Towards End-to-End Security

Thomas M. Chen
Dept. of Electrical Engineering
Southern Methodist University
PO Box 750338
Dallas, TX 75275-0338 USA
Tel: 214-768-8541
Fax: 214-768-3573
Email: tchen@engr.smu.edu
Web: www.engr.smu.edu/~tchen

1. Introduction

The Internet has been invaluable for facilitating the Web, e-mail, file sharing, and other communications services that we use every day. Unfortunately, the past few years of major virus and worm epidemics have demonstrated that ubiquitous Internet connectivity has a major downside. The communal nature of the Internet exposes organizations and home computer users to a multitude of malicious code threats. In the same way that germs are quickly shared among members of the same household, malicious code can spread rapidly among networked computers. Viruses and worms from Melissa in 1999 to Netsky and MyDoom in 2004 have become a prevalent problem for Internet users. In the 2004 CSI/FBI Computer Crime and Security Survey, 78 percent of the surveyed organizations reported being effected by viruses and worms [1].

A recent report found that 40 percent of Fortune 100 companies were infected by worms in the first half of 2004 [2]. The finding is significant because Fortune 100 companies are presumably well protected by firewalls, antivirus software, intrusion detection systems, and highly trained IT staff. After previous experiences with the devastating Slammer and Blaster worms in 2003, organizations have learned to guard against new worm outbreaks. However, current security measures are evidently not completely successful.

The inadequacy of current security is due in part to the difficulty of protecting hosts by antivirus software and software patching. New vulnerabilities in operating systems and applications are being continually discovered at an average rate of 48 per week [2]. Of these vulnerabilities, 46 percent are rated as highly severe (i.e., could lead to complete compromise of a computer), and 69 percent are considered easy to exploit. Moreover, an exploit appears within 5.8 days after a new vulnerability is announced. This means that organizations must patch frequently and apply new patches within a few days. The need for frequent patching places a considerable strain on large organizations with many computer users.

In addition, about 180 new Windows viruses and worms are being discovered each week on average [2]. Antivirus software signatures require constant updating to detect new viruses and worms. Many users neglect to update their antivirus software due to practical reasons. Antivirus software with outdated signatures may leave a computer vulnerable to new worm attacks.

Another reason for the inadequacy of current security is the ineffectiveness of firewalls after a network perimeter has been penetrated. Today, many avenues are possible to bypass perimeter defenses. For example, a notebook brought into an organization could be carrying malicious code. It takes only a single infection within an organization to effectively circumvent perimeter security.

The inadequacy of current security is worrisome for a number of reasons. First, the practical difficulties in constant patching and updating antivirus software means that organizations will likely contain a substantial number of inadequately protected computers. With weak defenses, organizations may experience a significant number of infections from new outbreaks. That could leave organizations with sizable expenses for recovery. In addition, even uninfected computers can be effected by the probing traffic or other side effects caused by a new worm outbreak. Therefore, it is not surprising that security experts are seeking better approaches to contain or quarantine new outbreaks.

2. End-to-End Security Approaches

Current security products - antivirus software, firewalls, intrusion detection systems - are designed as stand-alone pieces of equipment or software, for the host or the network. The basic idea of end-to-end security is to coordinate the security mechanisms on both hosts and network. The network and hosts are viewed as a single system to be secured, instead of separate systems. The network enforces a security policy by regulating network access of hosts; hosts provide information about themselves in terms of their security state, e.g., antivirus software state and operating system patch level. Through this coordination, hosts can verify they are secure and compliant with policy before they are allowed to access the network (with potential impact on other systems). This careful admission of trusted hosts helps to ensure that the overall health of the entire network is preserved.

End-to-end security may be viewed as automation of the manual process used by system administrators to identify and isolate vulnerable systems within their networks. System administrators can manually inspect the configuration of hosts to verify compliance with security policy. Hosts can also be tested using any number of vulnerability scanners to identify potential weaknesses. However, system administrators are challenged to keep track of mobile devices or home computers, or user installed software. For these reasons, automation of the process is highly desirable, especially for large organizations.

The first attempts at end-to-end security most likely began with organizations enforcing antivirus levels before logon. A host is initially restricted to a demilitarized zone (DMZ), a small subnetwork that sits between the trusted internal network and the Internet. During the quarantine, the host’s antivirus signature update level can be confirmed. If the antivirus is up to date, the restrictions on the host are lifted.

It is natural to extend this basic idea to scan hosts for more security information, such as operating system patch levels, known vulnerabilities, and malicious code infections. Until recently, security products were not designed to work together to accomplish this end-to-end security. A number of comprehensive approaches have appeared lately, including Cisco’s Network Admission Control and Microsoft’s Network Access Protection. They have attracted the most attention but are not the only strategies being offered. StillSecure’s Safe Access is an example of a simpler alternative approach. They are described below as examples of end-to-end security approaches without implying any endorsement for any commercial products.

3. StillSecure Safe Access

StillSecure’s Safe Access is aimed to protect the network by ensuring that hosts are free from threats and compliant with a security policy [3]. Safe Access systematically tests hosts for compliance with organization-defined security policies, quarantining non-compliant machines before they damage the network.

The process begins with definition of access policies that define which applications and services are permitted, as well as specify the actions to be taken for non-compliant devices. Safe Access automatically applies access policies to devices as they log onto the network.

Access policies actually consist of individual tests to evaluate the security status of hosts. The Safe Access appliance uses RPC (remote procedure call), a native Windows protocol, to interrogate machines trying to gain network access. Unlike some other approaches, it is not necessary to install a software agent on hosts. Safe Access can verify a number of antivirus and firewall programs are running, and includes an API to work with other programs as well. Specific tests assess operating systems; verify that key hotfixes and patches have been installed; ensure anti-virus and other security applications are up to date; detect the presence of malicious code (viruses, worms, Trojan horses); and check for potentially dangerous applications such as file sharing or spyware.

Based on test results, hosts are either granted or denied network access or quarantined to a specific part of the network. Hosts deemed compliant with the access policies are permitted access to the network. After the initial compliance tests, Safe Access continually tests hosts that have been granted access to ensure that real-time system changes do not violate access policies. Non-compliant hosts are either denied access or quarantined to give them a chance to verify their trustworthiness.

4. Cisco Network Admission Control

Cisco has enlisted antivirus companies McAfee, Symantec, and Trend Micro for its Network Admission Control (NAC) program up to October 2004. Naturally, Cisco's dominance of the data networking market gives it considerable leverage to promote acceptance of NAC.

The name "admission control" is not entirely appropriate. Admission control has traditionally been used in the context of traffic control to refer to the acceptance or rejection of new traffic to prevent congestion in the network. In NAC, traffic is accepted or rejected depending on compliance with security policies to protect the general population from vulnerable or non-compliant individuals.

As shown in Figure 1, NAC consists of Cisco Trust Agents; network access devices; and the Cisco Secure Access Control Server (ACS) [4].In contrast with StillSecure's agentless Safe Access, NAC requires every host to be running a Cisco Trust Agent. The Cisco Trust Agent is a small software program to communicate with the Cisco Secure ACS. On the host, the Cisco Trust Agent collects security state information from multiple security software clients, such as antivirus clients or the Cisco Security Agent. The Cisco Security Agent is a host-based intrusion prevention system and distributed firewall that identifies and blocks malicious behavior. The Cisco Security Agent is aware of the operating system version and patch level. After collecting security state information, the Cisco Trust Agent sends the information to network access devices, which forward it to the Cisco Secure ACS.

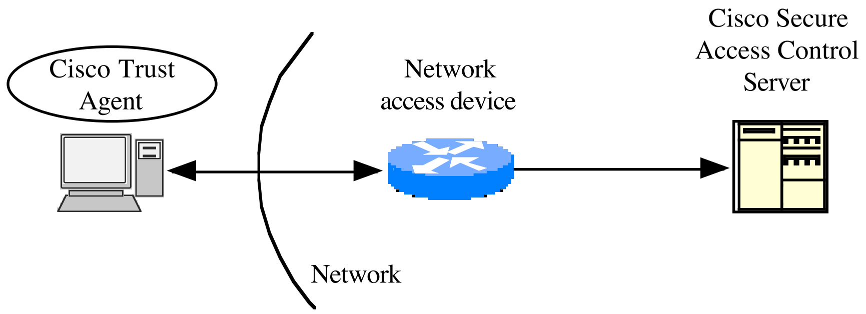

Fig. 1. Cisco NAP.

Network access devices include routers, switches, wireless access points, and security appliances. These devices enforce an admission control policy by requiring security "credentials" from hosts which are relayed to the policy server (Cisco Secure ACS). After the policy server makes an admission decision, the network devices enforce the control decision. As usual, routers use access control lists (ACLs) to restrict traffic.

The admission control decisions made by the Cisco Secure ACS can be: permit, deny, quarantine (temporarily to a special subnetwork), or restrict the network access (to a subset of services). Decisions are made by applying a security policy to the security state information collected from hosts.

In the future, hosts might be able to report misuse in real time, such that NAC could perform dynamic quarantining during an attack. However, there must be means to guarantee that intelligence communicated from hosts can be trusted. During an attack, it is possible that attacked hosts could be compromised and provide false intelligence. Because NAC depends on coordination between the network and hosts, it might fail if hosts are untrustworthy.

5. Microsoft Network Access Protection

Microsoft is planning Network Access Protection (NAP) as a new set of operating system components for Windows Server "Longhorn" that constitute a platform for protected access to private networks [5]. Microsoft has enlisted at least 27 partners for NAP, including antivirus companies, intrusion detection equipment vendors, endpoint policy management and enforcement companies, management and patch management organizations, network equipment vendors, and systems integrators.

The general idea of NAP is to detect the security state of a host attempting to connect to a network and restrict the access of the host until the policy requirements for connecting to the network have been met. NAP carries out this objective through a set of functions:

- inspection: when a host attempts to connect to a network, its "health state" is validated against the network access policies defined by system administrators

- isolation: non-compliant hosts are denied access or restricted to a quarantine network, depending on policies
- remediation: problems causing non-compliance are resolved by updating non-compliant hosts with the missing requirements.

The initial version of NAP will support enforcement of Dynamic Host Configuration Protocol (DHCP) address configuration and virtual private network (VPN)-based network connections. DHCP servers can enforce network access policies whenever a host attempts to lease or renew an IP address configuration on the network. VPN servers can enforce network access policies when hosts attempt to make a VPN connection to the network.

As shown in Figure 2, the elements in the NAP architecture include:

- VPN servers allowing VPN-based remote access connections to a private network
- DHCP servers providing automatic IPv4 address configuration to hosts
- IAS (Internet authentication service) server providing network access policy checking for DHCP or VPN clients (perhaps running on a VPN or DHCP server, or a separate server)
- Active Directory database storing user accounts and their credentials (network access properties)
- separate quarantine network for non-compliant hosts containing resources to remediate DHCP or VPN hosts
- NAP server enforcing restrictions on network access for DHCP or VPN clients
- policy server containing resources, such as antivirus signatures and software updates, to keep hosts compliant with security policies and provide remediation for non-compliant hosts.

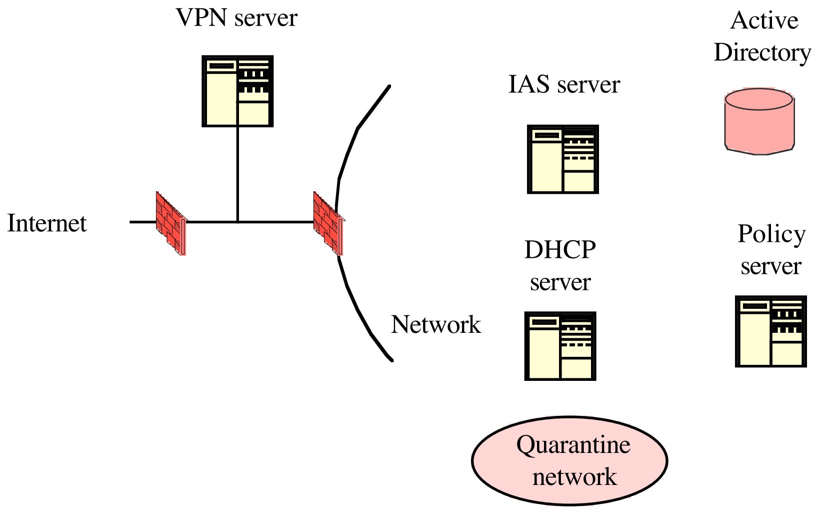

Fig. 2. Microsoft NAP.

A NAP host acting as a DHCP client uses DHCP messages to request a valid IP address configuration from the DHCP server. Its request should include an indication of its current system state called a statement of health (SoH). A host without a SoH is automatically deemed to be non-compliant. If the SoH is valid, the DHCP server assigns an appropriate IP address to the host for normal network access. If the SoH is not valid, the host is deemed non-compliant, and the DHCP server isolates the host into the quarantine network.

In the quarantine network, an isolated host reports its status to the policy server and requests updates for remediation. The policy server provisions the host with the required updates (e.g., antivirus signatures and software updates) to bring the host into compliance with policy. When the host updates its SoH, it can send another DHCP request to the DHCP server, which forwards the updated SoH to the IAS server. After the IAS server validates the SoH, the DHCP server grants normal access to the host.

A NAP host acting as a VPN client uses PPP (point-to-point protocol) messages to establish a VPN connection. The VPN client sends its authentication credentials to the VPN server using Protected Extensible Authentication Protocol (PEAP). If the authentication credentials are valid, the VPN server requests a SoH from the host. A host without a SoH is classified as non-compliant. If a host has a SoH, the SoH is passed to the VPN server which then forwards it to the IAS server. The IAS server communicates with the policy server to determine whether the SoH is valid. A SoH is valid if it matches the list of components and configurations that the policy server requires. If the SoH is valid, the VPN server completes the connection and grants the VPN client normal access to the network. If the SoH is not valid, the VPN server completes the connection but isolates the VPN client into the quarantine network.

An isolated host can send traffic only to the quarantine network, VPN server, and policy server. It reports its status to the policy server and requests updates for remediation. The policy server provides the VPN client with the required updates to bring the host into policy compliance. When the host updates its SoH, the host can send it to the VPN server through a PEAP exchange. After the IAS server validates the updated SoH, the VPN server grants normal network access to the VPN client.

The policy server is also used for compliant hosts. While a NAP host has normal access to the intranet, it accesses the policy server regularly to ensure that it remains healthy. For example, the NAP host periodically checks the policy server for the latest antivirus signature file or the latest operating system updates.

Finally, NAP is intended to be an extensible platform, rather than a comprehensive solution. In addition to operating system components, APIs allow expansion of functions and interoperability in the future through additional third-party components.

6. Conclusions and Open Issues

End-to-end security represents a new approach to coordinating host-based and network-based defenses in order to safeguard the integrity of the overall community. It is not intended to be a comprehensive security solution for all types of threats. Its sole objective is to identify non-compliant or vulnerable hosts and prevent them from effecting the shared network. Examples of end-to-end security approaches have been described only as examples and are not intended as endorsement of any products.

The practical issue of interoperability is important. If commercial products can not work together, system administrators will be challenged to choose among different vendors or manage a heterogeneous mixture of products. Significantly, Microsoft and Cisco have announced plans to work together starting October 2004. They will share APIs and develop protocols for interoperability between Microsoft's NAP and Cisco's NAC. Since both NAP and NAC are designed to allow organizations to enforce security policies on endpoint devices, they have been seen as competing approaches.

Another practical issue is dependence on honest cooperation of endpoint devices. None of the current end-to-end approaches try to address the potential problem of malicious users who may possess proper security credentials. Current proposed end-to-end approaches may fail if malicious users can present misleading credentials and be accepted by the network as trusted hosts.

References

[1] L. Gordon, M. Loeb, W. Lucyshyn, R. Richardson, "2004 CSI/FBI Computer crime and security survey," available at http://www.goscsi.com.

[2] D. Turner, et al., "Symantec Internet security threat report: trends for January 1, 2004 - June 30, 2004," available at http://www.symantec.com.

[3] StillSecure, "Safe Access," white paper available at http://www.stillsecure.com.

[4] Cisco Systems, "Cisco NAC: the development of the self-defending network," white paper available at http://www.cisco.com.

[5] Microsoft Corp., "Network Access Protection platform architecture," Microsoft Windows Server 2003 white paper available at http://www.microsoft.com.