CANDI: A System for Classifying the Security Risks in Network Appliances

Joseph G. Tront
Electrical & Computer Engineering
Virginia Tech
Blacksburg, VA USA 24061-0111
jgtront@vt.edu

Randy C. Marchany
Computing Center
Virginia Tech
Blacksburg, VA 24061
marchany@vt.edu

Abstract

Internet security is of utmost importance in today's e-commerce environment. Many different avenues are being taken in an attempt to secure the systems of both the end user as well as the server of information. The operating system, as well as application software, provide holes through which security is breached. One vulnerable part of the system that has not received much attention is the hardware. This study examines how various Internet appliances can be classified according to their vulnerabilities.

1. Introduction

Businesses as well as consumers are being directly impacted by breaches in network security measures. One of the major costs in operating a computing and communications environment involves the specification, installation and maintenance of effective security measures that preclude intrusion. Both commercial systems managers and consumers are continually seeking solutions that provide confidence in the safety measures that exclude network trespassers. These solutions must be relatively low cost, low maintenance, robust, and reliable. [1]

Most Internet security efforts focus on securing the operating system and the application software. While this type of effort is completely necessary, it leaves out one of the other critically vulnerable components of the system, the hardware. It further leaves out an entire class of computing/communication devices that do not contain operating systems or typical user application software. These devices are generally referred to as network appliances. Examples of a network appliance include simple devices like the parking gate that checks the users identity by reading an id card and verifying the validity of the information by sending it over the Internet to a central database system. See figure 1. The home lighting system that can be controlled remotely by sending a few commands over the Internet is also considered to be a network appliance. Other more recognizable devices like network routers, bridges, wireless hubs, etc. are also network appliances. Each of these pieces of equipment is potentially susceptible to attack. Each device may also be capable of being used to launch a malicious attack, unbeknown to the owner and normal user.

The CIRT Appliance & Network Defense Initiative (CANDI) is a term we use to describe a testing facility and the procedure followed to rate network appliances based on their susceptibility to intrusive and malicious activity. (CIRT is the Computer Incident Response Team). Persons responsible for purchasing network appliances can use the ratings to distinguish between more or less

Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)
0-7695-1874-5/03 $17.00 © 2002 IEEE

secure networks. The rating system includes various measures describing characteristics such as: out-of-box susceptibility, ease of maintaining the level of security, and expertise required for security initiation and maintenance. A date is associated with a rating indicating the relative validity of the evaluation.

Figure 1. A simple parking gate is a potential target for a security attack.

In order device a protection plan, the tactics used by the attacker must be understood. This paper will discuss some of the typical attack strategies followed by a description of the process planned to classify devices based on their ability to fend off known attack mechanisms. [2, 3]

2. Attack Strategy

In order to be able to classify the vulnerability of a particular device, it is first necessary to understand the possible strategies that may be used to attack the device.

Typically the first step in attacking a target network is the collection of information about the target device to form a *profile of the target.* Once the device is characterized, the attacker generally proceeds to *identify exploits* for the device. Generally, the final step of the strategy is to *perform the exploit* on the target. Sophisticated attack mechanisms may go one step further and analyze the progress of the attack and *report back* to the perpetrator for either self-satisfaction, material gain, or in order that future attacks can be improved. This last step occasionally proves to be the undoing of the attacker and is generally either very carefully done or is not done at all.

There are three major forms of attacks: logical, physical, and social engineering. These categories will be used to describe the formation of the profile, the exploits and the subsequent performance of the attack.

2.1 Profiling the Target

Logical profiles of a target are generally performed through the use of software capable of collecting critical information about an operating network device. For example, if the device is a general purpose computer, part of the profiling operation consists of collecting information describing the type of operating system, the revision level, the processes running on the system, etc. For devices not running an operating system, the profiling is slightly more complicated. In these cases the attacker generally tries to ascertain network address information, MAC address, open network ports, and any other system characteristics that the device is willing to report back.

The most useful tool to perform profiling is the port scanner. This software package is used to remotely analyze the responsiveness of a potential target and to initially identify possible vulnerabilities. Port scanner programs like *Nmap* can be used to identify unsecured ports, but also may be customized to probe further to gain details about operative processes, MAC address, etc. Programs such as *Nessus*, *Saint*, and *Sarah* not only perform port scanning, but

Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)
0-7695-1874-5/03 $17.00 © 2002 IEEE

they are also capable of reporting the best known exploit that can be used to attack the target system. Port scans are generally done in the open, but some port scanners attempt to glean information in stealth mode. This tactic somewhat complicates detection and defense mechanisms.

Physical profiling of a device can also occur wherein the perpetrator attempts to identify a physical vulnerability that allows the device to be disrupted or taken over for illicit purposes. An example might be a person identifying when an office is open in order to gain access to a persons unprotected computer keyboard or server console. Although this form of exploit can be effective, this paper will not delve into this aspect of security.

Another category of attack involves social engineering. In this type of profiling, the attacker identifies how a system owner or user may be manipulated so as to allow an incursion to take place. An example might be one in which an attacker identifies a situation where employees have not received raises for an extended period. The attacker leaves a floppy diskette labeled "Salary Information" in a conspicuous and unsecured location in the office. Employees are likely to insert the diskette in their system and unknowingly run a program that provides little interesting information to them, but does install a back door access to their system. Again, although these types of attacks are common and potentially devastating, they generally do not directly affect network appliances and will not be part of the focus of this paper.

2.2 Identify Exploits

Once the device is characterized, the attacker must decide on which exploits to attempt to use. Numerous exploits are published on various web sites and new ones are devised daily. Even those working to protect their systems provide exploit information for attackers as they publish their findings on attempts to breach their systems. Attackers can find database information relating vulnerabilities to specific programs and scripts that can be used to attack profiled systems. Sites such as Securityfocus.com and Packetstormsecurity.com can be used by the less experienced attacker to obtain pre-packaged attack mechanisms. More experienced attackers modify already written programs and generate their own code to perform particularly insidious assaults.

The database of attacks against network devices is not very extensive. However, the number of attacks on these devices continues to increase and it is very important that an effort be devised to deal with possible future attacks.

2.2 Perform Exploits

Performance of attack exploits can proceed in several different ways. For example, the attacker may perform a denial of service attack in which the network appliance is flooded with network interaction requests or commands. In this way, the appliance spends most of its time responding to inane or unrecognized communication and is unable to perform its normal operation. A similar attack may consist of an attacking device masquerading as the target device and providing equally inane responses to legitimate network traffic. This type of attack is a form of electronic vandalism and is popular among the less experienced attacker.

Another example of an exploit in action is a situation where an attacker is able to set up a backdoor access so that the invader is able to access a system and extract whatever information they like. This may consist of client files, credit card information, password information for other system, etc. In many cases this type of exploit is done in a stealthy manner and the system owner may never know they have been invaded. This form of attack

generally goes beyond vandalism and is more likely to result in some form of larceny.

A final example of an exploit is one in which the target machine is invaded and a program is placed on the system from which subsequent attacks are launched. The growth of this type of exploit can be exponential and can bring the Internet to a state of much degraded performance. This type of exploit has become popular these days and is quite difficult to terminate.

3 Identifying Risks

Network appliances encompass a growing group of devices that one way or another communicates via the Internet. Examples include: network printers and scanners, DSL modems, cable modems, routers, home environmental control systems, building security systems, personal identification systems, and many other devices that are being put to work in varied applications every day.

Although network appliances are less likely to be attacked than servers or client computers, there is still a significant potential for security problems in these devices. CANDI's approach to establish a program of safe use of network appliances is similar to the well-known UL approach to safety for electrical devices. The strategy involves classifying the individual network appliances according to their known susceptibilities and provide the user with a seal of approval (disapproval) that describes the level of safety that a device exhibits or has the potential to exhibit.

Evaluating network appliances against known security threats is a somewhat straightforward task. Devices, as delivered by the manufacturer, are checked logically to determine if they are susceptible to specific sets of threats. For example, a network router as delivered by the vendor may have no access password and a common well-known IP address. This constitutes a vulnerability that could allow an attacker access and control if not corrected. Devices with this characteristic are entered into a database and the threat type/level is noted.

Going further with the example above, some routers delivered with no password, or with well-known passwords, are not able to have their passwords changed. This is a more serious vulnerability that leads to a database entry noting this class of potential problem. Other routers that come with no passwords or well-known passwords and are able to have their passwords changed fall into a different category of vulnerability and are so classified in the database. The ultimate vulnerability of this last class of device will ultimately be determined by the user and their diligence in applying the needed security measures.

One of the best ways to identify risks is to measure devices against well-established points of reference. The CANDI project uses the following lists of known threats to test network devices:

1. Center for Internet Security (CIS) Security Benchmarks
2. SANS Top 10 Threats
3. SANS Top 20 Threats

These lists have become industry standards and are kept up-to-date making them a sound basis for testing. Devices are individually rated and a total risk score is established. The risk score can be employed by the user to make decisions on purchases or to better understand the security maintenance requirements of a device. We envision that eventually products will be branded with a CANDI security seal of approval much like the UL branding scheme. The seal of approval will convey three pieces of information: the date on which the testing

Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)
0-7695-1874-5/03 $17.00 © 2002 IEEE

established the device security status, the level of security, and the ease with which the device may be secured. The level of security branding includes the following information:

Level 1 – can affect security or operation of other sites
Level 2 – only a threat to the appliance itself
Level 3 – hardened against known threats
Level 4 – totally insensitive to all threats (not network accessible)

The ease of securing branding includes the following information:

A – easily secured/updated with regular software patch
B – patches more difficult to apply
C – system admin. knowledge required to update
D – history of poor update support or very difficult patching
F – updating not possible

These designations provide the user with a rich set of information that can be updated as the device continues in its life cycle. Shown in Figure 2 is a conceptual view of how the user will initially get a view of the CANDI security information. Details of the definitions of the CANDI ratings are placed on a web site. The web site also contains information that describes how the user may improve the security of the device for those devices whose security can be improved.

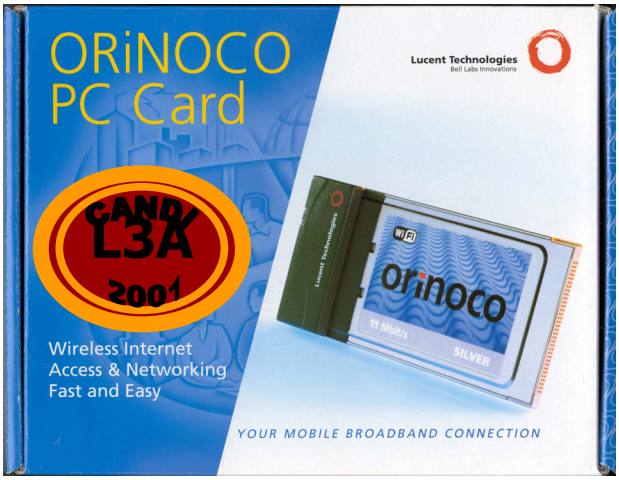

Figure 2 Branding scheme describing network appliance security status. Device is labeled as level 3A as of known threats in 2001.

3 Summary

The basis for CANDI has been established and devices are being examined and classified. Discussions are being held with various groups including security researchers, systems managers, manufacturers, and users to determined how this process can be improved.

4 References

[1] Householder, A., Houler, K., Dougherty, C., Computer Attack Trends Challenge Internet Security, *Security & Privacy*, IEEE Computer Society Press, June, 2002.
[2] Stajano, F., *Security for Ubiquitous Computing*, John Wiley & Sons, Chichester, UK, 2002.
[3] Sammes, A.J., Jenkinson, B., *Forensic Computing – A Practitioner's Guide*, Springer-Verlag, London, 2000.

Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)
0-7695-1874-5/03 $17.00 © 2002 IEEE