

Justification and Requirements for a National DDoS Defense Technology Evaluation Facility^{*†}

July 26, 2002

Wes Hardaker
Darrell Kindred
Ron Ostrenga
Dan Sterne
Roshan Thomas

Network Associates Laboratories

15204 Omega Drive
Rockville, Maryland 20850

* This research was supported by DARPA/SPAWAR contract N66001-01-C-8046.

† Approved for Public Release; Distribution Unlimited.

Executive Summary

The explosive growth of the Internet and its increasingly critical role in supporting electronic commerce, transportation, and communications, have brought an equally explosive growth in attacks on Internet infrastructure and services. Some of the most difficult attacks to defend against are the Distributed Denial of Service (DDoS) attacks, in which an overwhelming flood of network packets is generated from many different sources, with the intent of preventing legitimate use of services. Typically, DDoS attacks are directed at one or more targets, such as end-users, web servers, entire networks or parts of networks, or networking infrastructure components.

DDoS attacks pose a severe threat to the nation's ability to conduct business, defend itself, and provide vital government services to its citizens. Medium-scale DDoS attacks have been observed frequently during the past 2-3 years, and larger scale attacks are increasingly likely. For example, an amateur attacker disabled some of the world's largest web services (e.g., Yahoo!, CNN, Amazon, and Buy.com) for hours in February 2000. More recently, attacks against the CERT® Coordination Center [14] and edNET [15], a Scottish ISP, caused major disruptions in service. A determined enemy could perpetrate focused attacks that disable vital services at critical times, disrupt commerce, create uncertainty and panic among the public, and effectively prevent much of the electronic communication the U.S. Government relies on today. This serious national vulnerability can only be addressed through substantial and coordinated efforts by government and industry.

Although DDoS has recently drawn significant attention from the network security research community, no general approach to a solution has yet been identified. The nature and goals of effective DDoS solutions are not yet clear, and no broadly applicable, practical DDoS defense implementations have been produced. Even simple DDoS attacks cause significant disruptions today; more sophisticated attack techniques, undoubtedly being developed, will be difficult or impossible to counter with current defense technology. Research experts believe the DDoS problem to be fundamentally difficult for several reasons. The network traffic transmitted in a DDoS attack can be virtually indistinguishable from traffic for legitimate use of a service. DDoS attacks make use of forged source addresses, indirection, reflection, and other techniques to conceal the locations of the computers that are the real sources of the attacks. Finally, the attack sources themselves, typically commandeered from unwitting legitimate users, are widely distributed among different independent networks, so stopping an attack at its source presents technological and administrative challenges.

Developing DDoS defense solutions that will be effective against large, sophisticated attacks requires testing and experimentation in a realistic environment. Commercial and government entities developing or evaluating DDoS defense technology ordinarily have only small prototype test networks available to them. These small test networks do not accurately represent the Internet. Software is used to emulate hardware-based network routers, the available bandwidth is usually much lower than in the real Internet core, and the complex and dynamic structure of the Internet is not reflected. On the other hand, DDoS defense solutions cannot be effectively tested on the Internet because such testing would cause unacceptable disruptions of the Internet, which has become a critical national infrastructure component. A large-scale experimental network is thus the only environment in which DDoS defense technologies can be developed and evaluated effectively.

Construction and maintenance of a sufficiently large test facility for DDoS defense experimentation is too expensive for all but a few large companies. The largest Internet Service Providers (ISPs) and the large router manufacturers have built medium- to large-scale test networks, but these networks are not available to other organizations nor were they designed for DDoS defense research. More importantly, the companies that own large test networks are not motivated to invest aggressively in solving the DDoS defense problem. Router vendors have

little economic incentive for adding sophisticated DDoS defense features to their products. They assert that customer demand is insufficient, especially given their suspicions that DDoS defense features will be expensive to develop and test, and may reduce router throughput. Even if router vendors provide such features, ISPs may be hesitant to utilize them because ISPs want to avoid assuming any legal responsibility for policing the content of the network traffic they deliver. Finally, because a single DDoS attack typically travels over multiple ISPs and brands of network devices, responsibility for effective DDoS defense is diffuse; this reduces the potential impact gained by the introduction of advanced DDoS defense features by a single ISP or router vendor.

For these reasons, it is unlikely that industry, left to its own devices, will address the nation's DDoS vulnerabilities with the urgency required. To make rapid advances in DDoS defense, the United States must instead tap the energy and ideas of a broad cross-section of organizations—large and small, government and industry, research and operational—and promote coordination and collaboration on this critical problem. This in turn requires the existence of a large-scale, shared test facility that can support accurate evaluation of DDoS solutions as they are developed and provide experimental results to guide further research and product development. To date no such facility exists.

This report describes the results of a study funded by DARPA to determine whether a national facility for DDoS defense technology evaluation is needed and to identify its requirements. The study also lays the groundwork for follow-on activities, should they be initiated, such as obtaining funding to support coordination and planning among participating government agencies. The study's findings are based on interviews conducted with experts in relevant fields, including network security product vendors, computer security researchers, router manufacturers, network service providers, government organizations, network operators, and content providers.

The study envisions a National DDoS Defense Technology Evaluation Facility whose charter would be to provide a shared laboratory in which researchers, developers, and operators from government, industry, and academia can experiment with potential DDoS defense technologies under realistic conditions, with the aim of accelerating research, development, and deployment of effective DDoS defenses for the nation's computer networks. This facility would be a shared national asset, serving a wide range of clients attacking the DDoS problem. The following requirements were identified:

- The facility must realistically emulate conditions on the Internet. It must use hardware and software currently in use on the Internet, on a scale that partially represents the Internet's complex interactions.
- The network must be flexible and easily reconfigurable so that it can support experiments requiring wide variations in network topology and hardware configuration.
- The network must not be a production network. Network outages that would be unacceptable on a production network should be expected as a normal result of experimentation.
- The environment must provide realistic network traffic. One of the important criteria used in evaluating DDoS defense solutions is the ability of the solution to suppress attacks while allowing legitimate traffic to flow unimpeded.
- The environment must be sufficiently controllable to support repeatable experiments.
- All proposed uses of the facility must be reviewed to ensure consistent application of the facility's charter and usage priorities.
- The facility must have skilled, on-site technical staff that can help clients make efficient use of their time in the facility.

Other requirements concern physical location, security, operational requirements, service level agreements, data archiving, scheduling, staffing, and funding.

The study outlines and contrasts five design approaches for the facility, which vary in capability, complexity, and cost. The facility can be built in incremental phases, leveraging components, design experience, and other resources of existing test laboratories and network facilities. As it matures, the facility can provide DDoS defense evaluation, training, and consulting without bias towards specific vendors or technologies. The approximate cost of building and operating such a facility for the first year is estimated at \$18,000,000 to \$90,000,000, depending on a number of design parameters, with annual costs thereafter ranging from \$7,000,000 to \$55,000,000.

This study has confirmed that a facility meeting the requirements above is urgently needed to address the growing threat that DDoS attacks present to the nation. Such a facility can provide the test bed required for advancing the state-of-the-art in DDoS mitigation through careful study of DDoS attacks and defenses in a realistic, controlled environment. This facility can also foster, coordinate, and accelerate efforts by government and military agencies, researchers, and developers of DDoS defense products toward the common goal of protecting our critical national infrastructure from DDoS attacks.

Table of Contents

1	INTRODUCTION	6
1.1	MOTIVATION FOR THE STUDY	6
1.2	PURPOSE OF THE STUDY	8
1.3	SCOPE OF THE STUDY	8
1.4	METHODOLOGY	8
1.5	ORGANIZATION OF THIS REPORT	9
1.6	TERMINOLOGY USED IN THIS REPORT	10
2	DISTRIBUTED DENIAL OF SERVICE ATTACKS AND DEFENSES	11
2.1	CHARACTERIZING THE DDoS PROBLEM	11
2.2	DDoS DEFENSE METHODS	12
3	FACILITY CHARTER AND OBJECTIVES	14
3.1	POTENTIAL TYPES OF CLIENTS AND THEIR NEEDS	14
3.2	THE BENEFITS OF A DDoS TEST FACILITY	16
3.3	VALUE OF EACH BENEFIT TO EACH CLIENT TYPE	17
4	FACILITY REQUIREMENTS	19
4.1	PUBLIC AVAILABILITY	19
4.2	VENDOR NEUTRALITY	19
4.3	NETWORK TOPOLOGY	19
4.4	NON-PRODUCTION NETWORK	23
4.5	RAPID RECONFIGURATION	23
4.6	REALISTIC NETWORK TRAFFIC	25
4.7	DATA ARCHIVING	28
4.8	NETWORK MANAGEMENT, MONITORING, AND ANALYSIS	28
4.9	TECHNICAL STAFFING	28
4.10	USAGE ADMINISTRATION	29
4.11	PHYSICAL AND NETWORK ACCESS CONTROL	30
4.12	FINANCIAL REQUIREMENTS	31
4.13	LEGAL AND REGULATORY ISSUES	31
5	RECOMMENDATIONS	33
5.1	FACILITY DESIGN OPTIONS	33
5.2	ESTIMATED COSTS	48
5.3	FUNDING MODELS	51
5.4	MEASURES OF SUCCESS	54
5.5	OPERATIONAL RECOMMENDATIONS	55
6	FUTURE CONSIDERATIONS	57
6.1	FUTURE APPLICABILITY OF THE FACILITY	57
6.2	EVALUATION ORGANIZATION	57
6.3	BAKE-OFF HOSTING	57
6.4	DDoS SECURITY EXPERT CONFERENCE	58
7	SUMMARY AND CONCLUSIONS	59
	REFERENCES	61
	APPENDIX A. ACKNOWLEDGEMENTS	62
	APPENDIX B. ESTIMATED COSTS	63

Tables

TABLE 1. ESTIMATED VALUE TO EACH CATEGORY OF CLIENTS.....	18
TABLE 2. SELECTED CHARACTERISTICS OF THE DESIGN OPTIONS	46
TABLE 3. OPTION VALUE TO EACH CLIENT TYPE.....	48
TABLE 4. ESTIMATED COSTS FOR CONSTRUCTION AND OPERATION OF THE FACILITY (IN MILLIONS OF DOLLARS)	51
TABLE 5. ADVANTAGES AND DISADVANTAGES OF FUNDING AND REVENUE GENERATION MODELS.	54
TABLE 6. SUMMARY: ESTIMATED COSTS FOR CONSTRUCTION AND OPERATION OF THE FACILITY (IN MILLIONS OF DOLLARS)	67

FIGURES

FIGURE 1. NETWORK PEERING.....	20
FIGURE 2. ATTACHED NETWORKS AND SITES.....	21
FIGURE 3. AN EXAMPLE OF THE FACILITY’S BACKBONE	34
FIGURE 4. A COMPLEX EXAMPLE OF AN ATTACHED SITE	35
FIGURE 5. POSSIBLE STEPS IN A PHASED IMPLEMENTATION.....	36
FIGURE 6. OPTION 1: A SINGLE SITE FACILITY WITH NO ACCESS TO THE INTERNET FROM THE TEST DATA NETWORK.....	37
FIGURE 7. OPTION 2: A SINGLE SITE FACILITY WITH TEST DATA ACCESS TO THE INTERNET	40
FIGURE 8. OPTION 3: A MULTI-SITE FACILITY WITH NO ACCESS TO THE INTERNET FROM THE TEST DATA NETWORK.....	42
FIGURE 9. OPTION 4: A MULTI-SITE FACILITY WITH TEST DATA ACCESS TO THE INTERNET	43
FIGURE 10. OPTION 5: A MULTI-SITE FACILITY OFFERING ALTERNATIVE ROUTING	45
FIGURE 11. A SUGGESTED MODEL FOR FACILITY OPERATIONS	56

1 Introduction

With the increased availability of broadband access to the Internet and the lack of security associated with many university and home-user networks has come an increased proliferation of network-based attacks[1]. Compounding this problem is the increased reliance by the United States on the Internet as part of the critical infrastructure for electronic commerce and communications. Some of the most difficult network-based attacks to defend against are the Distributed Denial of Service (DDoS) attacks, in which an overwhelming flood of network packets is generated by many different sources, with the intent of preventing legitimate use of services. Typically, DDoS attacks are directed at one or more targets, such as end-users, web servers, entire networks or parts of networks, or networking infrastructure components (e.g., routers, communications links, load balancers, or firewalls).

DDoS attacks pose a severe threat to the nation's ability to conduct business, defend itself, and provide vital government services to its citizens. Medium-scale DDoS attacks have been observed frequently during the past 2-3 years [1], and larger scale attacks are increasingly likely. For example, an amateur attacker disabled some of the world's largest web services (e.g., Yahoo!¹, CNN, Amazon, and Buy.com) for hours in February 2000. More recently, attacks against the CERT Coordination Center [14] and edNET [15], a Scottish ISP, caused major disruptions in service. A determined enemy could perpetrate focused attacks that disable vital services at critical times, disrupt commerce, create uncertainty and panic among the public, and effectively prevent much of the electronic communication the U.S. Government relies on today. This serious national vulnerability can only be addressed through substantial and coordinated efforts by government and industry.

Although DDoS has recently drawn significant attention from the network security research community, no general approach to a solution has yet been identified. The nature and goals of effective DDoS solutions are not yet clear, and no broadly applicable, practical DDoS defense implementations have been produced. Even simple DDoS attacks cause significant disruptions today; more sophisticated attack techniques, undoubtedly being developed, will be difficult or impossible to counter with current defense technology.

Research experts believe the DDoS problem to be fundamentally difficult for several reasons. The network traffic transmitted in a DDoS attack can be virtually indistinguishable from traffic for legitimate use of a service. DDoS attacks make use of forged source addresses, indirection, reflection, and other techniques to conceal the locations of the computers that are the real sources of the attacks. Finally, the attack sources themselves, typically commandeered from unwitting legitimate users, are widely distributed among different independent networks, so stopping an attack at its source presents technological and administrative challenges.

1.1 Motivation for the Study

Developing DDoS defense solutions that will be effective against large, sophisticated attacks requires testing and experimentation in a realistic environment. Commercial and government entities developing or evaluating DDoS defense technology ordinarily have only small prototype test networks available to them. These small test networks do not accurately represent the Internet. They typically consist of a few dozen systems at most, connected via a handful of software-based routers with relatively low-bandwidth (100Mbps) network links. The Internet, by contrast, includes millions of systems (hundreds of thousands of which may be involved in a single DDoS attack [12]) connected via a backbone network comprising thousands of specialized hardware routers with network links of 2.5 Gbps or more. Therefore it is no surprise that the

¹ All registered and unregistered trademarks in this document are the sole property of their respective owners.

Internet behaves very differently, under both normal and attack conditions, than these test networks. The differences are not just ones of scale but include differences in structure, dynamism, and complexity as well. For example, DDoS attacks increasingly target network infrastructure components, including routers, and routing protocols [2]. Experiments on small test networks that lack these infrastructure components and use statically configured, software-based routers cannot accurately model such attacks or defend against them.

Given that small test networks are inadequate models of the Internet, one might suggest testing DDoS defenses on the Internet itself. This is impractical for several reasons. First, the Internet infrastructure has become a critical national resource; carrying out DDoS attacks or deploying experimental defense technologies in the Internet core would almost certainly cause unacceptable network disruptions. Second, the Internet environment is difficult, if not impossible, to control, and thus cannot support repeatable DDoS experiments. Third, privacy concerns regarding Internet communications would hamper effective data collection. A large-scale experimental network is thus the only environment in which DDoS defense technologies can be developed and evaluated effectively.

Construction and maintenance of a sufficiently large test facility for DDoS defense experimentation is too expensive for all but a few large companies. The largest Internet Service Providers (e.g., UU Net, AT&T, and Genuity) and the large router manufacturers (Cisco and Juniper) have built large-scale test networks, but these networks are not available to other companies, research organizations, or the government. These large corporate test facilities are usually reserved for corporate R&D, integration testing, and customer demonstrations. These facilities represent substantial investments and competitive advantage for their owners, so there is little incentive for their owners to permit use by others. Furthermore, these facilities tend to be homogeneous environments, tailored specifically to the needs of their owners, so they are not well-suited to general purpose DDoS defense experimentation and technology evaluations.

More importantly, the companies that own large test networks are not motivated to invest aggressively in solving the DDoS defense problem. Router vendors have little economic incentive for adding sophisticated DDoS defense features to their products. They assert that customer demand is insufficient, especially given their suspicions that DDoS defense features will be expensive to develop and test, and may reduce router throughput. Even if router vendors provide such features, ISPs may be hesitant to utilize them because ISPs want to avoid assuming any legal responsibility for policing the content of the network traffic they deliver. Finally, because a single DDoS attack typically travels over multiple ISPs and through multiple brands of network devices, responsibility for effective DDoS defense is diffuse; this reduces the potential impact gained by the introduction of advanced DDoS defense features by a single ISP or router vendor.

For these reasons, it is unlikely that industry, left to its own devices, will address the nation's DDoS vulnerabilities with the urgency required. To make rapid advances in DDoS defense, the United States must instead tap the energy and ideas of a broad cross-section of organizations—large and small, government and industry, research and operational—and promote coordination and collaboration on this critical problem. Creating a large-scale, shared test facility that can support accurate evaluation of DDoS solutions as they are developed and provide experimental results to guide further research and product development appears to be a crucial step in mobilizing these resources and increasing their effectiveness. Without such facilities, progress in DDoS defense will continue to be outpaced by improvements in DDoS attack technology. This in turn will further increase the likelihood of successful attacks that cause protracted network outages, disrupting communications and critical functions within both government and commercial sectors.

1.2 Purpose of the Study

This report describes the results of a DARPA-funded study to gather and analyze evidence concerning the need for a national DDoS defense technology evaluation facility and identify the facility's primary requirements. The study also lays the groundwork for follow-on activities, should they be initiated, such as obtaining funding to support coordination and planning among participating government agencies.

1.3 Scope of the Study

The scope of the study was limited to gathering and analyzing the needs of various classes of organizations concerned with DDoS defense, developing requirements for a DDoS defense evaluation facility, outlining alternative approaches for the design and evolution of the facility, estimating facility costs, and identifying potential funding models. Developing detailed design specifications and operational procedures for the facility is beyond the scope of the study.

1.4 Methodology

To ensure that the information gathered for this study properly reflects the views of the organizations concerned with DDoS defense technology, we solicited information from many sources. We interviewed experts in multiple fields of study and requested their input on the subjects of DDoS problems, DDoS solutions, technology testing, product evaluations, and general networking. These individuals are identified in Appendix A.

1.4.1 Types of Organizations Contacted

During the course of this study, we contacted many organizations and companies. We have identified the following categories of clients that can be served by the proposed facility:

- Government agencies (NASA, NSA, NRL, DARPA, DISA)
- Research labs
- Commercial network equipment vendors
- DDoS defense technology vendors
- Internet Service Providers (ISPs)
- Commercial users of DDoS defense technology (web hosting facilities, enterprise data centers, corporate networks, etc.)

We contacted a few organizations in each of the above categories. Some of the organizations above have a security research focus, while others are primarily interested in testing and piloting their networking and DDoS protection products.

1.4.2 Interviews and Tours

The requirements were gathered over a six-month period. Many on-site interviews, as well as telephone conference calls, were conducted to determine whether there was a need for a more realistic test environment in which research technologies and new DDoS defense security products could be validated, and to clarify what each participant needed to make the facility useful to them. Interviews generally lasted from one to two hours each, and interviewees were assured that any sensitive information (e.g., specific network and host configurations) would be kept confidential. On-site interviews were often accompanied by tours and demonstrations of the interviewees' facilities. The number of scheduled interviews grew as we interviewed organizations, as each participant typically suggested two to three more organizations we should contact.

In addition to the individual interviews, network operators were asked for their input at a DDoS Test Facility “Birds of a Feather” (BOF) session held at the 23rd North American Network Operators' Group (NANOG) meeting, October 21-23, 2001, in Oakland, California. Feedback was also solicited from computer security researchers at the DARPA Fault Tolerant Networks Principal Investigators meeting, January 15-18, 2002, in San Diego, California.

1.4.3 Types of Data Gathered

Most interviews were conducted in a free flowing format. Questions (topics) were focused based on the organization type that the interviewee represented. There were four major areas covered during the interviews:

1. General Questions
2. Researchers/Vendors Specific Questions
3. Government/ISPs Specific Questions
4. Test Facility Questions

1.5 Organization of this Report

This document begins with an introduction into the DDoS problem space in Section 2. This is followed by an overview of the proposed facility in Section 3. This overview includes detailed information about the facility's charter, the clients it would attract, and the benefits it would provide to those clients. Section 4 describes in detail the technical and administrative requirements that the facility must satisfy for it to be effective in helping solve DDoS-related problems. In Section 5 initial establishment of the facility is discussed. This includes details about how to conduct an incremental build-out, what the resulting facility might consist of, and a discussion about possible funding sources. This is followed by a discussion in Section 6 about future directions to consider once the facility has been established. Section 7 contains conclusions drawn from the study and summarizes the recommendations contained within the document.

1.6 Terminology Used in this Report

Denial of Service (DoS)	Attacks intended to disable or reduce the level of availability of a particular network service.
Distributed Denial of Service (DDoS)	A distributed denial of service is a class of denial of service attack that overwhelms a service by flooding it from a multitude of distributed locations.
Experiment	Any approved and scheduled use (e.g., DDoS defense product evaluation or testing of any DDoS technologies) of the testing facility.
Customer	A user of a vendor's product or service. Does not imply a relationship with the facility (see "Client").
Client	A user of the facility.
Researcher	A person, company, or organization attempting to develop solutions for the DDoS problem.
Facility	The testing environment for studying DDoS attacks and defense solutions. This paper outlines the requirements of such a facility.
Participant	A subject matter expert who was interviewed for this study.
Spoofed Address	A falsified Internet address used as the source address of a network packet. Typically, these addresses are randomly generated for each outgoing attack packet to hide the location of the real attacking source.
Internet Service Provider (ISP)	A company or an organization that provides or facilitates connections to the Internet.
Tier-1 ISP	A large ISP that provides central routing and backbone capacity to the Internet. Typically, they connect smaller ISPs together.

2 Distributed Denial of Service Attacks and Defenses

2.1 Characterizing the DDoS Problem

Distributed Denial of Service (DDoS) attacks have received much attention lately in the computing security community and in the industry at large. This can be attributed to the fact that the victims of these attacks have included well known web sites and electronic commerce companies. It is now estimated that the DDoS attacks in February 2000 on the CNN, Amazon, Buy.com, and Yahoo! Web sites caused millions of dollars in lost business [6].

Researchers and practitioners in the security community have long held that computer security has three primary objectives: confidentiality, integrity, and availability. A denial of service attack is fundamentally an attack on availability. The attacker seeks not to expose secrets or tamper with the victim's data, but to prevent the victim from effectively providing or using some service. DDoS attacks are a special class of denial of service attacks in which the attacker makes use of a large number of network-connected machines to carry out the attack.

The distributed denial of service problem is considered one of the most difficult security problems to solve. DDoS attacks are launched in a distributed and coordinated manner using automated agents on multiple machines. These agents are often difficult to locate as they may use spoofed source addresses. Many DDoS attack tools can be downloaded from well-known Internet hacker sites where new tools are being deployed at alarming rates. Accumulated experience by practitioners and researchers in dealing with denial of service (DoS) attacks has led to some consensus on the broad classification of these attacks. Attack classes include the following:

- **Bandwidth consumption.** These attacks consume all available bandwidth on one or more network links and thereby deny bandwidth to legitimate traffic. This may be accomplished in one of two ways. An attacker who has more available bandwidth than a victim's network can flood the victim's slower network connection. Alternatively, an attacker, even if using a slow network connection, can amplify the attack by using multiple sites to launch a distributed attack to flood the victim's network (see documentation on the Shaft tool [5]).
- **System resource starvation.** These attacks focus on consuming system resources such as CPU time, memory, and file-system usage quotas. By consuming these resources in an excessive manner, they are deprived for legitimate system and user needs.
- **Exceptional condition exploitation.** These attacks exploit design and programming flaws that result in the failure of an application, operating system, or hardware device to handle certain exceptional conditions. By inducing such conditions, the attack may slow down or disable the affected system. Some of the well-known attack techniques in this category involve sending malformed network packets to cause system crashes.
- **Routing and Domain Name Service (DNS) manipulation.** Routing-based DoS attacks involve malicious manipulation of routing table entries, causing network traffic to be improperly routed through the Internet. Attacks on DNS servers involve inducing these servers to cache bogus address information so that legitimate traffic is directed to the wrong Internet (IP) addresses. Either kind of attack may prevent the victim from properly sending or receiving network packets, or cause the victim to be flooded with packets misdirected to its network.

In principle, any of these denial of service attacks can be carried out in a distributed manner, as a DDoS attack, though distribution usually provides the most leverage in bandwidth consumption and system resource starvation attacks.

2.2 DDoS Defense Methods

Approaches to DDoS protection can be classified as network-based, source-based, or end-point-based, according to where the defenses are deployed. In this section, we describe these general classes and provide examples of defenses in each class that are currently in use or under investigation. Effective comprehensive DDoS defense will probably require implementing a combination of these methods.

2.2.1 Network-Based Defense

This approach takes a systemic view of DDoS protection. The goal is to protect as much of the network infrastructure as possible, by reducing congestion in communication links caused by attack traffic flows. These flows start as large numbers of relatively small flows from individual flooding agents. The small flows successively join into larger and larger flows as they approach the intended victim and in many cases overwhelm the capacities of one or more links along the way to the victim.

Intermediaries in the network, such as routers, switches, firewalls, proxies, and load-balancers, can be used to monitor network conditions as well as take defensive action when necessary. A complete network-based approach to DDoS prevention will require network operators and ISPs to have a thorough understanding of end-to-end network congestion and choke points, and to make a coordinated response to a DDoS attack. The ideal response will be to block traffic as close to the attackers as possible. Thus, the ability to locate the sources of attack traffic is a crucial component of effective network-based defense. Network-based solutions to DDoS prevention are most effective against bandwidth consumption and possibly network-wide routing attacks. The problem of coordination among independent administrative domains presents a significant obstacle to comprehensive network-based defense today.

Some limited, and not entirely effective, forms of network-based DDoS defense are in use today. *Rate limiting* and *quality-of-service* mechanisms are used to limit the bandwidth allocated to certain classes of traffic, such as the ICMP and UDP messages used in some DDoS attacks, at the expense of dropping some legitimate messages. *Ingress filtering*, in which any packet whose source addresses does not fit the network address of the interface it arrived on, can provide partial protection from some forms of source-address spoofing. A topic of current research is *traceback-and-block* mechanisms, in which attack traffic flows are traced to their sources through a variety of techniques and automatically blocked or constrained [3], [4], [11].

2.2.2 Source-Based Defense

Source-based DDoS defense approaches attempt to prevent attack traffic at or near its source, before it enters the Internet core. In many cases, it is easier to identify attack traffic near the source; for example, spoofed packet source addresses are easier to detect within the originating network since the range of legitimate addresses for outgoing traffic is known there. The primary shortcoming of this approach is that it relies on administrators for the many potential attack-source networks to deploy defenses. Those network administrators may have little motivation to stop outgoing DDoS attacks that do not directly affect their network, or they may lack the knowledge or resources required to establish effective defenses.

The most widely deployed form of source-based defense today is *egress filtering*, in which an organization blocks any outgoing traffic from its network whose source address does not match the organization's assigned address range. This is accomplished with access router and firewall filters. When deployed close to an attacker, this can greatly restrict the degree of source address

spoofing that is possible. Most security-conscious network operators do use egress filtering; unfortunately, most network operators today are not security-conscious. Since egress filtering is far from universal, attackers can selectively deploy their agents on the many networks known to have no egress filtering. Furthermore, in large networks, this filtering must be deployed at many points within the network to be most effective, and this imposes additional administrative costs.

2.2.3 End-Point-Based Defense

End-point-based solutions look at DDoS prevention primarily from the perspective of the end-points (servers) that need protection. The approach here is to pursue localized protection on the server itself, or at potential choke points (bottlenecks) in front of servers, without any network-wide or systemic collaboration. End-point-based DDoS solutions can be effective against localized bandwidth consumption attacks, as well as against attacks based on system resource starvation and exceptional condition exploitation of individual systems.

3 Facility Charter and Objectives

During the extensive interviews conducted during this study, the need for a national DDoS defense technology evaluation facility became clear. We were told repeatedly that DDoS attacks are a severe problem and that they are difficult to defend against. The network operators, government agencies, and commercial companies we contacted are desperately seeking a working solution to the DDoS problem. Defense research in this area is progressing slowly since realistically simulating the Internet is nearly impossible with the funding and equipment available to research and engineering teams.

The DDoS attack technologies, in contrast, are developing quickly. There are continuing increases in both frequency and size of DDoS attacks. The sites that are being targeted with DDoS attacks today are larger and more critical than those of even a year ago, yet the deployed defenses are frequently the same and are not succeeding in protecting these sites. If this progression continues at its current rate, it will not be long before critical Internet infrastructure components are successfully attacked. Attack technologies are mature enough that motivated attackers could cause prolonged outages of the country's government, financial, and commercial communication infrastructure. The United States has been fortunate that the attacks to date have merely caused local network outages and have not affected the country as a whole. Protecting the Internet infrastructure within the United States against DDoS attacks is becoming of critical importance, but no solution is on the horizon.

The lack of an available experimentation facility that can properly simulate complex DDoS attacks has greatly hampered defense research efforts to date. Government- and industry-sponsored research projects are forced to test their DDoS defense concepts on simple, small networks that are not representative of real, operational networking environments. In order to help accelerate the development of solutions for the DDoS problem, a common DDoS defense evaluation facility must be designed and constructed. This facility would provide a large-scale network infrastructure on which its clients could evaluate and study DDoS defenses. It would also provide a technical staff, skilled in configuring the infrastructure for effective experimentation. Clients of this facility would submit experimentation proposals, which would be reviewed and scheduled according to established facility priorities. The following charter statement is proposed for the facility:

The charter of the National DDoS Defense Technology Evaluation Facility is to provide a vendor-neutral shared laboratory in which researchers, developers, and operators from government, industry, and academia can evaluate potential DDoS defense technologies under realistic conditions, with the aim of accelerating research, development, and deployment of effective DDoS defenses for the nation's computer networks.

A national facility with this charter would provide many benefits to its clients. In this section, we briefly discuss the types of clients the facility will attract and their respective needs that the facility will fulfill. This is followed by a summary of how important these benefits are to each specific category of clients. In later sections, we will discuss requirements the facility must meet to satisfy this charter, and recommendations for the next steps toward the deployment of a facility that meets these requirements.

3.1 Potential Types of Clients and Their Needs

We envision that many types of organizations will wish to make use of this facility. They are categorized here and defined in terms of their needs.

3.1.1 Government Agencies

This facility would provide the Department of Defense (DoD) and other government agencies a means of evaluating emerging DDoS defense products prior to making purchasing decisions. The DoD could use this facility to run controlled testing of defensive technologies. For example, the facility would provide a controlled environment in which products or technologies from two or more vendors could be subjected to the same test scenarios to compare their relative effectiveness.

3.1.2 Research Labs

Some university- and government-sponsored research labs need commercial grade networking equipment and skilled network engineers to design and configure test scenarios to validate their research findings. Other researchers get equipment grants from network vendors, but lack real world experience in designing and configuring networks. This facility will have experienced network designers and technicians readily available to aid researchers. DDoS research requires a realistic mix of traffic traversing the test facility to evaluate the effectiveness of experimental DDoS defense technologies.

3.1.3 Network Device Vendors

Many networking device manufacturers have large test facilities. Their test facilities are used for network device integration and testing and not for solving the DDoS problem. Some vendors have expressed interest in a facility where their equipment can be tested in a heterogeneous environment containing equipment from other manufacturers, and where a realistic mix of network traffic and state-of-the-art DDoS attack traffic is available. Large networking companies could use this facility to demonstrate their products to potential customers, or arrange third-party evaluation and testing.

3.1.4 DDoS Defense Product Vendors

Many DDoS defense product vendors need larger scale facilities that can emulate large ISPs and content provider site. They need a realistic mix of traffic traversing the test facility to evaluate the effectiveness of DDoS defense technologies. They also need a large pool of attacking hosts. They need a facility where they could prove their technology to potential buyers, such as government agencies and corporations.

3.1.5 Internet Service Providers (ISPs)

ISPs will need to test DDoS defense technologies to see if they will work in their high-speed environments. They need a realistic mix of traffic traversing the test facility to evaluate the effectiveness of DDoS defense technologies. They also need to be able to test technologies that work across peering points. ISP's need to verify that DDoS defense products and solutions will not add additional points of failure to their networks. They will need to realistically emulate their own networking environments.

3.1.6 Other Commercial Entities

Companies with substantial network infrastructures, such as large corporations and content providers, need to verify DDoS defense technologies in environments that emulate their own. The facility must be able to simulate a typical content-serving farm with firewalls, load balancers, caches, and servers, and provide a realistic mix of network traffic.

3.2 The Benefits of a DDoS Test Facility

3.2.1 Realistic Generation and Understanding of Internet and Attack Traffic

One of the primary motivations leading to this study was the realization that most DDoS test facilities that exist today are unable to simulate realistic traffic conditions. These facilities are typically small laboratories or test environments with only a few hosts, limited network topologies and bandwidth, and a very small selection of networking equipment. The DDoS Test Facility could offer a superior test environment with the capability of generating realistic Internet traffic and attack conditions. For example, it should be possible to generate DDoS attack traffic originating from hundreds of hosts, coming to a victim from multiple ISPs, and capable of saturating gigabit speed connections. Thus, an essential benefit will be the ability to generate and simulate real Internet DDoS conditions. This facility will enable a realistic understanding of attack characteristics as well as potential DDoS solutions.

3.2.2 Testing and Experimentation with Complex, Realistic, and Scalable Network Topologies

The ability to generate real Internet traffic and attack conditions must go hand-in-hand with the capability to emulate a variety of complex network topologies that are representative of the topologies of client organizations. The value propositions derived from these capabilities are central to the viability and acceptance of a National DDoS Test Facility. Support for complex network topologies will allow for a realistic understanding of the spread and impact of DDoS attacks as well as the testing of potential DDoS solutions. The center must support topologies that utilize multiple peering points and ISPs as well as a variety of diverse equipment (routers, firewalls etc.) and network links with varying capacities.

3.2.3 Efficient Testing and Pilot Implementations through Rapid Reconfiguration and Adaptability

As mentioned earlier, there are a variety of clients that would use the features of the facility; therefore, an efficient time-sharing usage scheme must be put in place. However, such efficiencies can be realized for a multi-use environment only if the facility can be rapidly reconfigured with minimum downtime. Thus a key value proposition that the center will have to offer is rapid reconfiguration. Rapid and efficient reconfiguration will enable the network within the facility to be torn down after the conclusion of an experiment, reset to a neutral base, and to be quickly set up for a second client's experiments. It will also be possible to support multiple experiments concurrently. In general, reconfigurations should enable the center to experiment and support a variety of attack configurations as well as the configurations that mimic client networks. The latter will be a very compelling value proposition for the vendor of a DDoS product as it allows the vendor to demonstrate the applicability of the product in an environment that closely mirrors the potential customer's network. This may reduce the sales cycle.

3.2.4 Support for Large-Scale and Diverse Experimentation and Simulation

When compared to a small research lab or test facility, the DDoS Test Facility has to offer the capability to conduct large-scale and diverse experimentation and simulation. These capabilities include the:

- Ability to generate and emulate a variety of attack conditions;
- Ability to study attack impact on network performance and security;
- Testing of DDoS solutions to mitigate a variety of attack conditions;
- Testing of DDoS solutions in a multi-vendor environment; and
- Capacity planning, design, and testing of networks to withstand a certain level of attacks.

3.2.5 Availability of Consulting Services Through Highly Trained Engineers and Researchers

The facility will be staffed by or have ready access to affiliated highly trained researchers and engineers specialized in the DDoS defense field; this could be of tremendous value to many clients of the facility. The availability of such expertise will reduce the time taken to analyze, trace, and mitigate attacks. These experts could provide consulting and support services to various clients.

3.2.6 Encourages Collaboration Between Researchers and Engineers from Multiple Organizations

The Center will play a key role in helping researchers and engineers from multiple organizations collaborate on DDoS detection, prevention, research, and mitigation solutions. Such collaboration avoids duplication of effort by the security community in mitigating and researching DDoS attacks, and thus improves the efficiency and reduces the overall cost of these attacks.

3.2.7 Allows Interconnection Across Multiple Test Networks

Today, there exist a variety of test networks, scattered across a number of research labs and vendor facilities. From the business, organizational, and networking standpoint, no structure exists today to interconnect these test networks. However, such interconnections can help in harnessing the individual capabilities to collectively provide more powerful capabilities to analyze, test, and mitigate DDoS attacks. In particular, the amount of attacking power (number of attacking hosts, total available bandwidth, etc.) can be aggregated to conduct tests on a large scale. The DDoS Test Facility could play a pivotal role in interconnecting various test networks.

3.2.8 Encourages Communication and Cooperation among ISPs

All the big telecom carriers and ISPs today face DDoS attacks on their networks. DDoS attacks are increasing in sophistication and coordination, and often make use of the bandwidth available from multiple ISPs. Even though many ISPs are direct competitors, the ISPs may achieve greater efficiencies in DDoS recognition and mitigation if they share information on attacks. The DDoS Test Facility could play a key role in bringing about such sharing and cooperation.

3.2.9 Availability of Diverse Network Equipment and Software

The ability to support a diverse set of network equipment and software will be of great value to many clients. Many security vulnerabilities arise due to the combinations of hardware and software used and the subtle interactions of various constituent properties with the underlying network topologies. Availability of diverse network equipment and software is thus critical to understanding how attacks exploit various vulnerabilities, as well as testing the coverage provided by a particular DDoS defense solution. Most research labs can only provide a limited test environment in terms of the diversity of hardware and software.

3.3 Value of Each Benefit to Each Client Type

The table below gives a rough assessment as to how significant the above value propositions are to various kinds of clients. The study team made these qualitative assessments based on the aggregated comments and opinions expressed by interview subjects.

Value proposition	Government agencies	Research labs	Network vendors	DDoS defense vendors	ISPs	Commercial entities
Realistic generation of Internet and attack traffic	●	●	●	●	●	●
Testing and experimentation with complex topologies	●	●	○	●	○	●
Efficient testing and pilots through rapid reconfiguration	●	●	●	●	●	●
Large-scale and diverse experimentation and simulation	●	●	●	●	○	●
Availability of highly trained security engineers & researchers	●	●	●	●	●	●
Encourages research collaboration among organizations	●	●	●	●	●	○
Allows interconnection across multiple test networks	●	●	●	●	●	●
Encourages cooperation and interconnection between ISPs	●	○	●	○	●	●
Availability of diverse hardware and software platforms	●	●	○	●	○	●
Legend: ○ Low Value ● Medium Value ● High Value						

Table 1. Estimated value to each category of clients

4 Facility Requirements

Based on the information gathered through the interviews, tours, and meetings described in Section 1.4, we generated a set of requirements the facility should meet to satisfy the proposed charter presented in Section 3. These requirements and their justifications are presented below.

4.1 Public Availability

The facility must be accessible to a diverse set of organizations if it is to accelerate the development of DDoS solutions. Currently, DDoS technologies are typically tested in only small-scale test networks. Most research organizations, vendors, and small ISPs are simply unable to conduct truly valid tests within their own environments because of the prohibitive cost of deploying a test network that is representative of Internet conditions. DDoS technology vendors and research organizations frequently use testing facilities that number on the order of 10 to 20 network nodes. This size is simply not representative of large operational networks or the portion of the Internet that would be involved in a DDoS attack. Methods for generating background and attack traffic are typically primitive and do not adequately simulate a real networking environment. Making the facility publicly available will enable a broad range of researchers and vendors to make meaningful contributions toward solving the DDoS problem. Nearly all of the study participants we spoke with indicated a strong interest in making use of the facility, as their own test environments were inadequate.

To be permitted to use the facility, however, clients must be able to demonstrate that their proposed use of the facility will advance the state-of-the-practice, is scientifically valid, and is consistent with the facility's charter. It should not be open to simply anyone who wishes to make use of a complex network. The legitimacy of potential clients and the merits of their proposed experiments must be assessed before they are scheduled for access to the facility according to established priorities. This process is outlined in Section 4.10.

4.2 Vendor Neutrality

In order for the facility to be well respected with the DDoS technology experts, it must be controlled and operated in a vendor-neutral fashion. If any given company competing in the DDoS marketplace ran it, it would be perceived as a facility intended to benefit that company. Clients must be able to evaluate technologies without interference from unwanted marketing proposals, and without fear that the evaluation infrastructure is biased to favor certain products. The facility must be run by a vendor-neutral organization. Ideally, it should be administered directly by the government or by an organization that is contracted by the government solely for the purpose of running the facility.

4.3 Network Topology

The topology of a newly constructed facility must satisfy certain criteria to meet the needs of its prospective clients. The facility should be able to emulate some of the current Internet topology. It should be a scaled down but functionally accurate representation of the Internet. It minimally should be able to emulate multiple Tier-1 high-speed ISPs peering² with each other, which in total will form an Internet-like backbone. It should also be able to emulate various types of smaller attached networks at the edges of this core. The facility must be able to represent various security exposures that exist on the Internet. The facility should be able to emulate a large distributed pool of flooding hosts. The networking and hosting hardware must be representative of what is currently deployed on the Internet. The facility should have access to realistic traffic flows.

² Peering is the way in which ISPs share routing information (i.e., connectivity information).

4.3.1 Backbone Network

To properly represent the Internet, the design of the facility’s backbone network should be hierarchical. It must be able to emulate the functionality of the backbone networks of multiple ISPs, peering at various points. It should be able to represent at least four different ISP’s networks that exchange routes or peer with each other. Like the Internet, this facility should be able to run the Border Gateway Protocol (BGP), which is a protocol that facilitates the exchange of connectivity information between ISPs (see Figure 1). This facility may also need to be able to route live traffic across a wide area to properly represent Internet backbone routers. These routes would not be subject to any service level agreements³, but instead may be alternate routes to destinations that could be disrupted by the facility’s testing.

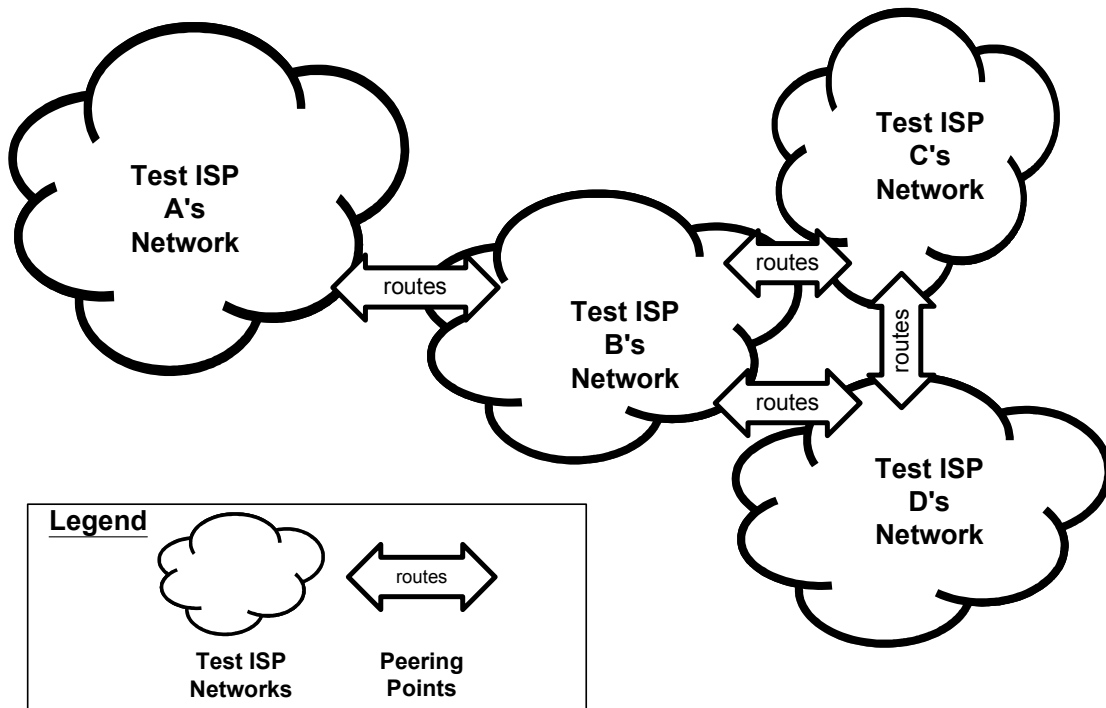


Figure 1. Network peering

The facility’s backbone network must have routers that are representative of those used by the Tier-1 service providers. During our interviews, it was determined that Cisco Systems and Juniper Networks manufacture the majority of the core backbone routers that are currently used by Tier-1 service providers.

4.3.2 Attached Networks and Sites

In a hierarchical network, many smaller, specialized networks will attach to one or more of the Internet Service Providers that make up the facility’s backbone network. Sites that connect to the Internet through multiple ISPs create interesting traceback, rate limiting, and filtering scenarios, which will need to be tested within the facility. Networks must be able to connect to the backbone at various speeds to emulate choke points or weakest links. The facility should be able to represent

³ A service level agreement is a legal or implied contract for the availability and/or quality of service. In this case it is for network connectivity.

several different types of attached networks, including broadband networks, corporate networks, web hosting networks, and university networks. See Figure 2.




Figure 2. Attached networks and sites

Broadband Access Networks

Broadband providers are becoming more prevalent on the Internet and need to be represented in the facility. The two most prevalent types of broadband access networks are Digital Subscriber Lines (DSL), which are offered by telephone companies, and cable networks, which are offered by cable television providers. Both services provide high speed Internet access to a large number of households across the country. Broadband Access Networks are a common source of DDoS attacks due to their high access bandwidth and usually low level of host security. The typical user attaches a personal computer that uses an operating system (typically a version of Microsoft Windows) that is neither kept up to date with security patches nor protected by a firewall that would help prevent attacks. Additionally, most of these computers are continuously attached to the Internet, which makes them always available for use as DDoS attack sources.

Government Agencies

Many government networks are connected to the Internet. Unfortunately, the security of some of these networks is inadequate because of poor funding and older equipment. Government hosts and networks are sometimes poorly configured, and may become unwitting attackers in a DDoS

attack. These networks may also propagate broadcast requests that make them amplifier networks for DDoS attacks. These sites and networks usually attach at line speeds of less than 1 Gbps, which make them attractive as potential flooding hosts as well as potential DDoS targets. Characteristically, these networks run at speeds of 1.5 Mbps to 43.2 Mbps and are usually a hub and spoke design. Most use firewalls and/or network address translation (NAT) to protect their internal sites from the Internet, but some only use filtering routers. Government agencies are comprised of hundreds of personal computers and workstations, which are attractive targets for hackers to compromise and make use of as potential DDoS flooding agents. Government agencies often contain servers that become the targets of DDoS attacks.

Web Hosting Sites

Web hosting sites are a source of normal Internet traffic. They are often connected to the Internet through multiple ISPs. Web servers are common targets of DDoS attacks. Connection line speeds can vary from 10 Mbps to 9 Gbps, depending on the site. The slower the connection speed to their ISP, the greater the chance of a DDoS attack succeeding. Web servers can also be used as attacking hosts if the servers themselves are compromised. Characteristically, hosting sites use filtering routers, firewalls, load balancers, and caches to help secure their sites from denial of service attacks and other nefarious activities. To accurately model a particular site, the facility may need to have specific network-support hardware and software such as network load balancers and various firewall packages.

Corporate Networks

Corporations often connect their internal networks to the Internet. The lack of security on a corporate network is often the result of poor budgeting for the IT division. Corporate hosts and networks are often poorly configured, and often become unwitting attackers in a DDoS attack. Many of these networks also propagate broadcast requests that make them amplifier networks for DDoS attacks. These networks usually attach at line speeds of less than 1 Gbps, which make them attractive as potential flooding hosts as well as potential DDoS targets. Characteristically these networks run at speeds of 1.5 Mbps to 43.2 Mbps, and they usually have a hub and spoke design. Most use firewalls and network address translation (NAT) to protect their internal networks from the Internet, but some only use filtering routers. Corporate networks usually include hundreds to thousands of personal computers running Microsoft operating systems, which are attractive targets for hackers to compromise and make use of as potential DDoS flooding agents.⁴

University Networks

Most universities connect their campus networks to the Internet. University networks are common sources of DDoS attacks. Each department within a university's network usually does its own system administration and security. This results in very uneven security across the university's network. Firewalls are not normally used, and most computers are connected to the Internet with fixed IP addresses. This makes them very attractive to hackers to use as DDoS agents since they are easily locatable (fixed IP addresses) and usually available (powered on and attached to the network) to be used as attackers. Many university computers may run a poorly configured version of UNIX, and support multiple user accounts that may have weak passwords. These networks usually attach at line speeds of up to 1 Gbps, which make them attractive as potential flooding hosts as well as potential DDoS targets.

⁴ The Nimda virus outbreak in the fall of 2001 had the effect of a DDoS attack against many corporate network infrastructures and the Internet's infrastructure due to the sheer number of compromised Microsoft Windows machines that flooded the Internet with traffic trying to seek out and infect other machines. <http://www.cert.org/advisories/CA-2001-26.html>

4.4 Non-Production Network

The facility must not be a “production” network that others rely on to carry non-experimental traffic. The nature of the testing and the amount of continuous reconfiguration of this facility will make this network facility very unstable. There should be no implied uptime guarantees. The only agreement about network availability is that the facility should be given to each group of experimenters in a default-working configuration at the start of their scheduled testing period. After the start of a DDoS technology experiment, each group of experimenters should have complete control over their allotted part of the facility, and they may, by design or accident, break connectivity within their allocated time slot and portion of the network.

4.4.1 “Breakable” Network

One of the expected results of a DDoS experiment is that some pieces of the network may be disrupted or disabled under the attack load. Network outages must be expected. This is the major reason that this facility cannot be created using existing high-speed networks. Each network has service level agreements with their users for availability. Most existing high-speed networks also have acceptable use policies that would not allow the use of malicious or destructive code.

4.4.2 Support for Repeatable Experiments

For an experiment to be valid, the conditions under which it was performed and its results should be reproducible. This facility must support highly controllable experiments to accomplish this goal.

Occasionally there may exist the need to temporarily join other networks and labs to the facility. These connections must be monitored closely and removed after the test has been completed.

The facility must also have policies on acceptable usage, software installation, and configuration.

4.5 Rapid Reconfiguration

When the test network is to be used by a new client, it is very likely that the physical and logical configuration of the network will need to be reset to a known neutral state and then changed to accommodate the needs of the new tests to be performed. It is also likely that each client will need to perform a series of experiments and each of these experiments will be based on a different network topology. Unfortunately, this setup work could easily consume a lot of time if the changes to be made are extensive. In order for the facility to be viable, a quick reconfiguration methodology must be deployed along with the facility itself. Without a fast turnover methodology in place, the number of evaluations and experiments that could be conducted within the facility over the course of the year would be greatly diminished.

The following types of configuration information will need to be identified as one experiment or evaluation team replaces another:

- Physical topology reconfiguration;
- Physical device model requirements;
- Node operating system type, version, and contained software;
- Network equipment software revisions (IOS versions);
- Experimentation data removal and restoring devices to a known operational state; and
- Management infrastructure type and configuration.

These types of configuration information are discussed in the sub-sections below.

4.5.1 Physical Topology Reconfiguration

A difficult problem is how to reconfigure the network's infrastructure if an upcoming test or evaluation requires that the current network topology be radically altered. The facility must deploy an infrastructure that solves this problem.

One method of ensuring that the facility is easily maintainable is to disallow changes in network topology. Mandating that all tests use one pre-defined network topology would certainly help reduce the amount of time needed for a client changeover, but would greatly limit the types of tests that could be conducted on the network. Different DDoS technologies instrument networks in different ways, and it is impossible for a single network topology to satisfy the testing requirements of every DDoS technology. Many network engineers will prefer to test technologies on networks that closely resemble their own. For example, some engineers would prefer to use Cisco routers to Juniper routers or would prefer to use Microsoft Windows operating systems instead of the Sun Solaris operating systems. Because of these needs, it is unlikely that a fixed topology will provide a suitable test environment for anyone. Therefore, the facility must not impose a fixed network topology upon the facility's clients.

In order to support rapid network reconfiguration, the facility must provide personnel who are experts in the usage and configuration of the equipment on hand. Rather than having the experimenters work with unfamiliar equipment, it would be more efficient to hire a team of skilled engineers who would be very familiar with the facility and could quickly configure the facility in advance of the experiment.

Of course, manual reconfiguration of such a large network carries several costs and risks, even when performed by experienced technicians. Reconfiguring many devices takes time, and could lead to significant delays between experiments that impede efficient utilization of the facility. Human error is unavoidable, and faulty configuration may lead to time-consuming diagnosis and invalidation of experimental results. Miscommunication between the clients and the facility's engineers could also lead to complications, problems, and delays.

Therefore, to the greatest extent practical, the network should support automated reconfiguration mechanisms. For example, switches with VLAN capabilities can be used to overlay a variety of virtual network topologies on a fixed physical network. Over-provisioning of network links can also reduce the need for physical reconfiguration; a fully connected set of routers (with links between each pair) can be made to emulate any network topology by disabling selected interfaces in software. However, care must be taken so that support for rapid reconfiguration does not substantially affect important characteristics of the network, as that could invalidate experimental results. Inserting extra bridges, routers, and other networking equipment into the topology can make reconfiguration easier, but may also introduce delays and unanticipated network behavior that would not exist in the simpler network being simulated. Specialized hardware such as physical layer switches can help address these problems, but manual and, in some cases, physical reconfiguration by the on-site technical staff will inevitably be required by some experiments.

To facilitate reconfiguration, the management network should be an out-of-band network, separate from the facility's experimentation network. An out-of-band management network will facilitate faster reconfiguration and setup between experiments. Most Tier-1 ISPs currently use an out-of-band management network to ensure accurate performance statistics collection, secure access to network hardware, and access to devices for problem determination during times of high network utilization (e.g., during a DDoS attack).

4.5.2 End Host Software Reconfiguration

End host devices, like web servers and desktop workstations, need to be re-configurable as well. Every device within the facility should be able to be reset to a known and trusted state, with no previous experimentation or configuration data left remaining on it. If a given product or evaluation needed a particular operating system type or version in order to function properly, they must be readily available. A system should be put in place that can reconfigure a large number of network clients as quickly as possible. It is imperative that the facility use techniques that automate software installation for any given network end-node.

4.5.3 Network Device Reconfiguration

Networking devices, such as routers, switches, encryption devices, and hubs, make up the Internet's infrastructure. These devices require updates to their operating systems, but they must be dealt with in a different manner than the end-host systems described above. The installation solutions will likely be different, as the devices run simpler operating systems, and the process of upgrading them is more difficult and requires a different installation tool set to be on-hand. Infrastructure must be developed that allows software on network devices to be quickly upgraded or downgraded. Within each network vendor's product line, the various operating system versions and patch levels should be available for use in the facility.

4.6 Realistic Network Traffic

The most difficult part of designing and conducting any DDoS evaluation is ensuring that the evaluation environment closely models the real world. To succeed in closely modeling the real Internet, researchers must carefully construct a simulation environment with appropriate network traffic traversing the test network. This traffic must include not only DDoS attack traffic for study, but typical day-to-day (legitimate use) traffic as well. The type of traffic needed to emulate a realistic environment will be specific to each client of the facility. For example, clients that need to emulate a web hosting facility environment will need different types of network traffic from those clients who need to emulate a university network. Traffic to and from web hosting facilities is almost entirely composed of World Wide Web traffic. Traffic to and from university environments, on the other hand, is frequently composed of a wide variety of traffic types that include not only web traffic, but also file sharing, music streaming, login connections, email transmission, and others. Ideally, the facility should be capable of representing network traffic needed by any client.

4.6.1 Statistical Traffic Generation

Many companies sell tools today that are capable of simulating traffic in a network environment by first statistically analyzing live data and building traffic profiles. They attempt to generate statistically similar, but artificial, traffic that models the same profiles. These traffic generation methods should reproduce not just the long-term average traffic characteristics, but also the highly bursty and self-similar nature of Internet traffic. Companies such as SmartBits currently offer products that do statistical traffic generation, though these products are often too expensive for smaller research organizations. Furthermore, the traffic these packages generate is based on random data and does not truly represent typical interactive session traffic. Additionally, these packages usually do not generate wide area network traffic; instead they generate traffic across a single network segment. Therefore, some researchers do not fully accept data from such a simulation environment. Still, most vendors and researchers do make use of these tools for preliminary analysis and testing since they are easy to find and simple to set up and use, and since they are often the only tools available. This facility must have these technologies readily available for its clients if it is to provide a traffic simulation environment.

4.6.2 Traffic Generation Using Network Applications and Attack Tools

The easiest and cheapest method of producing DDoS attack traffic is to simply run real DDoS attack tools on a variety of end-hosts and have them all attack a particular target within the test network. Researchers also need to produce background traffic by using networking software that can be automated (i.e., scripted) to simulate real users generating real network traffic. This approach works well in practice and is widely used by research organizations with small budgets. Unlike the statistical traffic generators discussed above, real DDoS tools and scripted network clients produce traffic with normal packet contents rather than random data. However, the generated traffic does not model the seeming random variations in traffic patterns produced by real human users sitting at networking consoles. It is still a useful technique, so the facility must have a library of network applications and DDoS attack tools on-hand. These tools may be brought in from the wild, but should be inspected to ensure that there are no backdoors or other unintended features present.

With today's fast processor systems, a small number of hosts can easily generate a lot of traffic using this technique. Many research groups that do not need to simulate complex networks rely on this type of traffic generation technique. For research groups that are studying how DDoS attacks affect more complex networks, this technique becomes quickly expensive if thousands of hosts and/or hundreds of attacking networks need to be simulated. Ideally, this facility should have a large number of end-host systems to enable simulations of complex attack topologies.

4.6.3 Simulating Network Delays and Congestion

When network traffic traverses a large disparate network, it is subjected to various delays in transit. These delays are a result of limited processing power of devices it must pass through as well as issues with crossing large physical distances. In a contained network facility existing solely at a single site, distance-related delays are mostly nonexistent. Some research environments, like Emulab [8], have solved this problem by carefully constructing delay devices that merely hold each packet for a period of time before letting it pass on. These delay devices should work well for small bandwidth connections, but simulating a delay at high speeds like an OC-192 link (near 4.8 Gbps) would be more difficult. The facility must support network-delay simulation for at least some portion of the network, even if simulating them at high speeds is not technically possible at first.

Network congestion also affects performance in a network and must be simulated in order to closely model a real world environment. It might be necessary, at times, to simply shrink the effective size of a given network link by pushing a fixed amount of traffic through it to decrease the link's available bandwidth. Fortunately, Emulab has addressed this problem as well and has developed techniques for applying a rate limit to a given connection by carefully adding in traffic generators and traffic sinks, which are designed to send a fixed stream of data across one or more network nodes. This type of technology must be present one way or another in any facility intending to simulate real world networks.

4.6.4 Capture and Replay of Real Traffic

Researchers who are studying network traffic need a method of testing their technologies on live network traffic without worrying about the possible ramifications of their experiments going awry and negatively affecting a production network. The best solution would be to capture network traffic in a real environment and then replay the traffic in the test facility environment, supplemented with traffic from developing DDoS technologies. Unfortunately, the technology available for capture and replay of traffic is currently limited in nature and is not suitable for use in a large-scale, complex networking environment.

The capture and replay of real traffic in a complex environment is a basic necessity that no existing technology satisfactorily provides. There is general agreement among network researchers that it is a hard, but solvable, problem. A research project should be undertaken to solve this problem; the results would greatly facilitate much of the research in not only the DDoS arena, but also in general network traffic research.

4.6.5 Using Live Network Traffic

Once DDoS defense technologies have been validated and it is believed they can be used safely on a live network, the next step is to deploy them in a test environment that can make use of a live network traffic stream. This stream should not be a true production network link, however, it should be subject to a degree of control and configurability that is not available in a fixed production network. The facility should be capable of mirroring external network traffic into the experimental network. For example, this could enable a client who operates a large web site to mirror some or all of its network traffic to the facility where that client could test DDoS defense mechanisms on the traffic. It should also be possible to operate portions of the facility as a test network that real traffic can be temporarily transmitted over.

Traffic mirroring has been implemented in the commercial world only in simple applications. Traffic traversing a device can often be routed to its real location as well as duplicated and sent to a monitoring port. This works for simple, single-point traffic-monitoring applications, but is unlikely to suffice for active, distributed DDoS defense mechanisms. Technology that could mirror traffic into a test network from different points and then absorb it after it finishes traversing the test network would be ideal for use in the facility, but does not exist today. It would have to support multiple input nodes and send traffic in both directions properly past the possibly numerous network nodes instrumented with DDoS defense technology. This is also believed to be a hard but technically solvable problem. If possible, a research project aimed at producing complex mirroring technologies should be undertaken. At some point, all DDoS technologies must be evaluated using live traffic to ensure that they are in fact viable technologies, as no simulation can exactly reproduce the non-deterministic behavior of real world networks. Because of this, it should be considered a requirement that the DDoS experimentation facility include the ability to make use of a live network stream.

Some legal issues arise when making use of live traffic. Specifically, personal and institutional privacy rights must be observed when dealing with traffic on a network. This is discussed in greater detail in Section 4.13

4.6.6 Standard Test Suites

To minimize setup time required by clients who merely need to quickly evaluate an existing technology, it should be easy and straightforward for them to make use of the facility to run a pre-defined sequence of DDoS tests. A default test suite is therefore needed, which can be used by clients who do not need specialized test scenarios.

When possible, these default simulations should address the needs of multiple client types. Some suite components could be designed to test DDoS defenses monitoring a high speed backbone, while others could be designed to simulate traffic over a very complex network topology with many DDoS defense technology instrumented nodes. These test suites should use all of the network traffic generation tools mentioned in this section. They should contain tests that use generated traffic as well as tests that make use of live or replayed traffic.

4.7 Data Archiving

In order for information to be imported and exported from the facility, a standardized method of data archiving and storage must be developed. It is likely that clients will wish to bring in DDoS technology software, captured network traffic, and network topology diagrams. It is also likely that they will need to extract any results from the test facility in a format they can take back with them. They will also need to obtain previous results to compare them against current results. Their experiment topology design may have been modified while they made use of the facility, and these changes must be extracted and stored for future reference. Additionally, it is likely that some clients will wish to leave data on the servers until the next time they return. Methods of storing, indexing, archiving, and transferring this data must be available to the clients. Those methods must consist of large capacity physical storage mechanisms and should include methods for transferring data to and from the facility over the Internet.

4.8 Network Management, Monitoring, and Analysis

The facility's equipment and software must run Simple Network Management Protocol (SNMP) services for device management, configuration, and problem determination. The facility must also be able to collect performance data from SNMP, cflowd, or Netflow. These are common network management mechanisms used on the Internet. Network management utilities that can analyze and generate reports from these data sources must be available for use. These management utilities should include network topology mapping tools. These tools can then be used to compare a deployed network against a desired network to ensure that the facility has been configured properly for the experiment.

The facility will require sophisticated network monitoring and traffic analysis tools for two purposes. First, these tools will support diagnosis and problem determination when the network does not behave as expected. Second, experimenters will use them to measure network conditions over the course of an experiment. For example, to determine the effects of a DDoS attack on legitimate network traffic when a specific DDoS defense technique is in use. For both purposes, the ability to aggregate and analyze results of a large number of monitors throughout the network will be crucial.

4.9 Technical Staffing

Other portions of this document have described the need for technical staff to be on hand to assist clients in making the most efficient use of their allocated time. Here we outline the type of staffing needed and the skill-sets they must possess. Without such staff on site, it would be difficult for clients to make efficient use of the facility.

4.9.1 Network Engineers

Many clients will not be able to use the site adequately if they are not familiar with the physical characteristics of certain types of networking components contained in the facility. It is unlikely that small research organizations will have proficient knowledge of high-speed backbone devices, such as large Cisco and Juniper routers with high-speed networking cards. To avoid the risks and liability associated with untrained personnel operating expensive equipment, an on-site technical staff team must be available to assist clients in setting up and configuring the network to their needs. They should be experts at using and configuring the wide variety of equipment found at the facility. They should also be able to provide assistance in diagnosing networking problems and utilizing the various features of the facility.

Nearly all of the study participants mentioned that having a team of experts on hand would greatly facilitate rapid experimentation. Access to the physical resources is only half of what

clients need. Personnel capable of running it as well as answering questions and helping the clients must be present if the clients are to optimize their time at the facility. Many network centers, such as Equinix, have on-site technical consulting staff. The availability of these consultants is well appreciated.

4.9.2 DDoS Experts

Many clients would benefit from having access to DDoS research experts who could help them conduct their experiments or evaluations. Clients who are not necessarily experts in the DDoS domain (e.g., some government organizations, ISPs, and other end users), but are using the facility to evaluate DDoS defense technologies, could make use of the DDoS expertise of the facility staff. Clients who themselves are experts in the DDoS domain (e.g., DDoS research institutions and DDoS defense product vendors) are likely not to need the expertise of on-site staff. Because the need is not evenly distributed across every client type, on-site DDoS experts are only a requirement if the clientele of the lab is to include groups who may not be experts in the field. If not every client will need to make use of the DDoS expertise found at the facility, a fee-for-consultation recharge scenario might make the most sense if the staff's salaries are to be recouped.

4.10 Usage Administration

To ensure effective and appropriate use of the facility, an oversight unit is needed to ensure long-term continued viability of the facility. Additionally, procedures must be developed for reviewing, prioritizing, and scheduling the proposed experiments. These requirements imply some further staffing requirements.

4.10.1 Steering Committee

To ensure that the facility continues to provide a beneficial evaluation environment, a steering committee will need to be formed to regularly evaluate the facility's current charter and perceived value, and weigh it against the needs of the client community. As DDoS technologies continue to evolve, it is likely that the facility will also need to evolve to accommodate the changing technological climate. Therefore, a steering committee should be formed to define the facility's immediate and long-term goals. This committee will need to take into account the viewpoints of the current technology experts, as well as the evaluation results of the previous experiments and should suggest any architectural or administrative changes that might improve upon the facility's usefulness. If the DDoS problem is solved at some future time, it is likely that the facility's charter will need to be reevaluated to determine if it might be useful as an evaluation facility for other national-scale cyber security problems.

4.10.2 Proposal Methodology

The first step any client must take if they wish to conduct experiments within the facility must be to contact the facility administrators to request access. Since it is likely that not all researchers will be granted access to the facility, a standard proposal form should be used to allow the approval process to be streamlined. This proposal form should include instructions for what type of information must be submitted before an approval can be granted. Many similar sites operate this way (e.g., Emulab and ATDnet, a government test network) and have found that having a standardized proposal mechanism allows them to more efficiently review proposals without having to sort through unfamiliar document structures looking for particular details about the experiment.

4.10.3 Prioritization

Proposals must be assigned priorities based on a variety of factors, including alignment with the facility's charter, relevance to near-term needs of government agencies, potential for impact on the state-of-the-art in DDoS defense, and scientific merit. Higher priority may also be assigned to certain clients, such as government or commercial organizations that help provide initial or operational funding for the facility. These priorities must be considered when selecting and scheduling experiments.

4.10.4 Scheduling

Scheduling of new experiments within the facility will require some detailed knowledge about how long it will take to switch from one client to the next. Once the facility is set up, it will likely take a while for the staff to get accustomed to reconfiguring the facility for a newly arriving client. Eventually, the technical and administrative staff should have a good feel for how long it will take to tear down one experiment's network and set up a new topology for an incoming client. Scheduling of a new client's arrival at the facility must take this reconfiguration time into account. Additionally, clients must stick to their defined schedule so they do not impact the experiments to follow by delaying the tear-down and rebuild portion of the changeover.

It is likely that the staff will be able to predict that certain experiments have similar set-ups, and grouping them together will result in a shorter changeover time. In addition to the client's needs, the scheduling of the various experiments must consider these experiment similarities. A method for estimating client changeover times must be developed. A method for determining how much time should be allotted to a given experiment must also be defined and adhered to.

4.10.5 Administrative Staffing

An administrative staff must be on-hand to handle the tasks outlined in this section.

- Distribute information that promotes the facility.
- Process proposal submissions.
- Schedule time for use of the facility.
- Request and receive physical or remote access to the facility equipment.
- Process and approve equipment that is to be brought into the facility.
- Take care of any financial obligations imposed on the client by the facility.

These kinds of tasks are important to meeting the operating requirements of the facility.

4.11 Physical and Network Access Control

The facility must be physically secure, with all sites having access controls and logging of visits. Access logs should be audited periodically for violations of policy and procedures. Any proposed introduction into the facility of hardware or communications links must go through a change control process and be approved prior to installation.

Access to administrative accounts for all hardware platforms must be tightly controlled, authenticated, and logged. In particular, strong cryptographic authentication must be required of each user remotely accessing the facility. All remote sessions must be encrypted and their occurrences logged. Remote access should require pre-approval and should only be granted to clients involved in experiments that are currently active.

If a facility site has been configured to allow exchange of network traffic with the Internet, then special precautions must be taken to ensure that attack traffic used within the facility cannot

spread to the Internet at large or disrupt other networks and systems. Consequently, all forms of network access into and out of the facility must be strictly controlled through the use of firewalls and routing protocol manipulation so that the type and amount of Internet traffic that can enter and exit the facility can be easily and consistently regulated.

4.12 Financial Requirements

The cost of this facility can be divided into two parts: the one-time cost of the initial design and construction, and the ongoing, maintenance-related costs of running the facility. Both of these cost sets need to be funded if the facility is to be built and remain viable.

4.12.1 Initial Design and Construction Costs

The initial design and construction of the facility is likely to consist of the following budgetary items:

- Initial equipment population;
- Technical engineering staffing to design the facility; and
- Physical construction costs, if any.

4.12.2 Recurring Operating Costs

In order to keep the facility running over the course of its lifetime, certain ongoing operational costs will require yearly funding. These year-to-year funding requirements will consist of at least the following budgetary items:

- Technical staffing salaries;
- Administrative staffing salaries;
- Housing space costs;
- Internet connectivity fees;
- Equipment replacement costs;
- Hardware and software maintenance contracts; and
- Other general administrative overhead costs.

4.13 Legal and Regulatory Issues

It is beyond the scope of this document to address all the legal issues that may affect the facility's operational capabilities. An in-depth study must be undertaken to ensure that the facility operates within the law. At a minimum, this study must address the legal and privacy issues involved with analyzing, collecting, and replaying live Internet traffic.

As described in Section 4.6, in order for the facility to be of the greatest value to its client base, the traffic flowing within it must be as realistic as possible. Section 4.6 outlines the need for evaluation of DDoS defense technologies in the context of live or previously recorded live traffic. The capture and use of any live or recorded traffic may subject to legal restrictions on wiretapping and tracing activities (e.g., 18 U.S.C. 3121 and 18 U.S.C. 2511). The legal issues relating to traffic data collection will involve both federal and state laws, and both areas of law must be adequately studied.

Some clients may be able to bring their own operational traffic into the facility legally by routing, mirroring or prior capture. It may be necessary for these clients to indemnify the facility and its

staff against any liability for the use or analysis of this data. Other clients may need to rely on traffic provided by the facility. It is important, therefore, that the facility find ways of obtaining and providing such traffic within the constraints of federal and state laws so that clients can be free to concentrate their efforts on technology development and evaluation.

5 Recommendations

This report has outlined a large number of requirements for the facility. Ideally, all of the requirements should be met; however, this may not be feasible due to economic, technical, legal, or logistical constraints.

The interviews revealed a seemingly equal need for a facility that can make use of live network traffic in a close-to-production manner and for an isolated facility that is composed of a large number of highly controllable networking devices and hosts. Some interviewees expressed a need for a combination of both of these options. A hybrid facility would allow data from the Internet to flow into or through the facility, and would allow some specific traffic to flow out of the facility into the Internet. In general, the research organizations we spoke with had a greater need for a highly controllable set of network devices. DDoS defense vendors and network device manufacturers, however, expressed more interest in modeling their test environments around a network infrastructure that is closely tied with a real production network and/or live network traffic.

The ideal choice would be to build a facility that housed a large collection of highly controllable networking devices and hosts, and that was additionally attached to the Internet for live traffic. Clients could then make use of the networking devices, hosts, and simulation tools available within the facility, or they could attach to the Internet for live traffic feeds. If the facility had both capabilities, clients could feed carefully controlled live traffic into a test network setup to study how network configurations performed under a real load with a more realistic traffic mix. In addition, assuming the proper controls and safeguards were in place, the facility could be used to perform DDoS vulnerability assessments of other sites on the Internet, though doing so would present a number of legal and liability concerns. Clients could also use the Internet connectivity as a way to attach the facility to their own testing environments for distributed, collaborative experiments.

If it is not possible to construct a facility capable of meeting all of these requirements, due to financial, administrative, or other constraints, a decision must be made as to which of the requirements will be met. The costs of meeting any particular requirement must be weighed against the perceived benefit from the functionality of that requirement.

5.1 Facility Design Options

In the sections below, we propose five high-level facility design options of varying functionality, complexity, and cost. All of the options share these common features. First, all options will have the ability to provide one or more core network backbones. A minimum of five core backbone routers is recommended as shown in Figure 3 since most basic routing topologies can be emulated with a combination of at least five routers. Second, each option provides the ability to represent multiple logically or physically distributed sites. Third, each site, whether geographically separated from other sites or collocated at one physical site, should provide a sufficient number of networking devices (e.g., routers, switches, and load balancers) and hosts (e.g., workstations, web servers, and e-mail servers) to represent one or more of the attached network topology components described in Section 4.3.2. Figure 4 depicts a complex example of such a site attached to a backbone router. Forth, all options have a secondary, lower-capacity management network for configuration, control, and monitoring. This management network is separate from the testing data network.




Figure 3. An example of the facility’s backbone

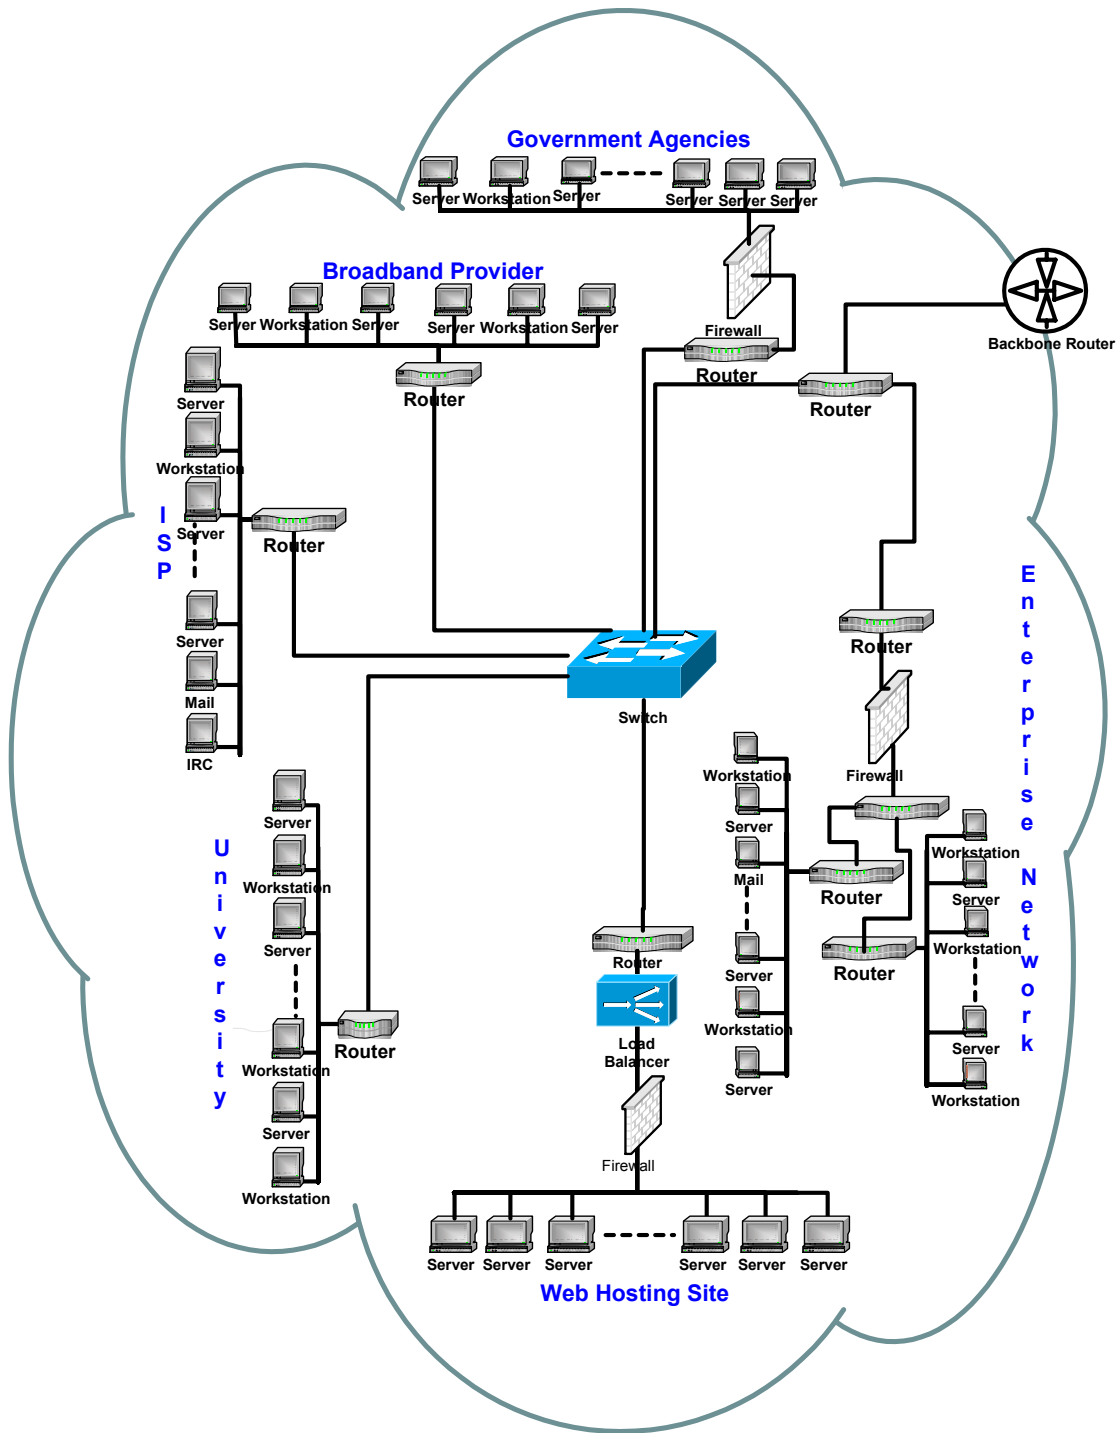


Figure 4. A complex example of an attached site

The following is an overview of the five proposed design options and the relationships between them. These options are described in greater detail in Sections 5.1.1 through 5.1.5.

- Option 1 is a single, self-contained data network isolated from the Internet and located at a single physical facility. It can be configured to emulate multiple logical sites and/or networks. Only the management network is accessible from the Internet; there is no direct Internet connectivity for the primary test traffic.
- Option 2 is similar to Option 1 with one major change: the facility's data network has direct connectivity to the Internet as well. This option will allow test traffic to be mixed

with some traffic from the Internet. Internet traffic in this option could be any of the following types: (1) destined from the Internet to a host within the facility; (2) originated from the network within the facility to destinations on the Internet; or (3) mirrored from the Internet into the facility.

- Option 3 is similar to Option 1, but the facility is distributed across multiple physically remote sites connected by private, point-to-point circuits or private virtual circuits over the Internet.
- Option 4 is similar to Option 2, but the facility is distributed across multiple remote sites connected by private, point-to-point circuits or private virtual circuits over the Internet. Option 4 is functionally a melding of Option 2 and Option 3.
- Option 5 is identical to Option 4 in physical topology, but in Option 5 there are agreements in place with one or more Internet Service Providers or with another backbone network such as Internet2 to route some traffic through the test facility’s backbone network.

The facility could be built in phases starting with any one of the five options and evolving toward others. Such a phased deployment would allow early experience with the facility to guide its future development. For example:

- Option 1 could be built first. Adding Internet traffic access to the test data network would transform it into Option 2; Option 1 could instead grow into Option 3 by adding additional physically remote sites without adding data network traffic access to the Internet.
- Option 2 could grow into Option 4 by adding additional physically remote sites.
- Option 3 could grow into Option 4 by adding Internet traffic access to the test data network.
- Option 4 could grow into Option 5 by adding peering agreements with other ISPs or backbone networks to route traffic through this facility.

These relationships are depicted in Figure 5.




Figure 5. Possible steps in a phased implementation

5.1.1 Option 1: A Single-Site Test Network with Firewalled Access from the Internet

This option would provide a highly controllable, easily changeable network in which DDoS test

traffic is completely isolated from the Internet. The facility would be constrained to a single physical site and have an out-of-band management network. Access from the Internet would be authenticated and allowed only to the management network for configuration changes and experimentation result analysis. The test network could be configured to represent multiple logical sites although the physical components would all be housed in one location. One realization of this option is illustrated in Figure 6. All of the circuits shown are directly connected circuits that are contained within the facility. Directly connected circuits are not provisioned by telephone companies, so there are no monthly fees associated with this type of connection. This should greatly reduce both the cost of deployment and the cost of ongoing maintenance in comparison to options in which the test network spans multiple physical locations. However, the facility would not support, as effectively, the evaluation of products that are designed to operate simultaneously at physically distant sites, which may be important to some large potential clients such as Tier-1 ISPs and router manufacturers.




Figure 6. Option 1: A single site facility with no access to the Internet from the test data network

Experimenters would have a plethora of networking devices and hosts to utilize. The facility should be able to represent most backbone topologies and attached sub-network configurations. The technology developed at the University of Utah’s Emulab could be extended and used to configure and manage these networking devices and hosts. The requirements for rapid reconfiguration (Section 4.5) and effective network management and monitoring tools (Section 4.8) could be more easily met since all equipment would be at a single location. The limited Internet connectivity would make effective network access control (Section 4.11) easier to achieve.

This design option is not as well suited as some others to meet the realistic network traffic requirements of Section 4.6. In particular, the lack of high-capacity Internet connectivity will prevent the facility from collecting live data, so it will have to rely on statistical traffic generation and replaying traffic captured elsewhere to meet these requirements. In addition, this option would not provide real distance-related network delays since the communication links would not traverse large distances. Network delays, however, could be simulated by the insertion of delay-emulating devices. A limitation of this option is that the use of a single physical site will make the facility less accessible to DDoS defense researchers, developers, and operators based in other parts of the country. Providing practical remote access mechanisms would thus be crucial to the facility's success.

A facility with this design could be used as the first building block for any of the other four options. This facility can be housed anywhere, but if there were plans to evolve toward options 2, 4, or 5, then it would likely be cheaper to house the facility at or near a large existing Internet exchange point (e.g., an Equinix [9] or PAIX [10] site). This would allow live Internet connections from Tier-1 providers, for example, to be connected quickly to facility's network with minimal cost. The cost of this option would be the lowest of all the proposed options since there would be fewer recurring circuit charges and only one physical site to equip. The lead-time to build it would be the shortest, because it would not require provisioning any circuits from telecommunications providers or identifying and configuring multiple sites.

Usage Examples:

- The network could be used as a “petri dish” for examining wild DDoS tools.
- The facility could be used as product evaluation environment for single-site products.
- The facility could be used for general DDoS defense research validation.

Advantages:

- The environment would be highly controllable.
- Access to the network from the Internet would be very limited, making it easier to secure.
- Start up costs would be low since it is an isolated single site.
- Recurring charges would be low because no cross-country circuits are required.
- Staffing costs would be lower because there is only one physical employment site.
- Experimentation setup would be easier because all of the equipment would be located at one physical site.

Disadvantages:

- Real Internet traffic could not flow through, into, or out of the facility, and thus would not meet the requirements outlined in section 4.6.5
- The facility would be less accessible to clients not based near the single physical site.
- Any network delays would have to be manufactured.
- It would be difficult to accurately emulate certain kinds of network topologies, in particular those of large ISPs or content providers.

5.1.2 Option 2: A Single-Site Test Network Fully Connected to the Internet

The second option would be similar to Option 1, but it would have the capability to route traffic to real sites on the Internet and to attract real traffic from the Internet to end points within this facility. This design could partially solve the realistic traffic problem (Section 4.6) by allowing some real traffic from the Internet to enter and exit the facility. Other research networks could be attached through virtual private network (VPN) tunnels. The facility could possibly be used to perform DDoS vulnerability assessments of other sites on the Internet, though doing so would present a number of legal and liability concerns (see Section 4.13).

This design would greatly increase the exposure of the facility to potential security risks since at least some portion of the facility would be directly addressable from the Internet. In addition, it presents the risk of DDoS attacks being carried out from the facility itself against sites on the Internet, due to accidental reconfiguration or compromise by malicious parties. To satisfy the access control requirements (Section 4.11), a “gate keeper device” (a specially configured access router and/or a firewall) would be used to control the amount and type of traffic that would be allowed into and out of the facility. Careful monitoring of the Internet connection for unintended uses would be required. Internet traffic should probably be allowed to pass to and from the experimental network only during experiments that require it, and measures should be taken to reduce the risk of hosts within the facility being compromised through Internet-originated attacks. After every experiment, care must be taken that the devices within the facility be cleansed and reconfigured to a known secure state.

Like Option 1, this design would have the advantages and disadvantages that result from housing the facility at a single location. Management and monitoring (Section 4.8) would be eased, and costs would be reduced. Effects of long-distance network connections would be absent except through emulation (Section 4.6.3), and the facility would be somewhat less accessible to a nationwide clientele.

It would likely be cheaper to house the facility at or near a large existing Internet exchange point (e.g., an Equinix or PAIX site) than elsewhere. Such a location would allow live Internet connections from Tier-1 providers to be linked quickly to the facility’s network with minimal cost. Figure 7 is an example representation of this design option.




Figure 7. Option 2: A single site facility with test data access to the Internet

The cost of this option would likely be higher than Option 1 since more security measures will be needed to safeguard the facility and to carefully control the traffic flowing to and from the Internet. In addition, recurring costs would be somewhat higher than Option 1 due to the connectivity charges. Like Option 1, the cost will rise if the size of the facility grows.

Usage Examples:

- The network could attract real Internet traffic by hosting informational servers or “honey pots” designed to attract DDoS attacks.
- The hosts in the facility could be used in a DDoS vulnerability assessment of a client’s external site.

Advantages:

- Real traffic could flow into and out of the facility.
- Traffic could be mirrored into the facility from the Internet for detailed analysis.
- Single location results in reduced management burden and lower cost.

Disadvantages:

- The facility would be more difficult to secure since traffic from the Internet would enter and leave the test network. In particular, if self-replicating DDoS toolkits are used within

the facility, special controls will be needed to ensure that such malicious software cannot escape from the facility.

- The facility would be less accessible to clients not based near the single physical site.

5.1.3 Option 3: A Distributed Multi-Site Test Network with Firewalled Access from the Internet

This option would be highly controllable, and all DDoS test traffic would be isolated from the Internet. It would have an out-of-band management network. Access from the Internet would be authenticated and allowed only to the management network for configuration changes and monitoring. It will have multiple physical sites that are connected via private network circuits. All of the circuits in Figure 8 are point-to-point leased lines that are provisioned by the telecommunications companies, so there will be recurring monthly fees associated with these connections. This option is the same as Option 1 but with physically disparate sites. These remote locations would add complexity to any network topology changes and device reconfiguration (Sections 4.5), as well as general network management and monitoring (Section 4.8). The cost of deployment and the cost of ongoing maintenance would be higher than Options 1 or 2, due to the physical distribution. The facility would have the ability to effectively evaluate products operating simultaneously at sites physically distant from each other; however, this does not offer any live Internet traffic, which makes providing a realistic traffic mix difficult. It would have a plethora of networking devices and hosts located at each physical site, and it should be able to represent most network topologies and attached site configurations. This option could be used to connect the various private research labs, government labs, and university labs together to test and share their DDoS defense technologies. The test network would support experiments requiring real distance-related network delays.

This facility design could be used as an intermediate step in building Options 4 or 5. The facility could be housed anywhere, but if there were plans to grow it into Options 4 or 5 then it would likely be cheaper in the long run to house the facility at or near a large existing Internet exchange point. This would allow live Internet connections from Tier-1 providers to be linked quickly to the facility's network with minimal cost.

This option would be more expensive than Options 1 and 2 since there would be many recurring circuit charges, and the lead time to build it might be long due to the number of sites to be established and the time required to provision circuits from telecommunications companies. However, the physical distribution would make it easier for clients to visit one site or another for experimentation, consultation, and collaboration.

If implemented, this option would enable more of the requirements outlined in this report to be met, and the facility could provide both a large collection of configurable networking equipment and easy accessibility to the facility from the Internet at multiple highly secure access points. This solution would provide easy, but controllable, access to the facility for all prospective clients of the facility. Figure 8 depicts an example facility configuration using this design option.

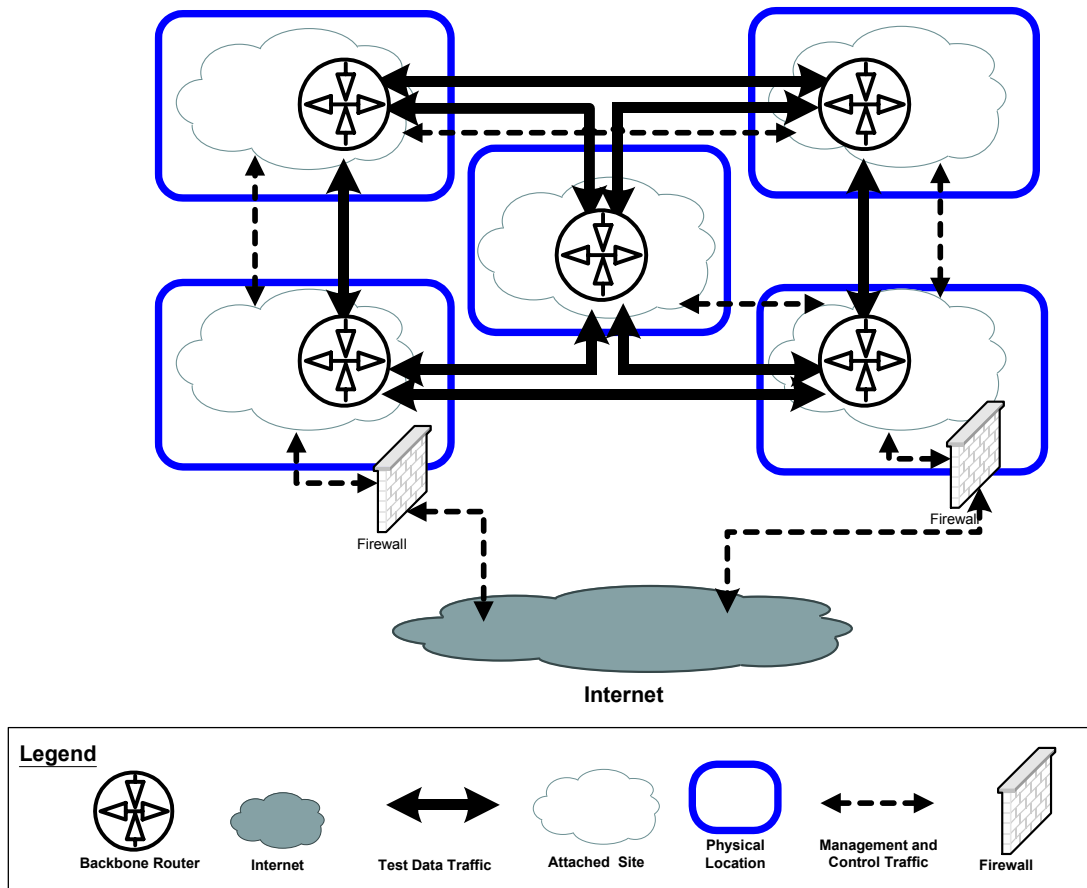


Figure 8. Option 3: A multi-site facility with no access to the Internet from the test data network

Usage Examples:

- The network could be used to connect the various government, private, and university research labs together.
- The network could be used to accurately emulate an existing Internet backbone.

Advantages:

- The environment would be highly controllable.
- The facility would be easier to secure than in Options 2, 4, or 5.
- The facility would have realistic distance-related network delays.
- A nationally distributed facility would be more accessible to clients who may need to visit a physical site to conduct experiments or consult with facility staff.
- Redundant access from the Internet to the management network would be available since each site could have an Internet access point.
- Could connect other research laboratories together using the facility’s backbone.
- The facility could support simultaneous experiments within separate network partitions (e.g., at different physical sites).

Disadvantages:

- Real Internet traffic could not flow to, from, or through the facility.
- Experiments spanning multiple sites might require physical configuration changes at remote sites, which would be difficult to coordinate.
- Management of the remote sites would be more difficult.
- The facility would need additional staff and other resources at each location.

5.1.4 Option 4: A Distributed Multi-Site Test Network Connected to the Internet

In this option, most of the requirements outlined in this report could be met, and the facility could provide both a large collection of configurable networking equipment and a high-speed connection to the Internet at multiple high-speed junction points. This solution would provide the ideal environment for many of the prospective clients of the facility.

As in Option 2, this option would offer the capability to route traffic to the Internet and to attract traffic from the Internet to destinations within this facility, but unlike Option 2, it would be connected to the Internet at multiple, geographically disparate sites. The same security considerations discussed under Option 2 would apply, including the need for gatekeeper devices to control traffic passing to and from the Internet. The existence of multiple Internet access points would make adequate network access control (Section 4.11) more challenging.

This design could partially address the realistic traffic problem (Section 4.6) by allowing some real traffic from the Internet to enter and exit the facility. Other research networks could be attached directly or through virtual private network (VPN) tunnels to the facility. As in Option 2, the facility could possibly be used to perform DDoS vulnerability assessments of other sites on the Internet. Figure 9 depicts an example topology for this design option.

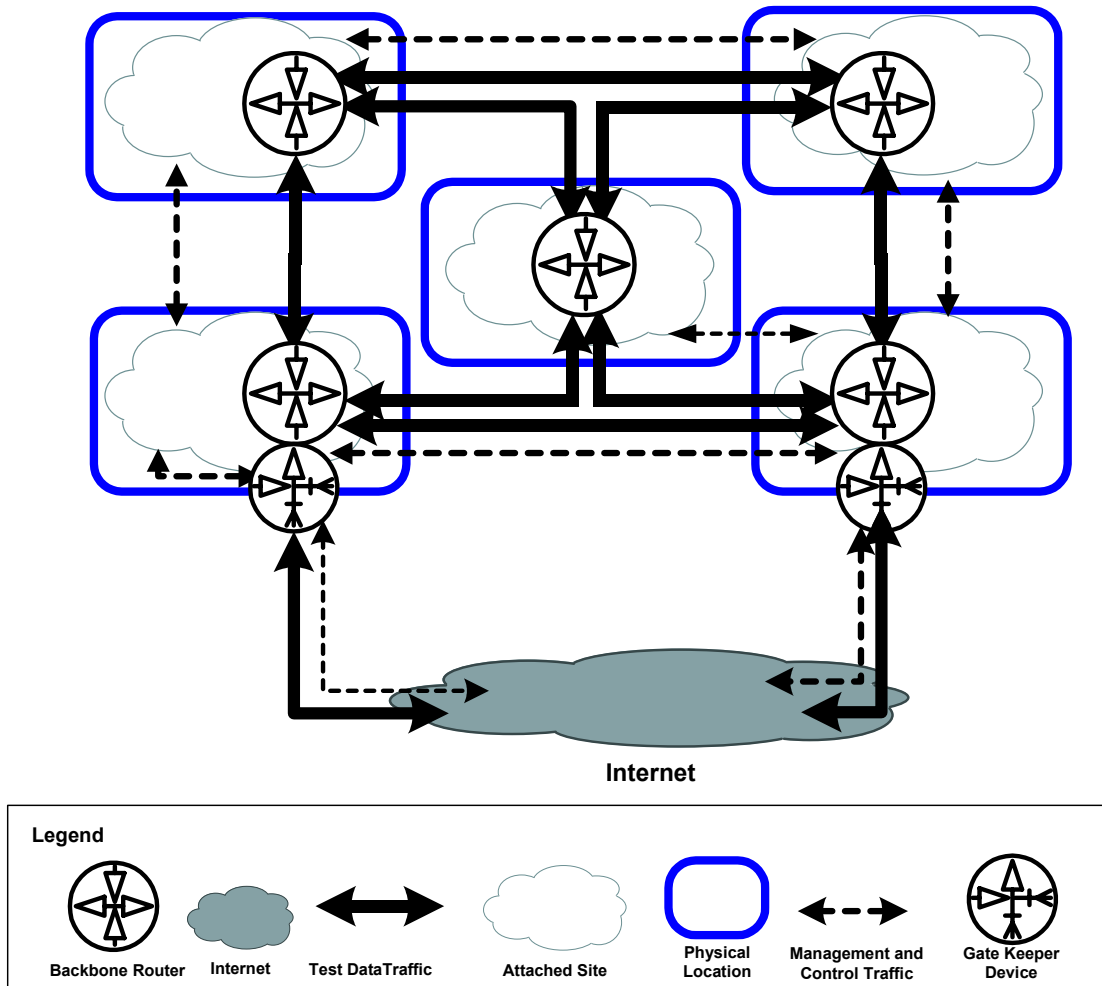


Figure 9. Option 4: A multi-site facility with test data access to the Internet

It would likely be cheapest to house the facility at or near a large existing Internet exchange point so that live Internet connections from Tier-1 providers could be linked quickly to the facility's network at minimal cost.

Usage Examples:

- The facility could simulate a complex distributed network.

Advantages:

- The facility would have realistic distance-related network delays.
- The facility could accommodate live traffic.
- Nationally distributed sites would offer greater accessibility for clients.
- The facility could offer connectivity between other research laboratories.
- Backbone traffic within the facility would be more realistic since it would have some live traffic interspersed with the test traffic.
- Redundant access from the Internet to the management and test networks could be made available since each site could have an Internet access point.

Disadvantages:

- Remote configuration changes would be more difficult to manage.
- Management of the multiple remote sites would be more difficult.
- The facility would need staffing present at each location.
- The facility would be more difficult to secure than in Options 1, 2, and 3 due to greater Internet connectivity.
- Higher recurring costs due to network connectivity expenses.

5.1.5 Option 5: A Distributed Multi-Site Test Network Providing Alternate Transit to the Internet or Another Network

Some DDoS defense technologies are focused on the networking backbone infrastructure [11]. In order to represent realistic backbone routing conditions and traffic mix, the facility would have to be able to route realistic traffic from the Internet or another live network through the facility's backbone. To accomplish realistic routing, Option 5 uses the same physical configuration as Option 4 but adds relationships with other Internet Service Providers or other backbone networks to offer alternate routes through the facility to hosts and networks that are outside of the facility. This option, like Option 4, would still provide both a large collection of configurable networking equipment and a high-speed connection to the Internet at multiple high-speed junction points. This solution would provide the ideal environment for all prospective clients of the facility as it would provide some aspect of what every client needs.

This option would have the capability to route traffic from the Internet or another network such as Internet2 into the facility and back out to its destination, making the backbone traffic more realistic. Legal agreements with other Internet Service Providers or other networks, also known as peering agreements, would have to be worked out in advance to allow traffic to traverse this facility. Routing such live traffic through the facility's network could present privacy concerns, depending on the kinds of monitoring used within the facility (see Section 4.13). Furthermore, these peering arrangements would have to be carefully managed so that live network traffic could be carried through the facility for only limited periods during experiments requiring it. If external parties grow to rely on the facility's network for carrying production data, it could dramatically reduce the facility's ability to support rapid reconfiguration and experimentation with less mature technologies (see Section 4.4).

The same security considerations discussed under Option 2 would apply, including the need for gatekeeper devices to control traffic passing to and from the Internet. In addition, further security

measures would be required in order to safely allow Internet traffic to pass through the facility without permitting, for example, DDoS floods originating in the facility to reach the Internet.

Of the five options, this design could best address the realistic traffic requirement (Section 4.6), by allowing real traffic from the Internet to traverse the facility. The facility could also be used as a “blackhole”⁵ network for real DDoS attacks.

This design would accommodate product evaluations that require operating simultaneously at sites physically distant from each other.

Figure 10 depicts an example of the data flows possible with this option. The cost of this option could be similar to that of Option 4, although higher-capacity site-to-site connections may be required.




Figure 10. Option 5: A multi-site facility offering alternative routing

Usage Examples:

- The network could offer transit routes to the Internet or another network (e.g., Internet2) that would generate real backbone traffic to evaluate the emerging DDoS defense technologies for backbone routers and backbone appliances.
- All of the examples for Option 4 would still apply.

⁵ A “blackhole network” diverts a DDoS attack from a victim’s network to a facility where the attack can be analyzed.

Advantages:

- The facility would have realistic backbone traffic.
- The facility would support realistic distance-based network delays.
- Nationally distributed facilities would be more accessible to clients.
- The facility could offer connectivity between other research laboratories.
- Backbone traffic within the facility would be more realistic since it would have some live traffic interspersed with the test traffic.
- Redundant access from the Internet to the management and test networks would be available since each site would have an Internet access point.

Disadvantages:

- Remote configuration changes would be more difficult to manage.
- Management of the multiple remote sites would be more difficult.
- The facility would need staffing present at each location.
- Very difficult to secure the facility since it would be a piece of the Internet or another network.
- There may be privacy issues if any real Internet traffic is analyzed as it traverses the facility.
- Reliance on the facility’s network could hamper reconfiguration and experimentation.
- Experiments would be less controllable and repeated due to the introduction of live traffic into the facility.

5.1.6 Comparison of the Major Characteristics of the Design Options

Table 2 lists some of the major differences between the five design options. The values “Yes” and “No” indicate whether the option provides or supports the characteristic. The value “Possible” indicates that an option provides or supports the characteristic under some conditions.

Characteristic	Design Option				
	1	2	3	4	5
	Single-site isolated	Single-site on the Internet	Multi-site isolated	Multi-site on the Internet	Multi-site offering routing on the Internet
Live Traffic	No	Yes	No	Yes	Yes
Repeatable Tests	Yes	Possible	Yes	Possible	Possible
Routes Non-Facility Traffic	No	No	No	No	Yes
Complex Peering Relationships	Possible	Possible	Possible	Possible	Yes
Breakable Network	Yes	Yes	Yes	Yes	Possible ⁶

Table 2. Selected Characteristics of the Design Options

5.1.7 Option Value to Client Types

Each of the above options will provide a different set of benefits to each facility client type. The options with the fewest features will please the fewest client types. Table 3 shows the study team’s qualitative assessments of how valuable each option is to each client type, based on the

⁶ Test facility network outages may effect other external networks if not carefully controlled.

interviews conducted. Only the fifth option, which can support all the requirements mentioned in this document, is equally beneficial to all types of organizations. The other options are more useful to a subset of the client types depending on how they might make use of the facility if it is constructed to meet only a subset of the requirements.

Option	Estimated Value to Client Type					
	Government Agencies	Research Labs	Network Device Vendors	DDoS Defense Vendors	ISPs	Commercial Clients
Option 1: isolated single site test network	●	●	○	●	○	○
Option 2: test network attached at one location	●	●	○	●	○	●
Option 3: isolated multi-site test network	●	●	●	●	●	●
Option 4: multi-location attached network of equipment	●	●	●	●	●	●
Option 5: multi-location attached network of equipment offering routing	●	●	●	●	●	●
Legend: ○ Low Value ● Medium Value ● High Value						

Table 3. Option Value to Each Client Type

5.1.8 Recommended Starting Point

The best choice for establishing a facility that meets these requirements would be to begin by building Option 1 or 2, depending on the funding available, and then migrate to Option 3, 4, or 5 over time as more funding becomes available and the facility becomes more heavily used. In order to make this happen as easily as possible, the facility should be designed from the beginning for migrating toward its final goal. It should be assumed that at some point in the future the facility will increase in size, functionality, and Internet connectivity. Planning for this expansion is critical for the facility’s success. If the needs of some clients are not adequately met by the current design, the facility could expand to encompass these needs in the future.

5.2 Estimated Costs

Any of the proposed design options will require a one-time initial capital expenditure plus recurring yearly costs. The one-time costs include initial equipment and software purchases, as well as initial facility design and setup expenditures. The recurring charges include networking connectivity charges, labor, rent, maintenance, upgrades, and insurance. The following sections identify the major budget items in each category and provide rough cost estimates for each of the design options described above. All costs given are based on advertised list prices. It is expected that donations or significantly lower prices may be made available for many of these items by vendors interested in supporting the facility’s goals and in encouraging use of their products through exposure to facility clients.

5.2.1 Equipment and Software

Networking equipment can be divided into three categories: network backbone equipment; intermediary equipment; and local area networking (LAN) equipment. The facility must be equipped with equipment that is representative of the equipment used on the Internet, as required in Section 4.3.

Host equipment can be divided into two categories: servers and end-user workstations. We recommend that all hosts within the test facility be rack mountable to efficiently utilize the facility's floor space.

Internet servers are typically either Intel-based processors or RISC-based processors. These servers will be used to emulate various functions including web-servers, databases, and business-to-business servers. Some of these machines will be used in experiments to emulate the defense environment, while others will be used in simulating DDoS attacks. We recommend that all of the hosts in this facility have at least high speed networking interface cards. Licenses for commonly used operating system software (e.g. Microsoft Windows, and Sun Solaris) must be available.

The estimated costs for equipment purchases ranges from \$11 to 32 million, depending on which optional features are selected.

5.2.2 Network Connectivity Costs

Most Tier-1 Internet service providers run their backbones at speeds of OC-48 (2.5 Gbps) or OC-192 (10 Gbps). We recommend that the test facility's backbone network should run at speeds of at least OC-12 (622 Mbps) and preferably be able to run at OC-48 speeds. Our research has indicated that OC-192 circuits are currently under-utilized in the Tier-1 ISP's network, and therefore, they are not usually the bottleneck in a DDoS attack. Because high speed networking cards for routers can be extremely expensive and because most current DDoS attacks can be adequately modeled at speeds of OC-48 or less, in the near term, facility will not need OC-192 networking. Circuit costs for high-speed links are a significant recurring cost, and they will greatly impact the total cost of operating the test facility.

The networking charges for operating the facility each year are expected to range from a negligible amount to \$47 million, depending on the level of internet connectivity is needed by the final facility design. Where the facility is located may significantly impact the telecommunication costs. For example, locating the facility at an Internet exchange point would offer lower connectivity costs since any required network connections can be over short distances within the exchange point.

5.2.3 Labor

The facility will require an administrative staff, an engineering staff, legal counsel, and a technical advisory panel.

The administrative staff would minimally include a director to run the facility, marketing personnel to publicize and promote the facility, clerical staff, and an accounting staff to pay expenses and collect fees from the clients.

A facility of this size would require an engineering staff to aid with reconfiguration and keep the facility operating. Our research estimates that we would need at least five (or more if the facility is geographically distributed) full time engineers/system administrators.

The legal counsel and the advisory panel could be kept on retainer to help with establishment of initial operating procedures and handle special circumstances as they arise.

Annual labor charges are expected to require \$2 to 3 million in funding.

5.2.4 Rent, Maintenance, Upgrades, and Insurance

This facility will take up a significant amount of floor space, and will have special electrical and cooling requirements. A large hosting complex, an Internet exchange point, a University's computer science department, or a government facility would be able to host this facility. Monthly rent may be expensive and all options should be considered to help minimize this cost.

Maintenance of all the hardware and software will incur additional annual costs. Depreciation and contingency accounts for new equipment should be setup and funded in advance to insure that the facility remains a state-of-the-art facility. We recommend that the facility budget 33% of the cost of the equipment per year to be used as a replacement and upgrade.

Annual charges for rent, maintenance and replacement costs are expected to fall between \$5 and 15 million. Insurance costs are not included in this estimate, but it is expected that the cost of insurance will be negligible in comparison to the rest of the charges.

5.2.5 Design Option Cost Comparison

Our research has indicated the cost to initially build and operate the facility for the first year will vary significantly due to differences in each of the five design options (e.g. one location versus multiple locations, local telecommunication circuits versus cross-country circuits). This section is intended to give rough cost estimates for each of the five design options. Detailed costing estimates are provided in **Error! Reference source not found.**

For comparison, we will use the following assumptions:

- The facility will consist of three *logical* sites.
- For Options 1 and 2 this means enough equipment to build three complete logical sites at a single physical location.
- For Options 3, 4, and 5 this means enough equipment and staff to equip three geographically dispersed test facility locations.
- All prices used are list prices and assume no discounts.
- The out-of-band management network will be 1.5 Mbps circuits.
- 160 Mbps circuits will be used for Internet access in options 2 and 4.
- Option 5 will use 2.6 Gbps circuits to connect the sites.
- All figures are rounded to the nearest million dollars.

	Option 1	Option 2	Option 3	Option 4	Option 5
One Time: Equipment, Software, Design, and Build-out	12	12	34	34	34
Recurring costs:					
Labor	1.7	1.7	2.7	2.7	2.7
Network Connectivity	0.02	0.5	0.5	1.4	37
Rent, Maintenance, and Replacement	5	5	15	15	15
Total for 1st Year	19	19	53	54	90
Annual Recurring Costs	7	7	19	19	55

Table 4. Estimated Costs for Construction and Operation of the Facility (in Millions of Dollars)

5.3 Funding Models

As discussed in Section 4.11, there are two financial budgetary concerns that must be addressed for the facility to be successfully deployed. The first is how the initial build-out will be funded, and the second is how the ongoing year-to-year costs will be funded. In this section, we examine a few funding possibilities for each of these concerns.

5.3.1 Funding for Initial Build-out and Set-up Investment

A facility of the magnitude envisioned in Options 3, 4, or 5, may require obtaining funding from multiple sources including the DoD, government agencies, as well as private sector organizations.

The DoD and government agencies will likely be the driving force for the initial build out and set up of the facility. It is expected that a large portion of the initial funding will need to come from the government organizations having the most urgent need to solve the DDoS problem.

With the right incentives and return on investment, it may be possible to attract some funds from the private sector. This may include endowments in the form of grants and donated equipment such as computers, routers, and switches. Currently, industry donates hardware and software to institutions that provide potential return on investment. The return on investment may include advertising, tax deductions, etc. For example, many commercial companies donate equipment to universities since they receive a tax deduction as well as the opportunity to expose students to their products. Network device vendors may see some benefit to donating equipment to the facility, especially if this can increase visibility and open potential sales opportunities to government and DoD customers. It is also likely that if future government purchases were to be made based on results of experiments conducted in this facility, product vendors would be more likely to help contribute its construction and operating costs.

One strategy of attracting donations would be to directly associate the facility with a public university or institution. This would provide both the exposure and tax benefits previously mentioned. Universities are also accepted as a vendor-neutral environment. Therefore, associating the facility with a university would help satisfy the public availability and vendor neutrality requirements described in Sections 4.1 and 4.2. Finally, associating the facility with a university whose faculty includes experts in DDoS research would also help fulfill the technical staffing requirements outlined in Section 4.9.2.

Investigations into the availability of the above two sources of funds should go hand-in-hand with explorations of various approaches to the actual construction and housing of the facility. A variety of possibilities exist, ranging from constructing the physical building and associated

network lines from scratch, to hosting the facility at one or more network exchange points that offer floor space in highly secure facilities that connect a variety of Internet Service Providers (ISPs) and Internet backbones together. This would allow the DDoS facility to be established in a shorter period of time and with lower construction and maintenance expenses. Network exchange points also offer a flexible physical footprint, allowing for easy expansion as the facility's spatial needs increase. A dialog with operators of these types of facilities (e.g., Equinix or PAIX) should be undertaken to determine their willingness to donate floor space, network connections, and other services to the facility.

5.3.2 Funding for Daily Operations, Maintenance, and Usage

Ongoing funding will also be required for staff salaries, training, sponsored clients, equipment replacement costs, etc. This raises a fundamental question: should the facility be self-sustaining or subsidized at least in part by government and other organizations? It is probably realistic to assume that the facility cannot be self-sustaining in its initial years of operation. However, during these initial years, a variety of revenue generation and funding models can be explored. These include the following:

- **Government sponsorships.** It is likely that the majority of the initial costs of the facility will need to be covered by the government. It is especially important that the government make the facility available to research institutions that may not have large expense budgets. A grant system should be set up to help these smaller institutions economically test their research within the facility.
- **Consortium-based memberships.** This is a multi-tiered membership structure along the lines of consortium memberships. Depending on the membership level and fees, different levels of access to the facility are guaranteed for a period of time for these members. The objective of this model is to be flexible and offer a variety of pricing plans to meet the needs and budgets of a diverse set of clients.
- **Sponsored research projects, experiments, and demonstrations.** The facility could be used by a variety of industry, government, and DoD research projects to host experiments and demonstrations. This could complement facilities such as DARPA's Technology Integration Center (TIC), which has been used to stage a variety of integration experiments, tests, and project demonstrations.
- **Consulting and research services.** These consulting services would be offered to facility clients by the facility's DDoS research and engineering staff. Affiliated university faculty or research staff members might also offer consulting services, thus enabling the facility to out-source some of the needed technical expertise. Fees would be charged to the clients to directly recover the cost of these staff or affiliate members.
- **Corporate sponsorships.** Companies may be willing to donate equipment, space, and networking to the facility if the donations provide a good return-on-investment. This is discussed in detail in Section 5.3.1.
- **Fee-for-time.** This would be the scheme preferred by small DDoS defense vendors and research organizations. Clients would "rent" the networks, equipment, and general usage of the facility on an as-needed basis for probably short periods of time. A variety of pricing plans and packages could be offered based on the level of access needed and the demand for the resources (such as high-end routers) to be utilized.
- **Training programs.** The facility could offer a variety of educational, awareness, and training programs to the networking and security community. Fees would be charged to people attending these training programs.

Each of the above funding and revenue generation models has its advantages and disadvantages, and no single model will be suitable for all clients. We discuss these aspects in the next section, along with a phased approach to exploring and implementing these models.

5.3.3 Funding Summary and Recommendations

Table 5 summarizes the advantages and disadvantages of the above-discussed revenue and funding models.

Funding model	Advantages	Disadvantages
Government sponsorships	<ul style="list-style-type: none"> • May provide a relatively steady source of funds. • Government is seen as vendor-neutral. • Government can provide financial assistance for small clients. • Promotes the long-term viability of the facility, as it is less prone to the revenue fluctuations of the corporate sponsors. 	<ul style="list-style-type: none"> • Difficulty in getting budgetary approval from various government agencies.
Consortium-based-memberships	<ul style="list-style-type: none"> • Most appropriate for clients who need to use the facility on a continual basis. • Allow clients to be guaranteed certain levels of access. • Multiple tiers provide flexible pricing and access guarantees. 	<ul style="list-style-type: none"> • Inappropriate for clients who need one-time/short-term use of the facility. • Small vendors may be reluctant to be tied down to yearly or long-term memberships.
Sponsored research projects, experiments, and demonstrations	<ul style="list-style-type: none"> • Promotes the facility and increases its name recognition and legitimacy 	<ul style="list-style-type: none"> • Sponsored research and experiments may tend to be lengthy and may require the aid of support staff.
Consulting and research services	<ul style="list-style-type: none"> • Could provide a very high profit-margin source of revenue. 	<ul style="list-style-type: none"> • There has to be enough demand to make consulting viable. • Added overhead and complexity in attracting, supporting, and retaining clients, consulting staff and infrastructure.
Corporate sponsorship	<ul style="list-style-type: none"> • Increases the chance of equipping the facility with state-of-the-art equipment. • Increases visibility and appeal to commercial clients. • Encourages vendors to implement DDoS solutions in their products. 	<ul style="list-style-type: none"> • Large commercial sponsors may want to exercise more control and this may result in the lowering of the vendor-neutral posture of the facility. • When corporations cut back on sponsorship in economically hard times, the facility may be denied the optimal mix of network equipment.
Fee for time	<ul style="list-style-type: none"> • Provides financial and short-term usage flexibility for small clients. 	<ul style="list-style-type: none"> • Will require publicly promoting the facility and its services. • A pure fee for time model provides no funding guarantees up front and makes budgeting more difficult.
Training programs	<ul style="list-style-type: none"> • Provides a complimentary and independent source of revenue. 	<ul style="list-style-type: none"> • Increased overhead in providing education and training (i.e., staff, class facilities, materials, etc.)

Table 5. Advantages and disadvantages of funding and revenue generation models.

As can be seen from the table, each funding model has various advantages and disadvantages. We expect these models to be adopted in a phased manner as the infrastructure and services offered by the facility evolve over time. Our recommendation for a phased approach would be as follows:

Phase 1: Initial build-out and operation through government sponsorships.

This phase will concern mostly the initial build-out and the fixed costs for the initial day-to-day operation of the facility, and we expect government sponsorships to be the predominant funding source. We also recommend that the government provide financial assistance to small research organizations. By increasing the availability of the facility to those research organizations that have promising ideas that merit larger scale testing and validation, the facility will gain visibility and credibility rapidly.

Phase 2: Expanding operations through consortium memberships.

We anticipate this phase to come after a year or two of initial operation. During this phase the facility will expand operations to encompass a wider array of clients and services. A variety of consortium memberships should be explored with various government, DoD, and commercial agencies and partners.

Phase 3: Full and mature operations.

Transition to this phase occurs when the facility has reached full maturity and offers a wide array of services to meet the demands of diverse clients. This phase will probably begin at least three or four years after initial start up. At this stage, we expect funding for the facility to be complemented by fee-for-time, consulting services, and corporate sponsorships.

In summary, we expect the funding models for daily operations and maintenance to be a combination of government and industry sponsorships, fee-for-time, staff consulting, and government and industry sponsorship. As discussed above, these funding models will likely be phased in over a few years as recognition of the facility and demands for its services grow. It is unlikely that any single one of the above funding possibilities will be capable of funding all aspects of the ongoing facility costs at any time.

5.4 Measures of Success

It is important that the facility be evaluated regularly to determine whether it is providing effective support for DDoS defense development and deployment, and how its role in finding a solution could be improved. The results of such evaluation can guide the further development of the facility and indicate when it has reached sufficient maturity to expand its mission or explore additional funding models. The high-level goals of the facility are expressed in the charter statement proposed in Section 3, but in order to objectively evaluate the facility's utility and effectiveness, more concrete metrics should be applied. A few such metrics are outlined here:

- **Client demand.** The number of clients seeking to make use of the facility's services, especially clients who return repeatedly, provides an indication of the perceived value of the facility.
- **Demand from non-government clients.** Commercial clients will have a different set of needs and cost/benefit tradeoffs from government agencies and military clients; the facility will ideally provide value to both client classes.
- **Buy decisions following facility evaluation.** Clients who use the facility for evaluation of DDoS defense products and subsequently buy one or more such products may be assumed to have found such evaluation convincing.

- **Product releases following facility testing.** Commercial products that are tested or developed within the facility and subsequently released for purchase suggest that the facility is effectively promoting development of new DDoS defense tools.
- **Publications.** The emergence of publications that include experimental results from studies conducted within the facility can indicate that the research community considers the facility a credible experimentation environment that fosters significant contributions to knowledge in the DDoS area.
- **Tools developed locally.** The facility can provide value to the research, development, and operational communities through publicly released tools for traffic generation, traffic capture, network management, network monitoring, automated reconfiguration, DDoS measurement, and other purposes that are developed by facility staff or by researchers using the facility. The utility and acceptance of such tools can be gauged by statistics including the number of times that a tool is downloaded from public repositories and the frequency with which such tools are cited in publications.
- **Client feedback.** Each facility client should be asked to complete a questionnaire following use of the facility's services. Responses to these questionnaires will provide information necessary for several of the measures above. In addition, direct evaluation by clients of the facility's strengths and weaknesses provide an important measure in itself.

5.5 Operational Recommendations

We recommend that the facility remain as vendor-independent as possible. To meet this goal, it should ideally be run by the government or housed within a neutral environment such as a university. We recommend that the operational procedures for running experiments in the facility be similar to the model depicted in Figure 11. A prospective client would likely begin by requesting information about the facility and would obtain the facility's charter along with a form to be used for submitting experimentation proposals. The client should then fill out the proposal form and indicate the following: why the use of the facility is sought; a description of the experiments to be performed; the length of time needed to conduct the experiments; the equipment needed; and any other pertinent details about the experiment. Upon submission of a proposal, the facility's steering committee should examine the submission and make a decision on whether the experiments to be performed are in line with the facility's charter and priorities (see Section 4.10.3). If the committee does not feel that the proposed experiments fall within the purpose of the facility, a reply indicating the reason for the rejection should be sent to the proposal's authors. If the proposal is approved, and if the client has requested financial aid for making use of the facility, a decision must be made regarding that topic as well. Finally, once approved, the client will coordinate with the administrative staff to schedule their experiment. When the client begins work within the facility, they must be shown how to make use of the facility's equipment, personnel and other resources. The network engineers and DDoS experts on-site should help the clients set up and control their experiment as well as help fulfill any other needs the client may have. Once finished, the client should fill out an evaluation form to provide feedback to the facility staff. It is likely that in many research cases the results of the experiments will lead to additional work. It is up to the steering committee and the clients to discuss how this new work is to be scheduled and whether a new proposal must be submitted.

It is likely that some of the facility's clients will need to bring in their own equipment when conducting experiments. It should be a goal of the facility to reduce the amount of equipment that must be brought in by each client, but it cannot be expected that every need of every client can be met by the facility itself. Therefore, a policy must be in place that allows clients to incorporate their own equipment into their experiments. This policy will need to take into consideration insurance liabilities of both parties.




Figure 11. A suggested model for facility operations

6 Future Considerations

To ensure long term viability of the facility, its expected usage over a long period of time must be considered. Although the initial purpose of the facility will be to accomplish the DDoS-specific charter objectives discussed in Section 3, it is likely that other critical network security threats will present themselves in the future and require national-level attention and a sophisticated research facility. In fact, individuals interviewed during this study stated repeatedly that a facility like the one proposed here could be extremely valuable for studying other network security problems and should not be exclusively reserved for DDoS defense experimentation. In this section, we outline some topics that should be considered in the future after the facility has reached some level of maturity.

6.1 Future Applicability of the Facility

Any facility to be designed and constructed from the recommendations outlined in this report must take into account how the facility might be useful after substantial progress has been made in developing effective DDoS defense solutions. The usefulness of this facility will need to extend beyond just the DDoS problem space. It should be possible to eventually use the facility for study of other networking problem spaces as well. Section 4.10.3, discusses the need to prioritize the facility's use. It should be noted that if other research work demonstrated a compelling need to make use of this facility, but outside of the DDoS research realm, its proposal should be considered and possibly allotted experimentation time. If a more serious problem surfaces in the future with an even more pressing need for a solution than DDoS attacks, then a modification of the facility's charter might be warranted.

As an example, many of the participants inquired as to whether the facility would be usable by research teams studying Internet worms, which continue to be a significant security problem. If a worm surfaces in the future that threaten the United State's security or infrastructure, this facility might prove to be an ideal place to study defenses to such a pressing problem.

6.2 Evaluation Organization

There is a need for a vendor-neutral evaluation organization that can be trusted to perform independent evaluation studies of DDoS defense technologies. The study participants expressed a desire for such an organization to act on behalf of the community-at-large. A published report on current technologies could be released by the organization describing the strengths and weaknesses of the evaluated technologies. A report produced by this organization should also identify the networking topologies in which the product is likely to perform best. It should indicate at what level of scalability the technology is sufficient for deployment. The organization producing these reports would likely need to make use of the facility in order to conduct its evaluations. Thus, it would be closely affiliated with the facility, although not directly tied to it. Like the facility itself, it is important that such an organization remain vendor-neutral. Vendors should be encouraged to help the evaluation organization to understand and utilize their DDoS defense products and technologies, and to help with setting up experiments within the facility. The organization and the reports they generate should strive to encourage deployment of any technologies that help mitigate the effects of the DDoS problem.

6.3 Bake-Off Hosting

Study participants also expressed a strong need for a vendor-neutral location where standardized "bake-off" testing could take place. As many DDoS defense technologies begin to emerge, sites that wish to deploy them need a side-by-side comparison to determine which technology is best suited for their environment. It would, therefore, be beneficial if a portion of the facility's time

could be devoted to holding a technology bake-off where different DDoS technologies could all be evaluated at once under the same networking conditions. DDoS product vendors and research experts should be invited to participate in the bake-off and should be encouraged to bring the latest versions of their products or technologies for evaluation. The results of these bake-offs should be collected and published in the form of a report, possibly from an organization like the evaluation organization described in the previous section. Hosting these bake-offs will provide the added benefit of promoting competition and increasing the speed at which DDoS defense technology development occurs.

6.4 DDoS Security Expert Conference

The need for two new conferences related to the DDoS problem space was also identified by this study.

Specifically, there is a need for a new DDoS-specific conference where network security officers could gather to compare notes and discuss solutions. This would be akin to existing conferences like the North America Network Operators Group [NANOG], but focused on current security concerns and practices. In particular, the DDoS conferences that currently exist are concerned more with research and development and there is a strong need for an operational forum as well.

There is also a need for a conference that specializes in topics relating to designing, running, and maintaining networking test infrastructures. Many small- to medium-scale test beds exist today, but the maintainers of these infrastructures do not typically have the opportunity to interact with one another in order to share ideas and techniques that can be used in emulating operational networking environments. It is likely that if such a conference were to take place it would result in an increase in the effectiveness of existing and future networking test infrastructures. It would be highly beneficial for the facility's own operational staff to attend so they can learn from other operators' techniques as well as share their own knowledge and experience from the facility.

7 Summary and Conclusions

The explosive growth of the Internet and its increasingly critical role in supporting electronic commerce, transportation, and communications, have brought an equally explosive growth in attacks on Internet infrastructure and services. Some of the most difficult attacks to defend against are the Distributed Denial of Service (DDoS) attacks, in which an overwhelming flood of network packets is generated by many different sources, with the intent of preventing legitimate use of a service. These attacks may target end-users, web servers, entire networks or parts of networks, or specific networking infrastructure components. It is now widely recognized that DDoS attacks pose a severe threat to the nation's ability to conduct business, defend itself, and provide vital government services to its citizens.

Researchers and vendors have begun to explore a variety of DDoS defense approaches. However, these development attempts today are hampered by the lack of realistic test environments in which to study sophisticated DDoS attacks as well as analyze, test, and develop DDoS defense solutions. Facilities such as the laboratories of academic, government, and commercial research groups, and the test networks of most vendors, have limited capabilities to simulate large-scale attacks and complex network topologies. While some large router vendors and ISPs have built medium to large test networks and laboratories, these are unsuitable for general DDoS experimentation as they only typically support the vendor's products and are not available for use by external organizations, vendors, and research groups.

More importantly, the owners of these large test networks have little economic incentive to add sophisticated DDoS defense features to their products and services, or invest aggressively in DDoS defense research. Consequently, it is unlikely that industry, left to its own devices, will address the nation's DDoS vulnerabilities with the urgency required. If the United States is to tap the energy and ideas of a broader cross-section of organizations to make rapid advances in DDoS defense, it must ensure that adequate testing and experimentation facilities are available. Without such facilities, progress in DDoS defense will continue to be outpaced by improvements in DDoS attack technology. This in turn will further increase the likelihood of successful attacks that cause protracted network outages, disrupting communications and critical functions within both government and commercial sectors.

The principal conclusion of this study is that a national facility for experimentation, testing, and evaluation of DDoS defense technologies is greatly needed. Specifically, the recommended charter for the National DDoS Defense Technology Evaluation Facility is to provide a shared laboratory in which researchers, developers, and operators from government, industry, and academia can evaluate potential DDoS defense technologies under realistic conditions, with the aim of accelerating research, development, and deployment of effective DDoS defenses for the nation's computer networks. The facility would be a public, national resource providing substantial infrastructure and equipment to support large-scale, reproducible DDoS defense testing and experimentation. Potential clients of the facility include government and DoD research labs and agencies, academic and industry research labs, network device vendors, DDoS defense product vendors, ISPs, and large enterprises that use and deploy DDoS defense technology. These clients would propose experiments that would be reviewed, approved, and prioritized by an oversight board according to established criteria. A skilled technical staff would help clients make efficient use of their time in the facility.

The facility should be built in phases in order to begin providing partial benefits as soon as possible, and to allow experience with early phases to guide further development. The networking infrastructure, staffing, funding, and revenue generation models should evolve over time to accommodate greater functionality, client diversity, and range of services. The networking infrastructure can start out as an isolated, single-site test network and evolve into a

multi-site test network attached to the Internet through multiple peering points and Tier-1 ISPs. The facility's offered services may be limited at first to basic DDoS defense experimentation support, while expanding later to include DDoS defense related external consulting and training. Initial facility development would rely on government funding and corporate equipment donations. Once the facility is more mature, a variety of other revenue generation strategies can be explored, including fee-for-time, consortium memberships, consulting services, sponsored research and experiments, and commercial sponsorships. It is important that the facility maintain a vendor-neutral orientation through every phase to promote broad participation.

Given the critical and growing threat posed by DDoS attacks, the U. S. Government must move forward decisively to establish a national facility that offers a large-scale, realistic environment for evaluation and analysis of DDoS defenses. This facility will allow a diverse collection of researchers to explore new ideas in DDoS defense. It will facilitate rapid testing and evolution of commercial DDoS defense products. It will enable government agencies and others to reliably evaluate potential DDoS solutions for their operational networks. By providing vital support for each stage from conception to deployment of DDoS defenses, the facility will significantly accelerate progress towards the common goal of protecting critical national infrastructure from DDoS attacks.

References

- [1] David Moore, Geoffrey Voelker, and Stefan Savage "Inferring Internet Denial of service Activity" Published in proceedings of the 2001 [USENIX Security Symposium](#). The full paper in PDF is located at <http://www.caida.org/outreach/papers/2001/BackScatter/usenixsecurity01.pdf>
- [2] CERT report on Trends in Denial of Service Attack Technology, October 2001 http://www.cert.org/archive/pdf/DoS_trends.pdf
- [3] Snoeren, A., Partridge, C., Sanchez, L., Jones, C., Tchakountio, F., Kent S., and Strayer, W., "Hash-Based IP Traceback," SIGCOMM 01, August 27-31, 2001, San Diego, California, USA.
- [4] S. Bellovin, M. Leech, "ICMP Trace-Back" Internet-Draft, July 2001.
- [5] S. Dietrich, N. Long and D. Dittrich, "Analyzing Distributed Denial of Service Tools: The Shaft Case", Proceedings of the LISA XIV, December 3-8, 2000, New Orleans, LA.
- [6] Cyber crime bleeds U.S. corporations, survey shows; financial losses for attacks climb for 3rd year in a row, April 7, 2002. <http://www.gocsi.com/prelea/000321.html>
- [7] D. Sterne, K. Djahandari, B. Wilson, B. Babson, D. Schnackenberg, H. Holliday, and T. Reid, "Autonomic Response to Distributed Denial of Service Attacks," 4th International Symposium on Recent Advances in Intrusion Detection (RAID 2001) held in Davis, CA on October 10-12, 2001.
- [8] Emulab: The Utah network testbed. <http://www.emulab.net/>
- [9] Equinix, Inc. <http://www.equinix.com/>
- [10] PAIX.net, Inc. <http://www.paix.net/>
- [11] John Ioannidis and Steven M. Bellovin, "Implementing Pushback: Router-Based Defense Against DDoS Attacks," AT&T Labs Research published in the proceedings from the 2002 Network and Distributed Systems Security Symposium February 6-8, 2002, San Diego, CA.
- [12] D. Moore, "CAIDA Analysis of Code-Red," <http://www.caida.org/analysis/security/code-red/>
- [13] W. Leland, M. Taquq, W. Willinger and D. Wilson. "On the self-similar nature of Ethernet traffic", Proceedings of ACM SIGCOMM, 1993.
- [14] CERT hit by DDoS attack for a third day, May 24, 2001. <http://www.itworld.com/Sec/3834/IDG010524CERT2/>
- [15] Scottish ISP floored as DDoS attacks escalate, September 4, 2002. <http://www.theregister.co.uk/content/6/24773.html>

Appendix A. Acknowledgements

Many individuals were interviewed during the data collection phase of the study and provided valuable insights and feedback:

- ❖ Jay Adelson, Founder and Chief Technology Officer, Equinix, Inc.
- ❖ Steve Bellovin, AT&T Research
- ❖ Matt Bishop, Department of Computer Science, University of California, Davis
- ❖ Sean Capshaw, Juniper Networks
- ❖ Mr. Kevin Charlow, Space and Naval Warfare Systems Center, Charleston SC
- ❖ Kelly J. Cooper, Genuity
- ❖ David Dittrich, Senior Security Engineer University of Washington
- ❖ Brian Gemberling, UUNET Technologies
- ❖ Andre Girona, operations.net
- ❖ Pat Hilton, AT&T Labs
- ❖ Jeff Ingle, Intelligence Community Chief Information Officer Staff
- ❖ Farnam Jahanian, Arbor Networks
- ❖ Jay Lepreau, Emulab at the University of Utah
- ❖ Thomas A. Longstaff, Ph.D. - Manager, Survivable Network Technology, Software Engineering Institute
- ❖ Rob Malan Arbor Networks
- ❖ David Moore, CAIDA/SDSC/UCSD
- ❖ Chris Morrow, UUNET Technologies
- ❖ Kihong Park, Purdue University
- ❖ Lane Patterson, Senior Research Engineer, Equinix Inc.
- ❖ Mr. John Peterson, Space and Naval Warfare Systems Center, Charleston SC
- ❖ Researchers at Indiana University's Advanced Network Management Lab
- ❖ Stuart Staniford, Silicon Defense
- ❖ LCDR Michael Shumberger, Space and Naval Warfare Systems Center, Charleston SC
- ❖ Mr. Robert Varnes, Scientific Research Corp.
- ❖ Mark M. Warren, Genuity
- ❖ Professor S. Felix Wu, Computer Science Department, University of California, Davis
- ❖ Marc Zwillinger

Appendix B. Estimated Costs

Each of the proposed design options will require a one-time initial capital expenditure plus recurring yearly costs. The one-time costs include initial equipment and software purchases, as well as initial facility design and setup expenditures. The recurring charges include networking connectivity charges, labor, rent, maintenance, upgrades, and insurance. The following sections identify the major budget items in each category and provide rough cost estimates for each of the design options described above. All costs given are based on advertised list prices. It is expected that donations or significant discounts may be made available for many of these items by vendors interested in supporting the facility's goals and in encouraging use of their products through exposure to facility clients. *These estimates are intended as rough guidelines only; a thorough enumeration of anticipated costs should be conducted as part of the facility design phase.*

B.1 Equipment and Software

Networking equipment can be divided into three categories: network backbone equipment, intermediary equipment, and local area networking (LAN) equipment. The facility must have equipment that is representative of that used on the Internet, as required in Section 4.3.

Host equipment can be divided into two categories, simulating servers and end-user workstations. We recommend that all hosts within the test facility be rack mountable to efficiently utilize the facility's floor space.

Internet servers are typically either x86-based machines using the standard PC architecture or RISC-based machines sold by operating system vendors (e.g., Sun, IBM, or HP⁷). These servers will be used to emulate various functions including web servers, databases, and business-to-business servers. Some of these machines will be used in experiments to simulate the defense environment, while others will be used in simulating DDoS attacks. We recommend that all of the hosts in this facility have high speed networking interface cards, since one host may simulate the output from a whole network in many experiments. Licenses for commonly used operating system software (e.g. Microsoft Windows and Solaris) must be available.

The estimated costs for initial equipment purchases range from \$11 million to \$32 million, depending on which optional features are selected.

B.2 Network Connectivity Costs

Most Tier-1 Internet service providers run their backbones at speeds of OC-48 (2.5 Gbps) or OC-192 (10 Gbps). We recommend that the test facility's backbone network run at speeds of at least OC-12 (622 Mbps) and preferably OC-48. Our research has indicated that OC-192 circuits are currently under-utilized in the Tier-1 ISPs' networks, and therefore, they are not usually the bottleneck in a DDoS attack. Because high speed networking cards for routers can be extremely expensive and because most current DDoS attacks can be adequately modeled at speeds of OC-48 or less, in the near term, the facility will not need OC-192 networking. Circuit costs for high-speed links can exceed \$800,000 per month. If design Option 1 were to be built, monthly circuit charges would be minimally about \$3,000 per month for T-1 access to the Internet. Design Option 5 with up to five geographically dispersed sites could cost in excess of \$4,000,000 per month based on list prices of OC-48 circuits. The facility's location may significantly impact the telecommunication costs. For example, locating the facility at an Internet exchange point would

⁷ All registered and unregistered trademarks in this document are the sole property of their respective owners.

offer lower connectivity costs since any required network connections would span only short distances within the exchange point.

B.3 Networking Equipment

Networking equipment can be divided into three categories: network backbone equipment, intermediary equipment, and local area networking (LAN) equipment. The facility must have equipment that is representative of that used on the Internet, as required in Section 4.3.1.

Tier-1 Internet service providers (ISPs) use almost exclusively Cisco and Juniper routers and switches in their core backbone networks. In the Cisco product line, the most commonly used backbone router models are the 12000 GSR and the 7000 series. In the Juniper product line, the most commonly used backbone router models are the M160 and M40. Each of the backbone routers will need to support multiple line cards at speeds up to OC-48. These high-speed line cards are often as expensive as the router itself. Our research has indicated that this facility would require at least five large backbone routers. These routers will be used to emulate one or more ISP backbones. The list price of this class of routers, equipped with two OC-48 line cards, is approximately \$300,000.

Intermediary equipment includes routers and large switches that are primarily used to aggregate smaller access routers into the core. These are typically Cisco 7000 and 3600 series routers, Cisco 6000 series switches, and Juniper M40 and M20 routers. These routers need to have multiple interface line cards at speeds of up to OC-12, at least one 100Mbps Ethernet interface, and possibly a gigabit Ethernet interface. This facility should have 15 – 25 intermediary routers and 5 large switches. These routers and switches will be used to emulate a variety of network environments, including ISP peering points, tier-2 and tier-3 ISPs, corporate backbones, university networks, and large hosting facilities. The list price of this class of routers is approximately \$75,000. Prices will vary depending on the number and speed of its line card interfaces. Large switches list for approximately \$150,000.

Access equipment usually consists of smaller routers or switches supporting relatively low bandwidth. These devices are used to connect remote sites to each other or to the Internet. They are typically Cisco 2600 series or Juniper M5 routers; however, IBM, Foundry, Lucent, and Nortel make up a smaller, but significant, share of the access router and switch market on the Internet. Load balancers, special switches that distribute incoming traffic across a number of servers, are a specialized class of access equipment that will be required for testing certain hosting scenarios. We recommend that this facility have 15 - 20 access routers or switches and at least three load balancers. A typical access router such as the Cisco 2600 series router has a list price of approximately \$20,000. Some commonly used load balancers are Foundry's ServerIron and Cisco's LocalDirector. They list for approximately \$15,000 per unit depending on the type of network interface.

Local area networking equipment includes intelligent hubs and switches. This facility will contain thousands of hosts connected by two disjoint networks. Our research indicates that the facility will minimally need 100 small VLAN capable switches and hubs. A small switch such as the Cisco 2850 lists for approximately \$5,000.

B.4 Host Equipment Costs

Host equipment can be divided into three general categories: servers, firewalls, and end-user workstations. The facility must be able to simulate an Internet server farm, a site or sites protected with firewalls, and end-user workstations.

Internet servers are typically either Intel based processors or RISC based processors. These servers will be used to emulate various functions including web-servers, databases, and B2B-servers (business to business). Note that some these machines can also be used as DDoS generating attack machines. These machines are usually rack mounted and have 100 Mbps or gigabit Ethernet network interface cards. Our research indicates that this facility will need 500 Intel servers running Microsoft, Linux, or BSD-based operating systems, and 100 RISC based servers running Unix operating systems built for that manufacturer's hardware (e.g., Solaris, HP-UX, or IBM's AIX). Appropriately configured, rack mountable Intel machines cost approximately \$4,000 per unit. The RISC servers cost \$10,000 to \$25,000 per unit.

Firewalls are devices which filter or proxy traffic from a less trusted network (e.g., the Internet) to a more trusted network or host. The facility will need at least three high-speed firewalls, such as a NetScreen 5000 series, and will need at least five smaller firewalls. High-speed firewalls typically cost \$100,000, which includes hardware, operating system, and firewall software. Smaller firewalls tend to cost around \$10,000, which usually includes the firewall software only.

End-user workstations will be used within this facility to emulate DDoS controlling agents and DDoS attacking hosts. These will typically be Intel machines running a Microsoft or Linux operating system. Our research has indicated that this facility should minimally have 500 end-user workstations. These typically cost around \$2,000 per unit without a monitor. (Note: This price is based on rack mountable systems. Efficient space utilization requires that this facility use rack mountable systems.)

B.5 Other Equipment

A requirement of this facility is the ability to generate large volumes of traffic (see Section 4.6.1). Traffic generation devices such as Spirent Communications' SmartBits cost around \$100,000 per unit and this facility will need at least two units.

To meet the requirement for rapid reconfiguration (see Section 4.5), intelligent patch panels are needed that can physically connect fiber from one device to another. APCON makes a physical layer switch that can switch OC-3 and OC-12 circuits. The cost of this device is about \$9,000 per unit.

If the facility requires that data and configurations be archived then the facility will need data storage devices such as tape backup systems or disk storage arrays. The cost of data archiving is not included in our estimates.

Other miscellaneous but costly equipment includes racks, monitors, intelligent patch panels, and KVM switches. A facility of this size will need approximately

- 50 racks at \$2,000 per unit,
- 100 monitors at \$300 per unit, and
- 50 KVM switches at \$700 per unit.

B.6 Labor

The facility will require an administrative staff, an engineering staff, legal counsel, and a technical advisory panel.

The administrative staff would minimally include a director to run the facility, marketing personnel to publicize and promote the facility, clerical staff, and an accounting staff to pay

expenses and collect fees from the clients. We estimate this cost at \$900,000 – 1,200,000 per year.

A facility of this size would require an engineering staff to aid with reconfiguration and keep the facility operating. Our research suggests that we would need four to eight full time engineers/system administrators that together will cost \$800,000 to 1,600,000 per year.

The legal counsel and the advisory panel could be kept on retainer to help with establishment of initial operating procedures and handle special circumstances as they arise.

Annual labor charges are expected to require \$1.5 to 3 million in funding.

B.7 Rent, Maintenance, and Replacement

This facility will occupy a significant amount of floor space, and will have special electrical and cooling requirements. A large hosting complex, an Internet exchange point, a University's computer science department, or a government facility would be able to host this facility. Monthly rent may be expensive and all options should be considered to help minimize this cost. We estimate this cost to be approximately \$1,000,000 per year.

Maintenance of all the hardware and software will incur additional annual costs. Depreciation and contingency accounts for new equipment should be setup and funded in advance to insure that the facility continues to provide an environment representing the state-of-the-art in networking. We recommend that the facility allocate 33% of the initial cost of the equipment per year to be used as a replacement and upgrade fund.

Annual charges for rent, maintenance and replacement costs are expected to fall between \$5 and 15 million. The cost of insurance against liability and equipment damage or loss is not included in these estimates, as it may or may not be relevant, depending on the facility's financial structure.

B.8 Design Option Cost Comparison

Our research has indicated the cost to initially build and operate the facility for the first year will vary significantly due to differences in each of the five design options (e.g., one location versus multiple locations, local telecommunication circuits versus cross-country circuits). This section is intended to give rough cost estimates for each of the five design options.

For comparison, we will use the following assumptions:

- Each facility location will have enough equipment to represent at least three *logical* sites.
- For Options 1 and 2, this means enough equipment to build three complete logical sites at a single physical location
- For Options 3, 4, and 5, the facility will have enough equipment and staff to equip three geographically dispersed test facility locations. In order to support independent experiments run at a single site as well as distributed, multi-site experiments, each of the three locations will be equipped with the same number of routers and hosts as the single site in Options 1 and 2.
- All prices used are estimated list prices and assume no discounts.
- The out-of-band management network will use 1.5 Mbps circuits for Internet or site-to-site connectivity.
- 160 Mbps circuits will be used for Internet access in Options 2 and 4.
- Option 5 will use 2.6 Gbps circuits to connect the sites. Option 5 will need to offer comparable throughput as a Tier-1 ISP if the facility is expected to attract live Internet traffic.

The table below summarizes the results of our calculations.

	Option 1	Option 2	Option 3	Option 4	Option 5
One Time: Equipment, Software, Design, and Build-out	12	12	34	34	34
Recurring costs:					
Labor	1.7	1.7	2.7	2.7	2.7
Network Connectivity	0.02	0.5	0.5	1.4	37
Rent, Maintenance, and Replacement	5	5	15	15	15
Total for 1st Year	19	19	53	54	90
Annual Recurring Costs	7	7	19	19	55

Table 6. Summary: Estimated Costs for Construction and Operation of the Facility (In Millions of Dollars)