THE BASTION NETWORK PROJECT

A Framework for Conducting Interscholastic Cyber-Exercises

J. D. Fulp

Naval Postgraduate School

Abstract: The Naval Postgraduate School's Center for Information Systems Security Studies and Research (CISR) has developed a small, but realistic network lab—the *Bastion Network*—that is dedicated to educating students in the myriad elements involved in the secure operation of a computer network. This paper describes the rationale for this network lab, and offers an overview of a simple framework that could accommodate educational network interaction with other schools that have similar IA educational goals, and that have, or may soon acquire, similarly designated labs. The framework describes the essential elements of a memorandum of understanding, and twelve suggested inter-network cyber-exercise scenarios.

Key words: Information Assurance, Network Security, Cyber-Exercise, Bastion Network, CISR

1. MOTIVATION

Operation of a network within acceptable levels of risk requires the proper installation, management, configuration, coordination and operation of all of its composite components; both hardware and software. The breadth, diversity and ever-changing design of these individual components, in addition to the additional complexity borne of their interconnected interaction, makes this a daunting task.

Though much of the knowledge required to perform risk mitigation of computer networks is liberally distributed in a host of texts and assorted online reference material; there is much less opportunity to employ and experiment with this knowledge in a practical environment. It is the purpose of this paper to describe a model framework for addressing this important Information Assurance (IA) education issue. This model program is called The Bastion Network Project (BNP). The name was chosen to convey the notion of a hardened, highly-defended, network; just as the term "bastion server" is used to convey the notion of an individually hardened server.

2. OVERVIEW

The BNP entails the construction and operation of a dedicated "exercise" network. Since a BNP network is devoid of operational information or resources, its attendants are free to experiment and learn without fear of adversely impacting their school's, business', or agency's

operational readiness. The BNP effectively provides a network security training "sandbox"; where devices, topologies, tools, configurations, applications, etc. can be added, modified, or deleted at will.

A BNP network can be as simple as one computer or as complex as resources and enthusiasm permit. Regardless of the initial size or topology of a BNP network, it serves as an accretion-like foundation for other technologies that a school may wish to add-on later. Such additions can be scaled and tailored to IA curriculum development and the availability of additional resources. Examples of such additions are PKI certificate servers, honeynets, wireless access points, and Kerberos security servers. In addition to the growth of an individual BNP network, adoption and participation in the BNP by other schools will enhance the quality and diversity of the educational experience for all participants.

The Naval Postgraduate School's Center for Information Systems Security Studies and Research (CISR), recently (June, 2004) completed the development of its own BNP network. I have had the pleasure of leading the development of this network, and hope, via this paper, to solicit the participation of other schools interested in a similar IA educational experience.

3. PREVIOUS RELATED WORK

In 2001, two of the U.S. service academies—the U.S. Military Academy and the U.S. Air Force Academy—and NPS participated in the first Cyber Defense Exercise (CDX) [1]. This exercise pitted roughly identically equipped networks at each of these three schools (Blue Teams) against assessment units (Red Teams) manned by select DOD agencies, in a network defense competition. The 2001 CDX was deemed such a success that it has been repeated in (so far) 2002, 2003, and 2004. Participants now include all of the U.S. service academies, the Coast Guard Academy, the Merchant Marine Academy, and the Air Force Institute of Technology (AFIT).

Though the CDX competitions have been a great success story, I believe that the competitive structure limits its viability for wider scale adoption and ease of year-round administration. In order to ensure fair and objective grading, the CDX organizers must require a fairly uniform network among all participants. This has been fairly easy for the service schools who have been the beneficiaries of DOD sponsored program money to provision their CDX networks with identical equipment. This will be virtually impossible; however, when trying to include other schools new to the cyber-exercise scene. This necessity for near-uniformity extends to the software as well. Generally, no school may install any software beyond the common installation unless it is free to the public. An additional hindrance to broader participation in a graded cyber-exercise program is the additional logistics required to obtain un-biased referees (White Teams) to be on hand to ensure adherence to all exercise rules of engagement (ROE). Competent Red Teams must also be identified and organized in such a manner as to ensure that equal assessment scrutiny is directed to each of the Blue Teams.

Another interesting form of cyber attack/defend exercise is the Capture the Flag (CTF) exercise [2] that is conducted at the annual DEFCON conference held in Las Vegas. Unlike the CDX, the CTF exercise has all participants engaged in both Red/attack and Blue/defense activities simultaneously. Also, very much unlike the CDX, the CTF style of competition gives each team no time to casually and methodically configure secure services on a secured network. Instead, participants are given an image with unknown services that must be installed, configured, de-bugged and patched while at the same time being potentially subjected to attack from other participating teams. The CTF might be thought of as the CDX in extreme fast forward with both attack and defend roles combined into one. Though this may well be the ideal venue to showcase existing IA skills, it is not the ideal venue for methodical learning by novices.

It is for the above reasons that the idea of the BNP came into existence. The BNP's merits are five-fold. 1) Virtually any school can participate, regardless of resources. 2) No school needs

to fret over the publishing of an embarrassingly low score. 3) Schools need not disrupt the configuration of any existing labs that they would like to utilize for BNP participation since virtually any configuration is acceptable. 4) As participation in the BNP grows, each school is rewarded with more diversity in the networks they can attack/evaluate, and in the diversity of received attacks they can observe. 5) Schools do not need to expend resources to arrange for a Red or White Teams.

4. ADMINISTRATIVE COORDINATION AND STANDARDIZATION

Once two or more BNP networks have been designated, it is a relatively simple matter to coordinate exercise details between the two participating schools. These details would include such items as: secure channel (i.e., VPN) setup, start/stop times and dates, roles (attacker or defender), ROE, and learning objectives. As an additional ease-of-administration feature; however, this author will suggest the drafting and publication to the Web of twelve pre-defined BNP scenarios and a memorandum of understanding (MOU). These twelve scenarios (described below) and MOU will serve as the necessary administrative coordination vehicle between participants. Though the exact content of these scenario descriptions and MOU have not been finalized at the time of this writing, the generalized contents and descriptions follow.

The MOU will contain four main elements. 1) A statement regarding the intent of BNP participation. 2) Elaboration regarding the mandatory implementation of a secure VPN tunnel between the participating BNP networks. 3) Delineation of cyber-exercise ethical conduct and ROE. 4) A statement indicating that each side has notified their local IT authorities regarding the exercise, and that each side has taken measures to ensure that their BNP network activities will not adversely hinder routine network operations at their school.

Regarding the VPN connection, the MOU will mandate that no split-tunneling is permitted. It will mandate that the connection utilize the ESP protocol in Tunnel mode, with DES or AES encryption mandatory. And it will mandate that the VPN gateway be the only access point into or out of the network.

Regarding ethical conduct and ROE, the MOU will (among other things) reiterate the importance of a cryptographically (i.e., within the VPN tunnel) and physically (i.e., no network connections aside from the VPN gateway) confined exercise. The MOU will proscribe the employment of worms or viruses of any kind, and will mandate immediate cessation of activities and removal of the BNP network from the larger internetwork in the case of an expected "spill" of exercise-related traffic outside the exercise boundary. The MOU will document both the IP addresses of each network's external (i.e., public-facing) VPN gateway interface, and at least one phone number for out-of-band exercise coordination.

5. TWELVE BNP CYBER-EXERCISE SCENARIOS

The following twelve scenarios were defined in order to provide a mix of exercises that include attack-only and defend-only scenarios, along with varying degrees of difficulty. Scenario #12 in not a specific scenario, but rather a place-holder that offers participants with the opportunity to agree upon their own personalized cyber-exercise agenda. The duration for any of these scenarios is at the discretion of the participants, and will be specified in the MOU. Suggested time frames are on the order of one to five days. The first eleven scenarios will be presented in order of the least to the most complex.

Scenario #1—**Attack with no perimeter to soft systems**: This scenario gives the attackers not only unfettered access to whatever systems may reside on the defender's network, but also presents them with intentionally un-patched and/or mis-configured systems that may ultimately

yield a successful penetration. The primary learning objective is to understand and employ the full range of hacker tools, practices and techniques.

Scenario #2—**Defense with no perimeter and soft systems**: This scenario represents the opposing side of scenario #1. The defenders will intentionally leave their systems exposed, and will intentionally leave several exploitable code flaws or configuration errors for the attackers to pursue. The intent is to observe the attack without responding to it. The primary learning objective is to practice intrusion analysis and post-attack recovery procedures.

Scenario #3—**Attack with no perimeter to hard systems**: This scenario presents the attackers with a much greater challenge: attacking hardened systems. No filtering is employed to restrict access to the systems, but the systems themselves have been hardened to the best of the defenders' capabilities. The most likely successful attack will arise from a recently announced flaw that the defenders have not yet corrected. Much more difficult, but also possible, is the development by the attackers of a "zero-day" exploit (i.e., a newly discovered, and thus not yet published exploitable code flaw). The primary learning objective is to understand and employ vulnerability assessment tools, practices and techniques.

Scenario #4—**Defense with no perimeter and hard systems**: This scenario represents the opposing side of scenario #3, and like scenario #2, the intent is for the defenders to simply observe the attack without responding to it. The primary learning objectives are to understand the process of system hardening, and to practice intrusion analysis.

Scenario #5—**Attack through perimeter to hard systems**: This scenario significantly increases the challenge for the attackers. It is expected that if the defenders do a fair job of configuring their perimeter defense (likely a firewall or filtering router), the attackers will have little success unless the scenario is allowed to run for several days. The attackers must essentially attack the target systems through tiny "holes" that they might identify in the perimeter defense. A denial of service (DOS) attack against the perimeter system is permitted so long as the attackers can perform it in such a way that they do not also prevent their own access to the protected systems. The primary learning objective is to understand and employ the full range of penetration testing and vulnerability assessment tools, practices and techniques.

Scenario #6—**Defense with perimeter and hard systems**: This scenario represents the opposing side of scenario #5. The primary learning objectives are to understand the process of system hardening and the employment of perimeter defense tools, practices and techniques.

Scenario #7—**DOS attack on hardened network**: This scenario greatly increases the range of tools available to the attacker, as attack methods intent on simply "breaking" systems or processes are far more common and accessible than those intent on more covert policy violations involving surreptitious data theft or modification, or obtaining root access. The primary learning objective is to understand and employ the full range of tools, practices and techniques that are known, or suspected, to cause disruption of normal system service.

Scenario #8—**DOS defense with hardened network**: This scenario represents the opposing side of scenario #7. Effectively the defenders are holding their breath and hoping for the best. The attackers have a wide variety of attack types and methods to employ, including protocol-violations, un-checked parameter entries, and excessive volume of network traffic. The primary learning objectives are to understand the process of system hardening and the employment of perimeter defense tools, practices and techniques.

Scenario #9—**Concurrent attack/defense with no perimeter**: This scenario effectively combines scenarios #1 and #2 (soft systems), or alternatively, scenarios #3 and #4 (hard systems). Both sides will agree on whether the systems will be hardened or not. The primary learning objectives are the same as for scenarios #1 and #2, or #3 and #4, as appropriate.

Scenario #10—**Concurrent attack/defense with perimeter**: This scenario effectively combines scenarios #5 and #6. The primary learning objectives are the same as for these two scenarios.

Scenario #11—**Concurrent DOS attack/defense**: This scenario effectively combines scenarios #7 and #8. The primary learning objectives are the same as for these two scenarios.

Scenario #12—**Ad Hoc**: This scenario is simply a place-holder to identify the scenario that will be completely defined, and agreed upon, by all of the BNP participants involved. This scenario can be tailor-made to fit any specific learning objectives that a school may wish to pursue. Some possibilities include combinations of the other scenarios. For example, schools X and Y may want to conduct a three day exercise. On day one, X will attempt to subvert Y's perimeter (scenario #5 for X, scenario #6 for Y). On day two, Y will remove the perimeter to give X greater access (scenario #3 for X, scenario #4 for Y). On day three, X will launch an all-out DOS attack on Y (scenario #7 for X, scenario #8 for Y). The primary learning objectives will be defined by the participants.

6. BNP NETWORK DESIGN

The last point of discussion is the design of the network itself. As previously mentioned, one of the merits of the BNP is that virtually any network design can be used. In this section, four basic building-block designs are briefly presented in order of simplest to most complex. At a very minimum a BNP network must have the ability to tunnel traffic in an IPSec-based VPN, and it must have at least one computer that can be used as a platform to host either attack (assessment) tools, or at least one service and/or operating system (OS) that can be targeted.

The simplest BNP network design (design #1) is one which consists of a single computer that is both the VPN endpoint and hosts either attack tools or a defended OS/service. This means that any school with at least one moderately capable computer can get participate in the BNP.

The next logical step (design #2), is to employ a dedicated VPN gateway in addition to a single computer that will host the attack tools and/or at least one OS/service. Depending on the operating system, CPU, and memory of this single computer, it could potentially host several services to defend, in addition to serving as the launch point for various attack/assessment tools.

Network design #3 introduces a LAN access device (hub or switch) and more computers hosting services. This permits the more secure deployment of services on separate computers (a security best practice outlined in [3]), and permits the operation of a designated IDS computer to monitor exercise traffic.

Network design #4 introduces a dedicated perimeter defense device, typically a packet-filtering router, a dedicated firewall appliance, or a general purpose computer that has two interface cards installed and is running a software-based firewall product. With this design, the arrangement would be as shown in figure 1 below. Note that the attack computer(s) is/are placed "outside" of the protected perimeter in order to preclude having to open holes in the firewall to support the various mischievous packets crafted by the attack tools. This design also raises an interesting IDS question: should the IDS sensor be placed in front of or behind the perimeter? Due to the relatively low volume of traffic expected during a typical exercise, it is recommended that the sensor be place outside of the perimeter in order to better understand what the attackers are attempting to accomplish. Otherwise, an effectively implemented firewall will block most, if not all, of the interesting traffic that schools will want to observe and study.

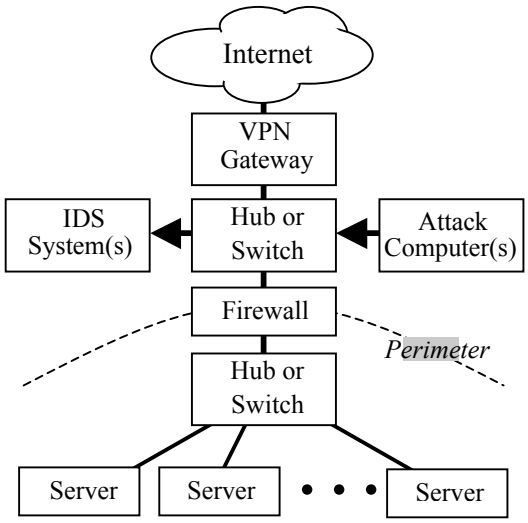

Figure 1. Basic BNP Network Design #4

Other network design issues include: IP space, network address translation (NAT) across the VPN gateway, VPN split-tunneling, and DNS (name service).

The simple solution for choosing IP space within a BNP network is to use any of the private IP address ranges specified in RFC1918 (e.g., 10.*any.any.any*, or 192.168.*any.any*). It is then only necessary that every BNP network utilize a separate "branch" of the private IP space tree. For example, the NPS BNP uses 10.1.*any.any*, so other schools might choose 10.>1.*any.any* for their networks.

With respect to NAT, any school that wants to connect their BNP network to a public network—perhaps during pre-exercise preparation—should enable NAT on their VPN gateway device, so as to convert their private (behind the gateway) IP addresses to static public IP addresses provided by their school's IT department or ISP as the case may be.

Split-tunneling; wherein some traffic leaving a network enters an encrypted tunnel, while some does not, should be disabled. This is to reduce the likelihood that malicious software might accidentally leave the confines of the exercise and send harmful packets to non-participating destinations. The details of how to disable split-tunneling will vary from device to device, but the basic remedy is to assign an ACL (access control list) rule to the VPN gateway device that will explicitly deny all IP traffic that does not have a destination network IP address equal to that of the co-participant's network.

Finally, for realism, each BNP network should install a DNS server in order to resolve local server names when queried by machines from outside the local network, and to cache name resolutions learned as the result of local queries to outside BNP network DNS servers. Since the BNP networks are connected within their own extranet, at least one of the BNP networks' DNS servers must be designated as the root DNS server. Currently, the NPS BNP is designated root for the <dot>bnp domain.

7. SUMMARY

The Bastion Network Project is intended to provide the framework for the safe, easy, and effective conduct of interscholastic cyber attack and defend exercises. The construction, defense,

and attacking of these exercise networks will provide real-world experience within a relatively small and contained environment. The non-competitive nature of the BNP frees participants from many of the resource and logistics constraints that might preclude them from participating in competitive cyber-exercises. A co-signed MOU, connectivity through a VPN, and a collection of generalized scenarios, should facilitate simplified administrative coordination between interested participants. The individual networks required to support participation in the BNP can be as simple as one computer, or as complex as each school desires. Readers interested in pursuing the BNP at their schools, should call 831-656-2280, or follow the *Bastion Network Project* link at www.cisr.nps.navy.mil.

REFERENCES

[1] Robert Lemos, "Training the Cyberwar Troops," in ZDNet, [online magazine] (2002 [cited 18 May 2004]); available from World Wide Web @ http://zdnet.com.com/2100-1105-893418.html

[2] Robert Lemos, "Hacking contest promotes security," in CNET, [online magazine] (2003 [cited 18 May 2004]); available from World Wide Web @ http://news.com.com/2100-1009_3-5059827.html

[3] National Security Agency, Systems and Network Attack Center, The 60 Minute Network Security Guide, [database online] (updated 12 July 2002 [cited 18 May 2004]); available from World Wide Web @ http://nsa1.www.conxion.com/support/guides/sd-7.pdf