

Volume I, Chapter 3 – Table of Contents

3. Conceptual ITI Network Architecture	3-1
3.1 ITI Network Architecture Approach	3-1
3.1.1 Geographic Considerations.....	3-2
3.1.2 Connectivity Grid	3-2
3.1.3 Autonomous Networks	3-3
3.1.4 Communities of Interest (CoI).....	3-3
3.1.5 Users	3-3
3.2 Layered Technical Model.....	3-4
3.3 ATM Connectivity Architecture Overview.....	3-5
3.3.1 DON ATM Architecture Strategy.....	3-5
3.3.2 ATM Connectivity Approaches.....	3-6
3.3.3 ATM Addressing Plan	3-8
3.3.4 ATM Routing Architecture.....	3-9
3.3.5 Rationale for ATM Technology in DON ITI Architecture.....	3-9
3.4 IP Connectivity Architecture Overview	3-11
3.4.1 Strategy.....	3-11
3.4.2 IP Connectivity	3-12
3.4.3 Autonomous Networks	3-13
3.4.4 Fleet Intranet.....	3-13
3.4.5 IP Addressing Plan	3-13
3.4.6 IP Routing Architecture	3-14
3.5 Network Connectivity Security Overview	3-15
3.6 ATM Connectivity - Detailed Architecture.....	3-17
3.6.1 ATM Planning and Implementation Constraints	3-17
3.6.2 ATM Architecture Design Factors	3-18
3.6.3 ATM Addressing and Routing.....	3-18
3.6.4 ATM Protocols	3-21
3.6.5 ATM Overlay Security	3-26
3.7 IP Connectivity Detailed Architecture	3-28
3.7.1 IP Connectivity Design Factors	3-29
3.7.2 Placement of Routers in IP Connectivity.....	3-29
3.7.3 IP Addressing	3-31
3.7.4 IP Routing.....	3-31
3.7.5 IP Overlay Security.....	3-32
3.7.6 Considerations for Connecting Contractors.....	3-33

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

- 3.8 DON ITI Architecture Plan of Action.....3-34
 - 3.8.1 Steps for Developing Detailed Enterprise ITI Architecture.....3-34
 - 3.8.2 Steps for Developing a MAN3-36
 - 3.8.3 Steps for Developing a CAN3-37

3. Conceptual ITI Network Architecture

Network connectivity is a fundamental requirement for a DON integrated enterprise information infrastructure. This chapter defines the ITI architecture that will support planning and implementation of the network connectivity segment of the infrastructure. The ITI planners address connectivity requirements by using architecture plans and templates that guide and constrain their solutions. The result is a set of network connectivity solutions that are consistent, complementary, and interoperable and that support the connectivity requirements of the DON functional missions.

3.1 ITI Network Architecture Approach

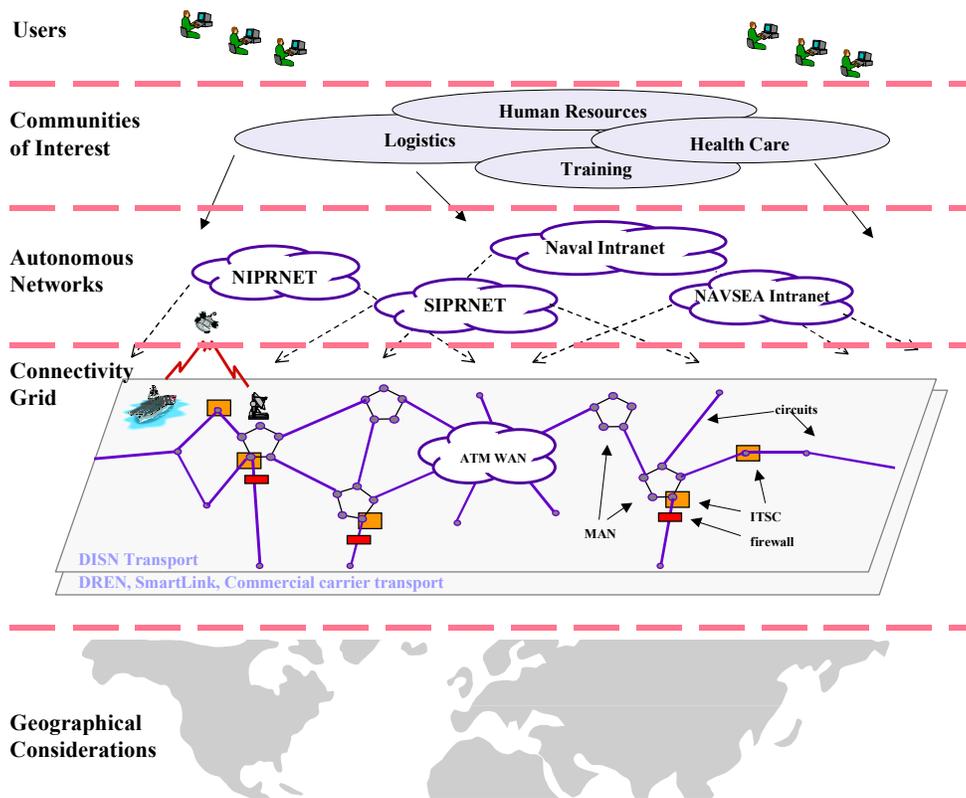


Figure 3-1. Conceptual Network Architecture

The development of the Naval ITI Network Architecture was predicated on a number of concepts, considerations, and doctrines associated with the Navy and Marine Corps mission, organizations, and operating relationships. Figure 3-1 is a composite of these factors and provides a simple view of the approach taken to develop this architecture. The five layers of the Conceptual ITI Network Architecture and the interrelationships of the layers are described in the following sections.

3.1.1 Geographic Considerations

Most of the Naval activities and personnel are concentrated at specific locations around the globe. These areas of concentration are illustrated in Figure 3-1.

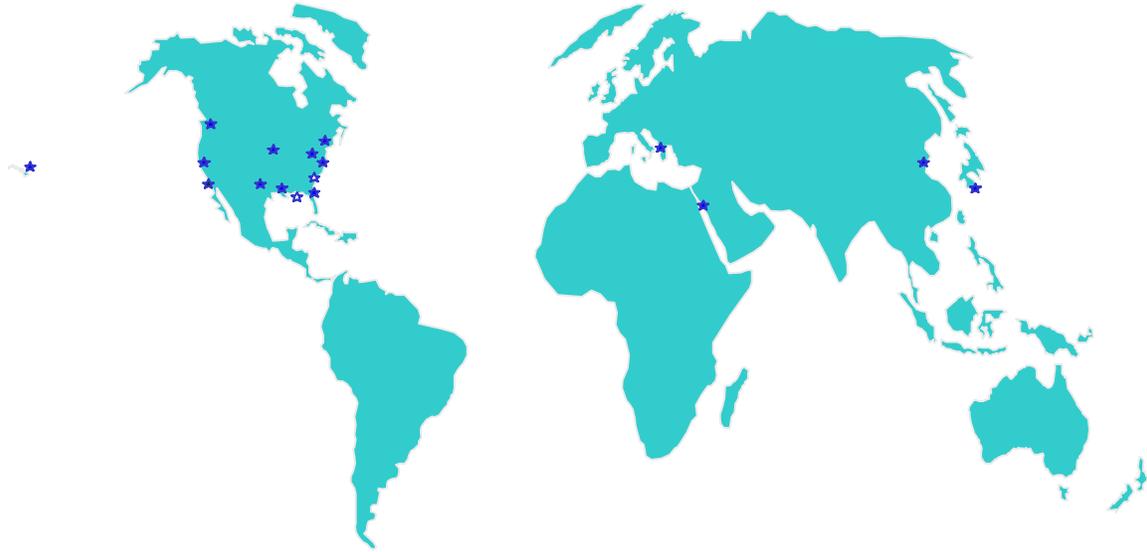


Figure 3-1. Naval Concentration Areas

These areas of concentration tend to align with metropolitan areas. Metropolitan areas typically have a high population density and an advanced telecommunications infrastructure. It is also predictable that they will have high bandwidth available and that it will be relatively inexpensive. A natural progression is to implement MANs in these areas to provide rich connectivity between Naval activities within the region.

Afloat units present a very different set of geographic considerations because of their dispersed, highly mobile, and wide-ranging deployment patterns. The satellite communications downlinks in support of afloat units are at well-defined locations around the world. When in port, shipboard units become aligned with metropolitan areas.

3.1.2 Connectivity Grid

In Figure 3-1, there is an underlying connectivity grid or common transmission fabric (includes the circuits, area networks, switches and routers, and RF links) that provides raw bandwidth, and on top of which are built all network services. The MANs are interconnected with wide area connectivity (networks or circuits). Campus networks outside the metropolitan area are linked via point-to-point links or other means. Each region has an Information Technology Support Center (ITSC) connected via the MAN. Fleet teleports are connected to the enterprise network in the same way that a base is connected (on the MAN or via a circuit). Piers are connected in a similar way, but are generally part of a base and are connected via the associated campus network.

Shipboard connectivity is provided via an RF communications infrastructure when afloat or via the pier when ashore.

The communications infrastructure is built primarily of components that are external to the Navy. In particular, the WAN and MAN connectivity is provided through a combination of telecommunications carriers and DoD- and/or Naval-owned devices. The campus connectivity, on the other hand, is primarily owned by the Navy and Marine Corps.

There are other global WANs to which the DON enterprise network must interconnect. Examples include NIPRnet, SIPRnet, and the Internet.

3.1.3 Autonomous Networks

Autonomous networks are built on top of the connectivity grid. Autonomous networks are generally independent of each other from the standpoint of media, technology, security, management, and other characteristics. They are more than virtual networks and include physical components that are unique to the particular autonomous network.

Autonomous network examples include the fleet intranet, the Naval Intranet, the SYSCOM networks, classified overlay networks, and voice and video networks.

3.1.4 Communities of Interest (CoI)

This layer represents a user-centric, geographically-dispersed grouping at or above the autonomous network layer. An example of a CoI is any functional area, such as logistics, where users need to exchange information relating to that functional area, and this exchange is across the user's normal organizational boundaries. The CoI, in cases such as logistics, will have extensive associated information management and systems applications that all DON users access to work in the functional area. A given CoI must have access to multiple autonomous networks to obtain the required connectivity services. This CoI capability forms a foundation requirement for enabling RBA and RMA.

Other CoI examples include intelligence, human resources, acquisition, training, and health care. Some, if not many, of these users may be members of multiple communities of interest.

3.1.5 Users

The users of the DON enterprise network include all Naval military and civilian personnel. It also includes contractors and other support personnel, as well as other parties that have some association with Naval activities. Because users are typically in multiple CoIs on physically diverse autonomous networks (as described above), Virtual Private Networks (VPNs) must be supported by the infrastructure.

3.2 Layered Technical Model

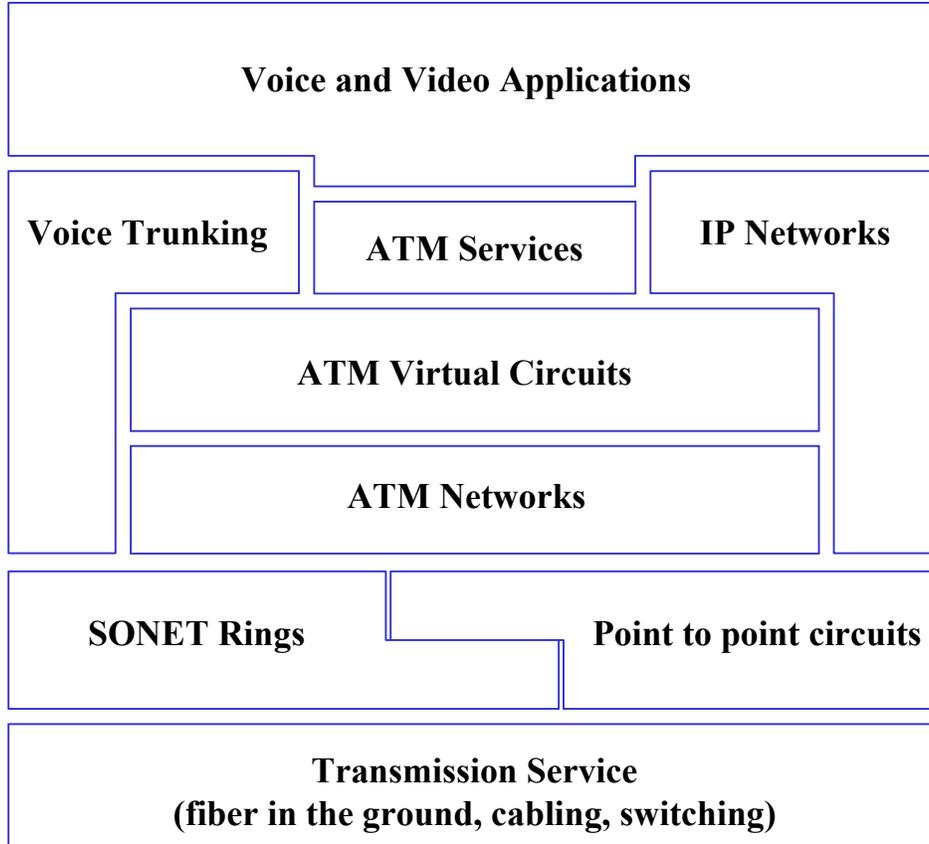


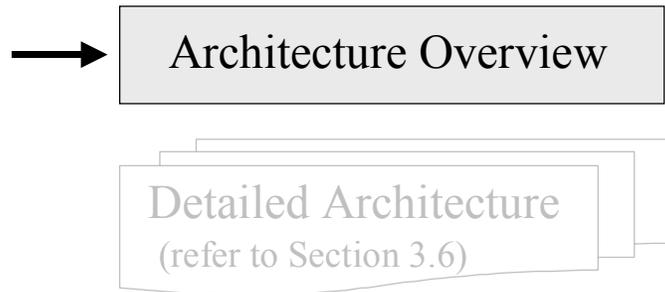
Figure 3-1. A Layered Approach to Networking Solutions

A layered technical model, depicted in Figure 3-1, was adopted as a basis for considering, organizing, developing, and presenting the ITI architecture. The IP network-centric view starts at the bottom with the transmission fabric (fiber, copper, multiplexers, etc.); on top of which are SONET rings or point-to-point circuits; on top of which are ATM networks; on top of which are ATM Virtual Circuits or links using Switched Virtual Circuits (SVCs), Permanent Virtual Paths (PVPs), or Permanent Virtual Circuits (PVCs); on top of which are IP networks. Some of the layers are optional, depending on the requirements of the autonomous networks or the communities of interest that use them.

3.3 ATM Connectivity Architecture Overview

This section presents the conceptual ATM connectivity architecture. Greater detail is provided for planners and implementers in Section 3.6. It is assumed that the reader has basic familiarity with ATM terminology, standards, and protocols.

3.3.1 DON ATM Architecture Strategy



The ATM architecture must support end-to-end ATM cell delivery and signaling. End-to-end is defined as from campus-to-campus or, in some cases, a campus may deliver ATM service deeper into the local infrastructure (in some cases to the desktop). In the latter cases, end-to-end service means supporting desktop-to-desktop ATM connectivity. To clarify, end-to-end service means that end-to-end signaling is supported to provide SVC service. In a large, complex network, this connectivity and service must be supported by a well-defined addressing and routing plan.

In the current ATM service provider market, each service provider would prefer that the customer use the service provider's addressing and routing architecture. This is acceptable except when customers are dual-homed to different service providers, such as in geographically-dispersed organizations. As a result, the customers will be compatible with one service provider but not others. This is particularly important in the DON because the entire DON enterprise network is comprised of many ATM service providers at both the MAN and WAN levels. It is a certainty that the DON will be inconsistent with some service provider's routing and addressing plans, no matter which plan is chosen.

The ATM architecture described here is independent of any specific ATM service provider and provides a mechanism to route and signal through any ATM network. The mechanism to accomplish this is "tunneling" through the various MAN and WAN ATM networks, which requires use of PVPs to create VPNs over the ATM service provider networks.

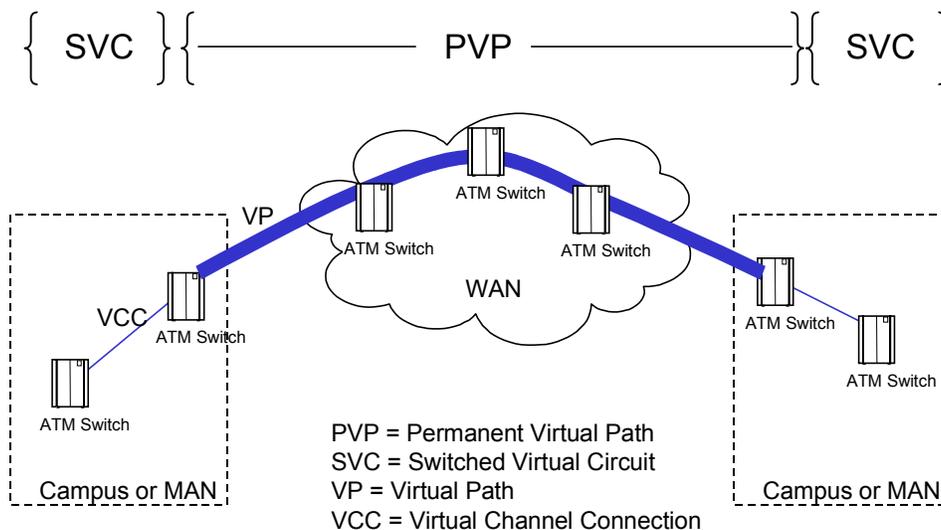


Figure 3-1. Combining PVPs and SVCs in an Architecture Solution

The VPN solution is depicted in Figure 3-1 and specifies a PVP mesh. It is for use at the levels of the network hierarchy in which the DON addressing and routing architecture are incompatible with that of the ATM service provider. The PVP mesh provides a set of fixed conduits through which the DON can establish a signaling and routing environment that supports SVCs and also support provisioning of PVCs. The advantages here are that Naval managers can still obtain SVC service from the PVP end points and they can easily add or modify PVCs as requirements change. With this approach, changes to the PVP mesh, and eventually the signaled environment, should have the least impact on campus and end users. The following PVP guidance applies.

- WANs. If the DON uses the DISA ATM addressing plan, and if the Wide Area Connectivity solution is the DISN ATM Service, then a PVP mesh will not be required for the general WAN connectivity case because the DON will be compatible with the provider's addressing conventions. However, in order to support a highly-mobile fleet autonomous ATM network, a PVP mesh will need to be established to support the fleet requirements.
- MANs. At the MAN level, some type of PVP mesh will be required as an overlay to any commercial ATM service offered as part of the MAN.
- CANs. At the campus level, the PVP solution should not be used. The campus networks are Naval owned and operated and therefore will comply with DON addressing and routing standards.

3.3.2 ATM Connectivity Approaches

There are two basic ATM connectivity cases that must be accommodated by this architecture. One is the general case of connecting the campus networks with the MANs with further connectivity to the WAN. This is a general (non-mobile) geographic situation. The other is the fleet case, which requires a high degree of mobility and must be supported by a dynamic routing architecture. The architecture presented here fully supports both of these environments. (The

intermediate case of “portable” users, in which service is not maintained during transit, is covered within.)

3.3.2.1 General (Non-Mobile) Case

The ATM switch that links the campus infrastructure to the external world (MAN) is referred to in the DON ITI architecture as a “premise” switch (or Primary Information Transfer Node (PITN) in BLII terminology).

The campus premise switch will connect to the MAN, preferably at the 155-Mbps OC-3c level of bandwidth (Figure 3-1 refers). Independent of the selected MAN technology (point-to-point, SONET, or ATM), the requirements for signaling and routing will drive the MAN architecture to a DON-controlled and -managed routing and signaling domain. For an ATM MAN, a PVP mesh overlay will be implemented to support the signaling domain between campuses.

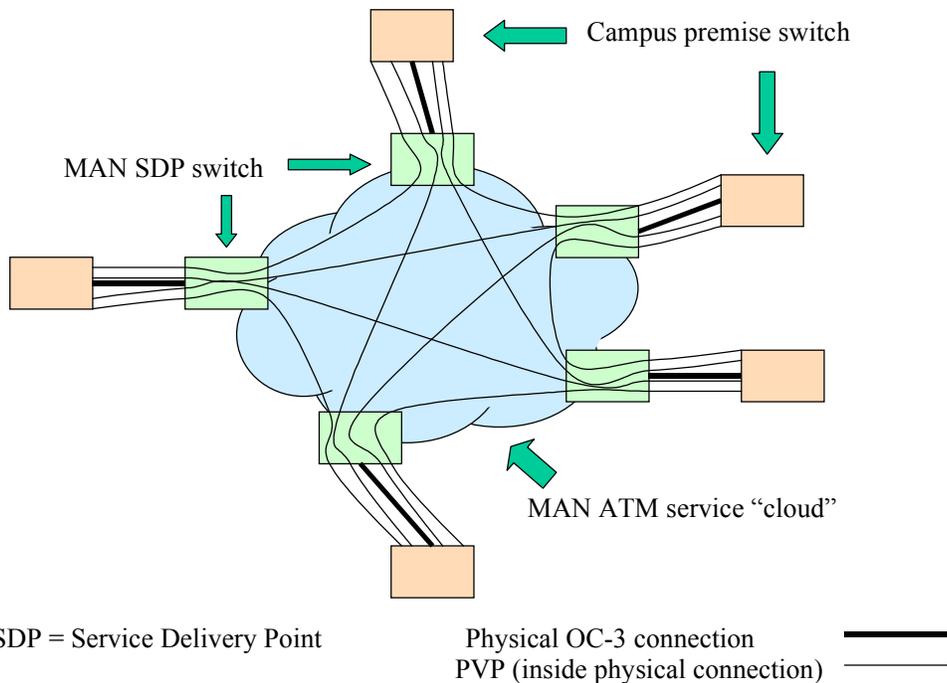


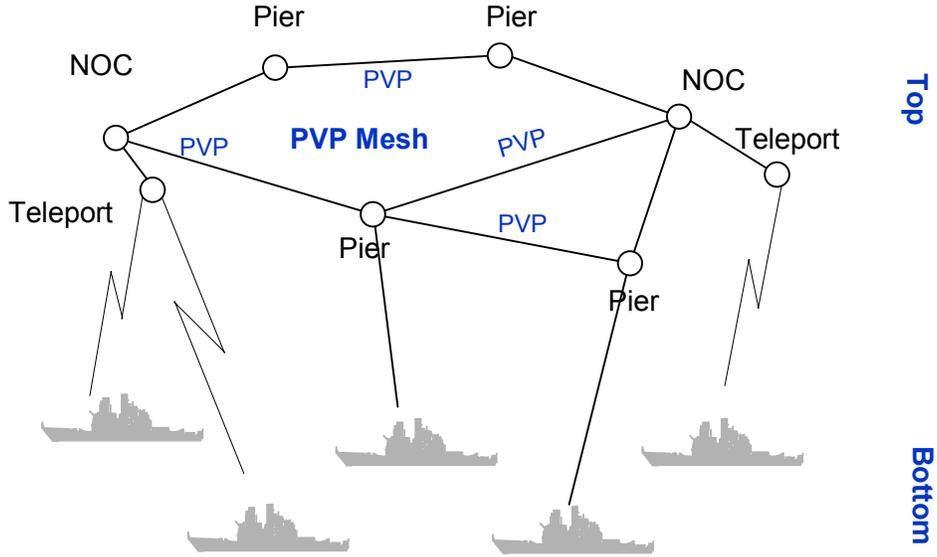
Figure 3-1. MAN Service delivered as ATM, with PVP Mesh to Campus Premise Switches

The ITSC premise switch normally will also have connectivity to the ATM WAN in order to establish connectivity with other regions. There are exceptions to this rule. For example, if the ITSC is at one campus and the WAN provider already has a POP at another campus, there may be no reason to install a new WAN POP at the ITSC if there is appropriate MAN access at the other campus premise switch.

3.3.2.2 Fleet (Mobile) Case

This section only applies to ATM-enabled ships that wish to have a native ATM connection to the rest of the fleet intranet. It does not apply to ships that have no external ATM connections - that is, those ships that are IP only.

In the fleet ATM connectivity case, our Naval mission requires full support for mobility. While fully-signaled mobile connections are desired, the standards (and existing implementations) are not yet developed. Until such time, there will be a separate fleet autonomous network based on ATM technology, with a PVP mesh linking the piers, Standard Tactical Entry Point (STEP) sites, and fleet NOCs. The mesh provides end-to-end signaling and cell delivery and allows easy reconfiguration as the mobile network moves. This does not require a full mesh because the supporting PVPs follow actual physical paths. As shown in Figure 3-1, routing is two-level PNNI with the ashore PNNI mesh and fleet switches at the top level, and the afloat units at the bottom. The boundary between the levels is at the RF links and the pier connections. This signaling domain does not extend to other Naval or even to DoD signaling domains. The IP connectivity must be separated by firewalls anyway. If signaling outside these constraints is required, it will be engineered on a case-by-case basis until a general solution becomes available.



PVP = Permanent Virtual Path
NOC = Network Operating Center

Figure 3-1. Two-level PNNI Hierarchy

Because the mobile platforms house other deployable forces such as MEUs, provisions for this multi-layer deployment must be addressed.

3.3.3 ATM Addressing Plan

The DON Connectivity Architecture incorporates a geographically-based addressing plan. To support DoD conventions and allow long-term interoperability within DoD, DON will use the

DISA ATM addressing plan whenever possible. For mobile users, DISA has offered Globally Unique Identifier (GUI) addresses that will follow the deployed forces. The GUI addresses do not meet the DON requirements and will not be used.

A MAN will obtain an ATM address block to support the MAN and all campuses within the region. Implementation will deviate slightly from the DISA address plan in that a single (or small) number of DISN nodes will be assigned in a Naval MAN. Addressing behind this node(s) will be performed at the discretion of the MAN architect. The hierarchy of the DISA plan will be maintained outside the MAN. The MAN will allocate address space from the DISA block to the campuses.

Fleet addresses will be allocated to ships based on their home port. A Naval base that includes piers should obtain enough campus network address space to allocate address prefixes to each ship that is homeported at that campus location.

Additional information on the DON ATM Addressing Plan is provided in Section 3.6.3.1.

3.3.4 ATM Routing Architecture

To exchange topology information for routing purposes, PNNI will be used whenever possible. Three levels of routing hierarchy are envisioned – campus, regional (MAN), and global. The manner in which this is achieved is dependent on the providers and technologies used at each of these levels.

From the MAN perspective, the MAN routing architecture must support routing between campuses. There will be limited prefix/masks per campus, but there may be multiple links between the campus and the MAN (for redundancy). PNNI will be the MAN routing protocol, and each campus will be viewed as a “logical peer group node.” The campus premise switch will “advertise” the appropriate campus prefix/mask to the MAN. When the campus ATM network has a rich topology, it may want to participate in the PNNI routing domain. In this case, it will be lower in the PNNI hierarchy. The choice to participate in the PNNI routing domain or to have prefixes/masks configured statically is determined by the MAN region/campus manager. In all cases, each campus will have a default route to the MAN, and each MAN will have a default route to the WAN.

Operation of the WAN requires that the prefix/mask information be known for each of the MANs and for any other entities to which it connects. This information can be inserted statically, or it can be derived through the normal operation of the PNNI protocol. Some type of dynamic scheme may be required because MANs need to have redundant paths to the WAN and only a dynamic routing protocol will allow re-routing in case of a connection failure.

Additional information on the ATM Routing Architecture is provided in Section 3.6.3.2.

3.3.5 Rationale for ATM Technology in DON ITI Architecture

Alternatives to ATM were considered during the selection of the networking technology; the ITI IPT selected ATM technology based upon a number of factors.

Today, time division multiplexing (TDM) technology is used in many Naval applications (using multiplexers) to provide multiple virtual connections over a single physical channel. Bandwidth allocation is fixed within each virtual channel, regardless of whether it is carrying any data, so it is impossible to “burst” to higher bandwidths on demand. ATM technology, on the other hand, provides the capability of multiplexing virtual channels and offers the efficiency of statistical multiplexing in which bandwidth is consumed only when needed and is shared on a demand basis across the entire physical channel. This allows bursting to the physical line rate or to whatever bandwidth is unused by other channels. As a result, ATM provides significant efficiency gains over TDM technology, even with ATM technology’s 10 percent overhead for header information.

DON requires that:

- ATM connectivity exists all the way to the desktop in some Naval network installations that have implemented IT21. This requires end-to-end ATM connectivity, even over the wide area.
- Multiple autonomous or virtual networks operate over the DON physical network infrastructure. ATM technology offers support for virtual channels with various classes of service and this provides the enabling capability for constructing the virtual networks. Alternatively, IP technology can be used to construct virtual networks on top of IP networks, but the IP approach is less efficient (it may increase overhead up to 50 percent) and fails to offer some required classes of service, such as circuit emulation.
- Networks provide a constant bit rate (CBR) or circuit emulation service for carrying synchronous trunks over a common communications fabric. Applications must include voice trunking between PBXs and bulk encrypted links using traditional type 1 encryption devices such as KG-84. ATM provides this class of service.
- Fast end-to-end key-agile (as opposed to bulk) type 1 encryption carry classified traffic over the common network. The one device that can achieve that operates over ATM networks (the KG-75).
- Significant scalability exists. ATM provides this without the need to replace the transport media (fiber).

The ATM decision is also examined in terms of other competing technologies. Compared to ATM, fast Ethernet and gigabit Ethernet technologies are better understood, appear to be scaleable, and are less expensive per drop. However, ATM was chosen for the “core” network technology, not necessarily for the “edge” networks. The appropriate technology solutions for connecting edge systems (e.g., hubs) are a separate issue. The advantages of ATM in the “core” network are exhibited on the backbone (wide area, metropolitan area, and some campuses) where scaleability is needed and many communications channels can be serviced over a single physical network. Additional advantages of ATM over the other technologies are addressed in the following.

- Ethernet Frame Size. The Ethernet class of protocols has a maximum frame size of only 1500 bytes. This frame size is too small to work well on networks near the border of the wide area and on the WAN itself. If the technology on or near the WAN allows a larger frame size,

then fragmentation of the packets will occur at the edge networks, and this will cause significant performance degradation on the WAN. The DON architecture must eliminate such problems because there is no way to engineer around these limitations. Hence, we have the goal of delivering ATM service to each campus from the wide area or metropolitan area networks.

- Ethernet Scaleability. With ATM, only the optics at the interface level must be upgraded to achieve a faster optical carrier (OC) rate. With Ethernet, different cabling may be needed to migrate to fast Ethernet, FDDI, or even gigabit Ethernet if the current copper cable plant is inadequate for the higher bandwidths.
- Packet over SONET (POS). POS or Packet over wave division multiplexing (WDM) have been the technologies of choice in a number of new high speed demonstration networks that are not using ATM. While these new transport protocols offer significant promise for the future, they currently lack the maturity as well as other features and robustness required to meet Naval requirements.

3.4 IP Connectivity Architecture Overview

The conceptual IP connectivity architecture is presented here. Greater detail useful for planners and implementers is described later in Section 3.7. It is assumed that the reader is familiar with IP terminology, standards, and protocols.

3.4.1 Strategy

This IP architecture must support ubiquitous, end-to-end IP connectivity within the enterprise network and must support full access to the global Internet. Determining the IP architecture is relatively straightforward because the DON has significant experience upon which to draw. Moreover, the conventions, protocols, and implementations are mature.

The IP challenge is to provide an architecture that accommodates the separate autonomous networks, particularly the “fleet intranet”. The underlying ATM infrastructure enables the construction of these autonomous networks using virtual circuits as links between the routers in a given autonomous network. Each of these autonomous networks is a separate routing domain unto itself, and the capability must exist to route IP traffic from one autonomous network to another and to the external world. An additional constraint is that some autonomous networks must be protected from other autonomous networks using firewalls or similar technology. The fleet intranet introduces the necessity and challenge of supporting network mobility. Ships constantly change geographic positions, and their connectivity is dynamic both in topology and in bandwidth. Ships also require a strong security perimeter. For these reasons, there will be a specific architecture discussion for the mobile networks.

In general, there will be a minimum of detail regarding the various autonomous networks. The emphasis is on enabling the implementation but not to dictate the internal architecture of each network. That is left up to those who assess the requirements and develop the specific design for the individual autonomous networks.

3.4.2 IP Connectivity

3.4.2.1 General Case

The campus router that connects to the MAN will be termed the “premise” router for the purposes of this discussion. Figure 3-1 shows the relationships of the premise router connecting an autonomous network to the MAN. In actuality, this router physically connects to the premise ATM switch over a direct link.

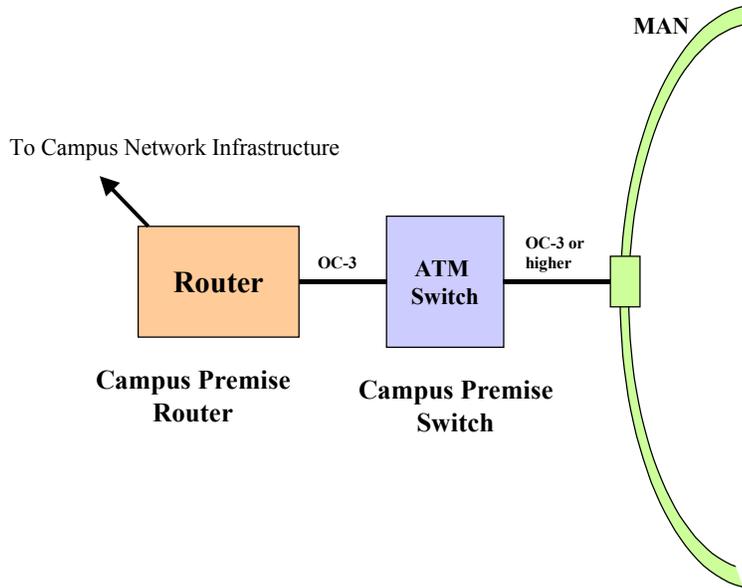


Figure 3-1. Campus Premise Router Connectivity to the MAN via Premise ATM Switch

The premise router needs connectivity to the other campus networks via the MAN. This should be done with a full Virtual Circuit (VC) mesh. The preferred choice is SVCs, which is accompanied by the challenge of figuring out how to perform the IP-to-NSAP mapping. Static mapping of the routers is one choice and probably the most stable, but another choice is RFC 1577 “classical IP.” The disadvantage of RFC 1577 is the instability of the Address Resolution Protocol (ARP) server as a single point of failure. Other choices are MPOA or IP Over NBMA (Non-Broadcast Multiple Access), which are both supported by a number of router vendors. Still another choice is to install a full PVC mesh as a single logical IP subnet (LIS) between the campus premise routers. This is a very stable configuration but is not very scaleable – the number of resulting PVCs is on the order of N^2 (where N is the number of participating routers) and each requires significant manual configuration.

As a means of summarizing the above, the ITI IPT recommends that the preferred campus-to-MAN architecture be full mesh connectivity to the MAN via SVCs, static mapping of routers, and MPOA. Additional IP connectivity guidance is discussed in section 3.7.

3.4.3 Autonomous Networks

From an IP perspective, the “autonomous networks” each stand alone from other IP networks in the Naval enterprise. They have their own routing architecture and have their individual Interior Gateway Protocol (IGP) operated as a separate routing domain. The IP routers in an autonomous network are managed separately from the rest of the enterprise.

Autonomous networks can be connected to each other by using an appropriate protocol, e.g., Border Gateway Protocol (BGP) 4. The users connected to one autonomous network should be able to communicate with the users on another autonomous network. The separate autonomous networks are all logically part of the global Internet.

Each autonomous network will have its own security profile and may include security controls at its perimeter. Network managers should employ the appropriate security mechanisms, i.e. zone 3 or zone 4 protection in the defense-in-depth model.

This architecture recommends that any autonomous network that has a presence in a given region should perform peering using an Exterior BGP (E-BGP) mesh. This is similar to what is currently being done in the Washington, D.C., area in the National Capitol Region Metropolitan Area Network (NCR MAN). Traffic from one Naval autonomous network to another should use “hot potato” routing in which a packet is moved to the destination network as soon as possible because the destination network can best move the packet to its ultimate destination.

3.4.4 Fleet Intranet

The fleet intranet, which is a special case of autonomous networks, must be operated as a separate routing domain because of the highly dynamic nature of the fleet environment. In the fleet environment, ship communication links change often and when they do change, automation must replace human intervention to produce fast convergence of the routing service.

This architecture specifies routing the fleet IP traffic using a single Open Shortest Path First (OSPF) domain. The fleet support infrastructure ashore (NOC, piers, STEP sites) is considered the “backbone area” in the OSPF sense (also known as area 0). There is a separate OSPF area for each ship. Each ship announces a single aggregated prefix and those of any MEUs, air wings, or other deployed forces that are onboard. Large routing updates to the ships are collapsed into a default advertisement to save bandwidth. If additional hierarchy is required, a ship could be a separate routing domain (OSPF with multiple areas) and use BGP to peer with the “backbone” for redistribution.

Connectivity to the external world is accomplished through the NOCs/NCTAMS. They announce the aggregate fleet address/mask from outside the fleet firewalls. From outside the fleet intranet, it is impossible to determine ship locations. Traffic originating from the external world destined to the fleet must first go to the nearest fleet border (NOC firewall or some pier firewalls), and then traverses the fleet intranet until it arrives at its ultimate destination.

3.4.5 IP Addressing Plan

The goals of the addressing plan are to aggregate address space and to use address space efficiently.

The most obvious places for aggregation are at the campus and MAN levels. Addresses on campus should all be summarized with a single address/mask, and addresses for all MAN customers should aggregate to a single address/mask. There may be exceptions (legacy installations, unexpected growth, etc.), but these should be minimized. To perform this aggregation, the MAN operators should obtain a Classless Inter-Domain Routing (CIDR) block large enough to support the region and allocate sufficient address space to the campuses to support the long-term needs of that campus. As a rule, a “class B” CIDR address block is adequate to support many campuses in a region and is inappropriately large for allocation to a single campus.

Unfortunately, large blocks of address space are difficult to acquire. An alternative is to use network address translation (NAT) technology at the campus or unit. With this alternative, non-unique address space is used internally but is translated into the globally unique space externally. This approach works quite well, and is sound from a security perspective. The NAT alternative is recommended whenever it is practical.

Many campuses in the DON already have allocated IP address space and do not want to change their IP addresses. Accordingly, these addresses in the near term will need to be globally routed.

A more prudent approach is for the Navy NIC to obtain a large block of address space and allocate it appropriately to each of the MANs. This would seem to be difficult initially, but once implemented would save significant management time.

To maximize efficiency of the address space usage, DON IT managers are encouraged to sub-allocate address space in appropriately sized blocks. For example, a point-to-point link should only be given a /30 (30 bits of mask, out of 32). A small IP subnet should only be given the amount it needs for the long term. Every entity should not automatically receive a “class C” chunk, (/24), unless it is justified to do so.

3.4.6 IP Routing Architecture

Figure 3-1 depicts the OSPF and BGP routing in the IP routing architecture. For the networks described, the preferred interior gateway protocol (IGP) is OSPF. Some autonomous networks may choose other protocols, but the MANs and campus networks will use OSPF. Each MAN will be configured as a “backbone area,” or area 0. Each campus on the MAN will be a separate OSPF area. The fleet intranet will also use OSPF.

BGP4 will be used between MANs and configured in a full E-BGP mesh. Because the WAN is implemented with a full VC mesh between the ITSC premise routers, there is no requirement for an additional IGP on the WAN. The BGP mesh is sufficient.

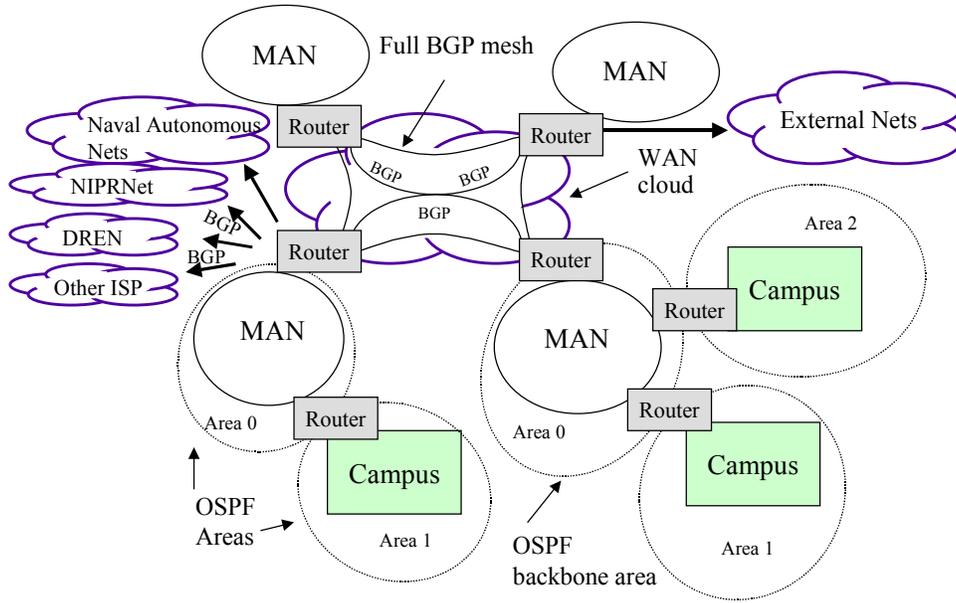


Figure 3-1. BGP and OSPF Routing in Enterprise Architecture

Between autonomous networks, BGP version 4 will be used. BGP4 will also be used to peer with the external world. Each autonomous network should only announce networks that are within its autonomous network. It should not announce networks that it learns from other BGP peers; it should not serve as a “transit” network. (There may be minor exceptions.)

Routing information between the IGP (OSPF) and the EGP (BGP4) should not be redistributed. Networks announced via BGP should be explicitly configured to ensure the highest network stability.

Naval autonomous IP networks should all peer at the ITSC to minimize routing distance.

3.5 Network Connectivity Security Overview

Defense in Depth is the preferred information protection approach for the DON. In Defense in Depth, information protection mechanisms are applied in multiple, complementary, and redundant locations to collectively form a system architecture. Defense in Depth is a layered approach analogous to the multiple zones around a carrier battle group. The layers provide maximum resistance to attack and minimize the likelihood that a single flaw can lead to a security compromise.

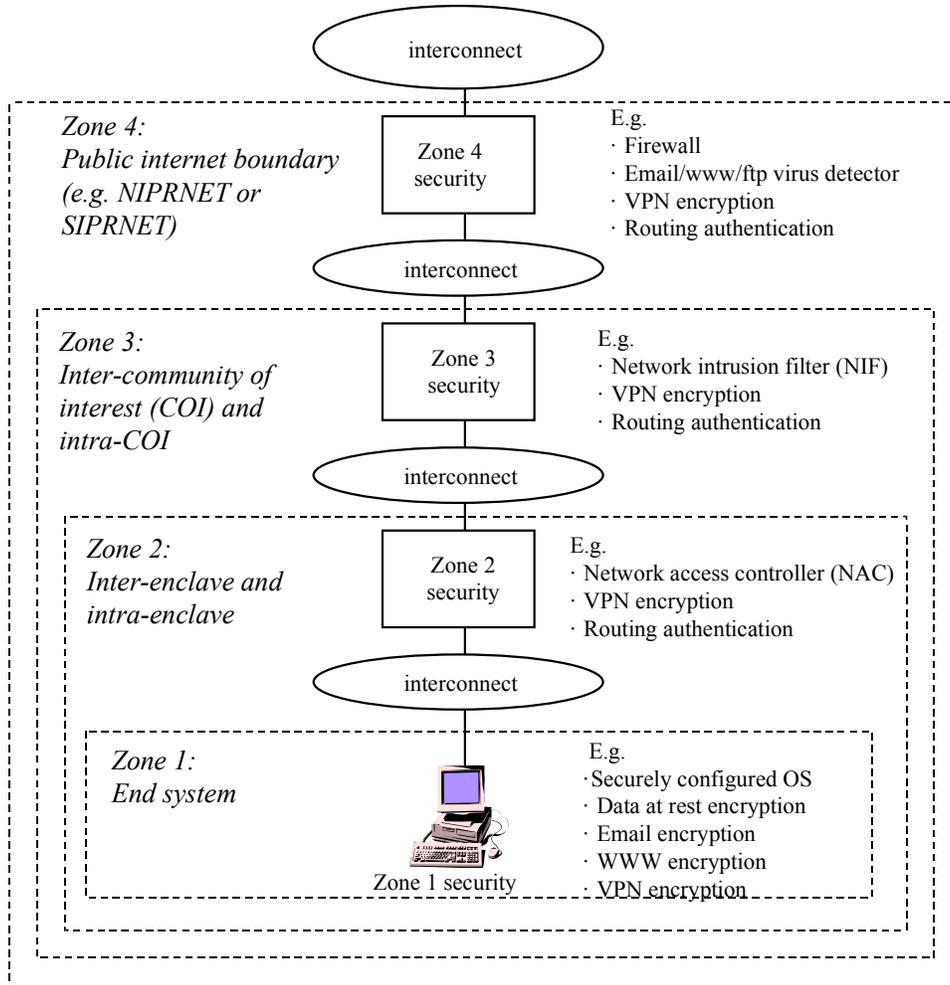


Figure 3-1. Generic Framework for Defense in Depth

A generic framework for defense in depth is illustrated in Figure 3-1. Four zones of defense are defined in this framework; these zones may be logical and not necessarily physically separate. Across the framework, numerous security mechanisms are used to protect information. In addition to the protection mechanisms, certain infrastructure components are required to build secure DON information systems. The most critical of these is a public key infrastructure (PKI) that provides identification and authentication mechanisms and encryption mechanisms for application across the various zones.

Appendix E provides a detailed account of the four security architecture zones and how the information protection components are applied to each zone and information dimension. Specific instances of how security is applied to area networks and network services is contained in each individual section of the document, for example, in this chapter – within ATM and IP, in Chapter 4 – within Directory and Domain Name System, and in the appendices – within each template.

3.6 ATM Connectivity - Detailed Architecture



As introduced in section 3.3.1, a successful implementation of ATM in the DON requires well-designed and coordinated planning across all Navy and Marine Corps organizations. Design factors to be resolved include address and routing strategies for interoperating in the joint DoD environment, service provider support of addressing schemes and ATM features/protocols such as Permanent Virtual Paths (PVPs) and Private Network-to-Network Interface (PNNI), and vendor support of ATM standards-based products. These and other implementation issues must be carefully evaluated in order to formulate an enterprise strategy that positions the DON to move toward the desired performance-oriented, cost savings potential of ATM.

The purpose of this section is two-fold. First, it outlines the ITI IPT's strategies and conclusions for choosing particular ATM architecture components. Second, it provides more detailed information on specific ATM protocols and addressing and routing plans.

3.6.1 ATM Planning and Implementation Constraints

The implementation is currently constrained by the following issues.

- Signaling (as opposed to tunneling) of virtual circuits over a public wide area carrier is difficult because service providers are not yet in full agreement on which standards to implement and when. For example, PNNI is not supported through the WAN but can be interfaced via Public UNI. Broadband InterCarrier Interconnect (B-ICI) is not implemented across service provider boundaries. Chosen solutions must be available to implement today and adaptable to future service-provider signaling offerings.
- ATM protocols are not only being introduced at an unprecedented pace but are still evolving. Chosen solutions must be fielded reliably today but yet take advantage of future, more robust protocols.
- Security solutions for ATM are still evolving. Basic NSA-approved Type 1 encryption exists today, e.g., the KG-75 Fastlane and the emerging KG-175 TACLANE, but traditional firewall-like devices are limited.

3.6.2 ATM Architecture Design Factors

Switched Circuits: End-to-end switched virtual circuits (SVCs) between any two DON ATM end systems is a requirement. Full signaling is sought from service providers but Permanent Virtual Paths (PVP) will suffice in the interim.

Ubiquitous ATM Service: Given the end-to-end SVC goal, it is desired that ATM service will be provided to every Naval base (campus), as appropriate.

Redundant MAN Links: No single switch or physical circuit will be a single point of failure for MAN access to the WAN.

Redundant CAN Links: No single switch or physical circuit should be a single point of failure for CAN access to the MAN (this requirement is driven by each particular site's mission).

3.6.3 ATM Addressing and Routing

3.6.3.1 Detailed DON ATM Addressing Plan

Implementation of the architecture described above requires a complete ATM addressing plan. The DON ATM Addressing Plan is based on the DoD ATM Addressing Plan as described in MIL-STD-188-176, Standardized Profile for Asynchronous Transfer Mode (ATM), dated 21 May 1996 and DISA implementation instructions contained in DoD ATM Addressing Plan, dated 17 April 1998.

The DON Network Information Center (NIC) is the single DON point of contact to the DoD NIC for both IP and Network Service Access Point (NSAP) registration. The NIC has established procedures to register ATM End System Addresses (AESAs).

DISA assigns the routing domain fields of the ATM address and the DON NIC assigns area field values to DON regions. Because DISA assigns the routing domain field (bytes 5 through 8) geographically, each Navy MAN will have a different address prefix.

For fixed sites, the NSAP addressing will be geographically based using the DISA NSAP addressing plan (as modified for DON MANs).

For the fleet autonomous networks and components, home-port geographically-based addresses will be used. A given autonomous network/component will aggregate to a single prefix/mask that will follow the unit; re-addressing of network devices is not acceptable. The network itself will adapt to the change in component location. Further, nesting of components will be supported, e.g., an MEU could deploy from its home base to a ship that in turn deploys to another area — requiring no change in addressing of either.

The System ID portion of the address range is unique within the MAN.

All addresses on a given campus will be summarized by a single prefix/mask and advertised to the MAN. Also, all addresses within a geographic area covered by a MAN and the campuses it

serves will be summarized into a single prefix/mask and advertised to the WAN. This hierarchy makes the routing problem more tractable.

An exception to the above is a campus that is dual-homed to a non-Naval ATM network. In this case, the campus may have to follow the addressing conventions of another and request that all providers accept those addresses for networks on which signaling support is required. This may be difficult when one service provider is unwilling to carry another service provider's NSAP address. Should this case arise, the dual-homed site must choose an NSAP address that is acceptable to all service providers. (DISA, for example, has agreed to carry non-DISA addresses on the DISN on a non-operational case-by-case basis.)

The scheme for fleet address advertisement is still being designed. Some of the constraints are to not require re-addressing of network components, minimize and hopefully eliminate the need to manually propagate static routes, not disclose the state or location of deployed forces, and allow for efficient routing between warfighting elements.

The DISA ATM addressing allocations provide a 24 bit (3 byte) field called "Area" for subscribers to use in assigning NSAP prefixes to switches. These would be applied at the MAN level. The MAN will then allocate address space down to the individual sites (campus networks). Figure 3-1 gives an example of the allocation hierarchy that can be used for a typical MAN.

Site (6)	classification (2)	Building (8)	Switch (8)
----------	--------------------	--------------	------------

Figure 3-1. Campus ATM Address Allocation (without Pier)

The allocated 24 bits are used to establish an addressing hierarchy in a region. Six bits allows 64 distinct sites to be connected to the MAN. Two bits of classification allows for four different classification levels at any given site. Eight bits allow for 256 buildings at a site. And finally, eight bits allow for 256 switches in a given building.

There is substantial latitude for variation in this allocation hierarchy. A given site can adjust its share of the suggested address space to meet local requirements. For example, if the MAN takes six bits for site designator, the site can allocate the remaining 18 bits, depending on local requirements. A site with a few large buildings may have a very different allocation scheme than a site with hundreds of small buildings spread over a large area.

Another variation is at the MAN level. If there is a small number of large sites in the region, along with a large number of small sites, then a dual allocation scheme could be used. One example is eight bits of site prefix for small sites and three bits for large sites. Thus, there would be $8 + 256 = 264$ sites total where a large site would have 21 bits to assign instead of 18.

Sites with piers for ship home-porting have additional requirements and are represented in Figure 3-2. Because the ship ATM prefixes are globally routed, each ship requires its own globally-unique prefix. From a campus network perspective, the ship can be considered to be a building. However, at the shipboard level there are multiple classification levels, so the bit allocation must be adjusted to meet that requirement.

Site (6)	classification (2)	Ship (8)	classification (2)	Switch (6)
----------	--------------------	----------	--------------------	------------

Figure 3-2. Campus ATM Address Allocation (with pier)

This allocation supports a site with up to 256 ships home-ported. This is more than enough, because the largest homeport is currently assigned 81 ships. This scheme also allows for four classification levels and up to 64 switches per classification level. Variations are supported as previously described.

3.6.3.2 Detailed DON ATM Routing Architecture

The MAN prefix/mask aggregate must be “advertised” within the WAN. Depending on the WAN provider, this might be done any number of ways. A fully-signaled, automatically-advertised solution such as full PNNI is strongly desired. Given the immobile nature of MANs, however, static addressing through Public UNI or a PVP mesh is acceptable in the short term.

For the fleet, the ashore components of the fleet autonomous ATM network will use PNNI routing and will be a peer group at the top of an independent routing hierarchy. These components will be part of a two-level PNNI hierarchy with ships at the bottom and the ashore infrastructure at the top. Each ship will be a logical group node, but each group can also be a logical node, which allows it all to collapse into a single level. Additional levels can be added at the appropriate time (battle group, theater). This architecture will be very simple and will only create additional hierarchy when it makes sense.

The simpler implementation is shown in the flat hierarchical routing depicted in Figure 3-1.

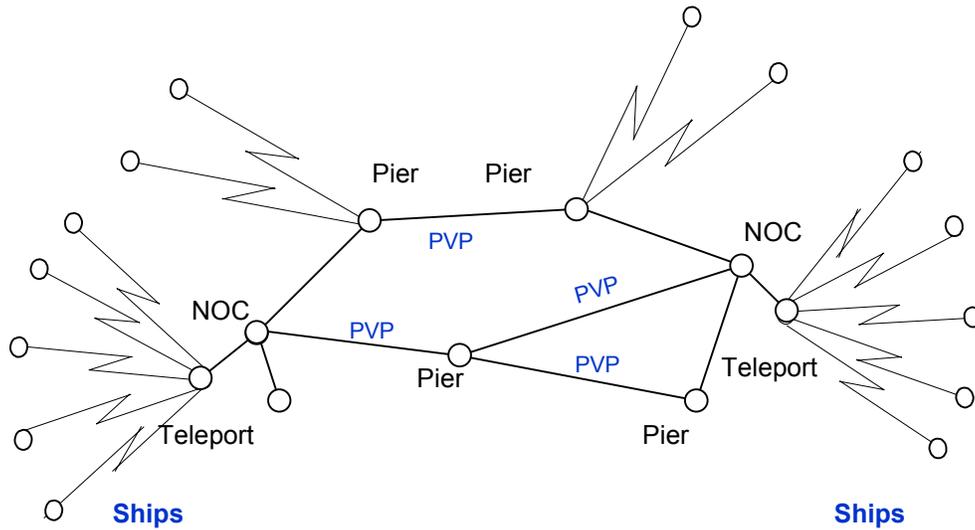


Figure 3-1. Flat PNNI Routing

A far more complex and long-term implementation is shown in the deep hierarchical routing depicted in Figure 3-2.

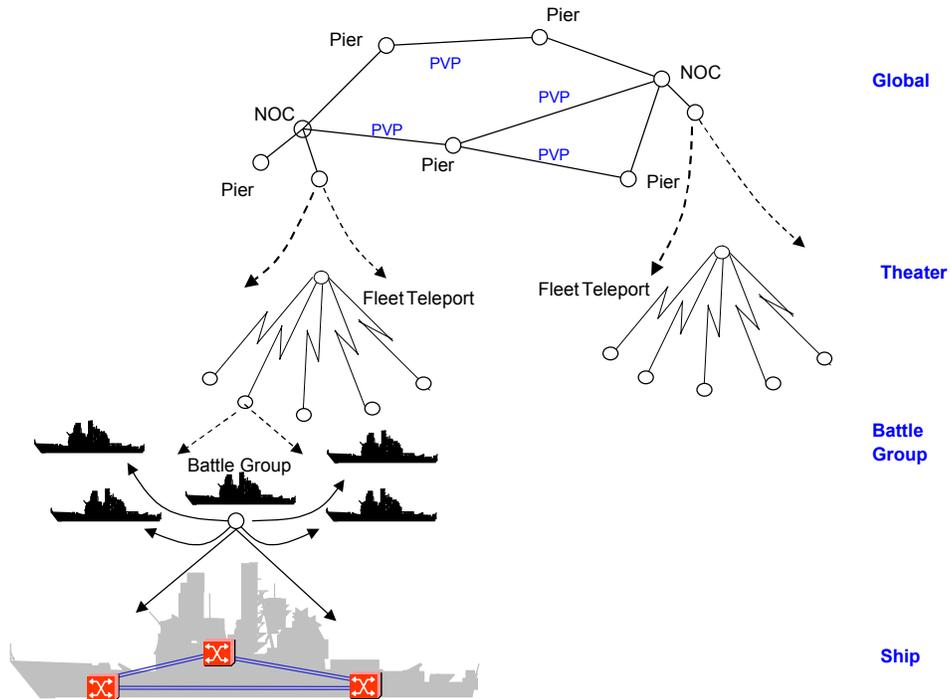


Figure 3-2. Deep PNNI Routing Hierarchy

While a deep routing hierarchy like this is possible, it is unclear whether it will be necessary or even beneficial. The advantage comes only if routing topologies can truly be summarized at the levels shown above. Unfortunately, the fleet is not static or predictable. Since the additional complexity offers little or no advantage, the simple flat hierarchy makes the most sense for the foreseeable future.

Whether a ship is connected via a pier or via an ATM-based satellite link, it will announce its ship prefix/mask via PNNI to the fleet autonomous network (ATM PVP overlay). The fleet network will then have full routing information for all ships, so signaling will take a direct path. When signaling originates outside the fleet network, the call will be routed to the region where the ship is home-ported. At that point, it will go directly to the ship if it is pier-side or will be routed to the fleet network if it is deployed.

3.6.4 ATM Protocols

A brief description of selected ATM protocols is contained in this section. More detailed information is available at the ATM Forum web site (<http://www.atmforum.com/>).

3.6.4.1 Private Network-to-Network Interface (PNNI)

This section on PNNI is provided for consideration by network planners. PNNI is critical to realizing the intended performance, interoperability, and ease-of-use of ATM.

PNNI is a dynamic routing protocol defined by the ATM Forum and specified for use between ATM switches and between groups of ATM switches. PNNI includes two categories of protocols.

- **Routing.** The ATM PNNI protocol allows each switch in the network to share topology information. This information is used to determine the best paths for an end-to-end route through the ATM network based on parameters such as Quality of Service (QoS) and advertised bandwidth use. A key feature of the PNNI mechanism is its ability to automatically configure itself in networks in which the ATM address structure reflects the hierarchical routing topology.
- **Signaling.** PNNI also defines a network-component protocol for signaling or message flows used to establish point-to-point connections across the ATM network. This protocol is based on the ATM Forum UNI signaling and uses the routing services of PNNI Routing to provide optimized end-to-end path selection for user calls.

A hierarchical relationship among ATM switch addresses is essential to successful operation of the DON Technology Infrastructure. Besides supporting a logical ATM addressing structure based on the network topology, the hierarchy makes ATM call routing simpler and more efficient. The hierarchy also enables scalability to a large world-wide ATM network.

The hierarchical addressing structure differs from the flat approach traditionally used in IP routers. In a flat address space, every router must maintain information about the topology of the entire network (or at least a “default” path to a more knowledgeable router). This results in excessively large routing tables and inefficiency. Networks based upon the first three bits of an IP address, together with a subnet mask, was presented as the original solution to this problem. Later, the notion of autonomous system numbers and the use of different routing protocols within IGP and between EGP networks was employed to contain the growing size of routing tables. Finally, classless interdomain routing was introduced primarily to conserve IP addressees and to further contain the growth of routing tables. In a hierarchical address structure such as PNNI, routing information is aggregated when detailed knowledge of the topology is not needed and a single routing protocol is used throughout the network.

The DON ATM hierarchy is created through the assignment of Network Services Access Point (NSAP) addresses. The network uses longer unique prefixes of the NSAP address for more detailed routing information (lower in the PNNI hierarchy). With the bits available in the NSAP address, there are levels in the hierarchy to enable efficient routing.

Each DON enterprise switch must have a unique address prefix/mask which is different from the address prefix/mask assigned to any other switch or device in the same network. All subordinate switches and ATM devices must have addresses based on the primary switch in the hierarchy.

Hierarchical addressing is the point of the PNNI Phase 1 (PNNI-1). The PNNI-1 protocol supports the concept of peer groups. Each peer group consists of multiple ATM switches which operate in the same hierarchical level. They communicate through a peer group leader who represents the peer group at the next higher layer of the hierarchy.

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

It may be necessary to establish PVPs between Navy and Marine Corps regions to facilitate top-layer PNNI routing within the DON. These PVPs might not be necessary if the architecture consisted of a collection of MANs using DISA (or an alternative WAN ATM service provider) as the communication infrastructure for interconnectivity between MANs. The goal of the DON enterprise backbone services is to provide a single PNNI peer group hierarchy.

3.6.4.2 Multi-Protocol over ATM (MPOA)

MPOA is the protocol for interoperability of both ATM and non-ATM devices in an ATM environment. While not emphasized in the early stages of this architecture strategy, MPOA will become increasingly important to improve DON network performance.

MPOA is the first standards-based protocol that allows routed networks to take advantage of the benefits of the ATM network (i.e., lower latency, performance, and scalability). It expands on schemes such as LAN Emulation, Classical IP/ATM (RFC 1577), and Next Hop Resolution Protocol (NHRP) to create a standardized notion of a virtual routing functionality integrated within a high-speed, dynamically-switched ATM network. MPOA, by allowing traffic to be forwarded to its destination over an ATM virtual circuit, reduces the cumulative latency in a multi-protocol routed network by reducing the number of intermediate points where packet processing must be performed. MPOA also allows non-ATM network layer protocols to take advantage of the QoS features of ATM.

The MPOA “switched-routing” methodology for IP consists of the following components:

- **ATM Network Cards/Drivers** allow ATM directly-connected hosts to send and receive traditional IP datagrams to and from an MPOA-capable network and interoperate with non-ATM hosts which are indirectly connected.
- **IP Switches** integrate routing (layer three) intelligence into an inherently switched (layer two) transport infrastructure (hence switched routing). The IP switches perform the packet-forwarding function on the non-ATM edge devices or hosts in an MPOA network.
- **Route Servers** maintain MAC address routing tables and act as distributed directories to translate the destination MAC addresses to the ATM address of the destination switch. Once the destination ATM address is known, the source device can establish a VCC directly to the destination device. This “cut-through routing” has the performance impact of eliminating all but a single IP “hop” from any source to any destination within the DON enterprise network. The resultant IP “diameter” of the network is 1, which is the smallest diameter possible. Participation in “cut through routing” is controllable, i.e., designers can choose which ATM devices are or are not allowed to accept “direct” VCCs through policy or flow patterns. The MPOA specification defines a virtual routing framework that separates the route calculation function from the actual layer 3 forwarding function. The MPOA capable edge device or host provides layer 3 forwarding of packets.

MPOA enables improved performance and interoperability of both ATM and non-ATM networks in wide area ATM network domains, not only in the MAN template environment but also in the WAN environment. It has even greater implications for the campus LAN environment.

MPOA is a logical evolution of ATM LANE and is upwardly compatible, i.e., a LANE host/edge device can participate in an MPOA-capable network.

3.6.4.2.1 Design Factors

MPOA is required for implementation in ATM networks based on this architecture guidance and the Information Technology Standards Guidance.

The WAN must support MPOA for TCP/IP.

3.6.4.3 LAN Emulation (LANE)

ATM LANs, non-ATM attached LANs, and accompanying end stations need to communicate over ATM networks. Prior to the deployment of MPOA, LANE was the acceptable technique to enable communication in this mixed environment. This is possible because the ATM network "emulates" the characteristics of broadcast LANs (e.g., Ethernet, FDDI, and Token Ring) and performs Media Access Control (MAC)-to-ATM address resolution. Implementation of LANE consists of a LAN Emulation Server (LES), a Broadcast and Unknown Server (BUS), and a LANE client (LEC). Each component resides in one or more ATM end systems or edge devices. Although it is possible to implement these functions in an ATM switch, this is usually avoided (as explained below). The LES and BUS work with LECs, typically Ethernet hubs with an ATM uplink, to provide layer 3 bridging functionality across the ATM network along with directly attached ATM hosts.

While it is possible to provide support for legacy networks such as Ethernet or Token Ring via LANE services in the core, most LAN emulation will be implemented in the local edge switches and ATM-connected hosts.

An edge switch is on the boundary of an ATM network. Typically it is an ATM switch which supports legacy networks via LANE services or an ATM end system workstation/host. For smaller networks, the edge switch can connect directly to the core.

The decision as to where to run LANE depends on the particular ATM devices that handle the LANE processes. For autonomous networks and/or communities of interest, the LES/BUS services should be implemented on ATM end systems (or on the edge switch) close to the units that are connecting to the various emulated LANs (e-LANs). (If desired, redundant LES/BUS servers should be employed.) In addition to required redundant hardware, there need only be one LANE Configuration Server (LECS) per ATM network.

Every e-LAN functions independently by using its own LES/BUS. Interconnecting multiple e-LANs requires a router typically called a one-armed router, named so because there is only one physical interface (usually 155-Mbps OC3 fiber) from the ATM switch to the router. Instead of routing between multiple physical interfaces, the one-arm router simply routes the layer 3 protocols onto multiple e-LANs connected through the same interface. The router can be a card that is inserted into an ATM switch or it can be a stand-alone router.

The natural upgrade path from LANE is MPOA.

3.6.4.4 Voice and Telephony over ATM (VTOA)

VTOA allows a traditional voice circuit to be dynamically signaled and mapped into an ATM virtual circuit (as opposed to circuit emulation in which voice circuit trunks are provisioned — not signaled — through ATM networks). VTOA provides improved bandwidth effectiveness because the voice calls are created and terminated on demand over the same network that is transporting data. VTOA also assures voice quality by employing voice adaptation techniques and Quality of Service support. VTOA reduces cost by consolidating networks for voice, data, video, and imagery via an ATM network.

WAN voice communications have traditionally used analog and digital leased lines between major locations supported by PBX switches or by public Centrex services using tandem PBXs to minimize facility costs. The design of PBX networks has changed little over time and produces poor efficiency (e.g., bandwidth use, compression degradation, blocked calls) and effectiveness. While circuit emulation can be used in the near term by provisioning traditional voice trunks through a network, end-to-end signaling is preferred. The DON enterprise architectural strategy is to provide MAN services that fully support the consolidation and implementation of voice. The general use of distributed sets of Remote Switching Modules (RSMs) that are homed to central office switches provides a Service Delivery Node. This supports use of existing infrastructure and centralization of egress points to minimize circuit costs, centralize trouble reporting operations, and expand the network to encompass additional users and locations.

The DON TI architecture for voice includes two options for the interworking between Defense Switch Network (DSN) and the ATM WAN:

- Based on Circuit Emulation Services (CES) in the ATM WAN that is transparent to the DSN.
- Based on PVC or SVC that map the DSN signaling and user traffic to and from ATM formats to provide cell-based direct interworking between the DSN and the ATM network. In a PVC architecture, the VTOA IWF and ATM infrastructure are transparent to the DSN switching systems. Calls are routed to the destination DSN switching systems over an ATM virtual circuit trunk carrying the narrowband traffic and associated signaling. In an SVC architecture, the VTOA IWF must process a subset of the DSN signaling messages to establish SVCs to carry the DSN calls.

3.6.5 ATM Overlay Security

The DON ATM overlay shall provide a private or virtual private network for the DON. The DON ATM overlay must satisfy the following security requirements:

- Confidentiality
- Integrity
- Reliability
- Robustness

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

The DON ATM overlay must provide data transport service at the secret (S) and sensitive but unclassified (SBU) system high levels.

It is assumed that CANs (switches, routers, fiber, etc.) are under positive DON control and are operated at the SBU level. It is also assumed that MANs and WANs may NOT be under positive DON control and are operated at the unclassified (U) level.

In order for the ATM overlay to satisfy confidentiality and integrity requirements, ATM cell encryption shall be used. For secret data transport services, information shall be encrypted using KG-75 Fastlanes before entering a DON CAN or Zone. For SBU data transport services, information shall be encrypted using KG-75 Fastlanes before entering the portion of the ATM overlay that is NOT under positive DON control. For example, if a MAN is constructed by leasing ATM service from a commercial ATM provider, this encryption would be applied at the CAN/MAN boundary.

In order for the ATM overlay to satisfy the requirement for reliability and robustness of data transport services, the network must provide redundancy, guarantee the minimum bandwidth required to support high priority data transport, and provide mechanisms for managing network bandwidth allocation (particularly when contention occurs).

Redundancy shall be such that no single failure of a network component or interconnection leads to the isolation of any CAN from the overall DON enterprise network. The careful satisfaction of the redundancy requirement will require detailed analysis of external vendor network configurations when commercial ATM services are used to implement MANs and WANs.

Guarantees of minimum bandwidth can be provided in a number of ways. These include physically dedicating bandwidth using either dark fiber (leased or owned) or SONET channels, and PVPs with appropriate committed information rates. In addition, if DISN ATM services are subscribed, it may be possible to use SVCs when MOAs guaranteeing minimum bandwidth can be established.

Management of bandwidth allocation within the DON ATM overlay must be provided. This management must provide authorized administrators with the capability to identify the priority of data transport requests and allocate network bandwidth to the highest priority transport requests when contention occurs. In addition, the ATM overlay must feature mechanisms to order and add additional bandwidth as required.

Finally, the ATM overlay must provide mechanisms to ensure that DON-controlled components of the overlay can only be managed by authorized administrators, are resistant to penetration attempts, and are resistant to ATM signaling-based denial of service (DoS) attacks. The use of KG-75 Fastlanes significantly reduces or eliminates the potential for unauthorized administration and successful penetration originating from outside the DON-controlled portion of the overlay. In order to reduce the potential of successful penetrations originating from inside the DON-controlled portion of the overlay, network components that are remotely managed must feature a non-spoofable authentication mechanism. Use of KG-75 Fastlanes may reduce the potential for successful ATM signaling DoS attacks originating from outside the DON portion of the overlay. However, due to the current lack of authentication in ATM signaling protocols, it is difficult to protect against attacks originating from inside the DON-controlled portion of the overlay.

A comprehensive treatment of ATM security objectives and threats is well documented in the ATM Security Framework 1.0 (ATM Forum AF-SEC-0096.000, Feb. 1998) and parallel those of IP networks. Efforts are underway at DoD to resolve these ATM security issues.

3.6.5.1 Design Factors

ATM implementations must reflect the most up-to-date ATM security guidance.

Encryption for ATM is provided through the use of Fastlanes or TACLANes. These devices encrypt the data payload of ATM cells but not the header. When traffic analysis is a concern, link encryption devices may also be used.

Remote management of ATM devices is a concern. SNMP should not be used for anything other than collecting status information. Remote login to these devices requires non-spoofable authentication such as token-based access control. Whenever possible, SNMP access and traffic should be limited to in-band, non-routable subnets.

Regarding denial of service, signaling in ATM is a special concern. ATM signaling does not have authentication. Guidance on safe ways to use signaling in ATM will be provided at a later date.

3.7 IP Connectivity Detailed Architecture

As introduced in section 3.3.2, the IP addressing architecture is focused on the myriad organizational and operational characteristics of the Naval enterprise. For example, commands vary in size, level of complexity, and networking knowledge and are distributed worldwide. Ships deploy and transit between theaters while maintaining network connectivity so that a ship's point of presence on the network changes constantly as its location changes. The latter condition is exacerbated by relatively low-bandwidth, high latency, and high error rates, which are all common with satellite and line-of-sight radio frequency (RF) voice and data links.

The network architecture comprises the connectivity upon which traffic is transported throughout the Navy and Marine Corps. The connectivity spans resources and activities across the physical, network, and application layers and uses a mixture of ATM and IP technologies. Combined, these two technologies allow for the creation of a high performance, flexible, and proven infrastructure.

This network connectivity uses ATM standards to create flexible, high-bandwidth, IP-based "autonomous networks" on top of the ATM "virtual circuits". This design supports flexibility to provision bandwidth where needed without being hampered by congested IP paths or too many router hops (i.e. network diameter). This services profile extends to all regional ITSCs and to each campus. In support of the fleet intranet, it extends to the piers, NCTAMS, and STEP sites.

At higher networking layers, voice, video, and data network services are provisioned. For data, this infrastructure is established through multiple IP subnets. These "intranets" consist of logically-separated autonomous networks that are part of the larger DON enterprise network and which, for security reasons, require some level of logical segmentation. This architecture provides for the creation of these virtual networks for both large commands and concentrations of geographically close smaller commands that obtain services from a MAN. The separate IP subnetworks are constructed on top of the ATM services provided in the network layer below

them. These IP subnetworks are aligned with separate security enclaves or autonomous networks that manage their own routing domains.

3.7.1 IP Connectivity Design Factors

The DON enterprise network does not carry IP traffic between other networks if neither the source nor the destination of the IP traffic is a DON enterprise subnetwork host. Exceptions will be made for legacy DON IP networks which will be treated at ingress and egress as external (i.e., Internet) addresses and will be subject to Internet access control policies.

Unrestricted IP traffic between regions (i.e., the associated IP networks) is supported on the backbone using methods described in the previous sections.

Access to DON enterprise network resources from legacy DON IP networks is controlled by Zone 4 security protection mechanisms.

In order to minimize the size of internal routing tables, regions obtain a CIDR block sufficiently large to provide IP service to all tenant organizations within the region. This is based upon the number of campuses and the population of each campus.

3.7.2 Placement of Routers in IP Connectivity

The connectivity of campus routers to a MAN via a full Virtual Circuit (VC) mesh was described in section 3.4.2.

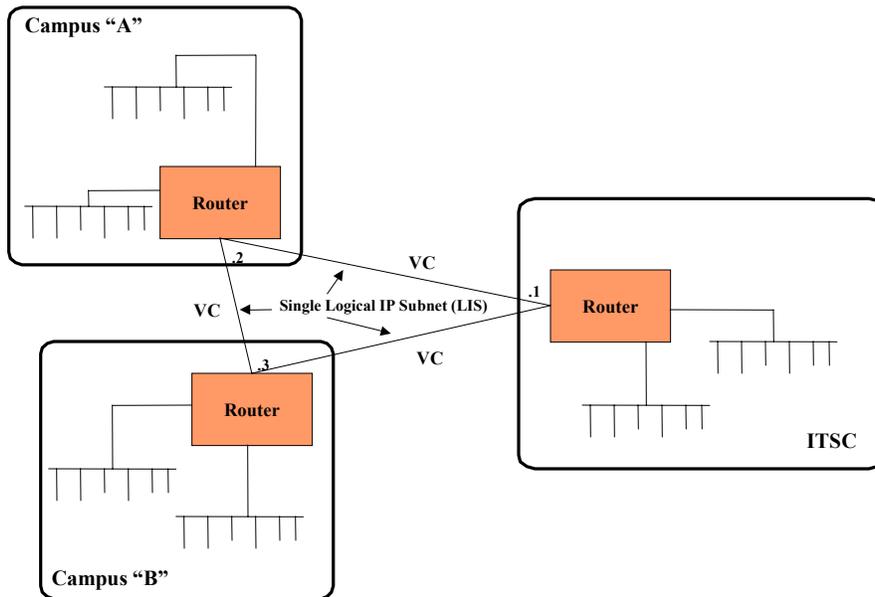


Figure 3-1. Multiple Campus Premise Routers, Showing VC Mesh Between Them, with Single LIS

One of the participating routers in a connectivity mesh will be the ITSC premise router. As shown in Figure 3-1, this router must include connectivity to the other Naval MANs and to the external world (NIPRnet, Internet, etc.) via the ITSC firewall. For reliability purposes, there should be two parallel routers; the figure is simplified for discussion purposes. To reach the other Naval MANs, there is a VC mesh over the WAN using the same connectivity principles as the MAN. This links all ITSC premise routers. The same recommendations and tradeoffs that are in the MAN case apply here.

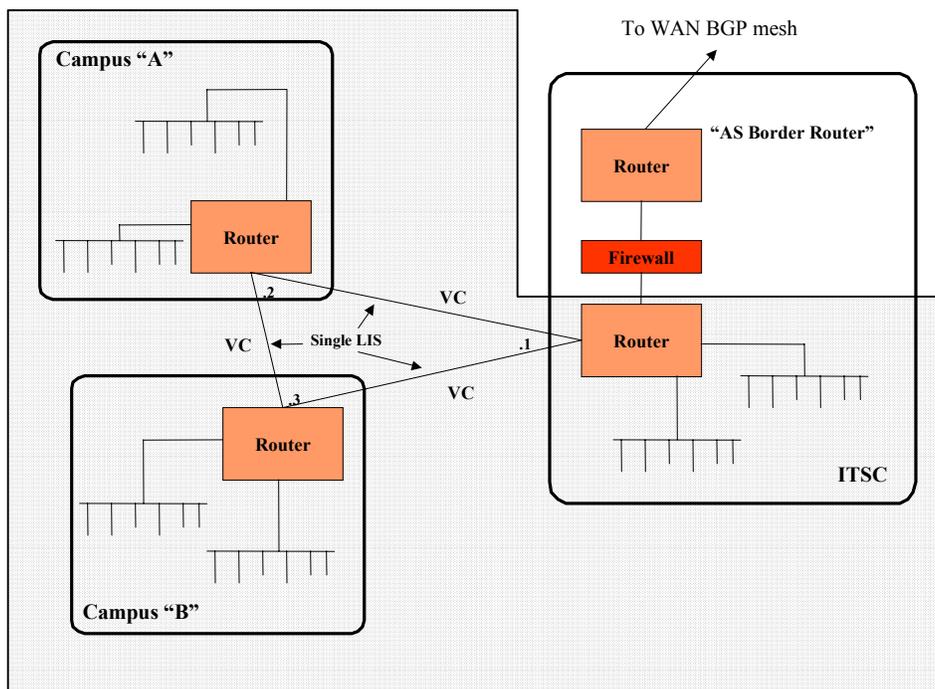


Figure 3-2. Multiple ITSC Premise Routers, Showing Full VC Mesh Between Them, Plus the ITSC Firewall, with Connection to the Outside World

There are cases in which a campus needs connectivity to the ITSC outside the firewall. Figure 3-2 shows that possible reasons for this may include an interim connection while back-door connections are eliminated. It may also be that the campus has a research network that needs to stay outside the firewall because of connections to external research nets. It may also be that the ITSC needs to peer with the external connections available at other campuses.

To accommodate these cases, it is clear that a separate VC mesh is required between routers that must peer outside the ITSC firewall. The architecture will be identical to the internal mesh described above.

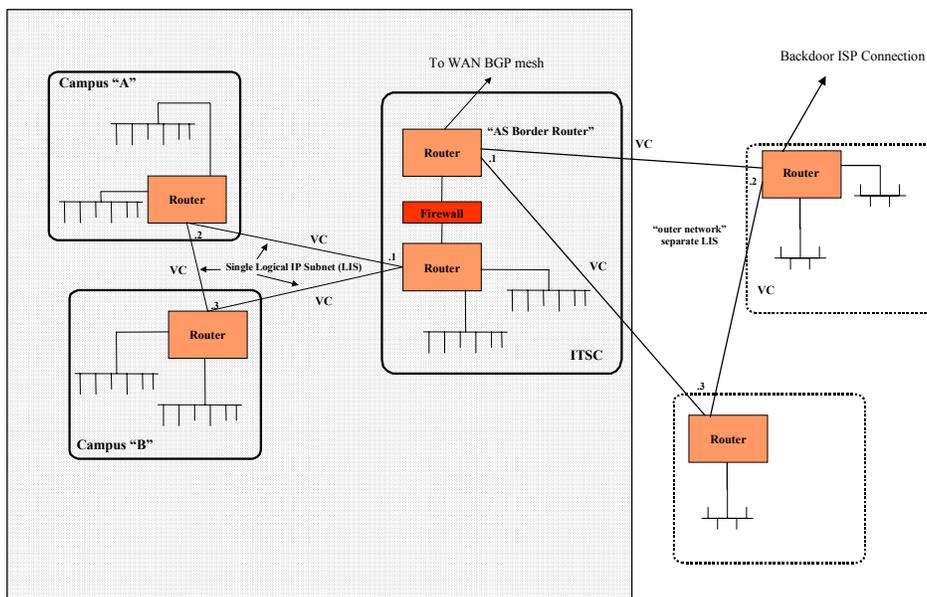


Figure 3-3. Picture of Multiple Campus Premise Routers, with VC Mesh, Connected Outside the Firewall

In order to minimize the number of routing hops between campuses and regions, the next step may be to implement MPOA to provide cut-through services for IP over ATM. It will not be necessary to do this right away, but this performance optimization can be designed and implemented at the appropriate time. Traditional router functionality, i.e., filtering, can be done on an MPOA-capable router. In this case, a cut-through virtual circuit will only be established if the source/destination IP addresses are "approved" to communicate.

The important point above is that all campus premise routers have direct access (zero hops) to the other campus premise routes on the MAN. Also, the ITSC premise routers all have direct access to each other via the WAN.

LANE technology should not be used across the MAN or WAN but should be restricted to the campus level.

3.7.3 IP Addressing

Regional commanders obtain IP network address space on behalf of ships and shore commands within their region from the Naval Computer and Telecommunications Station (NCTS) Pensacola, Florida, at (COMM) 850-452-3501, (DSN) 922-3501, or on-line at the Navy IP Network Number Registration page at <http://www.netreg.navy.mil>.

3.7.4 IP Routing

Establishing an appropriate routing architecture is critical to the success of the DON enterprise WAN. The connectivity, routing, and addressing schemas between ship, shore, MAN, and campus infrastructures are appropriately designed so that the enterprise network is scaleable and the number of hops between endpoints is minimized.

The routing architecture to address the fleet environment consists of a single routing domain for the entire fleet. Open Shortest Path First (OSPF) is used as the interior gateway protocol (IGP) for routing. The combination of dedicated and virtual circuits employed below layer 3 provides the networking infrastructure to support this design. The ITSCs and piers are included in this routing scheme. Ashore locations provide connectivity to the rest of the DON backbone network and the Internet and provide firewalls for Zones 2 and 4. Additionally, using OSPF allows fast convergence in response to topology changes that occur more often in the fleet environment.

Each of the MANs will need to obtain an Autonomous System (AS) number. This is required to support BGP peering with networks outside the region and for campus networks that already have their own ASN. Any given network should only be sourced from a single ASN.

The design factors for routing include determining required support of the following: LANE, Classical IP and ARP, and MPOA.

3.7.5 IP Overlay Security

The DON IP overlay shall provide a private or virtual private IP network for the DON and provide controlled interconnections to external IP networks. The DON IP overlay must satisfy the following security requirements:

- Confidentiality
- Integrity
- Reliability
- Robustness

The DON IP overlay must be provided at the secret and SBU system high levels.

The secret DON IP overlay will provide controlled interconnections to the SIPRNET and the SBU DON IP overlay will provide controlled interconnections to the NIPRNET. The interconnections must provide a high degree of isolation and penetration resistance from the NIPRNET and SIPRNET while allowing the flexible interchange of data between subscribers to the DON IP overlay. In addition, the DON IP overlays must continue to function in the event of SIPRNET or NIPRNET failure.

The DON IP overlay requirements for confidentiality and integrity shall be satisfied by using the DON ATM overlay at the appropriate system high classification level (secret or SBU). Also, the DON IP overlay requirement for reliability and robustness may be partially satisfied by virtue of the assured data transport provided by the DON ATM overlay.

The DON IP overlay requirements for privacy and controlled, secure interconnection to SIPRNET and NIPRNET (including continued operations of the DON IP overlay in the event of a NIPRNET or SIPRNET failure) shall be satisfied by employing a defense in depth approach to security. Key elements of this approach include network firewalls, secure domain name service (DNS), public key infrastructure (PKI), network intrusion detection, content security checking, secure configuration of workstations and servers, remote management security, and routing

protocol authentication. Optional elements of this approach include data at rest encryption, virtual private network (VPN) encryption, host-based intrusion detection, network intrusion filters (NIFs), and network access controllers (NACs). Detailed information on all of the elements can be found in appendix E.

From an architectural standpoint, network firewalls must be given particular attention in the DON IP overlay. Network firewalls shall be located between the DON IP overlay (Secret and SBU) and external IP networks (SIPRNET and NIPRNET). These firewalls, when combined with the other elements of the defense in depth approach, provide the controlled, secure interconnections to SIPRNET and NIPRNET. In addition, the combination of these firewalls and the use of the inherently robust/reliable DON ATM overlay provide a basis for satisfying the DON IP overlay requirement for robustness and reliability.

3.7.6 Considerations for Connecting Contractors

The customer frequently has a requirement to connect contractor and other non-Naval networks directly onto the Naval infrastructure. The reasons generally fall into two basic categories:

- Unsatisfactory performance of an Internet connection between a contractor site and a Naval site. This occurs because the contractor site receives connectivity via its own corporate intranet which has a gateway to the Internet. Traffic from the contractor site transits through their corporate infrastructure, through the network peering points, through the NIPRnet, and finally to the Naval site. The traffic may incur many router hops and congestion, even where the end points are at the same location. A direct connection between the two users would greatly improve performance.
- The contractor site is blocked access to the Naval site. The contractor appears as an Internet connection that cannot be trusted and is blocked for some access levels. To communicate, the contractor needs to bypass a security perimeter, such as a firewall or router filter, and establish direct connectivity “inside” the Naval infrastructure. To penetrate the security perimeter and gain access to systems on the protected side of the network, it must be done so that it does not compromise the security architecture.

If the main issue is performance and the connection to the Naval infrastructure is established outside the security perimeter (firewall or otherwise), then no additional security is required. If, however, the requirement is to bypass the security infrastructure, then it must be done in a manner that does not weaken or compromise the overall security architecture.

When non-Naval networks connect to the Naval Intranet inside the firewall, the following issues must be addressed:

- Does the contractor site connect to other networks? If so, these will constitute back-door connections and the connection should not be allowed.
- Can the contractor network be segmented with one segment being a separate network having no back-door connections? If so, this provides adequate protection and enables satisfactory connection to the Naval infrastructure. This segmentation may not be a desirable option for the contractor site.

- Does the contractor site have adequate controls of physical access? For example, are there locked doors? Who is allowed access to the protected network? What are the policies and procedures for establishing a connection to the protected network? Who can use the workstations that are connected to the physical network?

Extending the Naval Intranet to a contractor site opens up many new opportunities for security compromise because the contractor site is outside the Naval physical security perimeter. An acceptable solution must satisfactorily address the requirements and issues described.

The Naval Intranet will provide external access using Virtual Private Network (VPN) technology. This allows access to the Naval Intranet from anywhere on the global Internet by establishing a secure tunnel through the Internet to the Naval Intranet. Access can be controlled by use of an identity certificate provided by the Naval/DoD Public Key Infrastructure (PKI) solution described in the next chapter. Using PKI, access to the Naval Intranet is controlled on an individual basis instead of on the network or device level. Access to the Naval Intranet is granted only to those individuals who have a recognized need for network access. Authorized users must still be authenticated to gain access to systems and applications on the network, just like any other Naval user.

This VPN solution provides a general solution for contractor access, as well as for anyone (including Naval personnel that are travelers and telecommuters) who requires access from outside of the Naval Intranet. In this manner, direct connections can be established from outside the security perimeter while preserving the integrity of the security architecture.

3.8 DON ITI Architecture Plan of Action

This section reflects the priorities and steps that should be addressed in developing an ITI architecture and for planning and implementing selected components such as MANs and CANs.

3.8.1 Steps for Developing Detailed Enterprise ITI Architecture

The following outline reflects the steps that the ITI IPT will use to fully develop the ATM detailed architecture.

1. Identify the MANs and major sites in the Navy and Marine Corps that need ATM. The initial recommendation of the IPT is that every MAN and most campuses should have ATM-provisioned service.
2. Determine the list of components to be connected to ATM. A list of questions is relevant to making this determination:
 - Which of the components require ATM?
 - Which of the components can operate effectively with IP?
 - How much bandwidth is required?
 - What is the existing external IP and/or ATM connectivity?

These questions should be posed for the following:

- Major components
 - ◆ MANs
 - ◆ ITSCs
 - ◆ Piers
 - ◆ Teleports
 - Intermediate components
 - ◆ Large campuses not near a MAN
 - Small components
 - ◆ All the small bases
3. Develop a high level global wiring diagram that shows the world-wide connectivity required for the Naval enterprise network, including the above components.
 4. The conclusion of the IPT is that a multi-level PNNI hierarchy (WAN/MAN/Campus) is not a pragmatic solution for the near-term. This is in consideration of the need to provide a Naval enterprise solution for global routing and addressing across ashore and afloat platforms in view of the immaturity of standards and conventions for dynamic routing across public carriers. The near-term target architecture configuration for which the IPT will provide a clear definition during the coming months is as follows:
 - WANs: PVP service, end-to-end signaling
 - MANs: If SVC service available, fully supported signaling environment and PNNI; or if not available, PVP mesh overlay with required signaling domain between campuses
 - CANs: If ATM and SVC service, fully supported signaling environment and PNNI; or if not available, PVP mesh overlay with required signaling domain between campuses
 5. The IPT will develop the addressing scheme for DON, including a general case (showing all peer groups) and a fleet case (showing all peer groups). The scheme will include a diagram of the PNNI hierarchy. Additionally, the following detailed steps pertain.
 - Obtain an address block from DISA.
 - Support each MAN and all campuses within a region with the address block. Deploy the address block by region.
 - Determine how to allocate address space to the ships and ashore fleet infrastructure.
 6. Establish the MAN routing protocols using a 2-level PNNI hierarchy. Campuses will be able to participate as is appropriate. Show all peer groups for the general case. A similar hierarchy (with peer groups) will be developed for the fleet case. Both will be supported by graphic design guidance and supporting text.

7. Fully define LANE and the cases in which it should be used and not used. In general, LANE will be implemented only on campuses. The IPT will provide design guidance using graphics and supporting text.
8. Fully define MPOA and the cases in which it should be used. The guidance will present the current enterprise architecture for MPOA and provide specific guidance for both the general and fleet cases. It will provide graphic design guidance and supporting text, and diagram MPOA from campus to region to other regions.
9. Determine the additional routing architecture required for the DON enterprise network, including support of autonomous networks and fleet mobility requirements (for ships and for embarked staff, MEUs, air wings, etc.).
10. Develop a DON position for the establishment, configuration, and use of virtual private networks, including the N6 requested strategy and implementation issues (including interoperability, security, management, and performance), and coordinate with the joint services communities to reach a technically sound and implementable solution.
11. Determine a detailed DON connection architecture strategy for voice for both the ATM and IP connectivity environments.
12. Determine a detailed ATM security plan initially focusing on protecting the VPN.

3.8.2 Steps for Developing a MAN

1. Determine governance (including who builds, type of oversight, method of funding).
2. Understand customer base (including geography, population, mission, joint requirements).
3. Analyze requirements (current and projected, media, bandwidth, service delivery points, availability, alternate routes).
4. Develop MAN architecture (including security and management).
5. Identify provider alternatives.
6. Develop test plan.
7. Procure MAN.
8. Establish WAN connectivity.
9. Phase in campuses.
10. Connect outlying sites.

3.8.3 Steps for Developing a CAN

1. Perform inventory of circuits (voice, video, data) and cost of circuits that leave CAN.
2. Document routing architecture (routing domains, autonomous networks, Interior Gateway Protocol, External Gateway Protocol).
3. Determine performance requirements for connection to MAN.
4. Work with the MAN provider to establish Memoranda of Agreement and Service Level Agreements.
5. Determine connectivity strategy to MAN.
6. Determine management of network (if outsourced, turnover visibility and control are included in the MOAs and SLAs.).
7. Align the network topology along geographical versus organizational lines (use MAN to connect CANs in regions, use MAN to connect to WAN service delivery point).
8. Ensure that CAN accreditation is completed.
9. Align router and switching architecture with MAN standards.
10. Clean up local infrastructure to align with the DON ITI architecture.