications deployed within the ITI.</p>

Figure 4-1 COE Levels of Compliance

#### **4.14.6 Roles and Responsibilities**

Naval Systems commands and Program Executive Offices are responsible for ensuring that software developed under their cognizance meets the COE requirements in Figure 4-1.

Other Naval organizations which procure or develop non-Program of Record software for inclusion in the Naval Enterprise must abide by the COE requirements in Figure 4-1 whenever possible.

The ITSC planners and engineers should consider the DII COE compliance of software applications within their cognizance to ensure software interoperability within the unit, campus, regional, and global levels.