



SENTINEL RIDGE CONSULTING GROUP INC.

“YOUR COMPUTER SECURITY EXPERT”

Introduction to Security Technologies

Ian Parker, GCUX, GCFW, CCSA
President
Sentinel Ridge Consulting Group, Inc.

Introduction

The last few years has seen an explosion of network security technologies, as companies struggle to protect their information assets against an onslaught of new security threats. Often, these technologies are not well understood. To further confuse matters, several security technologies are often combined within a single device, giving rise to so called Unified Threat Management (UTM) devices.

The purpose of this, and future, articles in this series is to explain in relatively non-technical terms the basic purpose of some of the most popular security technologies, so that corporate customers have a better understanding of what kinds of technologies would be most appropriate for their own networks.

This first article will provide a brief introduction to the various technologies that are available. Future articles will delve more deeply into the features that each technology may incorporate.

Firewalls

Any discussion of security technologies would have to begin with one of the earliest technologies made available for securing corporate networks, namely, the firewall.

The basic function of a firewall is to control access between a secure, trusted network environment and an insecure, untrusted environment, such as the Internet. A firewall is an essential security component for any company with a connection to the Internet. However, in addition to Internet firewalls, many companies now realize the benefit of using additional firewalls to segment their internal network. An example would be a firewall that separates the servers running payroll applications from other areas of the network.

Firewall technology has advanced greatly in recent years. Early firewalls could perform only basic packet filtering and were unable to assess how individual packets of data related to prior and later packets that may belong to the same connection. By contrast, modern firewalls have the ability to perform what is known as deep packet inspection and can understand the relationship between packets, a technology known as stateful inspection. Some firewalls also have the ability to perform user authentication and can thereby control access based not only on the origin of the connection but also on who is making the connection attempt.

Proxy Servers

At one time, proxy firewalls, otherwise known as application-proxy gateways, competed in the marketplace with stateful inspection firewalls. The modern day proxy server preserves the proxy functionality but drops the firewalling capability. Hence, proxy servers are typically found alongside stateful inspection firewalls and are used primarily to mediate connections between HTTP and FTP clients and servers.

Outbound or forward proxy servers are typically used to regulate employee access to Web and FTP sites on the Internet. Outbound proxies are often used in conjunction with content filtering and anti-virus applications.

Inbound or reverse proxy servers are typically used to regulate access to a company's publicly accessible Web and FTP servers.

Content Filtering

Content filtering applications were originally designed to protect young children from exposure to inappropriate material on the Internet. However, companies quickly discovered that such applications were essential to prevent time wasting by employees and also to prevent possible legal implications should the company's network be used to host illicit material, such as pirated software or child pornography.

Content filtering is mainly applied to HTTP and SMTP traffic. In the former case, the filtering may simply take the form of a list of blocked URL's. These URL's can be a list maintained by the security administrator but more commonly is a list derived by a service provider that places sites into certain categories. The security administrator then allows or disallows access to each category of site according to company policy. More flexible applications will allow the administrator to use a combination of pre-defined categories and manually entered URL's.

As most spam that reaches a company's network is directed to invalid addresses, one of the most effective spam blocking methods is afforded by an e-mail content filtering application that can use a white list composed of all valid e-mail addresses used within the company. In addition to white listing, some applications can use a wide variety of validation checks of both e-mail headers and content. Most applications either have built-in virus scanning or can integrate with a third party anti-virus application.

Intrusion Detection and Prevention

The purpose of an intrusion detection system is to detect unauthorized activity either on the network or on a particular host. The former is a Network-based Intrusion Detection System (NIDS), while the latter is a Host-based Intrusion Detection System (HIDS).

A NIDS operates by sniffing traffic as it crosses the network and comparing that traffic to a database containing signatures that correspond to known malicious traffic. In this respect, the technology is very similar to anti-virus. The detection may also be performed by looking for traffic that constitutes anomalous behaviour, for example, a large amount of traffic directed at a particular host that rarely sees much traffic. Such detection is more problematic, as it can be difficult to ascertain what constitutes "normal" network traffic and this property can change over time.

A HIDS typically relies on an agent installed on the host machine itself. This agent looks for anomalous activity on the host, such as a large number of failed login attempts.

NIDS is the more popular of the two, as an entire network can be monitored using a single device, whereas the use of HIDS requires an agent on every host. This can be costly, both in terms of the security application itself and the administrative overhead required to maintain it.

Despite any marketing hype to the contrary, the only difference between intrusion detection and intrusion prevention is that the latter, in addition to detecting the event, also blocks the activity.

Network-based Intrusion Prevention Systems (NIPS) are gradually becoming integrated with the deep packet inspection capabilities of modern firewalls and may eventually disappear as a separate entity.

Often lumped into the general category of intrusion detection systems are file integrity checkers. These operate by creating a hash of all critical files on a host, periodically performing another hash and comparing the two hash values, to determine if a file has changed in any way.

Anti-Virus

Probably the oldest and most widely recognized security technology, virus scanning remains a critical component of any company's security infrastructure. Although modern day virus scanners have to deal with much more sophisticated malware than in the past, the basic mechanism for detecting malware remains virtually unchanged. For the most part, the scanner attempts to detect malware by pattern matching executable code against a database of known malware signatures. More recently, some virus scanners have also attempted to detect so called "zero-day attacks" for which no signature yet exists through the use of heuristic scanning. However, this type of scanning is far from perfect and is subject to a large number of false detections.

In keeping with the principle of defence in depth, virus scanning can and should be performed at several different locations on a corporate network. For example, scanners at the network perimeter can validate Web, FTP and e-mail traffic, while internal scanners can be used to protect servers and workstations. Ideally, the brand of scanner used at each location should be different, so that malware that goes undetected by one brand of scanner can be detected by a different scanner.

Virtual Private Networks

Virtual private networks (VPN's) have exploded in popularity in recent years, as companies have realized the enormous cost savings that can be realized by replacing expensive leased lines with connectivity over the Internet.

A VPN is a secure connection over an inherently insecure medium. The security is derived from encryption of the traffic and authentication of both parties involved in the exchange.

There are two basic types of VPN. A remote access VPN allows a remote client, such as an off-site employee, to securely access the corporate network. Until recently, most remote access VPN's used Internet Protocol Security (IPSec). However, more recently, Secure Socket Layer (SSL) VPN's have become very popular. A site-to-site VPN allows two networks, for example company headquarters and a branch office, to securely exchange data.

VPN functionality can be obtained from dedicated devices. However, the function is more often combined with firewall functionality.

Vulnerability Scanners

As the name suggests, vulnerability scanners look for security vulnerabilities associated with computer systems and applications. These vulnerabilities may take the form of missing security patches or mis-configuration. Some vulnerability scanners are designed to assess security weaknesses within the operating system, while other, more specialized, scanners target specific applications, such as Web servers and databases. Like the intrusion detection and prevention products discussed earlier, vulnerability scanners come in both network-based and host-based versions and, for the same reasons, the network-based scanners are more popular.

A network-based vulnerability scanner probes for security weaknesses using two basic methods. The first method is often referred to as "banner grabbing" and consists of querying the operating system or application and analyzing the response by looking for some tell-tale string of text that divulges which version of the operating system or application is running on the host. Based on this information, the scanner can report on what kinds of vulnerabilities the host may be subject to. The other method

involves actually launching an attack against a host and determining how the operating system or application handles the attack. Obviously, this method involves the risk of disrupting or even crashing the host and should never be attempted on production systems without advance notice of this possibility.

Ian Parker is the president of Sentinel Ridge Consulting Group, a firm specializing in the security of corporate computer networks. You can reach him at sentinelridgeconsulting.com.