THE PRESIDENT'S
NATIONAL SECURITY TELECOMMUNICATIONS
ADVISORY COMMITTEE

INTERNET SECURITY/ARCHITECTURE TASK FORCE REPORT

DEFINING THE EDGE OF THE INTERNET

June 25, 2003

TABLE OF CONTENTS

EXECUTIVE SUMMARY **ES-1**

1.0 INTRODUCTION AND CHARGE **1**

2.0 DISCUSSION **1**

2.1 Defending the National Edge Is Critical—Defining the National Edge Is Not 1
2.1.1 Non-Routing Systems 2
2.1.2 The Composition of Information Systems 2
2.1.3 Zones of Responsibility 3
2.2 Factors That Should Not Define the Edge 3
2.2.1 The Internet Edge Should Not Be Defined Geographically 3
2.3.2 The Internet Edge Should Not Be Defined By Applications 4
2.3.3 The Internet Edge Should Not Be Defined By The Terminal 4

3.0 DEVELOPING KEY WARNINGS AND INDICATORS **4**

3.1 Finding the Edge For Red Networks 5

4.0 CONCLUSIONS **5**

5.0 RECOMMENDATIONS **6**

APPENDIX A—TASK FORCE MEMBERS, OTHER PARTICIPANTS, AND BRIEFERS **A-1**

EXECUTIVE SUMMARY

During the President's National Security Telecommunications Advisory Committee (NSTAC) XXV Principals' Meeting of March 12–13, 2002, concern was expressed over the ability to protect the "edges" of the Internet against attack or exploitation. In response to these concerns, the NSTAC formed the Internet Security/Architecture Task Force (ISATF) and tasked the group to provide guidance to the President on how to define the edge of the Internet. The ISATF developed this report in response to the NSTAC tasking.

Through detailed analysis, the ISATF determined that because the Internet is not a single network but a network of interconnected networks, no single definition of the edge exists. Although an Internet service provider's (ISP) view of the edge may be the end-user customer at a home computer or the dial-up modem bank from which those end users gain access to the network, a backbone ISP may view the edge as the point where another ISP interconnects with the backbone network. The ISATF noted that three different, additional ways to define the edge include, but are not limited to, all systems that contain Internet Protocol (IP) addresses that do not route IP packets; the composition of information systems; and zones of responsibility for network operators versus end users. In addition, the group noted that the adoption of a single definition of the edge could prevent critical security precautions from being addressed in other areas.

The ISATF also identified several factors that should not be used to define the edge of the Internet: geographic, physical boundaries lacking physical connection to the networks that compose the Internet; applications; and the needs of home users and small businesses.

As there is no single definition of the edge of the Internet, the ISATF agreed that further discussion to define the edge is not critical. Instead, additional attention should be given to defending the Internet as a whole. The ISATF suggested that efforts be undertaken by Government to identify the critical NS/EP missions and functions supporting those missions that rely on the Internet and encourage the parties responsible for those missions to ensure that they are adequately protected through redundancy and alternate capabilities. Industry, standards bodies, software vendors, equipment vendors, network operators, and end users of all products and services that make up the Internet should ensure that these products have built-in baseline security features and that these features are appropriately configured and kept up to date. Finally, Government should work with industry to identify key warnings and indicators that service providers can use as a baseline to measure security threats and trigger notification processes to relevant stakeholders. On the basis of its analysis of issues related to defining the edge of the Internet, the NSTAC offers the following recommendations:

The NSTAC recommends that —

- The Government should continue its work to identify the critical NS/EP missions and functions supporting those missions that rely on the Internet and encourage the parties responsible for those missions to ensure that they are adequately protected through redundancy and alternate capabilities;

- Industry, standards bodies, software vendors, equipment vendors, network operators and end users of all products and services that make up the Internet should ensure that these products have built-in baseline security features and that these capabilities are appropriately configured and kept up to date; and

- The Government should work with Internet security experts and standards bodies to develop a standard set of "key warnings and indicators" that all service providers can use as a baseline to measure security threats.

1.0 INTRODUCTION AND CHARGE

During the President's National Security Telecommunications Advisory Committee (NSTAC) XXV Principals' Meeting of March 12–13, 2002, concern was expressed over the ability to protect the "edges" of the Internet against attack or exploitation. To defend the edge, it is first necessary to define the edge and to understand the relevant perimeter that is needed to protect the Nation's critical infrastructures. However, there is no universal definition for the edge of the Internet. Many consider the Internet borderless or without any edge. Some consider the edge of the Internet to be the end-user terminal level—the millions and millions of desktops in homes and small businesses around the Nation and around the world—while others believe it to be the point at which an Internet service provider (ISP) interconnects with the backbone network. Consequently, the task force concluded that to protect the critical infrastructure of the United States, the focus should involve all elements of the Internet rather than attempt to focus security efforts at the edge. Security practices throughout the Internet must be as fluid and dynamic as the Internet itself.

2.0 DISCUSSION

2.1 Defending the National Edge Is Critical—Defining the National Edge Is Not

During the NSTAC Principals' meeting, participants sought a definition for the boundaries and edge of the Internet, and the means by which the boundaries and edge can be monitored. There are many different interpretations and definitions of the edge of the Internet. The concept, however, is clouded by the fact that the Internet is not a single, tangible thing, but a fabric of separate and distinct networks that are interconnected. Instead of a single edge, there is an ever-changing series of concentric circles that make up the Internet, extending all the way to the end device or user. This dynamic network of networks results in a perspective of the domain under one's management control based on where that individual sits in the overall architecture at any point in time. For instance, the edge concept as defined by the operators of the backbone networks—the large, long-haul fibers that help bring hundreds or thousands of networks together—may differ from that of the ISP network operators, which can either own a network or lease fiber from backbone operators or from network operators at large corporations or universities that can also act as de facto ISPs.

The definition of the edge may even vary within the ISP and backbone operator communities themselves. For ISPs, the end-user customer could be considered the edge, as could the dial-up modem bank accessed by those end users. For backbone providers, the edge could be the point at which an ISP interconnects with the backbone, or it could be the outermost edge of the downstream ISP. Both definitions are equally plausible.

Other methods of defining the edge include the following:

- Non-routing systems—all the systems with Internet Protocol (IP) addresses that do not route IP packets, e.g., leaf nodes;[1]
- The composition of information systems; or
- Zones of responsibility—the space between the network operator's zone of responsibility and the end-user's zone of responsibility. This could be considered the boundary between a customer edge and a provider edge.[2]

2.1.1 Non-Routing Systems

As noted in Request for Comment (RFC) 1958, the most robust end-to-end network functions are realized by end-to-end protocols. By keeping information about the end-to-end communications at the edge of the network, the network becomes more secure. Information about the state of end-to-end communications "should be maintained only in the endpoints, in such a way that the [end-to-end communications] can only be destroyed when the endpoint itself breaks (known as fate-sharing)."[3]

2.1.2 The Composition of Information Systems

The composition of information systems is important to the edge discussion because the Internet must accomplish business functions, e.g., e-commerce, e-government, to be considered useful. The components of an information system are as follow:

- The infrastructure, e.g., the Internet and the local access (including wireless systems) right out to the wall plug (utility model);
- The performing (user) computing elements or the "seats"; and
- The software applications.

In this case, the edge is the wall plug because users do not want to "own" the seats or take responsibility for the applications when viewed from an Internet provider perspective. It is necessary to monitor the security and performance of the edge to maintain the ultimate business applications/continuity.

---

1 This concept is consistent with the Internet Engineering Task Force's (IETF) RFC 1958, *Architectural Principles of the Internet.* RFC 1958 defines an end-to-end argument for implementing complex functions. One could interpret "end" as "edge" from this context, which would assert that "certain required end-to-end functions can only be performed correctly by the end-systems themselves." RFC 1958 at 2.3, p. 2.

2 The IETF's Provider Provisioned VPN Working Group is currently studying this issue. The customer edge, or customer edge device, "faces the users at a customer site that has an access connection to a provider edge (PE) device. It may be a router switching-router, or a switch that allows users at a customer site to communicate over the access network with other sites in the VPN." Provider edge devices are defined as "facing the provider network on one side and attached via an access connection over one or more access networks to one or more consumer electronic devices. It may be a router or a switching router."

3 RFC 1958 at 2.3, p.2-3.

2.1.3 Zones of Responsibility

Each network service provider or ISP is responsible for the equipment it needs to operate and maintain its network. This responsibility can include equipment that it owns and operates on the premises of one of its customers as well as equipment that it leases to a customer and maintains under a service agreement. If a piece of equipment fails or a group of customers experience latency or packet loss issues, the network operator or ISP responds accordingly. In this instance, the edge is the point at which the network service provider or ISP no longer has responsibility to act to protect its own assets, or by contract, the assets of its customers.

Nevertheless, instead of defining the edge, emphasis should be placed on defending the edge because the adoption of a single definition of the edge may prevent critical security precautions from being addressed in other areas. The concept of a single edge is impractical because the Internet is not a single network. Defensive measures focused only at the edge of a network are less effective than security measures employed in every software application and network element that constitutes the Internet. All network operators have responsibility for their networks and all of their components that make up the network. By encouraging network operators to regularly scan, monitor, and maintain not only the perimeter but the interior of the network, the Internet as a fabric becomes stronger. Treating each network as a separate and distinct component of the Internet allows for the reasonable identification of areas of responsibility for network security as defined and controlled by the individual network administrators. Responsibility for security can be extended to the users of the network through customer agreements, acceptable use policies, or terms and conditions.

2.2 Factors That Should Not Define the Edge

The Internet edge should not be defined using any of the following factors:

- Geographically by physical boundaries lacking physical connection to the networks that compose the Internet as a whole;
- By applications, such as e-mail or the World Wide Web, neither of which should be considered an acceptable boundary for the Internet; or
- By the needs of home users and small business users.

2.2.1 The Internet Edge Should Not Be Defined Geographically

Due to its continually evolving configuration, the Internet should not be defined geographically or geo-politically. The very word "Internet" belies its intangible nature. The Internet cannot be correlated with geography because permanent geographical features that pose physical barriers to the movement of people and goods are neutralized by the Internet's multiple technological pathways. Examples of this are the submarine fiber optic cables that crisscross under oceans and rivers, the fiber optic cables that pass through mountain passes and mountain tunnels, and the communications satellites that fly over these physical obstructions altogether. Nor can the Internet be defined geo-politically. As recent developments have shown, geopolitical boundaries are porous to the movement of people and are even more porous to streams of data packets.

Authoritarian governments have become increasingly aware of this phenomenon as they have sought to "turn off" the diverse pathways of ideas and information.

2.3.2 The Internet Edge Should Not Be Defined By Applications

Some may choose to define the edge of the Internet by examining the applications that touch all persons and working to ensure the security of those applications. E-mail, the World Wide Web, and the "walled gardens" of an ISP's closed community all provide areas of focus and interest in defining a national perimeter. Applications also exist that are not oriented towards the end user, but are machine-to-machine applications. However, while these applications have some common elements, they appear to be too generic to provide significant benefit (e.g., both classes can run on end-user terminals, some common protocols, or software platforms). Given that Americans do not use a standard suite of applications, the "edge" becomes too jagged to use applications as a defining factor for the Internet's edge.

2.3.3 The Internet Edge Should Not Be Defined By The Terminal

Another suggested method is to define the edge desktop by desktop. If all Americans were to install a base level of security on their computers, conceptually, it would become more challenging for that computer to be hacked and used as an attack node against the core network. However, defining the edge by end user presents significant security challenges and risks. Requiring secure terminals for every home user will automatically increase the cost of ISP service. ISPs simply do not have the resources or ability to proactively monitor home users' machines—especially when the purpose of monitoring is to ensure the currency of another company's software products. Liability issues also make this avenue significantly less inviting for the ISPs.

3.0 DEVELOPING KEY WARNINGS AND INDICATORS

In today's technology environment, monitoring networks for denial of service (DOS) attacks is extremely challenging. Many in the Government are calling for the development of network monitoring capabilities that will allow the detection of DOS attacks as they are launched. This will, someday, be a possibility; however, without significant advancements in router functionality, the processing power needed to examine each and every packet as it travels across the Internet will slow the Internet to a crawl or shut it down entirely.

That does not mean that the Nation is defenseless against cyber-attacks. From any perspective, crucial steps exist that must be taken to improve the security of the Internet as a whole. Companies and home users alike can implement best practices in their networks or home computers and can install and update anti-virus and firewall software to protect their systems from hackers and other intruders.

The Government can and should work with Internet security experts and standards bodies to develop a standard set of "key warnings and indicators" that all service providers can use as a baseline to measure security threats. Service providers and network operators are in the best position to monitor their networks. Although certain service providers can meet and exceed those baselines as a matter of customer service, standardized key indicators and warnings will provide earlier notice of national security level network events to downstream customers,

upstream providers, and relevant security groups or Government organizations that may then react accordingly.

Several existing organizations are able to develop key warnings and indicators. The Network Reliability and Interoperability Council's (NRIC) Focus Group 2 is involved in studying voluntary outage reporting for both Internet network operators and ISPs. The NSTAC Network Security Information Exchange (NSIE) is also equipped with the membership and technical expertise necessary to initiate publication of warning and indicator guidelines. Additional efforts to develop key warnings and indicators are included in the President's *National Strategy for Cyberspace Security*. These initiatives should be encouraged to combine with those efforts currently under way within the National Communications System (now part of the Department of Homeland Security) and privately, within organizations such as Dshield.org, Incidents.org, and others.

3.1 Finding the Edge For Red Networks

Historically, the Federal Government established a "red network" program for the transportation and containment of classified information. In essence, the edge of a red network exists at the point where the classified or sensitive data leaves the "red" (classified information) side and mixes into the "black" (unclassified information) side. These networks attempt to bound the traffic by a range of techniques ranging from "air gapping" through protective hardware and software elements to establishing laws and policies. Although these boundaries are the most easily identifiable, logical places to monitor the integrity of the edge of the red networks, these boundaries and protective measures institute a protected edge. Prudently, an operator of a red network should monitor traffic both well outside and inside the boundary between the red and black networks. The network operator can enhance protection and protect the network by blocking sensitive traffic, or reroute it somewhere else to allow the traffic to proceed. Ensuring the integrity of red or classified networks is an increasingly complex task as these red networks become more complex and intertwined with the rest of the public network. Thus, even in the red network program, defining the "edge" of the red network is becoming even more difficult. Even though the red network program is a good model to evaluate for protection of key assets within the Internet, various factors prevent its application to the entire Internet.

4.0 CONCLUSIONS

Defending—not defining—the national edge of the Internet is most important. Yet, defense of the Internet is a concept that is almost impossible, given that it implies that the Internet is defensible everywhere it touches—across every border around the world. The concept of a secure Internet will remain a global work in progress as it addresses a global dynamic problem. Neither hard nor soft technologies are currently at the level that allows for a totally secure Internet worldwide. Until these technologies become available, the United States Government should continue to focus its attention on critical NS/EP missions and functions supporting those missions that rely on the Internet and encourage the parties responsible for those missions to ensure that they are adequately protected through redundancy and alternate capabilities. Industry, standards bodies, software vendors, equipment vendors, network operators, and end

users of all products and services that make up the Internet should ensure that these products have built-in baseline security features and that these features are appropriately configured and kept up to date. Finally, Government should work with industry to identify key warnings and indicators that service providers can use as a baseline to measure security threats and trigger notification processes to relevant stakeholders.

5.0 RECOMMENDATIONS

- The Government should continue its work to identify the critical NS/EP missions and functions supporting those missions that rely on the Internet and encourage the parties responsible for those missions to ensure that they are adequately protected through redundancy and alternate capabilities;

- Industry, standards bodies, software vendors, equipment vendors, network operators and end users of all products and services that make up the Internet should ensure that these products have built-in baseline security features and that these capabilities are appropriately configured and kept up to date; and

- The Government should work with Internet security experts and standards bodies to develop a standard set of "key warnings and indicators" that all service providers can use as a baseline to measure security threats.

APPENDIX A—TASK FORCE MEMBERS, OTHER PARTICIPANTS, AND BRIEFERS

TASK FORCE MEMBERS

MCI, Inc.	Ms. Joan Grewe, Chair
Raytheon Company	Mr. Sebastian Taphanel, Vice-Chair
AT&T Corporation	Mr. Harry Underhill
Bank of America Corporation	Mr. Roger Callahan
BellSouth Corporation	Mr. Shawn Cochran
The Boeing Company	Mr. Robert Steele
Cisco Systems	Mr. James Massa
Computer Sciences Corporation	Mr. Guy Copeland
Lockheed Martin Corporation	Mr. Daniel Tolley
Lucent Technologies	Mr. David Massarik
Nortel Networks	Dr. Jack Edwards
Northrop Grumman Corporation	Mr. Scott Freber
Science Applications International Corporation	Mr. Hank Kluepfel
SBC Communications, Inc.	Ms. Rosemary Leffler
VeriSign, Inc.	Mr. Ken Silva
Verizon Communications, Inc.	Mr. James Bean

OTHER PARTICIPANTS

Cisco Systems, Inc.	Mr. Brian O'Connor
George Washington University	Dr. Jack Oslund
Juniper Networks, Inc.	Mr. Pejhan Peymani
Microsoft Corporation	Mr. Sean Finnegan
National Security Council	Mr. Marcus Sachs
Raytheon Company	Mr. James Craft
SBC Communications, Inc.	Mr. Paul Hart
VeriSign, Inc.	Mr. Michael Aisenberg
MCI, Inc.	Ms. Cristin Flynn

BRIEFERS

CanSecWest	Mr. Dragos Ruiu
Equinix, Inc.	Mr. Jay Adelson
The Honeynet Project	Mr. Shane Macaulay
National Security Council	Ms. Marjorie Gilbert
Sun Microsystems, Inc.	Mr. Lance Spitzner
University of Washington	Mr. David Dittrich