

**INTERNET/PSN INTERCONNECTIVITY AND
VULNERABILITY REPORT**



December 1996

**Office of the Manager
National Communications System
701 South Courthouse Road
Arlington, VA 22204-2198**

INTERNET/PSN INTERCONNECTIVITY AND VULNERABILITY REPORT

December 1996

Prepared by:
Booz, Allen and Hamilton
8283 Greensboro Drive
McLean, VA 22102

Prepared for:
Office of the Manager,
National Communications System
Under Contract: DCA100-95-C-0113
Optional Task Orders, (NS/EP Telecommunications
Performance Analysis Task, Network Security Support)
CDRL Item L001 and DI-MISC-80711

DI-MISC-80711	Internet/PSN Interconnectivity and Vulnerability Report Section
3.2.5 Contents	Table of Contents
3.2.6 Figures and Tables	List of Exhibits
3.2.9 Symbols, Abbreviations, and Acronyms	List of Acronyms
3.3.1 - 3.3.6 Summary and Body	Sections 1 - 5
3.3.7 References	References

TABLE OF CONTENTS

EXECUTIVE SUMMARY.....	vi
1. INTRODUCTION	1-1
1.1 BACKGROUND	1-1
1.2 SCOPE	1-1
1.3 ORGANIZATION	1-2
2. HISTORY OF THE INTERNET.....	2-1
3. INTERNET DEFINITION	3-1
3.1 INTERNET SERVICE PROVIDERS	3-1
3.1.1 National Service Providers	3-2
3.1.2 Regional Service Providers	3-4
3.1.3 Resellers	3-5
3.2 INTEREXCHANGE POINTS.....	3-6
3.2.1 IXP Functionality and Architecture	3-8
3.2.2 IXP Peering Agreements	3-10
3.2.3 National-scope IXP Architecture Example	3-11
3.2.4 Metropolitan IXP Architecture Example	3-12
3.3 INTERNET ROUTING PROTOCOLS	3-13
3.3.1 Routing Information Protocol	3-14
3.3.2 Open Shortest Path First.....	3-15
3.3.3 Border Gateway Protocol Version 4	3-16
3.4 INTERNET ACCESS	3-17
3.4.1 Business Access.....	3-17
3.4.2 Residential Access	3-17
4. INTERNET ANALYSIS.....	4-1
4.1 INTERNET ANALYSIS TOOL FUNCTIONALITY	4-1
4.2 INTERNET ANALYSIS TOOL IMPLEMENTATION	4-3
4.3 INTERNET ANALYSIS RESULTS	4-4
4.3.1 Internet Analysis Methodology	4-4
4.3.2 Internet Analysis Results.....	4-6
5. VULNERABILITIES.....	5-1

5.1 INTERNET SERVICE PROVIDERS	5-1
5.1.1 National Service Providers	5-1
5.1.2 Regional Service Providers	5-3
5.1.3 Resellers	5-4
5.2 INTEREXCHANGE POINTS	5-5
5.3 INTERNET ACCESS	5-5

APPENDIX A

APPENDIX B

LIST OF ACRONYMS

REFERENCES

LIST OF EXHIBITS

Exhibit 2-1	Internet Timeline.....	2-1
Exhibit 2-2	Original NSFNET Backbone.....	2-4
Exhibit 2-3	NSFNET Three Tier Infrastructure (1986-1995).....	2-4
Exhibit 2-4	1988 T1 NSFNET Backbone.....	2-5
Exhibit 2-5	1992 T3 NSFNET Backbone.....	2-7
Exhibit 2-6	The National Science Foundation vBNS Network.....	2-9
Exhibit 2-7	Countries and Networks Connected to NSFNET as of April 1995.....	2-11
Exhibit 3-1	Representative NSP Backbone Network.....	3-3
Exhibit 3-2	NorthWestNet Backbone Network.....	3-6
Exhibit 3-3	CERFnet Backbone Network.....	3-7
Exhibit 3-4	Selected Major IXP Locations.....	3-8
Exhibit 3-5	Typical National-scope IXP Configurations.....	3-9
Exhibit 3-6	PacBell San Francisco NAP ATM/FDDI Hybrid Architecture.....	3-12
Exhibit 3-7	Analog Modem and ISDN Characteristics.....	3-18
Exhibit 3-8	Asymmetric Internet Access Characteristics.....	3-18
Exhibit 4-1	Sample Output From the IAT.....	4-2
Exhibit 4-2	IAT Site Locations.....	4-3
Exhibit 4-3	Internet Analysis Methodology.....	4-5
Exhibit 4-4	Status of IAT Traces.....	4-6
Exhibit 4-5	Categorization of Unsuccessful IAT Traces.....	4-7
Exhibit 4-6	Average Round Trip Time Versus Time of Day.....	4-8
Exhibit 4-7	Typical Traffic Patterns at MAE-EAST.....	4-9
Exhibit 4-8	Typical Traffic Patterns at MAE-WEST.....	4-9
Exhibit 4-9	Average Number of Hops Versus Time of Day.....	4-10
Exhibit 4-10	Top 50 Routers' Normalized Frequency of Use.....	4-11
Exhibit 4-11	Normalized Frequency of ISP Network Use.....	4-12
Exhibit 4-12	Booz • Allen's Critical ISP Networks.....	4-13
Exhibit 4-13	Proxima's Critical ISP Networks.....	4-13
Exhibit 4-14	Shared Critical Nodes.....	4-14
Exhibit 5-1	PN Three Tier Restoration Architecture.....	5-2
Exhibit 5-2	Internet Architecture Vulnerabilities.....	5-4

EXECUTIVE SUMMARY

Background

The Office of the Manager, National Communications System (OMNCS) performs a broad range of activities in fulfilling its mission. These activities include analyzing communications networks that support national security and emergency preparedness (NS/EP) communications. As more businesses, government organizations, and the public use the Internet for their daily activities, it has become more important for the OMNCS and its constituents to understand the operation of the Internet and its dependence on the existing communications infrastructure.

The phenomenal growth of the Internet has been one of the most significant technological events of the last several years. As an instrument for sharing and distributing information, the Internet will be judged one of the major milestones of the latter part of the 20th century. The exponential growth in Internet traffic has fostered the concept of the "Internet" as the ubiquitous tool for sharing information. However, the accessibility and availability of the Internet depend on a physical infrastructure of software, routers, and transmission media. It is commonly perceived that the Internet and the public telephone networks in the United States are two separate and distinct systems. Although this is true to a certain extent, most data networks, including the Internet, rely on the public networks (PN) to transport their traffic.

Internet Definition

At the highest level, the current Internet consists of multiple national and regional Internet Service Providers (ISP) and interconnection points where the ISPs meet and exchange traffic. This infrastructure is similar to that of the old National Science Foundation (NSF) network, NSFNET, which consisted of a three-tier structure:

- Backbone network
- Regional networks
- Local/campus networks.

The NSFNET was decommissioned in 1995. In its place are multiple nationwide networks similar to the original NSFNET backbone network. Regional networks still aggregate their traffic and hand it off to the nationwide backbone networks to which they are connected. Interexchange points (IXP) are located nationwide to facilitate the exchange of traffic between national and regional ISPs.

National Service Providers (NSP) provide national backbone service. This type of service provider owns or leases its own backbone network and has a nationwide customer base. Additionally, NSPs are generally connected to all the major IXPs and

have peering agreements with other major NSPs at these exchange points. Traffic originating with a customer on an NSP that is destined for a customer on another NSP is transferred from the originating NSP's network to the terminating NSP's network at an IXP

Regional Service Providers (RSP) are similar to the NSPs in that they own or lease their backbone network, but they are much smaller in scale. Their networks encompass a single region and usually have a regional customer base. RSPs have peering agreements with NSPs to transfer traffic over the Internet. RSPs either connect directly to the NSP or connect to an IXP where they transfer traffic to the NSP network.

With the dissolution of the NSFNET backbone, the NSF sponsored three primary and one secondary Network Access Points (NAP). The NSF's concern was that without the sponsorship of a core set of exchange points, the commercial backbone providers would set up a conglomeration of bilateral connection points that would potentially result in routing chaos.

Each NAP operator provides the exchange facility while the ISP that connects to the NAP establishes peering agreements with the other ISPs connecting to the same NAP. The purpose of a peering agreement is to ensure that traffic from one ISP can reach all the customers on another ISP by exchanging routing information between the two ISPs.

The current number of IXPs on the Internet far exceeds the original four NAPs sponsored by NSF. The term "NAP" is applied only to the NSF-sponsored IXPs, whereas all IXPs provide the same functionality, which is a common place for ISPs to exchange data.

Analysis

The Internet is a very dynamic entity in that it is constantly evolving and growing. Therefore, it is impossible to accurately identify all components of the current Internet. To develop the data for this report, the Internet was analyzed to identify key components used to transmit network traffic across the Internet. To achieve this purpose, a software tool, called the Internet Analysis Tool (IAT) was used to automatically trace the routes used to send traffic between two hosts on the Internet. The IAT collects data from the set of routers an Internet packet traverses on its path from one host to another. The analysis of the routes identified by the IAT yields traffic trends and identifies key components in the Internet infrastructure.

For this analysis, two IAT source sites were chosen:

- Booz • Allen & Hamilton, McLean, Virginia, on the PSINet network
- Proxima, Inc., McLean, Virginia, on the MCI Network.

The tool collected routes from each of these two sites to 105 other sites located across the United States. The type of Web sites chosen for this analysis were the following:

- 23 NCS Member Organizations' Web sites
- 50 State Web sites
- Major university Web sites
- Popular commercial Web sites.

The output from an IAT execution is the set of routers in the path between two hosts. For each router, three datagrams were sent at different times of the day, and the round trip time from the originating host and the router was collected.

Analysis Results

Traces performed throughout the test period indicated high success rates averaging between 87 and 89 percent. Of the unsuccessful trace attempts, most resulted from an unreachable node (i.e., a router or the destination server) in the path that was probably either shut down or incompatible with the IAT software.

Internet use is highest during mid-to-late afternoon business hours. Based on the round trip time for packets to traverse the network, congestion peaks between the hours of 12:00 noon and 4:00 p.m. eastern time.

This analysis indicated that the number of router hops did not vary in accord with the time of day or the day of the week. Thus, the predictability of Internet routing, along with an increasing dependency on this communications medium, renders it vulnerable to targeted and intended network disruptions.

Routers appear to share a somewhat balanced traffic load within the backbone networks (excluding those routers closest to the two sources). As expected, a high number of router "visits" occurred in the initial hops of the traces. These initial routers are critical to the sources; however, they are not necessarily critical to the entire Internet. As the trace moved away from the source and into the backbone networks, the number of visits per router stabilized. Therefore, a single critical router could not be identified, however, it could be determined which networks were more heavily traversed. For this analysis, MCI's network was traversed most frequently and was, therefore, critical to the success of the traces.

Vulnerabilities

The Internet can provide service in a volatile, unreliable network environment. But, like the PN, the Internet has vulnerabilities that can severely degrade its level of service. Because the Internet relies on PN packet and circuit switched networks, it is vulnerable to the same cable cuts and other damage that can affect the PN. In addition, some restoration techniques used by the PN carriers for circuit switched traffic cannot be used on the Internet's packet switched traffic.

National IXPs are critical to the operation of the U.S. portion of the Internet. An IXP failure could greatly reduce the Internet's ability to transport traffic nationwide or even worldwide. Congestion at these IXPs has also convinced ISPs that it is necessary to establish secondary means of interconnecting with one another.

Network routing protocols dictate how traffic is directed through the network in that they determine the paths that should be taken through the network to avoid congestion and network outages. Some Internet routers are vulnerable to "thrashing" – the optimal path through the network changes so frequently that the router spends more time computing these paths than actually routing users' data.

In summary, the initial analysis has determined that the Internet's physical vulnerabilities are consistent with the vulnerabilities of other large communication networks, most notably "last mile" issues and loss of backbone transport. Additional vulnerabilities exist that are distinct to the Internet – congestion (exponential growth in traffic), routing software, and network server management issues.

1. INTRODUCTION

The National Communications System (NCS) is a federation of 23 federal departments, agencies, and organizations that are responsible for the survivability and interoperability of various components of government communications supporting national security and emergency preparedness (NS/EP) activities. The Office of the Manager, National Communications System (OMNCS) is the planning and operational element of the NCS. The OMNCS performs a broad range of initiatives in fulfilling its mission, including analyzing communications networks that support NS/EP communications. The analysis process utilizes a standard OMNCS modeling methodology that incorporates OMNCS and commercial-off-the-shelf models, as well as public and proprietary data.

1.1 BACKGROUND

The phenomenal growth of the Internet has been one of the most significant technological events of the last several years. As a instrument for sharing and distributing information, the Internet will be judged one of the major milestones of the latter part of the 20th century. The introduction of Web browsers, dial-up communications protocols (i.e., Point-to-Point Protocol (PPP), Serial Line Interface Protocol (SLIP), and WinSock), and the increased efficiency of routers have made Internet access possible and cost effective even for small-business and at-home personal computer (PC) users. The exponential increase in Internet traffic has fostered the concept of the "Internet" as the ubiquitous tool for sharing information. However, the accessibility and availability of the Internet depend on a physical infrastructure of software, routers, and transmission media. As more businesses, government organizations, and the public use the Internet for their daily activities, it becomes more important to understand the operation of the Internet and the reliance of the Internet on the existing communications infrastructure.

The infrastructure that supports the Internet has evolved from mainframes and large minicomputers using dedicated transmission lines to low-cost routers and dial-up access from modems on PCs. Additionally, a growing support industry is providing Internet services, software, and content. As the Internet continues to evolve, its users will increasingly be dependent on not only the physical infrastructure but also the supporting services that have allowed the Internet to become an unparalleled information sharing tool.

1.2 SCOPE

This report describes the Internet by tracing its growth and development over the last three decades. It is difficult to provide a detailed, definitive history of the Internet

because much of its history has incorporated computer folklore and anecdotes. However, the major Internet milestones have been captured and serve as a baseline for its future growth. In context of the current description of the Internet and the Public Networks (PN), this document addresses several key vulnerabilities. These vulnerabilities are quantified using a simple route tracing tool that determines the physical path of Internet traffic. The Internet routes are then overlaid onto the PN infrastructure to illustrate the interdependence of the PN and the Internet.


1.3 ORGANIZATION

This document is organized into five sections. Section 1, Introduction, provides the background and scope of the Internet/PSN Interconnectivity and Vulnerability Assessment. Section 2, the History of the Internet, provides a detailed description of the history of the Internet from its earliest inception in 1969 up to the dissolution of the National Science Foundation's (NSF) backbone network in 1995. Section 3, Internet Definition, presents a breakdown of the different types of service providers, a description of the Internet infrastructure at a high level, and a discussion of the relationship of the Internet infrastructure to the PN infrastructure. Section 4, Internet Analysis, describes the Internet Analysis Tool (IAT) functionality and implementation. This section also presents the analysis methodology and results from the IAT. Finally, Section 5, Vulnerabilities, analyzes the current infrastructure of the Internet and discusses its major vulnerabilities.

2. HISTORY OF THE INTERNET

The Internet is a very complex entity of more than 10 million hosts connecting over 95,000 networks. To fully describe what the Internet consists of today, it is necessary to look at how the Internet began and evolved to its current state. The roots of the technology employed by today's Internet are found by analyzing its evolution. This section provides a detailed description of the history of the Internet beginning with the initial work performed by the Defense Advanced Research Projects Agency (DARPA) in 1969 to the recent commercialization of the Internet and the dissolution of the National Science Foundation Network (NSFNET) backbone in 1995. Exhibit 2-1 shows a timeline of the history of the Internet that this section will discuss in detail.

Exhibit 2-1
Internet Timeline



The inception of the Internet can be traced to 1969 when DARPA was commissioned by the United States Department of Defense (DoD) to develop a communications system that would be survivable in the face of enemy attacks including nuclear war. In addition, the network should allow military and academic researchers to collaborate on research projects and share computer processors across the country. In response to this

direction, DARPA, later renamed ARPA, set up a network consisting of the following four nodes:

- University of California at Los Angeles
- Stanford Research Institute
- University of California at Santa Barbara
- University of Utah.

ARPA used this four-node network, referred to as ARPANET, to experiment with the linkage to be used between DoD and military research contractors.

In 1970, ARPA began researching packet switched technology. The goal of this technology was to decentralize the network by giving all nodes on the network equal authority to transmit and receive packets across the network. The route each packet took to its destination was unimportant as long as it reached its destination. Thus, packet switching technology was effective when network connections were unreliable. This packet switching technology, employed by ARPA during the seventies, was known as the Network Control Protocol (NCP). By the end of 1971, there were 15 nodes connecting 23 hosts to ARPANET.

In 1973, ARPA began the "Interneting" project. The goal of this project was to develop a protocol that could seamlessly pass information between different networks. This project culminated in 1977 in a demonstration of networking through various media including satellite, radio, telephone and Ethernet. The protocol developed in this project formed the basis for the Transmission Control Protocol and Internet Protocol (TCP/IP), where IP handles the addressing of the individual packets while TCP coordinates the proper transmission of information.

By the end of 1982, ARPA established TCP/IP as the protocol suite for the ARPANET, requiring that all nodes connecting to ARPANET use TCP/IP. Additionally, DoD declared that TCP/IP was to be its standard protocol. The official cutover from NCP to TCP/IP was executed on January 1, 1983. Aiding this transition was the incorporation of TCP/IP into Version 4.2 of Berkeley Standard Distribution of UNIX. This version of the UNIX operating system was free to anyone who wanted it, thus ensuring a wide deployment for TCP/IP. The marriage of TCP/IP and UNIX began a long-standing affiliation between the Internet and the UNIX operating system that continues today.

Another major event in 1983 was the division of ARPANET into two networks: ARPANET and MILNET. MILNET was to be used for military specific communications, whereas ARPANET was to continue its research and development in networking computers. MILNET was integrated with the Defense Data Network created in 1982. The funding for ARPANET was provided by Defense Advanced

Research Projects Agency (DARPA). By 1984, the number of hosts connecting to the ARPANET was more than 1,000.

While the ARPANET was undergoing major changes, another significant event in the history of the Internet occurred. In 1979, representatives from DARPA and the NSF and computer scientists from several universities met to establish a Computer Science Department research computer network (CSNET). One of the driving forces for the establishment of CSNET was the concern that computing facilities located at universities not connected to ARPANET did not have the same advantages in research and staff and student recruitment as those who were connected. In 1981, CSNET was fully operational through money granted by NSF.

Although designed initially to be a standalone network, CSNET later incorporated a gateway connection to the ARPANET. In the summer of 1980, a DARPA scientist proposed the interconnection of the not yet established CSNET and ARPANET using protocols that would provide services and the seamless transmission of information between users regardless of the type of network. This set of protocols was TCP/IP. The gateway connection between the two networks was established in 1983.

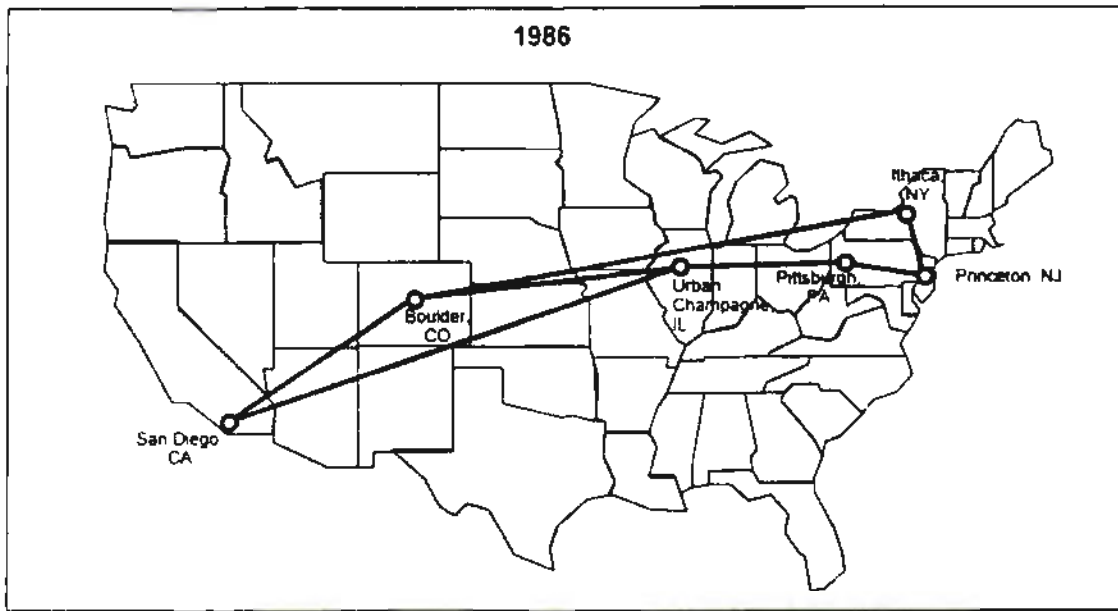
In 1986, the NSF created the NSFNET. The purpose of this network was to provide high-speed communications links between five major supercomputer centers located across the United States. Although ARPANET was flourishing, its 56-Kbps backbone and network topology could not fulfill the demand for high-speed networking required by multiple research projects. The goal of the NSFNET was to provide a reliable environment for the U.S. research and education community and access to the major supercomputing centers. The NSFNET essentially duplicated the functionality of the ARPANET. NSF chose TCP/IP as the standard protocol for its new network. This new network ultimately led to the downfall of ARPANET. In 1990, ARPANET was formally retired.

The infrastructure of the NSFNET was a three-tier hierarchical structure:

- National backbone
- Regional networks
- Local area networks (LAN).


The original backbone of the NSFNET, depicted in Exhibit 2-2, consisted of a 56-Kbps network. This backbone network is considered the basis of what is now called the Internet. Regional networks hung off the backbone network and provided services to LANs at education and research facilities. Universities and research associations combined to form the regional networks, which in turn would aggregate their traffic and "hand it off" to the NSFNET backbone. Exhibit 2-3 depicts the three-tier structure implemented in the NSFNET throughout its existence.

Exhibit 2-2
Original NSFNET Backbone



Source NSF


Exhibit 2-3
NSFNET Three Tier Infrastructure (1986 - 1995)



Because NSFNET's primary focus was for nonprofit research and development by universities and research groups, NSF instituted an "acceptable use" policy that restricted the use of the NSFNET to noncommercial activities. Additionally, NSF offered financial help to those regional networks, composed of university and research facility LANs, who wished to connect to the NSFNET backbone.

By 1987, the NSFNET outgrew its existing capacity. NSF awarded a five-year contract to Merit, the Michigan state networking organization, with MCI and IBM. The purpose of this contract was to transition the NSFNET backbone to T1 links and provide several access points around the country. Merit's role was to manage the backbone including routing, whereas IBM provided the routing equipment and MCI provided the trunk lines. The transition to a T1 backbone was completed in 1988. By the end of the 1980s, more than 100,000 hosts from 17 countries worldwide were connecting to the NSFNET. Exhibit 2-4 depicts the T1 backbone of the NSFNET in 1988.

Exhibit 2-4
1988 T1 NSFNET Backbone



Source: NSF

As the NSFNET grew, some organizations realized that providing services and functionality similar to that of the NSFNET without the access restrictions was a golden business opportunity. These organizations, experienced in providing regional network operations, seized the opportunity to set up their own nationwide backbone networks. Thus, the first commercial Internet service providers were created. These providers included Performance Systems (PSINet), and Altnet, which was generated from


UUNet Technologies. The main focus of these networks was to provide the same functionality as the NSFNET, over their own networks, but without any access restrictions.

In 1991, the fourth year of the five-year contract, Merit, IBM and MCI formed a new **nonprofit corporation**, Advanced Networks and Services (ANS), which was given the operational responsibilities of the NSFNET. In June 1991, ANS announced it would provide commercial access to the Internet, thus nullifying the acceptable use policy. By broadening access to the Internet, ANS increased its efforts to expand connectivity and make the Internet a more powerful tool. The new evolving private commercial networks were hindering research, forcing researchers to spend time accessing several networks all in the name of science. With expanded commercial providers on the Internet, there was a single common network that increased a researcher's ability to find any information needed and focus on the research at hand. When NSF lifted its access restrictions in 1991, allowing commercial traffic on the NSFNET, ANS formed a for-profit subsidiary, ANS CO+RE (Commercial + Research & Education), to provide full commercial traffic across the backbone. Once the "acceptable use" policy had been abolished, PSINet, UUNet Technologies, and General Atomics (CERFnet) created the Commercial Internet Exchange (CIX). CIX was a traffic exchange point between the NSFNET and the commercial Internet service providers networks.

The other major event that occurred in 1991 was the transition of the NSFNET backbone from T1 links to T3. This transition, like the initial transition from 56-Kbps to T1 links in 1988, was because of the capacity of the backbone network could not meet the traffic loads. Although this transition required new routing equipment and interfaces and, at times, proved to be technically challenging, it was accomplished with relative ease. This was due to the fact that the same organizations who were managing the old T1 backbone were responsible for implementing and overseeing the new T3 backbone network. Additionally, the T1 backbone still existed as a backup if the new network failed. Exhibit 2-5 depicts the T3 NSFNET backbone as of 1992.

In 1992, Vice President Al Gore drafted legislation that proposed a National Research and Education Network (NREN). This new network would consist of T3 links (separate from those making up the NSFNET backbone) and would connect all schools, libraries, etc., for a cost of over \$2 billion. Even though the legislation was passed, no new network ever came into existence. The NREN effort did, however, succeed in sparking a greater interest in the Internet.

Exhibit 2-5
1992 T3 NSFNET Backbone



Source: NSF

The new public Internet coincided with the release of the first Microsoft Windows version of Mosaic in 1993. Mosaic, developed by the University of Illinois at Urbana-Champaign, was an X-Windows interface to the World Wide Web (WWW). The concept of the WWW was started in 1989 in Switzerland as a means to easily share information among researchers in high-energy particle physics. In 1991, the first WWW server came into existence, but without any client software. The introduction of the first interface to WWW included the capability to navigate through the Web via the mouse. Today's Web browsers, such as Netscape, include File Transfer Protocol (FTP), E-mail, Telnet, and many more capabilities. The use of a graphical interface to access the Internet has played a significant role in the popularity growth of the network because it allowed access to the Internet without having knowledge or possession of the UNIX operating system.

While the look and feel of the Internet was undergoing changes, NSF, in 1992, began to question its role in the network. NSF observed that its backbone network was operating in conjunction with several commercial nationwide backbone networks. Essentially, NSF was paying for users to access its network, and thus the Internet, whereas the other commercial service providers were being paid for access to theirs. Although in 1991 the NSF had notified the regional networks that they would have to become self-sustaining, it was 1992 before the NSF took action. The NSF began considering ways in which it could successfully pull out of the Internet arena with little

disruption to the Internet while continuing its commitment to the education and research community.

With the five-year contract between the NSF and Merit drawing to a close, Merit was granted an 18-month extension (beyond the original October 1992 expiration date) to allow the NSF time to work out how to transition its backbone network into a new structure. This work culminated in a solicitation for proposals (Solicitation 93-52) in the following four areas that compose the new national Internet structure:

- Network Access Points (NAP)
- Routing Arbiter
- Regional network provider awards
- A very high-speed Backbone Network Service (vBNS).

The NAPs act as interconnection points where commercial Internet service providers can meet and exchange traffic. The NSF believed that without such interconnect points, backbone providers would likely establish their own independent bilateral connect points that would stifle the NSF's plan for full connectivity for the research and education community. The NAP manager contracts were awarded to the following:

- Sprint, for a New York NAP
- Metropolitan Fiber Systems (MFS) Datanet for a Washington DC NAP
- Bellcore and Ameritech for a Chicago NAP
- Bellcore and Pacific Bell for a California NAP.

The Routing Arbiter is an independent group that operates route servers at each NAP. The transfer of traffic among the backbone providers that meet at the NAPs is facilitated by route databases contained in the route servers. These databases contain routing information and policy requirements for each backbone provider and therefore indicate to which provider the incoming information should be sent. This contract was awarded to Merit and the Information Sciences Institute (ISI) at the University of Southern California, which together make up the Routing Arbiter group.

With the dissolution of the NSFNET, and the introduction of NAPs and commercial traffic, access to the Internet by the NSF subsidized regional networks was no longer free. The commercial backbone providers were now paying a fee to interconnect with the NAPs and passing these charges to their users – the regional network providers. Therefore, the NSF decided to create the regional network provider contracts to alleviate the regional networks' initial shock of having to pay for Internet access. The awards provided the regional networks with annual NSF funding, with the funding declining to zero over a four-year period. The regional network providers would use the subsidy to pay the commercial Internet providers who were in turn required to connect to the NAPs. There were 17 contracts awarded to regional network providers for interregional connectivity.


The NSF also proposed to sponsor a new backbone, the vBNS, operating at a minimum speed of OC-3 (155 Mbps), to link the following five NSF supercomputer centers:

- Cornell Theory Center
- National Center for Atmospheric Research
- National Center for Supercomputing Applications
- Pittsburgh Supercomputing Center
- San Diego Supercomputing Center.

Unlike the general purpose NSFNET infrastructure, the vBNS functions as an advanced research laboratory, allowing research, development, and integration of new networking requirements, using technology beyond just IP routing. There is a strict acceptable use policy: the vBNS may only be used for meritorious high-bandwidth research activities and it may not be used for general Internet traffic. NSF entered into a five-year agreement with MCI to provide the vBNS.

During this five-year agreement, MCI is expected to participate in the development and use of advanced Internet routing technologies. At the end of the agreement, it is anticipated that technology will exist that will increase the transmission speeds beyond 2.2 Gbps. Additionally, the vBNS will act as an experimental platform for the development and testing of broadband Internet services and equipment. Exhibit 2-6 depicts the NSF's vBNS network.

Exhibit 2-6
The National Science Foundation vBNS Network



Source: MCI

The result of NSF's solicitation for proposals was a new Internet structure. In April 1995, the NSFNET backbone was formally retired. At that time, 93 countries and more than 50,000 networks were connected by the NSFNET backbone. Exhibit 2-7 details the number of networks, by country, connected to the NSFNET backbone by the end of the project.

NSF's original task was to improve the previous NSFNET backbone, push the technology to new heights, and implement it on a national level. It was hoped that this would place a powerful tool in the hands of the research and education community, and create innovative use and applications. Its goals were accomplished: the NSFNET backbone connected most of the higher research and education community to a robust and reliable high-speed network and it served as the sole player in making the Internet industry.

The NSF's task will continue to evolve in two directions: 1) providing support for the research and education community by guaranteeing the availability of services, resources, and tools to keep the Internet connected, and 2) by continuing to push networking technology using the vBNS.

Exhibit 2-7
Countries and Networks Connected to NSFNET as of April 1995

Country	Total Networks	Country	Total Networks	Country	Total Networks
Algeria	3	Greece	105	Norway	214
Argentina	27	Guam	5	Panama	1
Armenia	3	Hong Kong	95	Peru	44
Australia	1875	Hungary	164	Philippines	46
Austria	408	Iceland	31	Poland	131
Belarus	1	India	13	Portugal	92
Belgium	138	Indonesia	46	Puerto Rico	9
Bermuda	20	Ireland	168	Romania	26
Brazil	165	Israel	217	Russia	405
Bulgaria	9	Italy	506	Senegal	11
Burkina Faso	2	Jamaica	16	Singapore	107
Cameroon	1	Japan	1847	Slovakia	69
Canada	4795	Kazakhstan	2	Slovenia	46
Chile	102	Kenya	1	South Africa	419
China	8	Korea, South	476	Spain	257
Colombia	5	Kuwait	8	Swaziland	1
Costa Rica	6	Latvia	22	Sweden	415
Croatia	31	Lebanon	1	Switzerland	324
Cyprus	25	Liechtenstein	3	Taiwan	575
Czech Rep.	459	Lithuania	1	Thailand	107
Denmark	48	Luxembourg	59	Tunisia	19
Dominican Rep.	1	Macau	1	Turkey	97
Ecuador	85	Malaysia	6	Ukraine	60
Egypt	7	Mexico	126	Unit Arab Emirates	3
Estonia	49	Morocco	1	U. K.	1436
Fiji	1	Mozambique	6	United States	28470
Finland	643	Netherlands	406	Uruguay	1
France	2003	New Caledonia	1	Usbekistan	1
French Polynesia	1	New Zealand	356	Venezuela	11
Germany	1750	Nicaragua	1	Vietnam	1
Ghana	1	Niger	1	Virgin Islands	4

Source: Merit Network, Inc.

3. INTERNET DEFINITION

At the highest level, today's Internet consists of multiple national and regional Internet Service Providers (ISP) and interconnection points where the ISPs meet and exchange traffic. This infrastructure is similar to that of the old NSFNET, which consisted of a three-tier structure:

- Backbone network
- Regional networks
- Local/campus networks.

On the NSFNET, regional networks would aggregate their traffic and "hand it off" to the NSFNET backbone. The regional networks comprised multiple local business and campus networks. Although there were many regional and local networks, there was only one backbone network.

As mentioned in Section 2, the NSFNET has been decommissioned. In its place are multiple nationwide networks, which are similar to the NSFNET backbone network. Regional networks still aggregate their traffic and hand it off to the nationwide backbone network to which they are connected. Interexchange points (IXP) are located around the country where traffic is exchanged between national and regional ISPs. Peering agreements are used between the ISPs connected at an IXP to determine how traffic is routed. These service providers and interexchange centers are the main components of the U.S. Internet. This section will describe different elements of the Internet architecture and the different routing protocols used on today's Internet.

3.1 INTERNET SERVICE PROVIDERS

ISPs are classified according to their network and customer base. The network classification refers to whether or not the ISP owns or leases its network. An ISP that does not own or lease its network is referred to as a reseller. The customer base classification refers to an ISP's type of customers, national or regional. A particular ISP may have national and regional customers, but generally it has more of one type than another. There are three types of ISPs:

- National Service Providers (NSP)
- Regional Service Providers (RSP)
- Resellers.

The following sections provide further detail for each type of ISP.

3.1.1 National Service Providers

The first category of ISPs is NSP, which provide national backbone service. This type of service provider owns or leases its own backbone network and has a nationwide customer base. Additionally, NSPs are generally connected to all the major IXPs and have peering agreements with other major NSPs at these exchange points. Traffic originating with a customer on an NSP that is destined for a customer on another NSP is transferred from the originating NSP's network to the terminating NSP's network at an IXP. The NSP's network infrastructure consists of routers (network layer) and switches (data link layer) that are owned by the NSP. The following are examples of NSPs:

- ANS
- BBN
- MCI
- PSINet
- Sprint
- UUNet.

Of the NSPs, MCI and Sprint are the only two that own their entire network. Other NSPs may own small parts of their networks, but most of their networks consist of circuits leased from the PN providers. Most of these circuits are leased from the large Interexchange Carriers (IEC)¹. However, some circuits are also leased from the Local Exchange Carriers (LEC) (e.g., Bell Atlantic), Competitive Access Providers (CAP) (e.g., Metropolitan Fiber Systems), and smaller IECs (e.g., LDDS).


Exhibit 3-1 depicts a representative backbone network for UUNet, one of the NSPs mentioned above. As shown in the exhibit, UUNet (like most NSPs) has redundant connectivity between each switching node on its backbone network.

NSPs rarely sell directly to small consumers (e.g., small businesses and residential customers) because of the added "customer handholding" required by smaller, less experienced users. Instead, NSPs sell their services to large businesses and resellers. Resellers in-turn resell Internet service to small business and residential customers. It is important to note that not all NSPs resell their networks, e.g. PSINet.

The architecture of an NSP's network may be separated into access and transport. Access refers to the customer's connection to the NSP, whereas transport refers to the backbone of the NSP's network. Customers connect to NSPs via leased and dial-up lines. Typical leased lines are 56-Kbps or T1 and usually terminate at an NSP's point of presence (POP).

¹ MCI advertises that 40% of all Internet traffic travels over MCI circuits. This includes traffic on MCI's NSP and traffic on other NSPs that use MCI leased lines.

**Exhibit 3-1
Representative NSP Backbone Network**



Source: UUNET

For dial-up customers, the NSP usually has digital and/or analog modem banks terminating from its POP into the local central office using T1s. Because NSPs have national presence and reach, once a customer's traffic reaches an NSP's POP, it has essentially reached the Internet. The typical backbone of an NSP comprises routers and switches connected by T1, T3, or even OC-level circuits. These circuits may be leased from one or more IEC. One NSP, PSINet, leases backbone circuits from five different IECs.

The NSP market has not escaped the notice of existing PN providers anxious to get involved in the growth of the Internet. In the short term, PN providers have chosen to partner with NSP providers for Internet backbone transport instead of developing NSP expertise in-house. For example, GTE recently announced a partnership with UUNet to provide Internet access under the GTE name to customers in 46 U.S. states². Cross PN-NSP service agreements also exist between Pacific Bell and America On-Line (which owns ANS), and AT&T and BBN.

² Washington Post, July 11, 1996, Page D9.

The recently announced merger between UUNet and MFS may be a harbinger of future mergers between NSPs and PN providers. PN providers own the data links necessary to run an NSP and have the marketing savvy to sell Internet service to business and residential customers. NSPs, on the other hand, have the in-house technical expertise to manage the switches, routers, and interconnection arrangements necessary to make the NSP backbone work.

Other future developments in the NSP market will include service differentiation to target selected customer markets. For example, MCI and BBN have announced services that provide a higher quality of service to business customers who subscribe to their NSP. BBN provides priority treatment to business customers through Internet Protocol version 6 (IPv6) priority service protocols. MCI provides a separate network for its business subscribers' Internet traffic. This separate network includes locally hosted mirror sites from popular Web sites on other NSP networks and in the future will include IPv6 priority treatment.

3.1.2 Regional Service Providers

The second category of ISPs are the RSPs. These service providers are similar to the NSPs in that they own or lease their backbone network but are much smaller in scale. Their networks encompass a single region and usually have a regional customer base. RSPs have peering agreements with NSPs to transfer traffic over the Internet.

RSPs either connect directly to the NSP or connect to an IXP where they transfer traffic to the NSP network. NorthWestNet is an example of an RSP that connects directly to an NSP. NorthWestNet, which provides service to customers in Washington, Oregon, and Idaho, has direct connections to both MCI and Sprint's NSP networks. Erols is an example of a network with a direct connection to an IXP. Erols, which provides service to customers in the metropolitan Washington, DC area, is connected to the Metropolitan Area Ethernet-East (MAE-EAST) IXP where it can transfer traffic to most of the larger NSPs and several smaller RSPs.

RSP service is an attractive option for residential and small business customers. Because of the small customer base, RSPs can offer more "hands-on" assistance in the form of customer training and help desk operators trained to assist less knowledgeable users.

Like NSP networks, the RSP's network architecture may be separated into access and transport portions, though with different meanings. In the RSP scenario, access refers not only to the customer connecting to the RSP but also the RSP connecting, if at all, to the Internet. Transport refers to the backbone of the RSP's network. As in the NSP scenario, customers connect to RSPs via leased and dial-up lines. Typical leased lines are 56-Kbps or T1 and usually terminate at an RSP's POP. For dial-up customers, the

RSP usually has digital and/or analog modem banks terminating from its POP into the local central office using T1s.

An RSP's backbone is typically restricted to a region, as opposed to NSPs who have a national presence and whose backbone spans the entire United States. Transport on an RSP's network, or backbone, comprises T1 and T3 circuits that connect their POPs and customers in a particular region. These circuits are leased from LECs, CAPs, and IECs. As noted above, RSPs' customers are primarily small business and residential subscribers. In the coming years, new companies will enter this market. Most notable are the Internet service offerings from the IECs and the Region Bell Operating Companies (RBOC). This increased competition may cause some consolidation of the RSP market when smaller RSPs go out of business or are bought out by larger firms. The remaining RSPs will survive by targeting market niches, such as high volume residential users or businesses new to the Internet.

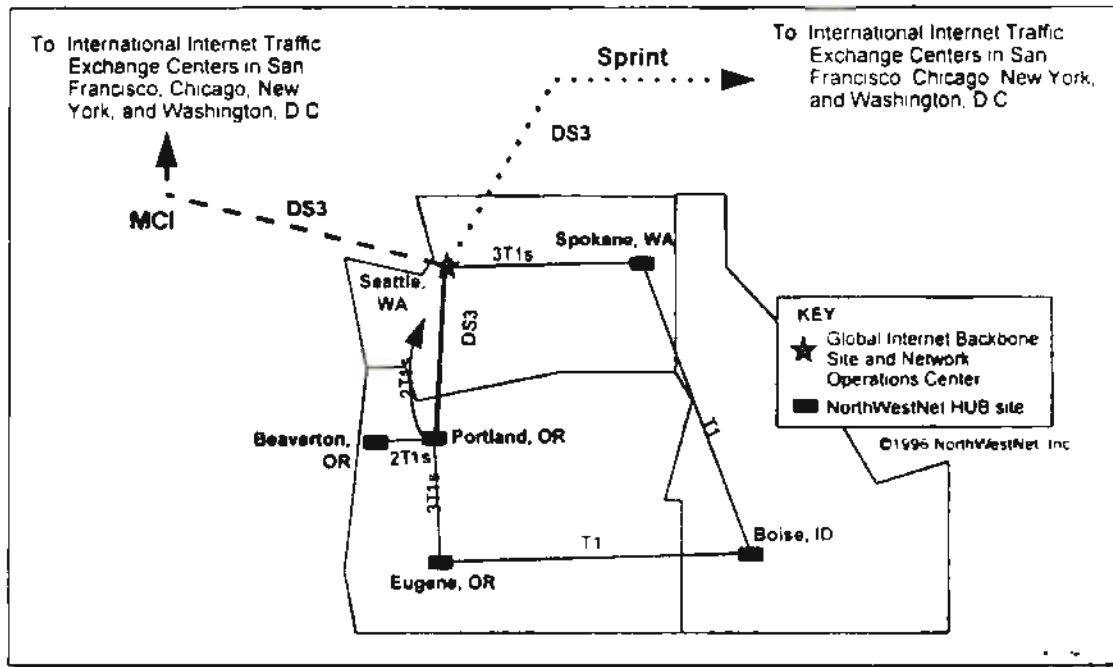
The exhibits below show two example RSP network backbones. Exhibit 3-2 shows NorthWestNet's backbone and Exhibit 3-3 shows CERFnet's backbone. Note that NorthWestNet has redundant connections to Sprint and MCI to transfer traffic, whereas CERFnet connects directly to NAPs to share traffic.

3.1.3 Resellers

Resellers are another member of the Internet provider family. Resellers purchase service from NSPs or RSPs and resell this service to small business and residential customers. Resellers are differentiated from RSP because resellers do not own or lease a network infrastructure. Instead resellers typically operate out of a single site with a modem bank for customer access and a T1 connection to transfer traffic to the NSP/RSP network.

There are approximately 1,400 Internet resellers in the United States, most of which base their business on monthly subscriptions to Internet service. As the Internet market matures, monthly Internet service is becoming a commodity. This trend has been furthered by the entry of the RBOCs and IECs into the residential Internet service market. Typically, unlimited access is provided on a monthly basis for a flat-rate fee or a combination of flat-rate and usage-based pricing.

**Exhibit 3-2
NorthWestNet Backbone Network**



Source: NorthWestNet, Inc.

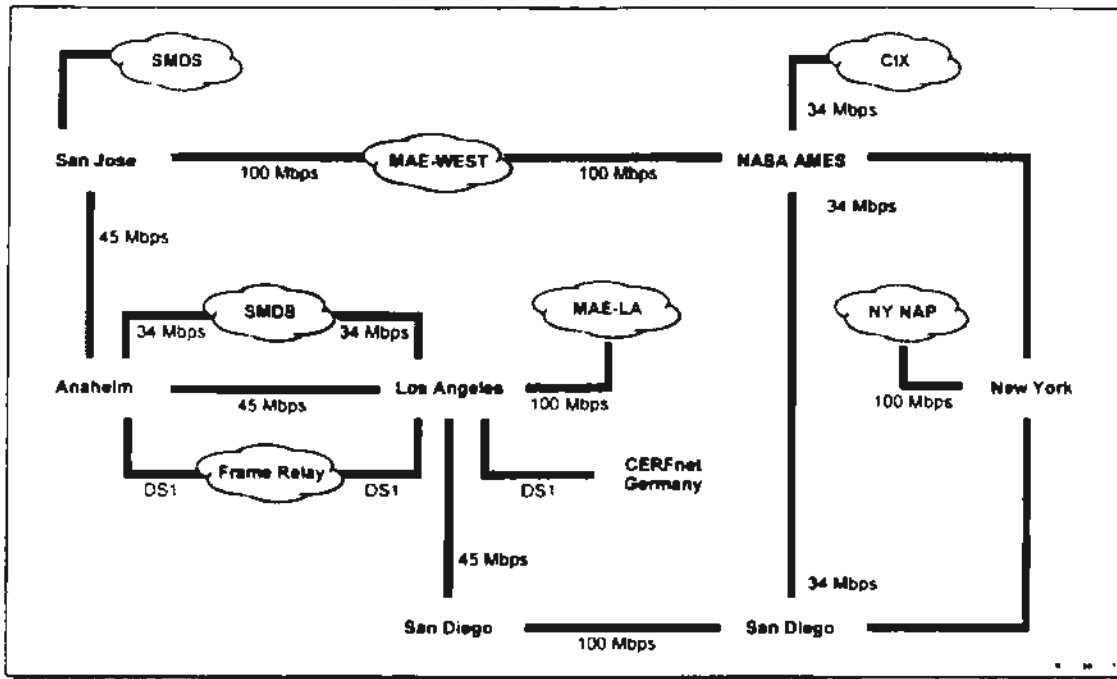
Because Internet service has significant economies of scale, the market favors the larger providers who can spread their fixed costs over a larger customer base. Because of this, many experts predict that the number of Internet resellers will decrease dramatically in the next few years. The Yankee Group predicts that there will only be 200 resellers left in business by 2000.

The remaining resellers may survive by looking for market niches. For example, instead of simply providing monthly Internet subscriptions, resellers are already starting to provide value-added services such as Web page hosting, Web page development, security management, and electronic commerce consulting. In these areas, a reseller may be able to provide better service to small businesses than a larger NSP or RSP company.

3.2 INTEREXCHANGE POINTS

With the dissolution of the NSFNET backbone, the NSF was concerned with maintaining connectivity between the commercial ISP networks and the research and education community. To address this issue, the NSF sponsored three primary and one secondary NAPs. Without the sponsorship of a core set of exchange points, the NSF feared that the commercial backbone providers would likely setup a hodgepodge of bilateral connect points potentially resulting in routing chaos.

Exhibit 3-3
CERFnet Backbone Network




Source: CERFnet

Under the NSF model, each NAP operator provides the exchange facility while the ISP that connects to the NAP establishes the exchange agreements, also known as peering agreements, with the other ISPs connecting to the same NAP. The purpose of a peering agreement is to ensure that traffic from one ISP can reach all the customers on another ISP by exchanging routing information of the two ISPs.

Today there are many more IXP centers on the Internet other than the original four sponsored by NSF. The term NAP is applied only to the NSF sponsored IXPs, whereas all IXPs provide the same functionality, a common place for ISPs to exchange data. Various cities and organizations have used different names for the exchange point, e.g., NAP, MAE, CIX, Federal Internet Exchange (FIX). Exhibit 3-4 presents a snapshot of several of the larger IXPs in the United States.

It is important to note that an IXP does not have to serve the national ISPs. There are metropolitan exchange points (MXP) used today, which are similar in structure to the NAPs, but service only local and regional traffic. This means that traffic originating and terminating in a single region would not traverse any of the national ISPs' backbones, thus removing some of the burden on these networks. The remainder of this section describes the structure of an IXP and details the different types of peering agreements used by the ISPs at an IXP.

**Exhibit 3-4
Selected Major IXP Locations**



3.2.1 IXP Functionality and Architecture

The large, national-scope IXPs, such as the NAPs or MAEs, interconnect numerous national ISPs and may exchange data requiring large amounts of bandwidth. The smaller regional or metropolitan IXPs will have fewer interconnects and require much less bandwidth. The IXP structure is similar regardless of the size of the IXP or the technical architecture used to exchange the traffic.

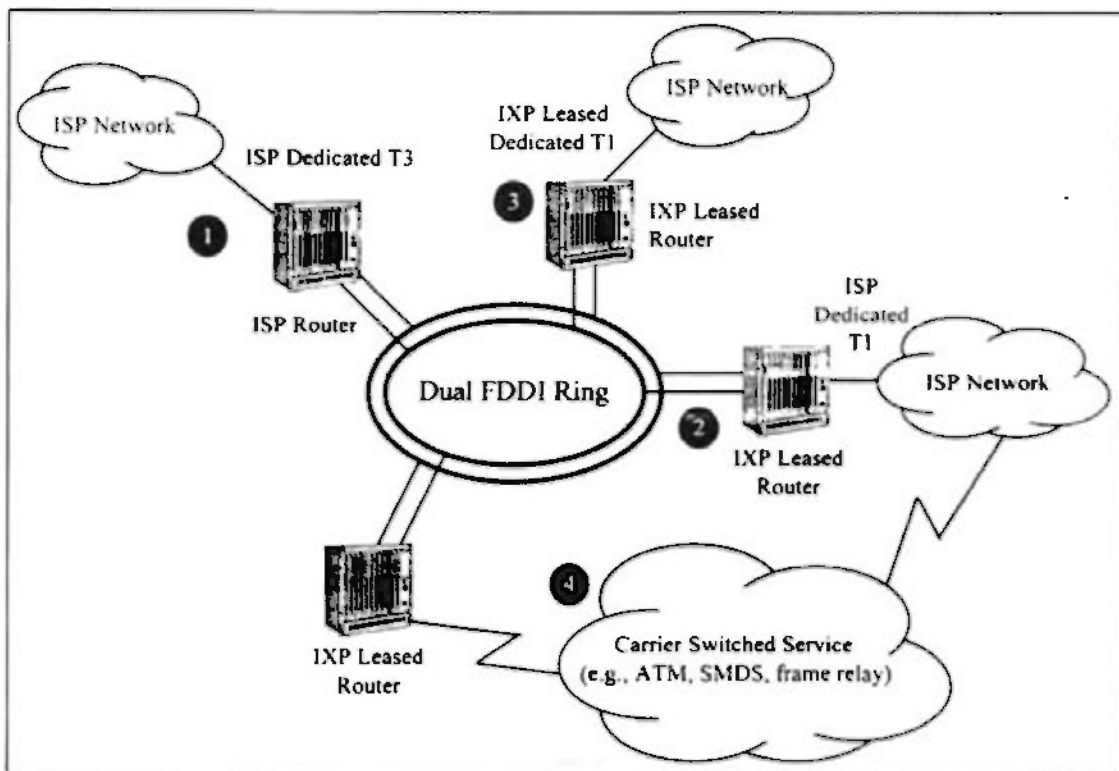
IXP facilities generally consist of a high-speed LAN or metropolitan area network (MAN) architecture capable of interconnecting various wide area network (WAN) technologies. ISPs connect to the IXP LAN via either a high-speed router or an asynchronous transfer mode (ATM) switch capable of connecting to the IXP architecture. Each of the connecting ISPs must negotiate bilateral or multilateral peering agreements with other ISPs interconnecting at the IXP. The Routing Arbiter administers the traffic routing resulting from these peering agreements. This traffic routing and addressing information is provided to each ISP's router by a route server within the IXP LAN. Incoming packets are routed to the high-speed LAN ring where the route server indicates the possible routes available to the packet.

The most common NAP architecture is a Fiber Distributed Data Interface (FDDI) dual ring backbone LAN running at 100 Mbps. Routers for each ISP are homed to the dual ring bus in the various access configurations discussed below:

1. The ISP provides and manages its own router collocated at the IXP facility. The ISP would have dedicated access to this router via its own dedicated line (typically a T1 or T3). This option may not be available at all IXPs because of space limitations.
2. The ISP leases an IXP provided router located at the IXP. The ISP has dedicated access to the IXP router via its own dedicated line.
3. The ISP leases the dedicated connection and the router from the IXP.
4. The ISP leases switched access service to the IXP facility from the IXP or another provider. Switched access may include ATM, Switched Multimegabit Data Service (SMDS), and frame relay.

These access configurations are shown in Exhibit 3-5. For each of the access configurations, all equipment is located in a single facility.

Exhibit 3-5
Typical National-scope IXP Configurations



Source: Sprint

Other IXP architectures that have been used include SMDS and ATM networks. Lower bandwidth solutions such as SMDS may be more commonplace in regional or metropolitan IXPs.

All IXPs are privately owned and administered by IECs, Incumbent Local Exchange Carriers (ILEC), Competitive Local Exchange Carriers (CLEC), or ISPs. The four NSF-sponsored NAPs are owned by Sprint, MFS, Pacific Bell, and Ameritech. Regional and metropolitan IXPs may also be owned by ISPs, e.g., the SMDS Washington Area Bypass (SWAB) is operated by PSINet. The IXPs normally charge flat interconnection fees and usage based fees to the interconnecting ISPs.

Large IECs, ILECs, and CLECs can provide network management for their IXPs from their PN network management centers. Most IXP operators will ensure reliability of service and mean time to repair, and provide maintenance for collocated equipment. The dual ring FDDI buses used in many large IXPs are also very robust to a single line fiber cut. A single dedicated connection from the ISP network to the IXP router will pose the greatest vulnerability in the IXP architecture. Redundant connections to the IXP should be used by regional ISPs that do not have presence at multiple IXPs.

3.2.2 IXP Peering Agreements

The policies for data exchange at an IXP are set forth by the parties involved. Just because an ISP connects to a particular IXP does not guarantee that that ISP can exchange traffic with every other ISP connected to that exchange point. Agreements that specify how traffic is carried and transferred, and how billing is handled have to be established and maintained between the ISPs on an IXP. Any ISP can connect to an IXP as long as the ISP agrees to the predefined policies. Currently, there are three different types of exchange policies:

- Bilateral
- Multilateral
- Multi-party bilateral.

A bilateral agreement is between only two ISPs at an exchange center. A multilateral agreement is between many ISPs at an exchange center. A multi-party bilateral agreement is between a small ISP and a large ISP to carry the small ISP's traffic to other ISPs. The more IXPs a single ISP connects to the better the performance and reliability of the ISP's service. Each IXP has its own procedures for establishing peering agreements among the IXP-attached ISPs.

A peering agreement is defined as the advertising of routes via a routing protocol for customers of the IXP participants. Specifically, the ISP is obligated to advertise all its customer's routes to all other participating ISPs and accept routes from the customer's

routes advertised by the ISP. ISPs are required to peer with the IXP's route server which facilitates the routing exchange between the ISPs routers. The route server gathers the routing information from each ISP's router, processes the information based on the ISP's routing policy requirements, and passes the processed routing information to each of the IXP-attached ISPs. Currently, ISI handles the work done on the routing management system, while Merit implements and maintains the route servers and route server databases.

3.2.3 National-scope IXP Architecture Example

Pacific Bell's NAP, located in San Francisco, California, is fairly typical of national-scope IXPs. PacBell's NAP is an ATM/FDDI hybrid LAN, whereas other national-scope IXPs may be straight FDDI design or an FDDI/Ethernet hybrid. PacBell's use of ATM makes it one of the fastest IXPs, capable of up to 139 Mbps for OC-3 access. PacBell's FasTrakSM ATM Cell Relay Service offering is being rolled out in phases, first utilizing Permanent Virtual Circuits (PVC) and in the future, Switched Virtual Circuits (SVC). As the ATM technology matures and becomes more of an industry and user standard, PacBell and other IXP operators will migrate to fully switched ATM IXP backbones.


SF NAP consists of ATM switching sites in the San Francisco area connected by OC-3 Synchronous Optical Network (SONET) links. Participants can access the NAP network using an ADC Kentrox ADSU and a Cisco 7000 or 7010 router. Access speeds reach 36.8 Mbps for DS-3 access and 139 Mbps for OC-3 access.

In addition to the ATM network, the NAP includes an interconnected FDDI dual-ring LAN. The FDDI LAN provides service to customers that require bandwidth less than 30 Mbps. The FDDI LAN was added when PacBell tests indicated that the ATM network was dropping cells at speeds between 20 Mbps and 30 Mbps. ISPs provide or lease dedicated T1 or T3 connections to PacBell DSUs and Cisco 7000 routers connected to the FDDI backbone. Exhibit 3-6 depicts PacBell's San Francisco NAP ATM/FDDI hybrid network architecture.

3.2.3.1 Routing

Each participating ISP must negotiate bilateral peering agreements with other ISPs before connecting with PacBell's San Francisco NAP. Routing on the FDDI ring is accomplished via the route server database maintained by the Routing Arbiter. On request, PacBell will provide NAP clients with a PVC to the Routing Arbiter route server database to receive and provide routing updates. Routing among peered ISPs may also be accomplished by direct PVC connections between the ISPs at the NAP, without regards to the route server database.

Exhibit 3-6
PacBell San Francisco NAP ATM/FDDI Hybrid Architecture



Source: PacBell

3.2.3.2 National ISP Clients

The San Francisco NAP interconnects numerous national and regional ISPs. National ISPs include ANS, MCI, and Sprint.

3.2.4 Metropolitan IXP Architecture Example

PSI, Inc. manages a metropolitan IXP in the Washington, DC area. PSI established the SWAB as an alternative IXP to the MAE-EAST NAP. SWAB operates nearly identically to the national-scope IXPs, requiring participating ISPs to negotiate peering agreements. Unlike the NAPs, the SWAB network is not facilities-based. Instead, each interconnecting ISP subscribes to Bell Atlantic's SMDS service over which the TCP/IP is routed.

Each participating ISP must subscribe to Bell Atlantic's SMDS service at a specified access class (speed). SMDS may be accessed at up to 34 Mbps, making it a lower bandwidth solution than FDDI or ATM. The ISP must supply its own dedicated access (either T1 or T3) to the SMDS service. To route TCP/IP over SMDS, the ISP must also provide an SMDS capable CSU/DSU and an IP router that supports SMDS encapsulation at the SWAB interface.

SWAB provides broadcast capabilities by use of SMDS address groups. The SWAB participants can have their SMDS address included in the SWAB SMDS address group for broadcast purposes.

3.2.4.1 Routing

The functionality of the Routing Arbiter's route server database is provided using SMDS address screening. Address screening is used to filter out SMDS addresses from the SMDS connection, analogous to how the SS7 network can screen calls from a voice line. An ISP's screen accepts packets from peered ISPs, while refusing packets from other ISPs. Each ISP must request that Bell Atlantic screen SMDS addresses from their SWAB interface.

3.2.4.2 National ISP Clients

Currently, PSINet and UUNet are the only national ISPs interconnected at the SWAB.

3.3 INTERNET ROUTING PROTOCOLS

The Internet, as previously described, is a collection of networks that allows communications between research institutions, universities, and many other organizations worldwide. These networks are connected by routers. A router is connected to two or more networks, appearing to each of these networks as a connected host. Forwarding an IP datagram generally requires the router to choose the address of the next router in the path or, for the final hop, the destination host. This choice, called routing, depends on a routing database located within the router. The routing database is also known as a routing table or forwarding table. The routing database should be maintained dynamically to reflect the current topology of the Internet. A router normally accomplishes this by participating in distributed routing and least-cost routing algorithms with other routers.

Routers within the Internet are organized hierarchically. Some routers are used to move information through one particular group of networks under the same administrative authority and control, known as an autonomous system (AS). Routers used for this purpose are called interior routers and they use a variety of Interior

Gateway Protocols (IGP). Routers that move information between ASs are called exterior routers and they use Exterior Gateway Protocols (EGP).

There is no standard protocol for either IGP or EGP. However, there are three protocols that are used by the ISPs and at the IXPs on the Internet. Generally, ISPs use the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol.

Most IXPs use the Border Gateway Protocol Version 4 (BGP4) as their routing protocol. All three protocols are dynamic in that the routers interact with adjacent routers to learn which networks each router is currently connected. The IGP protocols, RIP and OSPF, are detailed in Section 3.3.1 and 3.3.2, respectively. BGP4 is presented in Section 3.3.3.

3.3.1 Routing Information Protocol

RIP was developed by the Xerox Corporation in the early 1980s for use in Xerox Network Systems (XNS) networks. RIP is a dynamic protocol that continually updates its routing table based on information received from its adjacent routers. RIP is a distance-vector protocol, meaning that each router maintains a table of distances (hop counts) from itself to each other router in the system. These routing tables are updated based on RIP messages from adjacent routers.

RIP performs five basic operations:

- Initialization
- Request received
- Response received
- Regular routing updates
- Triggered updates.

On execution, RIP determines which of the routers interfaces are up and sends a request packet out on each interface. The purpose of this request packet is to ask each of its adjacent routers for their entire routing table.

A request received operation occurs when a router receives a packet from one of its adjacent routers asking for all or part of the router's routing table. The router will process the request and reply by sending the requested data.

A response received operation occurs when a router receives a response to its request for all or part of its adjacent routers' routing table. When a response is received, the router must validate the response and update its routing table.

Regular routing updates occur every 30 seconds. A router sends either all or part of its routing table to all of its adjacent routers. This ensures that each router on the network consistently has an accurate routing table.

Finally, a triggered update occurs when a router notices that one of its routes has changed. The router sends all routes from its routing table which are affected by the changed route, which may or may not be the entire table.

Although RIP appears to be a very simple protocol, it does have serious limitations. First, as shown by the series of operations in RIP, the protocol propagates either all or part of a router's routing table every 30 seconds in addition to any triggered updates. Subsequently, the protocol is very slow to stabilize when network failures or routing errors occur.

Second, RIP limits the number of hops between any two hosts on the network to 16. This means that hosts that are more than 15 hops apart within a single AS will not be able to communicate with one another. As a result, RIP is not well suited for large internetworks and works best in small environments.

Finally, when faced with multiple routes between a router and a network, RIP always chooses the path with smallest number of hops. This choice does not consider other cost factors such as line speed and line utilization, which are important when choosing a path between two nodes. Although RIP is still a very popular protocol, many companies are moving toward its replacement, OSPF.

3.3.2 Open Shortest Path First

OSPF was developed by the Internet Engineering Task Force (IETF) as a replacement for RIP. OSPF is designed to overcome the limitations of RIP and is supported by all major routing vendors. OSPF uses IP and its own protocol and the transport layer, not UDP or TCP. OSPF is a dynamic link-state protocol, unlike RIP, which is a distance-vector protocol. In a link-state protocol, a router does not exchange distances with its neighbors. Instead, each router tests the status of its links with its neighbors and sends this information to each adjacent router. Routers using OSPF are able to build an entire routing table based on the link-state information received from each of its neighbors.

In contrast to RIP, OSPF does not make its routing decisions based on the number of hops to a destination. Instead, OSPF assigns a dimensionless cost to each of interfaces of the router. This cost is not based on hop count, but on throughput, round trip time, reliability, etc. When the router is faced with multiple paths for a particular route, the routing decision is made using this cost. If two routes exist with the same cost, OSPF distributes the traffic equally among the routes. Additionally, OSPF allows multiple routes to a destination based on the IP type of service, e.g., Telnet, FTP, SMTP. This

means that a router can choose the best route for outgoing packets based on the type of traffic contained within the packet.

As described in Section 3.3.1, RIP is not well suited for larger internetworks because of its functionality. OSPF however, is designed for larger networks and stabilizes much faster when network failures or routing errors occur. OSPF also does not impose limitations on the number of hops between any two hosts because it does not use this metric when making routing decisions. Although RIP is still very popular, OSPF will ultimately replace RIP as the Internet grows.

3.3.3 Border Gateway Protocol Version 4

The primary routing protocol used on the Internet is BGP4. This protocol is used on Internet core (high level) routers to dynamically learn network reachability, respond to outages, and avoid routing loops in interconnected networks. Although RIP and OSPF are IGP's, BGP4 is an EGP used to pass traffic between different autonomous systems. BGP4 uses the TCP protocol to communicate routing information with its BGP4 peers.

Routers using BGP4 classify traffic as either local traffic or transit traffic. Local traffic is traffic that either originates or terminates in the router's AS. All other traffic is classified as transit traffic. The goal of BGP4 is to reduce the amount of transit traffic on the Internet.

The BGP4 system exchanges network reachability information with other BGP4 systems. This information includes the full path of autonomous systems that traffic must transit to reach the destination. The network reachability information is used by the router to construct a graph of AS connectivity. Once constructed, routing loops can be removed from the AS connectivity graph and routing policy decisions can be enforced.

BGP4 peers initially exchange their full routing tables. From then on incremental updates are sent as the routing tables change. BGP4 assigns a version number to the routing table and all adjacent routers will have the same version number for their routing tables. This version number changes whenever the routing table is updated as a result of routing information changes. To ensure that each adjacent router is alive, keepalive³ packets are sent between BGP4 peers whereas notification packets are sent in response to errors or other special conditions.

After a router using BGP4 receives routing updates, the protocol decides which paths to choose to reach a specific destination. Like RIP, BGP4 is a distance-vector protocol that allows only a single path to a destination. However, BGP4 does not impose a limit on the number of hops between two hosts and stabilizes quickly after network failures or

³ The BGP4 keepalive operation is independent from the TCP version of keepalive.

telephone lines. However, ISDN is gaining popularity with residential users as ISDN equipment and service prices drop. Both ISDN and analog modem connections use PN switched connections. The characteristics of analog modem and ISDN connections are described in Exhibit 3-7 below.

The bandwidth allocation for ISDN and analog modems is symmetric, meaning that there is an equal amount of inbound and outbound bandwidth. Unfortunately, many traffic applications are asymmetric, whereby the user receives far more inbound traffic than he or she generates. Examples of asymmetric applications include video-on-demand (small request to access a movie results in many gigabits of high resolution video) and Internet access (small request to access a Web page results in many megabits of text and images from the Web page).

**Exhibit 3-7
Analog Modem and ISDN Characteristics**

Characteristics	Analog Modem	ISDN
Speed	2.4 to 33.6 Kbps	64 to 128 Kbps
Equipment Cost	\$100 to \$150	\$300 to \$400
Representative Service Cost ⁴	Flat Rate Monthly (\$40/month)	Monthly Plus Usage (\$100/month + \$0.02/minute)

ILECs, cable companies, and direct satellite companies are testing and deploying several asymmetric access technologies (see Exhibit 3-8 below). These technologies have up to 30 Mbps of inbound bandwidth and up to 2 Mbps of outbound bandwidth.

**Exhibit 3-8
Asymmetric Internet Access Characteristics**

Characteristics	Direct Broadcast Satellite	ADSL	Cable Modems
Service Provider	DirecTV Satellite	ILECs	Cable Companies
Inbound Speed	400 Kbps	1.544 to 6 Mbps	10 to 30 Mbps
Outbound Speed	28.8 Kbps (over analog phone lines)	16 to 512 Kbps	768 Kbps to 2 Mbps
Equipment Cost	\$1,700	\$1,000	\$500
Service Cost	\$40/month	\$60 to \$100/month	\$40/month
Status	Deployed	In trial	In trial

⁴ Includes the cost of service from the LEC and the cost of access from the ISP.

routing errors occur. The decision process is based on different factors, including next hop, path length, route origin, local preference. BGP4 always propagates the best path to its adjacent routers. Currently, BGP4 is used by most IXPs on the Internet but is not defined as the standard EGP.

3.4 INTERNET ACCESS

The last (and in some ways the most vulnerable) component of the Internet architecture is the link between the service provider and customer. This access connection is typically a single dedicated or switched line over PN facilities.

Because access is provided over a single PN line, the connection is vulnerable to outages. This situation is identical to the "last mile" vulnerability of the PN architecture. Most other parts of the Internet architecture can use redundant links to route around outages. However, the access link is typically a single point of failure for an end user's connection to the Internet.

Internet access can be divided into two broad categories: business access and residential access. These categories are described separately below.

3.4.1 Business Access

Large and medium-size businesses use dedicated lines to connect their enterprise LAN/WAN to the Internet. These lines are either bundled with the ISP's service or leased separately by the company. In either case, the connection travels over PN facilities.

Most large businesses use T1 (1.544 Mbps) or higher connection speeds. Medium-size businesses use T1 or fractional T1 speeds (i.e., 128 Kbps to 768 Kbps) depending on their traffic requirements. Small businesses (10 to 50 employee sites) may be able to get by with a 56 Kbps leased line or a 128 Kbps Integrated Services Digital Network (ISDN) connection.

Leased line connections are available from ILECs and in metropolitan areas from CLECs. Today, CLEC companies include CAPs (e.g., Metropolitan Fiber Systems, Teleport Communications Group) and in many cases, IECs (e.g., LDDS, AT&T, MCIMetro). As legislation opens the local exchange to increased competition, leased lines may be available from utility companies, cable companies, or other providers.

3.4.2 Residential Access

Residential access connects a single user's computer to an ISP, reseller, or on-line provider. Most residential access is through modem connections over a LEC analog

The direct broadcast satellite offering is the only one of the three that is currently in widespread distribution. Direct broadcast satellite allows a user to receive inbound traffic over a 1-meter satellite dish and transmit outbound traffic over a standard analog modem line.

Asymmetric Digital Subscriber Line (ADSL) is a technology developed by the RBOCs to provide high bandwidth asymmetric connections over standard copper twisted pair wire. ADSL was originally developed exclusively for the home entertainment market (e.g., video-on-demand, interactive cable). However, as residential Internet access has grown in popularity, the LECs have added Internet access to their ADSL marketing efforts. ADSL is popular with LECs because copper cable is the basis for almost every residential phone installation. ADSL has a head start over its rival technologies because of the widespread deployment of copper wire (which reaches 98 percent of U.S. homes compared to 60 percent for cable). However, ADSL does have several drawbacks:

- Installation costs are high to upgrade existing copper cable to carry ADSL signals.
- Subscribers must be within 10,000 feet of the central office to reliably receive ADSL signals.
- Strong local AM stations can interfere with ADSL signals.
- The bandwidth available for communication is far less than the bandwidth available over cable modems.

Cable modems have the highest inbound and outbound bandwidth, but also have the most obstacles to widespread deployment. Cable modems depend on a two-way communication path between the cable operator and the subscriber. Almost every cable installation is designed to provide only a one-way path for video. To facilitate Internet access over cable plant, cable operators must upgrade their coaxial cable networks to two-way operation. Once upgraded, cable operators may have additional problems with the reliability of their plant, e.g., cable wires are installed only several inches below ground level and are highly susceptible to outages due to unintentional cable cuts. Once these issues are addressed, cable modems may easily fill a niche in the new market of Internet-enabled television (i.e., WebTV). Currently, access for these devices is provided using analog modems over dial-up lines.

4. INTERNET ANALYSIS

As described in Section 3, the Internet can be viewed as an interconnection of national and regional networks, end-users and organizations, and interexchange points. The Internet is a very dynamic entity that is constantly evolving and growing. Therefore, it is impossible to identify all of the components of today's Internet. For this report, the Internet is analyzed to identify key components used to transmit network traffic across the Internet. To achieve this purpose, a software tool, referred to as the IAT, was used to automatically trace the routes used to send traffic between two hosts on the Internet. The tool collects the set of routers an IP packet traverses on its path from one host to another. The analysis of these routes will identify traffic trends and key components in the Internet infrastructure.

This section provides an in-depth description of the IAT and analysis results. Section 4.1 details the functionality of the IAT. Section 4.2 details the implementation of the tool, including the set of hosts that was analyzed. Section 4.3 presents the analysis methodology and the results of the analysis.

4.1 INTERNET ANALYSIS TOOL FUNCTIONALITY

The purpose of the IAT is to collect the routes traveled by IP packets from one host to another. Because it is impossible to collect and analyze routes between every host on the Internet, a subset was chosen to provide an accurate sample of U.S. Internet traffic. Section 4.2 details the sites chosen for this analysis.

The IAT utilizes a UNIX utility, *traceroute*, to record the different routers a packet traverses once it is sent from the originating host to the destination host. The *traceroute* application is available with all UNIX and UNIX-variation operating systems. *traceroute* uses the Time To Live (TTL) field in the IP packet header to determine the routers in a particular path. The purpose of the TTL field is to ensure that packets do not stay on the Internet for an infinite amount of time (e.g., as a result of a routing loop). Each router that receives an IP packet is required to decrement the TTL field in the IP header by the number of seconds the router holds onto the datagram. Because most routers process a datagram in less than one second, the TTL field effectively becomes a hop counter that is decremented by one by each router.

IP packets are usually transmitted with a TTL of 60 by the originating host. When a router has an IP datagram with a TTL of one, the router decrements the TTL to zero, discards the packet, and returns an error message to the originating host. This error message is an Internet Control Message Protocol (ICMP) packet

that identifies the router that sent the error message and indicates that the time has been exceeded on the datagram.

The basic operation of the IAT is to send out *traceroute* IP datagrams beginning with a TTL of one, then a TTL of two, and so on, until the entire route between two hosts is determined. The router receiving the first IP datagram with a TTL of one will decrement the TTL and return an ICMP message to the originating host. This identifies the first router in the path. The IAT will then send out a second *traceroute* IP datagram with a TTL of two. The first router decrements the TTL to one, and sends the datagram to the next router in the path. The second router will decrement the TTL to zero and return the ICMP message. This continues until enough datagrams have been sent to have one of them reach the destination host. The destination will not discard the *traceroute* IP datagram, even though it will have a TTL of one because the datagram is addressed to that host.

For the IAT to determine that a datagram has reached its destination (because it has not received the final ICMP message), the IAT sends UDP datagrams to the destination host using a very high destination port number. The destination host will not respond to incoming packets on this port number; thus, the destination host will send back an ICMP "port unreachable" error to the IAT. The IAT differentiates between the time exceeded and port unreachable errors to determine when the route has been fully traced.

The output from an IAT execution is the set of routers in the path between two hosts. For each router, three datagrams are sent, and the round trip time from the originating host and the router is collected. Exhibit 4-1 depicts the sample output from a source host to the destination, www.disa.mil.

Exhibit 4-1 Sample Output From the IAT

```
Traceroute to www.disa.mil
 1 Cisco-AGS.dcmetro.bah.com (156.80.1.1)  2 ms  2 ms  2 ms
 2 fr.herndon.va.psi.net (38.2.104.1) 128 ms 227 ms 47 ms
 3 38.1.2.19 (38.1.2.19) 45 ms 77 ms 138 ms
 4 mae-east.ddn.mil (192.41.177.130) 68 ms 54 ms 41 ms
 5 137.209.1.2 (137.209.1.2) 176 ms 168 ms 204 ms
 6 198.26.127.10 (198.26.127.10) 179 ms 132 ms 114 ms
 7 164.117.2.13 (164.117.2.13) 134 ms 127 ms 134 ms
 8 164.117.1.1 (164.117.1.1) 143 ms 135 ms 125 ms
 9 www.disa.mil (164.117.147.116) 135 ms 147 ms 176 ms
```

4.2 INTERNET ANALYSIS TOOL IMPLEMENTATION

This section details the implementation of the IAT described in Section 4.1. For this analysis, two sites were chosen as source sites:

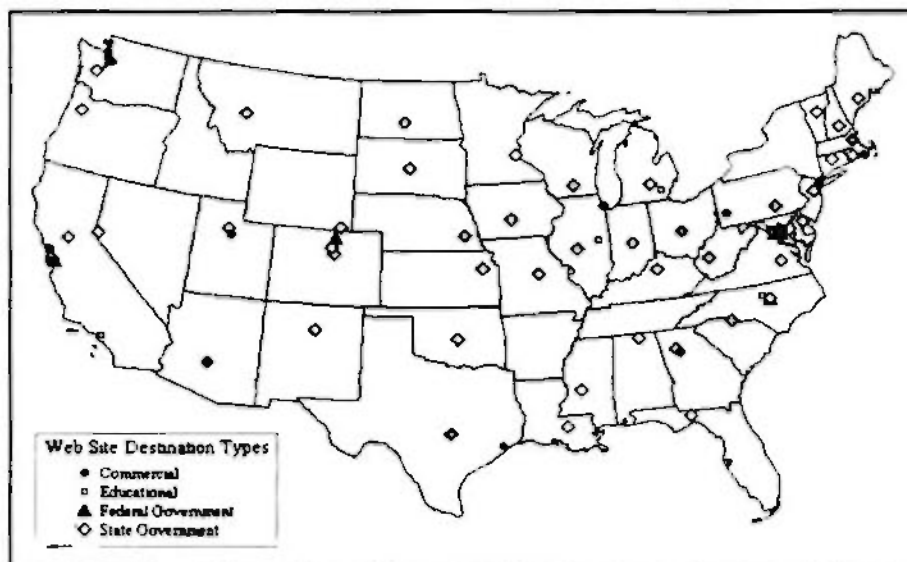
- Booz-Allen & Hamilton, McLean, Virginia, on the PSINet network
- Proxima, Inc., McLean, Virginia, on the MCI Network.

The tool collected routes from each of these two sites to 105 other sites located across the United States. The Web sites chosen for this analysis included the following:

- 23 NCS Member Organizations Web sites
- 50 State Web sites
- Major university Web sites
- Popular commercial Web sites.

Appendix A provides the entire list of Web sites used in this analysis. Exhibit 4-2 also shows the geographic locations of these sites. The IAT, which collects the routes from the two source locations to all 105 sites, is executed six times daily, every four hours beginning at midnight. This results in a sample of Internet traffic throughout the day. The output from this tool is formatted and loaded into an Oracle database where the analysis on the collection of routes is performed. Section 4.3 presents the analysis methodology followed by the IAT study.

Exhibit 4-2
IAT Site Locations



4.3 INTERNET ANALYSIS RESULTS

The data collected using the IAT represents a general picture of Internet connectivity. The destination Web sites used in the analysis were selected to provide both a United States and NCS specific view of the Internet's topology.

4.3.1 Internet Analysis Methodology

An in-depth analysis of the physical topology of the Internet would be an incredibly complex and difficult task. Because of the number of national backbones and regional distribution networks spanning multiple carriers, the Internet's topology is an amalgamation of CLEC, ILEC, and IEC networks. Determining the entire physical topology of the Internet may well be impossible without the cooperation of these PN carriers.

An analysis methodology was developed to provide the most complete and valuable view of the Internet and its topology. The methodology defines the steps used to evaluate the data obtained from the IAT. A description of the methodology is presented below.

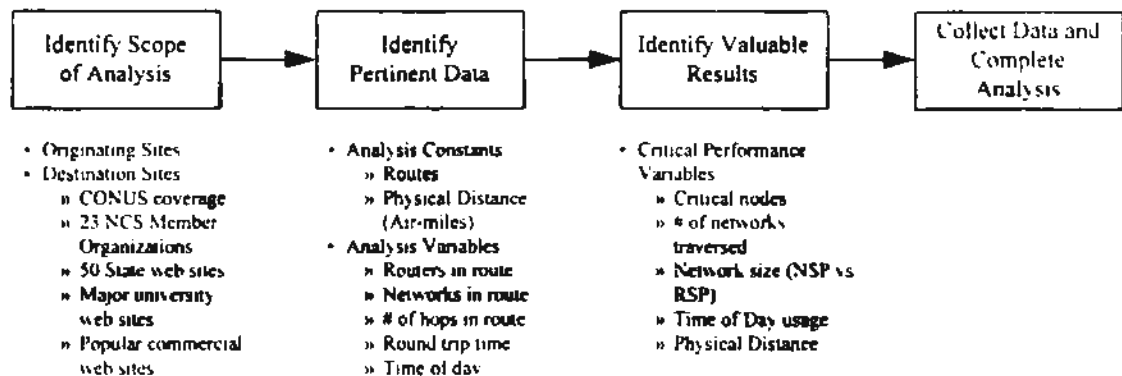
- *Identify Scope of Analysis.* Although the Internet is too large and complex to be handled in its entirety, the scope of the analysis was selected to provide a representative view of the Internet. The IAT is most useful in analyzing single, specific routes, not large network topologies. Given enough representative routes, the collective results of the IAT can provide a view of portions of the larger Internet. By collecting data at various times of day from multiple routes, the IAT provides a representative set of data. The originating and destination sites selected for this analysis provide a distribution of sites across the United States. The inclusion of the NCS Member Organizations provides a capability to capture and analyze data specifically for the NCS community.
- *Identify Pertinent Data.* The data used in the analysis must provide a complete and accurate picture of how IP packet traffic will be routed over the Internet. Variables such as the distance traveled, number of networks traversed, and the congestion of the network will affect how packets traverse the network. The IAT provides a host of data that is used to analyze our representative Internet routes. The data used in this analysis includes the following:
 - Origin to destination route
 - Physical distance of route in air-miles
 - Time of day

- Round trip time
 - Number of hops in route
 - Routers in route
 - Networks in route.
- *Identify Valuable Results.* The purpose of the analysis is to identify the discriminating variables that affect the Internet's performance. Using the available data (i.e., round trip time and number of hops in route), the analysis should indicate differences in performance based on the following variables:
 - Critical nodes
 - Physical distance between hosts
 - Time of day congestion
 - Number of networks traversed
 - Relative size of networks traversed (NSP versus RSP).

These results will provide input to the analysis of the vulnerabilities of the Internet. They may also identify how the OMNCS and NCS Member Organizations can improve Internet reliability by choosing certain ISPs, mirroring important Web sites, or performing downloads in off-peak hours.

Exhibit 4-3 illustrates the Internet analysis methodology.

**Exhibit 4-3
Internet Analysis Methodology**



4.3.2 Internet Analysis Results

The IAT analysis focuses on identifying the path that is critical for transmitting data across the Internet's regional and national backbone networks. The data provided by the IAT was analyzed to trace the paths through the Internet and to identify how Internet data traffic is affected by daily traffic surges and congestion and network outages. This analysis is intended to provide an estimate of the performance characteristics of a portion of the U.S.-based Internet. However, the results presented here cannot be assumed to represent the entire Internet, or even the entire U.S.-based network. This is because of the limited scope of the data, and the sheer size of the Internet in terms of routers and hosts. More thorough analyses of the entire Internet are planned as a follow-up to this initial analysis. We chose to use two source hosts for this analysis, one based on Booz • Allen & Hamilton's network and the other on Proxima, Inc.'s network. Booz • Allen and Proxima, Inc. receive Internet service from two of the six NSPs, PSINet and MCI, respectively. Therefore, this analysis may primarily represent the characteristics of these two networks.

The data provided by the IAT traces is the basis of a statistical analysis of the number of hops and round trip time for the 210 source and destination pairs (2 sources and 105 destinations). Traces were given a status of either "successful" or "unsuccessful." A successful trace was one in which the IAT packets generated reached the destination router address and an unsuccessful trace was one in which they did not. Exhibit 4-4 shows the number of successful traces per source and the percentage of the total traces performed.

Exhibit 4-4
Status of IAT Traces

<i>Source</i>	<i>Total Traces</i>	<i>Successful Traces</i>	<i>Unsuccessful Traces</i>
Booz • Allen	5134	4468 (87 %)	666 (13 %)
Proxima, Inc.	9128	8098 (88.7 %)	1030 (11.3 %)

A small percentage of the traces for both sources was determined unsuccessful. An unsuccessful trace could typically be attributed to one of the following reasons:

- The destination name server entry could not be resolved and therefore the trace never began
- An initial router of the ISP could not be reached
- A router or gateway in the path of the trace was unreachable

- The destination server was unreachable, most likely due to it being shut down
- The host's network might use code that is incompatible with the IAT testing protocol. That might have resulted in a router not returning the ICMP messages required for the operation of the IAT.

Exhibit 4-5 illustrates an approximate categorization of reasons why traces were unsuccessful. The percentages of those due to an unreachable path router, an unreachable destination server, or incompatible network code were combined. A hop-by-hop analysis of all unsuccessful traces, comprising nearly 45,000 hops, would be required to determine the component percentages.

**Exhibit 4-5
Categorization of Unsuccessful IAT Traces**

	Booz • Allen	Proxima
Unresolved host name:	2.6 %	0 %
ISP unreachable:	0.2 %	0.8 %
Router or destination machine unavailable:	10.2 %	10.5 %
Total Unsuccessful:	13.0 %	11.3 %

The results described in the remainder of this analysis are solely based on successful traces.

4.3.2.1 Traffic Congestion

Internet traffic encounters congestion due to surges in its use in daylight hours. Traffic surges occur during working hours, and most notably between noon and 6:00 p.m. Weekend traffic should not be as susceptible to Internet congestion because of the reduced number of business users. Our analysis assumes the effects of congestion will become manifest in the response time for data traveling over the Internet.

The IAT collects the round trip time for a single datagram to travel to and from each of the destinations. For each destination, three datagrams are sent, and the total travel time is recorded for each. The average travel time versus time of day for these datagrams is shown in Exhibit 4-6. As expected, these results appear to coincide with traffic patterns for a typical east coast IXP, MFS's MAE-EAST. The additional traffic on the Internet results in a proportional increase in the delay time. Representative weekday and weekend data for MAE-EAST and MAE-WEST are shown in Exhibit 4-7 and Exhibit 4-8, respectively. The traffic

increase between 12:00 noon and 4:00 p.m. shown in the MAE-EAST traffic profile is similar to that of our round trip time results. Note that traffic on MAE-EAST, located in Washington, DC, and MAE-WEST, located in San Jose, CA, are nearly identical for the time of day, based on eastern standard time (EST). Because of the large amount of traffic traveling between the east and west coasts, these two IXPs are interdependent. The traffic generated on the east coast between 12:00 noon and 4:00 p.m. eastern time affects the west coast traffic patterns between 9:00 a.m. and 1:00 p.m. Pacific time.

Exhibit 4-6
Average Round Trip Time Versus Time of Day

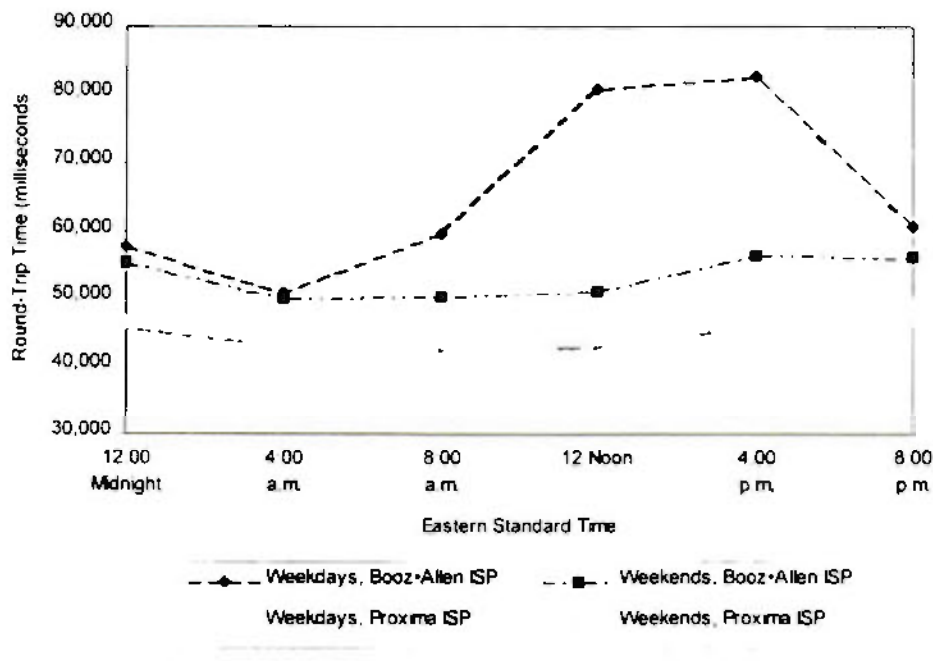
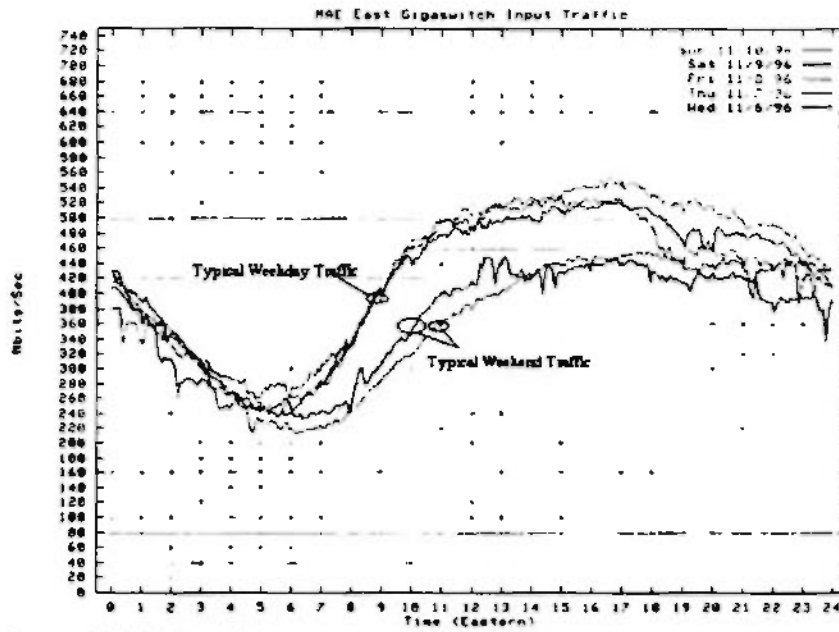
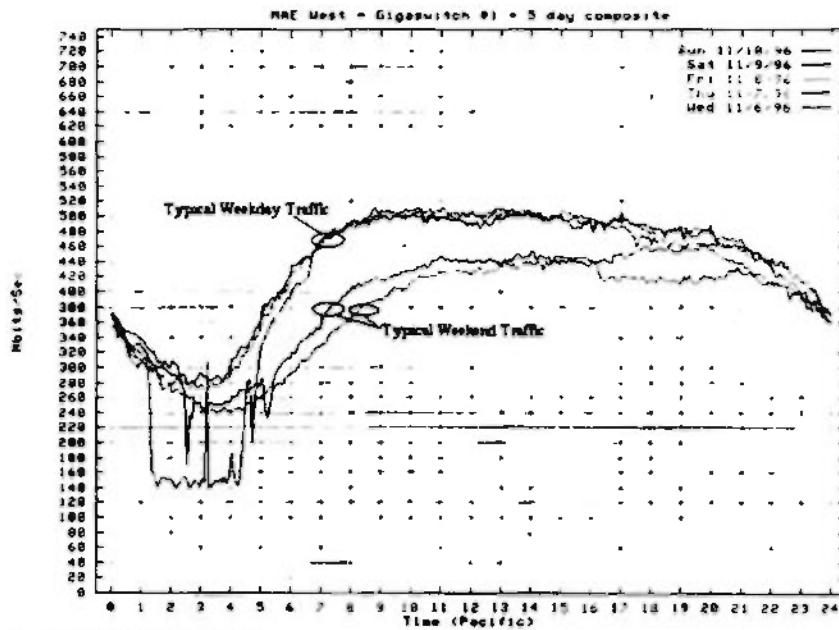


Exhibit 4-7 Typical Traffic Patterns at MAE-EAST



Source: MFS Datnet, Inc.

Exhibit 4-8 Typical Traffic Patterns at MAE-WEST



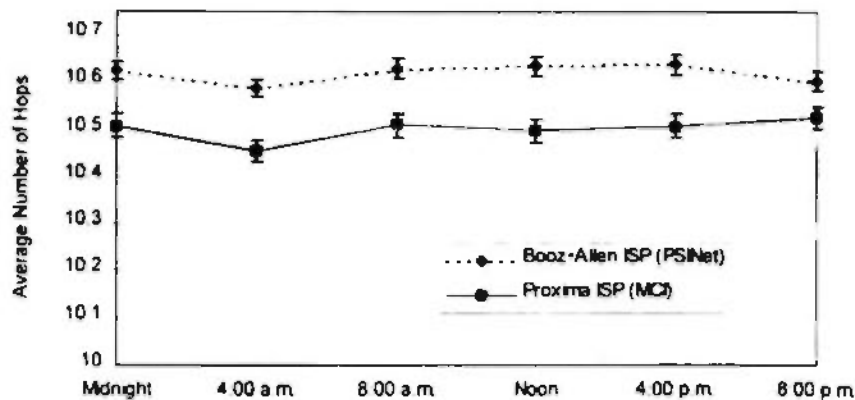
Source: MFS Datnet, Inc.

4.3.2.2 Network Outages

Network outages are the most disruptive of the Internet's vulnerabilities. In the case of critical nodes, a network outage can preclude access or egress from the network (as in the case of an isolated regional or local network) or severely hamper the flow of traffic (as in the case of a NAP or IXP failure). Network outages will occur with much less frequency than network congestion, but they may result in a significant reduction of network capacity and availability depending on their severity.

We determined the number of hops in each successful IAT trace from source to destination. Exhibit 4-9 compares the average number of hops with the time of day for each source network. It is clear that the number of hops does not depend on the time of day. This indicates that the path taken from source to destination does not change frequently due to outages or routing around network congestion. This is because Internet routing tables are generally static. Routing tables are meant to change during a disruption in service and in the event of network congestion. Although some routing algorithms will route around link congestion, this analysis indicates this is uncommon, because the number of hops does not depend on the time of day, while congestion does. Creating large Internet routing tables requires expensive processing power. This process can result in more route "thrashing" than actual routing. In fact, routing tables will normally only be recreated when links become disrupted, or when a network administrator manually replaces the routing table.

Exhibit 4-9
Average Number of Hops Versus Time of Day




We hypothesize that an outage in a critical network node, such as a national IXP, would greatly reduce, but not eliminate, the ability of the Internet to route traffic quickly nationwide.

4.3.2.3 Critical Network Nodes

As explained in Section 3.2, the Internet relies primarily on the national IXPs to route and exchange traffic. Exhibit 3-4 shows the locations of the major IXPs across the United States. These high-speed LANs provide the majority of the routing among the backbone NSPs and the RSPs. Additionally, traffic is exchanged at private direct connects between ISPs' networks. Private direct connect exchange points of this kind are becoming more common due to congestion at the IXPs. ISPs are establishing private direct connects to avoid congestion problems and improve routing redundancy.

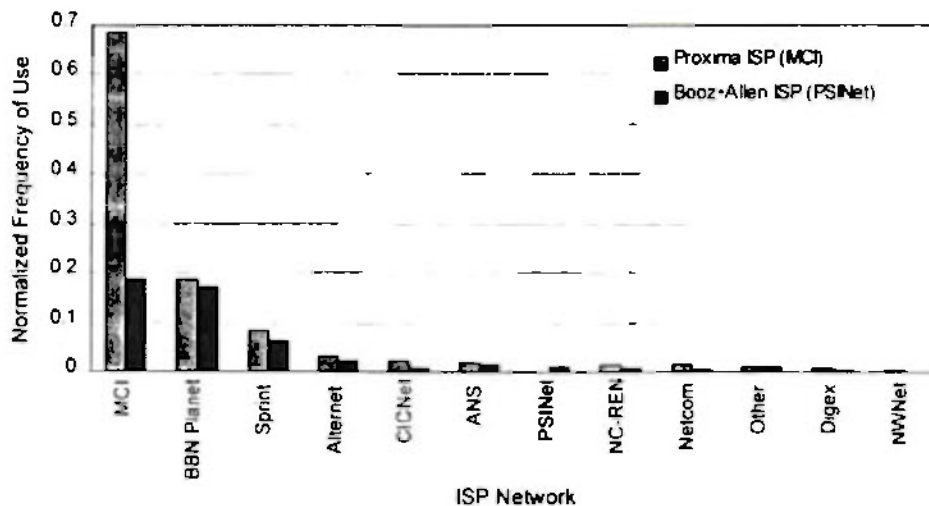
The IAT output provides the IP address of each of the routers traversed in the Internet traces. Using this data, we compiled lists of the most commonly visited routers for our two source networks. Exhibit 4-10 shows the distribution of the normalized frequency of use for the top 50 routers for both sources. The normalized frequency was obtained by dividing the number of hits on any router by the total number of hits recorded by the IAT for that source. This allows a direct comparison between the two sources. The first, second, and third router in each trace is considered to be specific to the source. These three routers show a very high frequency of use for our sources, and they are therefore critical to these sources, but do not fairly represent the remainder of the Internet. These three routers have been eliminated from the remainder of this analysis.

Exhibit 4-10
Top 50 Routers' Normalized Frequency of Use



The network domain names provided by the IAT output identify the router's owner. Using these domain names, we identified the relative importance of ISP networks to the two sources. The normalized frequency of use for each network is shown in Exhibit 4-11. "Other" networks were those networks that did not individually represent a large portion of the total frequency of use or that were not identified by a domain name.

Exhibit 4-11
Normalized Frequency of ISP Network Use




The critical network nodes had the highest frequency of use. Network nodes were considered critical if:


- They had a high frequency of use
- They were not too specific to the source routes, i.e., the top three routers
- They were not too specific to the destination routes.

All routers with normalized frequency greater than 0.004 were considered critical. Each of the sources shows a dependence on multiple critical network routers to trace a path to the destinations. Exhibits 4-12 and 4-13 show the critical ISP networks for the Booz • Allen and Proxima ISPs, respectively.

**Exhibit 4-12
Booz • Allen's Critical ISP Networks**




**Exhibit 4-13
Proxima's Critical ISP Networks**



Some of the critical nodes for one source were also critical to the other source. These nodes become our most critical nodes, which we can then identify as critical to the Internet based on our study. Exhibit 4-14 shows the distribution of these critical nodes to ISP networks.

Exhibit 4-14
Shared Critical Nodes



4.3.2.4 Conclusions

Traces performed throughout the test period indicated high success rates averaging between 87 and 89 percent. Of the unsuccessful trace attempts, most were due to an unreachable node (i.e., a router or the destination server) in the path that was probably either shutdown or incompatible with the IAT software.

Internet use is highest during mid-to-late afternoon business hours. Based on the round trip time for packets to traverse the network, congestion peaks between the hours of 12:00 noon and 4:00 p.m. eastern time. However, the dependence of businesses on the Internet could not be determined, i.e., the analysis did not determine whether the Internet was used to conduct critical business communications and research, or simply for personal use.

This analysis indicated that the number of hops did not depend on the time of day or the day of the week. Generally, routing tables are rarely modified to route around network congestion. Unlike switched traffic, the routes of Internet connections were somewhat "predictable." Therefore, the predictability of Internet routing, along with an increasing dependency on this communications media, renders it vulnerable to targeted and intended network disruptions.

Routers appear to share a somewhat balanced traffic load within the backbone networks (excluding those routers closest to the two sources). As expected, a high number of router "visits" occurred in the initial hops of the traces. These initial routers are critical to the sources, however, they are not necessarily critical to the entire Internet. As the trace moved away from the source and into the

backbone networks, the number of visits per router stabilized. Therefore, a single critical router could not be identified, however, it could be determined which networks were more heavily traversed. For this analysis, MCI's network was traversed most frequently and was therefore critical to the success of the traces.

5. VULNERABILITIES

This section addresses connectivity vulnerabilities that are inherent in the architecture of the Internet. These systemic vulnerabilities result from the utilization of the current PN infrastructure by the Internet composite networks. The vulnerabilities include second order effects such as availability and reliability due to outages on critical links and routing database errors. Security issues and vulnerabilities from outside influences, such as hackers, are not addressed. Internet vulnerabilities from hackers are addressed in the *Electronic Threat Intrusion Report*. The vulnerabilities associated with the ISPs, IXPs and Internet access connections are discussed below.

5.1 INTERNET SERVICE PROVIDERS

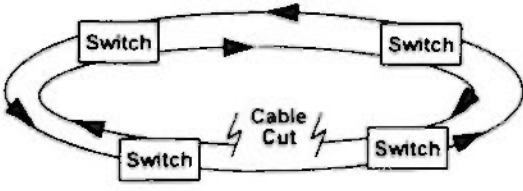
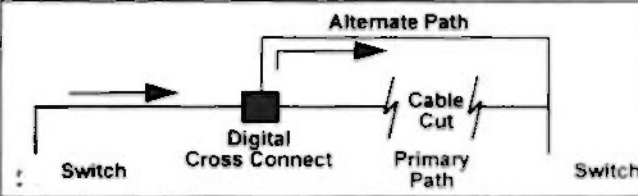
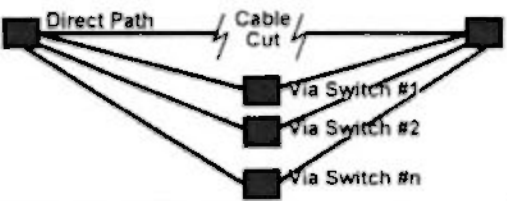
As introduced in previous sections, the ISPs provide the basic backbone architecture of the Internet. The ISPs can be divided into three categories: NSPs, RSPs, and resellers. Internet vulnerabilities that are unique to each category are detailed in the following sections.

5.1.1 National Service Providers

The majority of the NSP links travel over dedicated lines leased from the PN carriers. PN dedicated lines travel in the same conduit as other switched PN lines. Thus, the NSP links have a physical reliability comparable to that of the carrier's network. The IEC maintain their high reliability standards through a three tier restoration architecture. This architecture is based on protocols, physical diversity, and switching algorithms. Figure 5-1 details this tiered architecture.

The PN providers' current restoration techniques for cable cuts, the most frequent cause of outages, are not available for the dedicated lines used in the Internet. The switched-based mechanisms are not available because of the fundamental differences between switched voice and data communications. The protocol- and physical-based restoration mechanisms, however, could be employed for dedicated line failures. Each NSP needs to work closely with the PN providers to ensure that their dedicated lines are afforded these restoration techniques. For example, SONET rings are currently being deployed to increase the reliability of communications links. Traffic on a SONET ring automatically reverses its direction as a result of a cable cut. However, a PN provider may impose additional charges to add dedicated lines to a SONET ring if there are unused non-SONET protected lines available. Thus, the primary alternate routing schemes used to ensure connectivity is dependent on the NSP's routers, routing protocol, and restoration plans.

**Exhibit 5-1
PN Three Tier Restoration Architecture**

Mechanism	Basis	Description
SONET	Protocol Based	
Digital Cross Connects	Physical Based	
Dynamically Controlled Routing	Switching Based	

NSPs connect to multiple IXPs nationwide. Typically, an NSP's connection at each of these IXPs is non-redundant. If this connection is lost, the NSP will lose its connectivity to the IXP and the ability to exchange traffic with the other interconnected NSPs. However, if the NSP has connections to other IXPs, either regional or national, the NSP can still exchange traffic with the IXP-attached NSPs. The loss of the connection between an NSP and an IXP is critical only if the NSP does not have connections to multiple IXPs.

NSP networks are also susceptible to routing problems, such as slow convergence and routing loops. The three routing protocols discussed – BGP4, OSPF, and RIP – can affect routing within and between ISP networks. Because BGP4 is an external protocol, it can affect routing between ISPs. RIP and OSPF, which are internal protocols, will only affect an ISP's internal network.

RIP, the oldest of the three routing protocols discussed, has particular vulnerabilities that have been addressed by the newer protocols. RIP is a distance vector protocol based on hop count to the destination node. RIP routing tables contain only the single best route from origin to destination; when a better route is present, it replaces the old

route. When determining the best route available, RIP only considers the hop count and not other important factors such as bandwidth and line utilization.

Additionally, RIP is very slow to converge after a network failure or routing error has occurred. If a link in the route path is disrupted, RIP may not settle on the new best route for several minutes. During those minutes, service between those particular nodes is disrupted.

RIP is also susceptible to routing loops. In the minutes that it takes RIP to converge after a failure, routing loops may develop that will cause packets to route endlessly over the network until their TTL expires. Although there are modifications to the implementation of the RIP protocol that will help to avoid routing loops, they are subtle and may not be present in every network using RIP.

Finally, because RIP propagates its routing table to each of its neighbors every 30 seconds, RIP networks that are already congested by user traffic will be congested further by these routing tables.

The OSPF routing protocol overcomes RIP's shortfalls. The link state vector characteristic of OSPF allows each router in the network to have complete routing tables with multiple paths to destination. This greatly improves convergence time during a network failure and eliminates the chance of routing loops. OSPF routinely propagates route advertisements every half hour. OSPF also uses IP's multicasting capability to reduce the bandwidth requirement for these advertisements. This reduces the overall bandwidth overhead on the network attributed to the routing protocol. In time, OSPF will replace RIP as the standard internal routing protocol on the Internet.

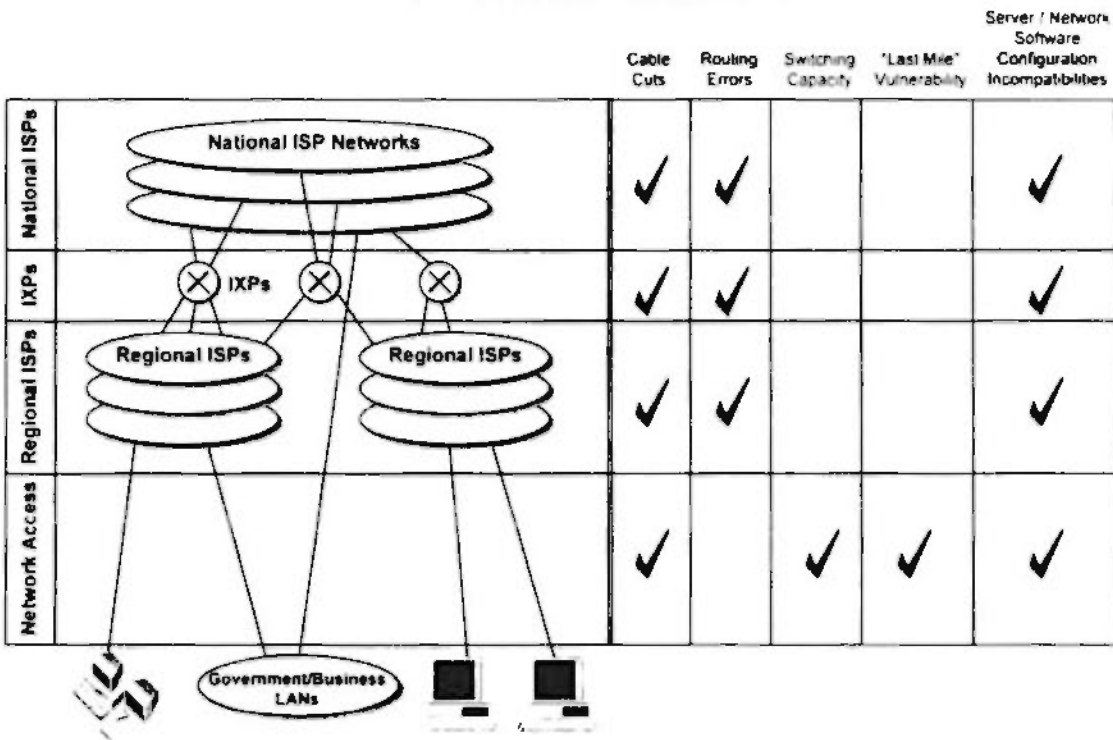
Exhibit 5-2 summarizes the vulnerabilities of the NSP networks, RSP networks, IXPs, and the access portion of the Internet architecture. These vulnerabilities are described in greater detail in the following sections.

5.1.2 Regional Service Providers

RSPs have similar vulnerabilities to those of the NSPs. These vulnerabilities may be compounded since RSP's smaller geographic scale limits the availability of physical diverse paths and their choice of a PN provider. This increases the possibility of isolation of the RSPs.

RSPs usually have fewer connections to IXPs. These connections may also be limited to the region that the RSP services. If one or more of an RSP's IXP connections is disrupted, the RSP's service will suffer greater degradation than an NSP. RSP service could be seriously affected by a regional natural or man-made disaster.

**Exhibit 5-2
Internet Architecture Vulnerabilities**



Because of an RSP's smaller geographic coverage, traffic will be carried over fewer links. If a major link fails because of a cable cut, it can have a large effect on the traffic within the RSP's network. For example, in NorthWestNet's backbone shown in Exhibit 3-2, the DS-3 circuit between Seattle, WA, and Portland, OR, is a critical high-bandwidth link. If that link fails, Portland's bandwidth to the national Internet connectivity provided at Seattle will fall from DS-3 (45 Mbps) to 2 T1s (3.088 Mbps) – a possible 93 percent drop in speed and bandwidth.

Since RSPs have smaller networks, much of their traffic is transmitted over other ISP's networks. Thus the effect of EGP routing, the IXP connection, and bilateral NSP network connection failures are more pronounced in an RSP's network. The RSP's traffic will also encounter the vulnerabilities of the NSP network carrying its traffic, including the reliability problems encountered due to routing errors.

5.1.3 Resellers

Resellers depend on their "host" ISP network to provide reliable and responsive service. Resellers typically have a single dedicated connection between their distribution facilities and its ISP. This connection typically travels over PN dedicated

lines. A failure in the dedicated line will result in a loss of service for the users homed to that distribution facility.

A reseller's network may become a congestion bottleneck when multiple customers access a single distribution facility with dedicated lines. If a reseller has not engineered the network connection for sufficient bandwidth to support dedicated and dial-up users, congestion may occur. This problem may occur in some reseller networks more than others.

Network availability is also a concern for dial-up customers of reseller networks. The ratio of customers to reseller modems may vary from 5 to more than 15. During high congestion periods, customers may be unable to gain access to the Internet. Higher ratio resellers have a greater potential for customer blocking.

5.2 INTEREXCHANGE POINTS

The interexchange point is the central location where ISPs meet to exchange network traffic. Recall from section 3, that all the necessary switching and routing equipment for all IXP-attached ISPs and for the IXP are physically located within a single facility. Subsequently, any disruption or disaster encountered at that facility could result in the loss of service at the IXP. For most NSPs the loss of one IXPs is not critical because NSPs generally have connections to multiple IXPs nationwide. However, for RSPs, the loss of an IXP is more critical, specifically if the RSP is connected to a single IXP.

In addition to the physical vulnerabilities, IXPs are susceptible to routing problems between the various interconnected ISPs. Routing problems could come from EGP protocol faults or invalid IXP routing tables. IXP operators attempt to eliminate routing problems by requiring a single EGP protocol at the IXP.

5.3 INTERNET ACCESS

The Internet access connection is the most vulnerable aspect of the Internet with respect to business and residential end users. Business connections are typically single, non-redundant connections from the business' LAN to the ISP. Like all critical single lines, if the connection is lost, the company loses Internet connectivity. Large companies with advanced nationwide WANs (e.g., GE, IBM, and Boeing) may employ redundant connections to the Internet for reliability. A business' Web page will also be vulnerable to a cut in the Internet access link. However, businesses may have their Web pages hosted on an ISP Web server instead of hosting them on their own network. This practice reduces LAN traffic and provides those Web pages with the additional reliability provided by the ISP network.

Residential access to the Internet is provided almost exclusively through analog modem or ISDN dial-up access. Both connections are over single connections and are a single point of failure for the residential connection. However, overall reliability of the PN remains very high. Reliability will drop when users access the Internet using alternate schemes, such as cable, which are not built to telephone industry standards.

Flat-rate pricing for Internet service has also introduced new availability issues for LEC PN networks. These networks' demand and pricing models were designed based on a 5-minute voice call, whereas Internet data calls can last hours. During times of crisis when voice and Internet traffic surge, long dial-up data calls may reduce the availability of the voice network using the same end-office switching capacity. Continued growth in the use of alternative access techniques such as cable modems and DirectPC satellites should eventually reduce these switching issues in PN carrier networks.

Some Internet users connect over direct broadcast satellite services, such as DirecPC. DirecPC uses an inbound satellite connection over a 1-meter dish and an outbound connection over an analog modem. If either leg of this connection fails, the entire connection will be lost. The reliability of the analog modem link will be the same as described above. The reliability of the satellite link will depend on the satellite terminal at the residential location and the satellite company's downlink location.

ADSL will be comparable in reliability to other LEC access technologies (e.g., analog modem, and ISDN). However, ADSL has limitations to where it can be installed. ADSL cannot be installed near a strong AM radio station because of AM frequency interference on the ADSL signal. Additionally, only homes within 10,000 feet of the LEC central office may be serviced by ADSL.

In the short term, cable modem reliability is close to that of the cable television provider. Cable modem service poses special reliability concerns because the cable industry, unlike the voice telephone industry, has not been required or expected to have the degree of reliability of phone service because it is not considered essential to public welfare (e.g., 911 emergency access). Typically, the cable has not been installed to telephone industry standards and has been installed in shallow trenches (typically less than 6 inches deep). Additionally, cable providers do not employ the restoration mechanisms of the traditional carriers. These factors make the cable facility, and ultimately the cable modem connection, very vulnerable to cable cuts and outages.

**APPENDIX A
INTERNET ANALYSIS TOOL SITES**

Organization	Web Site
Central Intelligence Agency	www.odci.gov
Department of Commerce	www.doc.gov
Department of Defense	www.dtic.dla.mil
Department of Health and Human Services	www.dhhs.gov
Department of Energy	www.doc.gov
Department of the Interior	www.doi.gov
Department of Justice	www.usdoj.gov
Department of State	www.state.gov
Department of Transportation	www.dot.gov
Department of the Treasury	www.ustreas.gov
Department of Veteran Affairs	www.va.gov
Federal Communications Commission	www.fcc.gov
Federal Emergency Management Agency	www.fema.gov
General Services Administration	www.gsa.gov
Joint Staff	www.dtic.dla.mil
National Aeronautics and Space Administration	www.nasa.gov
National Communication System	www.disa.mil
Nuclear Regulatory Commission	www.nrc.gov
United States Department of Agriculture	www.usda.gov
United States Information Agency	www.usia.gov
United States Postal Service	www.usps.gov
FedWorld Information Network	www.fedworld.gov
Library of Congress	www.loc.gov
Alabama	www.asc.edu
Alaska	www.state.ak.us
Arizona	www.state.az.us
Arkansas	www.state.ar.us
California	www.state.ca.us
Colorado	www.state.co.us
Connecticut	www.state.ct.us
Delaware	www.state.de.us
Florida	ww.state.fl.us
Georgia	www.state.ga.us
Hawaii	www.hawaii.gov
Idaho	www.state.id.us
Illinois	www.state.il.us
Indiana	www.state.in.us
Iowa	www.state.ia.us

Kansas	www.state.ks.us
Kentucky	www.state.ky.us
Louisiana	www.state.la.us
Maine	www.state.me.us
Maryland	www.mdarchives.state.md.us
Massachusetts	www.state.ma.us
Michigan	www.state.mi.us
Minnesota	www.state.mn.us
Mississippi	www.state.ms.us
Missouri	www.state.mo.us
Montana	nris.mls.mt.gov
Nebraska	www.state.ne.us
Nevada	www.state.nv.us
New Hampshire	www.state.nh.us
New Jersey	www-ns.rutgers.edu
New Mexico	www.state.nm.us
New York	www.state.ny.us
North Carolina	www.state.nc.us
North Dakota	www.state.nd.us
Ohio	www.ohio.gov
Oklahoma	www.oklaosf.state.ok.us
Oregon	www.state.or.us
Pennsylvania	www.state.pa.us
Rhode Island	www.state.ri.us
South Carolina	www.state.sc.us
South Dakota	www.state.sd.us
Tennessee	www.state.tn.us
Texas	www.texas.gov
Utah	www.state.ut.us
Vermont	www.state.vt.us
Virginia	www.state.va.us
Washington	www.wa.gov
West Virginia	www.slate.wv.us
Wisconsin	www.state.wi.us
Wyoming	www.state.wy.us
Alta Vista	altavista.digital.com
America Online	www.aol.com
Apple Computer, Inc.	www.apple.com
Computer Network (cnet)	www.cnet.com
CNN	www.cnn.com
CompuServe	www.compuserve.com
Digex (ISP)	www.digex.com
Interport (ISP)	www.interport.com

Lycos, Inc.	www.lycos.com
Macro Computer Systems, Inc. (ISP)	www.mcs.com
Microsoft, Inc.	www.microsoft.com
MTV	www.mtv.com
NetCom (ISP)	www.netcom.com
Netscape	www.netscape.com
Olympics	www.atlanta.olympic.org
Oracle Corporation	www.oracle.com
Primenet (ISP)	www.primenet.com
Sun Microsystems, Inc.	www.sun.com
USA Today	www.usatoday.com
WebCrawler	www.webcrawler.com
Windows95 Home Page	www.windows95.com
Word Magazine (on-line)	www.word.com
World Wide Web Consortium	www.w3.com
Yahoo	www.yahoo.com
Massachusetts Institute of Technology	www.mit.edu
Ohio State University	www.osu.edu
Stanford University	www.stanford.edu
University of California, Los Angeles	www.ucla.edu
University of Illinois, Urbana-Champaign	www.uiuc.edu
University of Michigan	www.umich.edu
University of North Carolina	www.unc.edu
University of Texas	www.utexas.edu

sl-dc-8-F0/0 sprintlink.net	186	(144.228.20.8)	Sprint
borderx1-fddi-1.WillowSprings.mci.net	180	(204.70.104.52)	MCI
core1.SanFrancisco.mci.net	176	(204.70.4.169)	MCI
ntis.bbnplanet.net	174	(192.221.253.22)	BBN Planet
Fddi0-0.CR2.DCA1.Alter.Net	173	(137.39.33.131)	Alternet
border2-fddi-0.Boston.mci.net	172	(204.70.3.34)	MCI
rtp1-gw.ncren.net	172	(128.109.70.248)	NC-REN
border1-fddi-0.Greensboro.mci.net	171	(204.70.80.18)	MCI
nc-research-net.Greensboro.mci.net	171	(204.70.81.6)	MCI
atlanta2-cr99.bbnplanet.net	170	(192.221.25.1)	BBN Planet
t3-2.was-dc-gw1.netcom.net	170	(163.179.220.181)	Netcom
rtp5-gw.ncren.net	170	(128.109.32.2)	NC-REN
mae-east.digex.net	170	(192.41.177.115)	Digex
fddi.mae-east.netcom.net	170	(192.41.177.210)	Netcom
f0.cnss61.Washington-DC.t3.ans.net	170	(140.222.56.197)	ANS
core1-aip-4.Greensboro.mci.net	170	(204.70.1.21)	MCI
border2-fddi-0.Denver.mci.net	170	(204.70.3.114)	MCI
atlanta2-cr99.bbnplanet.net	170	(192.221.25.230)	BBN Planet
atlanta3-cr1.bbnplanet.net	164	(192.221.42.1)	BBN Planet
border1-fddi-0.KansasCity.mci.net	163	(204.70.2.66)	MCI
midnet.KansasCity.mci.net	163	(204.70.40.6)	MCI
StLouis-StLouis2-f30.gi.net	162	(192.35.171.35)	Other
ut8-h1-0.the.net	162	(129.117.16.241)	Other
borderx1-fddi-0.WillowSprings.mci.net	161	(204.70.104.20)	MCI
cambridge2-cr3.bbnplanet.net	160	(192.233.33.10)	BBN Planet
cambridge1-cr1.bbnplanet.net	154	(192.233.149.201)	BBN Planet
cambridge2-cr2.bbnplanet.net	154	(192.233.33.2)	BBN Planet
cambridge2-cr2.bbnplanet.net	154	(199.92.129.2)	BBN Planet
sl-chi-15-H2/0-T3.sprintlink.net	153	(144.228.10.69)	Sprint
sl-kc-2-F0/0.sprintlink.net	152	(144.224.20.2)	Sprint
jackson-cr1.bbnplanet.net	148	(192.221.5.17)	BBN Planet
border1-fddi-0.Chicago.mci.net	147	(204.70.2.82)	MCI
merit-michnet-ds3.Chicago.mci.net	146	(204.70.24.6)	MCI
Hssi2-0.Vienna6.VA.Alter.Net	136	(137.39.100.78)	Alternet
sprint-nap.WestOrange.mci.net	135	(204.70.1.210)	MCI
Fddi0-0.SR1.TCO1.ALTER.NET	132	(137.39.11.22)	Alternet
borderx2-fddi-1.Seattle.mci.net	127	(204.70.203.68)	MCI
seabr1-gw.nwnet.net	127	(192.147.179.5)	NWNet
nwnet.Seattle.mci.net	127	(204.70.203.118)	MCI
core2.Seattle.mci.net	126	(204.70.4.33)	MCI
144.228.135.34	123	(144.228.135.34)	Sprint
core-hssi-3.Boston.mci.net	122	(204.70.1.2)	MCI
sl-atl-1-F0/0.sprintlink.net	119	(144.228.80.1)	Sprint
sl-fw-6-H3/0-T3.sprintlink.net	119	(144.228.10.86)	Sprint
dgc-fddi5-0.chicago.cic.net	118	(131.103.1.18)	CICNet
sl-fw-15-F0/0.sprintlink.net	116	(144.228.30.15)	Sprint
dgb-fddi5-0.chicago.cic.net	106	(131.103.1.17)	CICNet
f11-0.t56-0.Washington-DC.t3.ans.net	97	(140.222.56.66)	ANS

LIST OF ACRONYMS

ADSL	Asymmetric Digital Subscriber Line
ANS	Advanced Networks and Services
ANS CO+RE	ANS Commercial + Research and Education
ARPA	Advanced Research Projects Agency
AS	Autonomous System
ATM	Asynchronous Transfer Mode
BGP4	Border Gateway Protocol Version 4
CAP	Competitive Access Provider
CIX	Commercial Internet Exchange
CLEC	Competitive Local Exchange Carrier
DARPA	Defense Advanced Research Projects Agency
DoD	Department Of Defense
EGP	Exterior Gateway Protocol
EST	Eastern Standard Time
FDDI	Fiber Distributed Data Interface
FIX	Federal Internet Exchange
FTP	File Transfer Protocol
IAT	Internet Analysis Tool
ICMP	Internet Control Message Protocol
IEC	Interexchange Carrier
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
ILEC	Incumbent Local Exchange Carrier
IPv6	Internet Protocol Version Six
ISDN	Integrated Services Digital Network
ISI	Information Sciences Institute
ISP	Internet Service Provider
IXP	Interexchange Point
LAN	Local Area Network
LEC	Local Exchange Carrier
MAE	Metropolitan Area Ethernet
MAN	Metropolitan Area Network
MFS	Metropolitan Fiber Systems
MXP	Metropolitan Exchange Point
NAP	Network Access Point
NCP	Network Control Protocol
NCS	National Communication System
NREN	National Research and Education Network
NS/EP	National Security Emergency Preparedness
NSF	National Science Foundation
NSFNET	National Science Foundation Network

NSP	National Service Provider
OMNCS	Office of the Manager, NCS
OSPF	Open Shortest Path First
PC	Personal Computer
PN	Public Network
POP	Point of Presence
PPP	Point-to-Point Protocol
PVC	Permanent Virtual Circuit
RBOC	Regional Bell Operating Company
RIP	Routing Information Protocol
RSP	Regional Service Provider
SLIP	Serial Line Interface Protocol
SMDS	Switched Multimegabit Data Service
SONET	Synchronous Optical Network
SVC	Switched Virtual Circuit
SWAB	SMDS Washington Area Bypass
TCP/IP	Transmission Control Protocol/Internet Protocol
TTL	Time To Live
UDP	User Datagram Protocol
vBNS	Very High Speed Backbone Network Service
WAN	Wide Area Network
WWW	World Wide Web
XNS	Xerox Network Systems

SECTION 2 REFERENCES

1. Cerf, Vinton G., "Computer Networking: Global Infrastructure for the 21st Century," World Wide Web, www.cs.washington.edu/cra/networks.html, 1995.
2. CERFnet, "Network Service Provider Interconnections and Exchange Points," World Wide Web, www.cerf.net/cerfnet/interconnects.html
3. Cooper, Lane F., "The Commercialization of the Internet," *Communications Week*, April 1, 1996, pp. 135-139.
4. Fazio, Dennis, "Hang Onto Your Packets: The Information Super Highway Heads to Valleyfair or Building a High Performance Computer System Without Reading the Instructions," World Wide Web, www.mr.net/announcements/valleyfair.html, March 14, 1995.
5. Frazer, Karen D., "The NSFNET Phenomenon," World Wide Web, www.merit.edu/nsfnet/final.report/phenom.html.
6. Hardy, Henry Edward, "A Short History of the Net," World Wide Web, www.ocean.ic.net/ftp/doc/snethisthew.html, 1995.
7. MCI Telecommunications Corporation, "The vBNS Network," World Wide Web, www.government.com/vBNS/network_map.html, 1995.
8. Merit Network, Inc., "NSFNET: Transition to T3," World Wide Web, www.merit.edu/nsfnet/final.report/transition.html.
9. Merit Network, Inc., "Router Server Technical Overview," World Wide Web, www.ra.net.routing.arbiter/RA/.rs.overview.html.
10. National Laboratory for Applied Network Research, "Background Information," World Wide Web, www.nlanr.net/VBNS/background.html.
11. National Laboratory for Applied Network Research, "Collaboration on the Very High Speed Backbone Network Services (vBNS)," World Wide Web, www.nlanr.net/VBNS.
12. National Laboratory for Applied Network Research, "NSFNET -- The National Science Foundation Network," World Wide Web, www.nlanr.net/INFRA/NSFNET.html, November 23, 1995.

13. National Science Foundation, "NSF 93-52 - Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for NSFNET and the NREN(SM) Program - Program Solicitation," May 6, 1993.
14. Quarterman, John, "What is the Internet Anyway?," World Wide Web, gopher.tic.com/00/matrix/news/v4/what.408, 1994.
15. Rietz, Randy, Lewis, Will, "History of the Internet," World Wide Web, falcon.cc.ukans.edu/~wlewis/project/history.html, July 31, 1995.
16. Sprint, "Network Access Point Handbook," October 25, 1994.
17. Sprint, "SprintLink Customer Handbook 2.1," Sprint Document #5953-2, October 11, 1995.
18. Zakon, Robert Hobbes, "Hobbes' Internet Timeline v2.4a," World Wide Web, info.isoc.org/guest/zakon/Internet/History/HIT.html, 1996

SECTION 3 REFERENCES

1. Ameritech, "The Chicago NAP," World Wide Web, www.ameritech.com/products/data/map/The_Chicago_NAP.html, 1995.
2. Associated Press, "Computer Network Weathers Big Jolt: Internet Users Swap News, Worries After Quake Hits," Associated Press, January 18, 1994.
3. Bickel, Robert, "Building Intranets," *Internet World*, March 1996, p. 73.
4. CERFnet, "CERFnet T3 Backbone and Interconnectivity," World Wide Web, www.cerf.net/cerfnet/about/T3-map.html, June 7, 1996.
5. Cisco Systems, "BGP4 Case Studies Tutorial Section 1," World Wide Web, ciso.cisco.com/warp/public/459/13.html.
6. Cisco Systems, "Protocol Brief," 1994.
7. Cortese, Amy, "Here Comes the Intranet," *Business Week*, February 26, 1996, p. 76.
8. Coy, Peter, Judge, Paul, "Limo Service for Cruising the Net: MCI and BT Will Help Business Surfers Go First Class—for a Price," *Business Week*, June 24, 1996, p. 46.
9. Detroit MXP, "What is an MXP?," World Wide Web, www.mai.net/mxp/detroit/bckgrnd.htm.
10. Eng, Paul M, "War of the Web: Commercial Online Service Providers, Upstart Companies and Telecommunications Companies All Fighting for Internet Market," *Business Week*, March 4, 1996, p. 71.
11. Finneran, Michael, "Cable Modem Madness," *Business Communications Review*, March 1996, p. 68.
12. Holmes, Allan, "Flood Data Rides Internet Wave," *Federal Computer Week*, February 5, 1996, p. 1.
13. IITF, Reliability and Vulnerability of the National Information Infrastructure (NII), Information Infrastructure Task Force (IITF), August 17, 1995.
14. Loeb, Larry, "The Stage is SET: The SET Agreement Between MasterCard and Visa Could Pave the Way for Widespread E-commerce," *Internet World*, August 1996, p. 54.

15. MacKie-Mason, Jeffrey, Varian, Hal R., "Pricing the Internet," World Wide Web, gopher.econ.lsa.mich.edu, April 1993.
16. MacKie-Mason, Jeffrey, Varian, Hal R., "Economic FAQs About the Internet," World Wide Web, gopher.econ.lsa.umich.edu, August 21, 1994.
17. MacKie-Mason, Jeffrey, Varian, Hal R., "Some FAQs About Usage-Based Pricing," World Wide Web, gopher.econ.lsa.umich.edu, November 4, 1994.
18. Mendes, Gerald H, "Next-Generation IP Takes Shape," *Business Communications Review*, March 1996, p. 49.
19. Mills, Mike, "MCI Offers Customers Free Internet Access," *The Washington Post*, March 19, 1996, p. C1.
20. Netscape, "Netscape Announces New Real-time Audio and Video Framework for Internet Applications," Netscape Press Release, January 31, 1996.
21. Pacific Bell, "Multi-Lateral Peering Agreements Pacific Bell Network Access Point," World Wide Web, www.pacbell.com/Products/NAP/mlpa.html, August 14, 1995.
22. Pacific Bell, "Pacific Bell Network Access Point," World Wide Web, www.pacbell.com/products/business/fastrak/networking/nap/features.html.
23. *PC Week*, "UUNet puts ADSL on trial," *PC Week*, June 17, 1996, p. 3.
24. PSINet, "PSINet Technology and Infrastructure," World Wide Web, www.psi.new/psi-tech/psi-tech.shtml, 1995.
25. PSINet, "SWAB - SMDS Washington Area Bypass," World Wide Web, www.psi.net:80/misc/swab.html.
26. PSINet, "Typical POP Design," World Wide Web, www.psi.net/psi-tech/pop.html
27. Reilly, Patrick, "More Publishers Charging for Web Services," *Wall Street Journal*, May 8, 1996, p. B8.
28. Rigdon, Joan E, "Blurring the Line: New Technology Aims to Make the Web Look and Act More Like Television," *Wall Street Journal*, March 28, 1996, p. R5.
29. Sandberg, Jared, "Making the Sale: The Allure of On-Line Commerce, Its Proponents Argue, Will Eventually Prove Overwhelming," *Wall Street Journal*, June 17, 1996, p. R6.

30. Scott, D.F., "The Underground Internet: Through the MBONE, the Internet May Become the World's Largest Broadcast Service," *Computer Shopper*, March 1996, p. 548.
31. Sprint, "Network Access Point Handbook," October 25, 1994.
32. Stevens, Richard, "TCP/IP Illustrated Volume 1, The Protocols," Addison-Wesley Publishing, 1994, Chapter 10, pp. 97-110.
33. Swisher, Kara, "By the Sweat of Their Browser: District Entrepreneurs Turn a Web Search Idea Into a \$38 Million Deal," *The Washington Post*, June 4, 1996, p. C1.
34. Vaughan-Nichols, Steven J., "Radio Comes to Cyberspace," *Byte*, October 1995, p. 46.
35. Verity, John W., "Invoice, What's An Invoice: Electronic Commerce Will Soon Radically Alter the Way Business Buys and Sells," *Business Week*, June 10, 1996, p. 110.
36. Wingfield, Nick, "RSA to Connect Virtual Private Networks," *InfoWorld*, January 15, 1996, p. 47.
37. UUNET, "The UUNET Network Backbone," World Wide Web, www.uu.net.bbone.html.
38. Ziegler, Bart, "Up and Running: Why Did the Web Replace Interactive TV as the New Mantra? A Simple Reason: It's Here," *Wall Street Journal*, March 28, 1996, p. R6.

SECTION 4 REFERENCES

1. Asif, "U.S. Federal and State Government WWW Sites," World Wide Web, www.ilinks.net/~ace/html/government/html#sh6.
2. Bruno, Charles, "Internet Health Report: Condition Serious," Network World, September 16, 1996, pp. 1, 104-111.
3. InterNIC, "InterNIC Whois Service," World Wide Web, www.internic.net/wp/whois.html.
4. MFS Datanet, "MAE East Statistics," World Wide Web, ext2.mfsdatanet.com/MAE/east.stats.html.
5. MFS Datanet, "MAE West Statistics," World Wide Web, ext2.mfsdatanet.com/MAE/west.stats.html.
6. University of Illinois at Urbana-Champaign, "Host Name to Latitude/Longitude," World Wide Web, cello.cs.uiuc.edu/cgi-bin/slamm/ip2ll, June 19, 1995.
7. Stevens, Richard, "TCP/IP Illustrated Volume 1, The Protocols," Addison-Wesley Publishing, 1994, Chapter 10, pp. 97-110.