

Computer and Network Security Analysis

A WHITE PAPER



Computer and Network Security Analysis

With businesses everywhere interconnecting more of their mission-critical computers to the Internet—in fact many staking their very existence on online commerce—the dangers of malicious intruders invading computer networks becomes a threat of enormous significance.

The objective of this paper is to describe the various approaches and products/services available as tools to assist Information Technology managers in addressing the growing threats of unauthorized entry to computer networks.

It should be noted with emphasis at the outset that protecting computer networks against virus attacks and other forms of unauthorized access to secure data is a multi-faceted and dynamic process. That is to say, different techniques for breaking into secure systems are constantly being developed as networks' scope and complexity evolve, and, more importantly, there is no one 'silver bullet' that will assure complete security under all circumstances.

The best that organizations can do is zealously engage all prudent precautions and maintain continuous surveillance, using the most current techniques available to stay on top of a constantly changing and expanding threat.

That said, following is a summary of the principal techniques for addressing network security and the evolution of delivery modes from packaged software to the new self-administered around-the-clock online interactive penetration testing services.

Two Principal Approaches to Security Assessment/Detection

The two principal security assessment techniques are *Intrusion Detection* and *Penetration Testing*. Intrusion Detection is a "defensive" approach insofar as it provides warning that unauthorized entry has occurred. Penetration Testing, on the other hand, is an "offensive" approach designed to identify areas in which computer network firewalls are vulnerable to unauthorized entry and thereby enable companies to prevent invasion before it occurs.

Intrusion Detection Systems

As the name implies, Intrusion Detection Systems (IDS) attempt to detect, identify and often isolate attempts to "intrude" or make inappropriate unauthorized use of and/or entry into computer network systems.

Attacks can originate either by an external network connection or from within an organization's network. Systems that are usually targeted are service or workstation systems; however, attackers many times also focus on network devices such as managed hubs, routers and switches when attempting illegal entry. An IDS can help identify the fact that attacks may be occurring. It may also be able to detect attacks that other security components don't see and help collect data and/or evidence which can be used to identify intruders. Many current IDS products use a passive approach to collecting data by protocol analysis gathered by watching traffic on a network segment. Thus, when this is the case, there may be many IDS products on a network, one for each segment that the company wants to monitor. In this application, the IDS gets copies of its segment's traffic to inspect by listening in promiscuous mode (much the same way a snoop session works on a UNIX-based system) and having its network interface card bring in a copy of every packet it sees. The IDS examines these packets and attempts to determine whether they represent an intrusion attempt. This is done by seeing if the contents of packets contain the signature of a known attack method; that is, whether it contains a string of characters that matches a specified pattern, or otherwise fits profiles that define known attack methods.

Intrusion detection products are based on the assumption that an intruder can be detected through an examination of network traffic and of various system events such as CPU utilization, system calls, user location, and various file activities. Network sensor and system monitors convert observed events into chronologically sorted records of system activities. These activities are commonly referred to as audit trails. The records are then analyzed by an IDS product for unusual or suspicious behavior. Figure 1 shows a basic IDS implementation architecture.




Figure 1

There are two principal approaches to Intrusion Detection:

- *Signature-based*
- *Statistical-based*

Signature-based Intrusion Detection is designed on the assumption that intrusion attempts can be characterized by the comparison of user activities against a database of known attacks that lead to compromised system status. Most commercial IDS products perform signature-based intrusion detection against properties that initiate rules when audit records or system status information begin to indicate illegal activity. These predefined rules typically look for high-level status change patterns observed in the audit data compared to pre-defined penetrate status change scenarios. In general, a signature can be associated with a process or an event.

Statistical-based Intrusion Detection systems seek to identify abusive behavior by noting and analyzing data that deviates from a predicted norm. Statistical-based Intrusion Detection Systems are founded on the premise that intrusions can be detected by inspecting a system's audit trail for "out of the ordinary" activity, and that an intruder's behavior will be noticeably different from that of a legitimate system user. Any sequence of system events deviating from the expected profile by a significant amount is flagged as a potential intrusion attempt.

Analysis of Intrusion Detection

Today's first generation Intrusion Detection Systems were designed to protect networks by attempting to watch all traffic for signs of attack—a typical "first line of defense" type of system. Although the network-based approach to intrusion detection showed a lot of initial promise, it is now running up against some critical limitations. Additionally, within the past year, several new types of attack techniques have been identified that are beyond the detection capabilities of the IDS products.

Following is a summary of seven key issues that illustrate the limitations of current IDS products and the flawed underlying pattern-matching methodology. It recognizes that it is extremely difficult—perhaps impossible—to produce a single tool that will detect *all* types of potential attack scenarios, and the corresponding probability that many network-based attacks go undetected by the IDS products.

1. The inability to see all traffic on the network.

Many companies are beginning to use “switched Ethernet” technology to architect their LANs. Since current network IDS technology can’t reliably watch this traffic on the wire, another method employed is to connect the network IDS to the spanning port of the switch. Alas, this unfortunately has not been the answer either. Implementing a network IDS on a switch will considerably degrade the packet throughput, inasmuch as the hub has to wait for the spanning port to catch up before sensing the packet. This means that implementing a network IDS attached to the spanning port of a switched hub would defeat the purpose of a switched high-speed network in the first place.

2. Too much traffic to watch efficiently.

As networks get faster, network-based IDS products simply cannot keep pace. Examining the contents of each packet, and seeing if it matches any of the known signatures and rules takes time and resources. For example, a network IDS can examine all of the traffic on a 10 megabits per second (Mbps) LAN and check it for several dozen signatures. However, many of today’s LANs are running at far higher speeds—100 or even 1000 Mbps. And network speeds are virtually certain to grow faster than the technology for high-speed packet signature analysis. Today’s network IDS can’t reliably watch more than 10 to 20 Mbps of traffic; above these speeds you run the risk of losing data or being unable to watch for more than just a few signatures. Although IDS technology will keep improving and speeding up, the speed of networks will grow faster, given the obvious business-driven incentives.

3. Inability to evaluate impact of suspect packets.

A network-based IDS can’t predict whether a given destination machine will see a suspect packet (for example an Internet Protocol packet with a bad UDP checksum) and, if seen, whether it would be processed as expected by the network IDS. This means the IDS must inspect *all* packets, which in turn may overload it. Obviously, in network communications, packets can be sent unreliably or duplicates may be sent. When network protocols such as IP receive duplicate data, however, the IDS must choose either the old or the new data. Different implementations of IP from different OS vendors make different decisions on this issue, adding yet another dimension of uncertainty.

4. Fail-open architectures.

Certain types of network sensor-based security systems, when they fail due to overload, crashes, or denial-of-service attacks, leave the network they were guarding “open,” often without notification of the problem to the central console. The only other option for a sensor is to “fail closed,” in effect shutting down essential network services until the sensor is brought back online. Fail-closed architectures are possible on host-based systems (i.e., a fundamental practice in firewall technology), but only if the authentication module of the OS integrates with the host-based IDS and ensures that the last set of ‘rules’ implemented stays in effect until the administrator resets the system locally. This is, in effect, a ‘personal firewall’ for critical systems requiring the utmost security and data integrity.

5. Not enough information is known.

A network IDS can’t predict the implication of a packet just by looking at it. It also has to have information about the network segments, the end systems, etc., none of which is provided by a simple packet capture. Trying to provide this information may be too much work for many organizations, and trying to include it in high-speed attack monitoring may not be possible. Also, because a network IDS is usually on a dedicated machine rather than on one it is watching, differences between the network IDS host and protected hosts in hardware, drivers, etc. can lead to discrepancies in what may or may not work as a type of an attack signature.

6. Failure to detect certain types of attack.

Perhaps most significantly, as pointed out in a paper entitled *Insertion, Evasion and Denial of Service: Eluding Network Detection*, by Ptacek and Newsham (1998), there are at least 26 different techniques for executing a single attack that may elude the detection of current IDS products. Three classes of attacks have been found which exploit the fundamental nature of a network IDS, based on IP and TCP protocol analysis. These attacks either evade signature recognition, or consume enough resources to disrupt/disable the network IDS and result in what is commonly referred to as a Denial of Service (DoS) attack. In trials run by many of the leading product testing labs, all network IDS products on the market today proved vulnerable to each type of attack. The implication is that, short of some fundamental redesign, today's network IDS systems cannot claim to offer significant intrusion protection.

For example, one form of insertion attack inserts extra characters into the packet stream, which keeps the contents of the packets from matching an attack signature, i.e., data is sent one character per packet, and the "exploit" string "GET/cgi-bin/phf" is masked by inserting padding characters, so the network IDS sees a pattern such as "GET/cgi-bin/phearf." The end node discards the padded data due to processing and the assumption that the additional characters are caused by inherent noise on the line, resulting in the target system "recreating" the original exploit string.

Similarly, since TCP/IP reassembles data streams using sequencing numbers in the packets, one evasion technique is for an attacker's packets to be sent out of sequence. So, again, what the network IDS sensor sees doesn't look like an attack, but when these packets are reassembled by the target system, they contain the original attack. Additionally, the paper entitled *Insertion, Evasion and Denial of Service: Eluding Network Detection* identified numerous Denial of Service (DoS) attacks against either the network IDS device itself, or a host that cannot be detected or prevented by a network IDS. DoS attacks are intended to compromise a device's availability by making them too busy, crashing them, etc. Common network DoS attacks include mail bombs, ping floods and attack or exploit problems known as software bugs. Moreover, reports show these types of attacks occurring with greater frequency, thus confirming that intrusion detection analysis is incapable of ensuring protection or safety from network-based attacks.

7. Too many false positive responses.

Still another problem with network IDS products is that they are prone to perceive threats that appear to the IDS to be real, but in fact are just normal transactions. Such incorrect diagnosis is known as "false positive responses," and can be quite disruptive. It is easy for a network IDS, for example, to produce an alert when a "ping flooding" attack occurs. If the network IDS produces an alert every time it sees any kind of ping, however, it may produce the opposite effect. In fact, knowledgeable hackers have been known to create a "boy who cried wolf" scenario by generating so many alerts that appear to be false positive responses that network administrators simply filter out or ignore these alerts, allowing serious attacks to go unnoticed. There is currently no benchmark for evaluating the "signal-to-noise" ratio produced by network IDS products.

Penetration Testing

Following the philosophy of Vince Lombardi, legendary coach of the Green Bay Packers, many Information Technology professionals believe "the best defense is a strong offense." Penetration Testing (PENTEST) is the preferred preventive monitoring tool. Today, most Information Technology professionals consider PENTEST to be basic to firewall surveillance and responsible network security management. Even though it is axiomatic that no system can be 100% secure, it is equally sure that it is simply not practical to maintain network security today without some form of PENTEST tool.

In Penetration Testing, companies' firewalls and defense devices are actually "invaded" to determine where and how they may be vulnerable to attack. As organizations' use of the Internet for secure information exchange and e-Commerce has increased, so too has the incentive to breach security and invade networks. In the process, companies have come to realize that firewalls, while critically important, are by no means an impenetrable deterrent to invasion. In fact, a recent study by the United States CIA

concluded that more than 50% of security breaches were from the *inside* of the network. Thus, testing for vulnerability to unauthorized entry has evolved into the preferred mode of network security preventive monitoring.

Through Penetration Testing, companies unleash a proactive approach to network security, as opposed to relying on Intrusion Detection to issue an alert when a security breach has occurred. Just as specific information is known about computer viruses, there is a similar body of knowledge regarding known security vulnerabilities and attack methods, many of which are utilized by hackers, “wannabies,” and/or other potential network intruders, including company employees. These weaknesses may be as simple as whether the password on a default or field service maintenance account has been left unset or with the commonly known default password, or as complex as what level of a certain type of attack a network-based IDS can sustain before failing. Categories of typical vulnerabilities include misconfigurations, security policy violations, service exploitations, software updates, security holes and known bugs. Given that this knowledge exists, it makes sense for a company to focus *first on known vulnerabilities* by regularly testing its firewalls, routers, hosts and other devices. Once these vulnerabilities are identified, an organization can then make decisions what to do, rather than wait for attackers to identify and take advantage of weaknesses, with potential disabling results.

Penetration Testing Delivery Systems: Standard Software to Online Interactive Services

There are three types of Penetration Testing systems: Standard Software Packages for in-house installation, Managed Services and, recently for the first time, Online Interactive Service over the Internet.

1. Standard Software Packages.

This was the first form of Penetration Testing and one still in widespread use. The advantage to in-house installed software is that Information Technology specialists have a tool at their disposal for penetration testing at will. The disadvantages are (1) the high cost of hardware, installation, system operation and maintenance, and (2) the problem of software obsolescence as new penetration techniques are developed.

2. Managed Service Providers.

Managed Service Providers provide remote penetration testing, receiving orders for tests via the Internet and confirming the completion of testing via e-mail. The advantages of Managed Services are reduced cost and a greater likelihood that testing protocol will be up-to-date. The disadvantages are that the testing protocol is one-size-fits-all, that is, no matter what a companies' network configuration may be, it gets the identical regimen of scans for penetration vulnerability, which may or may not be applicable.

3. Online Application Service Providers.

In 2000, Network Security Systems introduced iNETPATROL™, the first ASP-model online interactive Penetration Testing service. Now, for the first time, Information Technology computer network security specialists can self-administer penetration testing without the expense of in-house software installation.

Planning a Comprehensive Security Strategy

Using an online scanning service like iNETPATROL™, which has a database of vulnerabilities to test against and is continually updated, an organization can conduct its own routine checks of routers, firewalls, network devices, operating systems, web servers, applications, workstations and even intrusion detection systems. This provides a company with a more proactive approach to security, an attractive alternative to simply waiting for potential hackers or intruders to invade the network.

The iNETPATROL™ application service currently documents hundreds of tests, and, unlike other Penetration Testing Systems, tailors the testing to the particular network being tested.

The patent-pending scanning engine for iNETPATROL™ and its sister product LANPATROL™, a black box appliance for in-house penetration monitoring, includes the techniques practiced today to attack first-generation network intrusion detection systems. Based on experience, it may take only a matter of seconds or minutes to conduct a full set of the appropriate set of tests against a single type of system architecture. However, experience also shows that given the varying characteristics of the Internet and traffic on any given day, what takes seconds or minutes one day may take much more time on another day. So, by running multiple scanner sessions in parallel, tests of larger networks and larger numbers of hosts can be done even more efficiently. It may make sense to approach scanning as a multi-step process, especially given that an organization is typically bringing in new machines and upgrading or replacing existing ones weekly, daily or even hourly. This is one of the primary reasons the iNETPATROL™ service is available on a 24x7 unlimited basis. Upon using the tool, organizations can then focus on what to fix, based on many factors, including how difficult and/or expensive it may be to resolve the problem, how vulnerable the system may be if the problem isn't fixed, or how important the system, application and/or data is to the organization. Figure 2 represents the iNETPATROL and LANPATROL architecture.

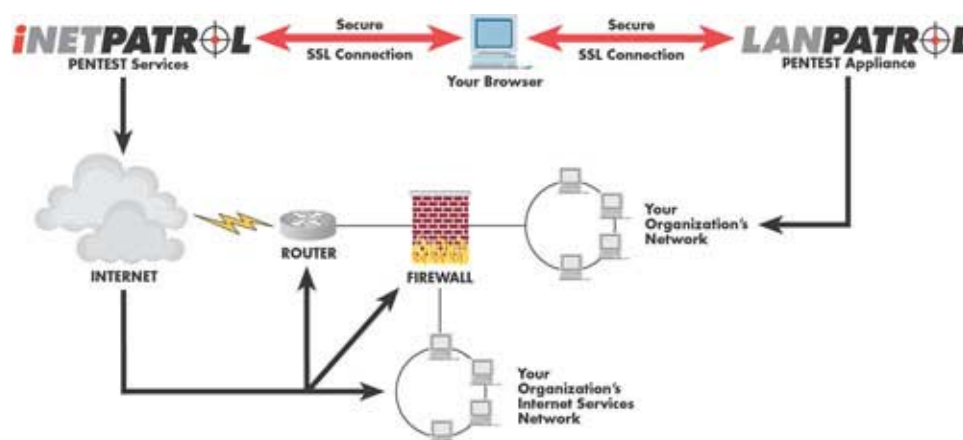


Figure 2

Vulnerability Analysis and Management

Of course to be meaningful at most enterprise levels, vulnerability and exploit testing needs to be conducted on a regular basis. Tests should be initiated after fixes, patches or other operating system modifications are applied, and after any meaningful change or upgrade is done to hardware or software, as new vulnerabilities are discovered, and on a frequently scheduled basis. Ongoing testing should be viewed as one aspect of the vulnerability and exploit management process that organizations institute for mission-critical resources and services, including networks, systems, manufacturing and support facilities. Assuming the vulnerability database is comprehensive, scanning followed by applying the recommended actions will close a significant percentage of security vulnerabilities, and provide organizations with a good sense of what can and can't practically be remedied in terms of vulnerability to invasion.

There is always the possibility a security vulnerability will be diagnosed but can't be remedied. A few examples are system design issues, employee misuse and unsecured architectures. Since these potential weaknesses can be "closed," they must be dealt with in other ways. In many cases, the unsecured resource will be one that can't simply be removed. In these cases, host-based security monitoring and network analysis tools are advised to assure that attack attempts are detected quickly should they occur.

Host-based Security Monitoring and Network Analysis

If organizations want the assurance that they know what is happening in real-time to all systems on the network, there is no substitute for monitoring them *directly*. In addition to the many shortcomings of network IDS products, some attack behavior can only be seen from inside the host and viewed directly from the operating system and/or kernel rather than from a network session.

An ideal solution is to deploy some type of intrusion detection agent directly on each of the computers to be watched, and have some form of online network analysis also available. However, given large heterogeneous networks with many hosts this can be a very difficult process to manage. But because the agent is running on the host, it can observe events and system behaviors, including some which are difficult or impossible to see from a telnet, login or UNIX shell session. This in turn makes it possible for the agent to watch and analyze events like logins and also watch different types of behavior on the system.

By comparing behavior against a database of rules such as IP addresses or services to block and thresholds such as a number of failed login attempts, the agent can identify possible intrusion attempts. The response may be any combination of logging, alerting or other activities. For example, iNETPATROL™ and LANPATROL™ contain a system profiler, vulnerability scanner and network analyzer, but do not include host-based intrusion detection agents, due mainly to the management and update complexity. As operating systems continue to mature, it is believed that host-based intrusion detection functionality will in time become part of the base operating system. Even now, it is important that system administrators use the features in current operating systems to setup and monitor the appropriate processes and logs.

Discovering and Profiling the System

There are currently thousands of types of vulnerabilities to attack at the packet and operating system levels. However, a good vulnerability scanner should be able to check and discover any or all of them, and report on recommendations for corrective action when a particular vulnerability or exploit is found. In addition a good vulnerability scanner should be able to discover what nodes are on a network and profile what types of systems those nodes entail. For example, if a network is addressed 172.16.1.0, the scanner should be able to find all of the nodes in the range of 172.16.1.1 through 172.16.1.254 and detect what type of device it appears to be communicating with and what operating system that device is running. This is called “the discovery and profiling process,” and both iNETPATROL™ and LANPATROL™ contain this functionality. In addition, this information can be used to provide an inventory assessment of what devices are actually up and running on the network versus what devices are thought to be up and running on the network.

Network Security Solutions for Effectively ‘Patrolling’ the System

To address the foregoing challenges and requirements in network security penetration and vulnerability testing, Network Security Systems provides a comprehensive solution.

iNETPATROL™, the online application service, provides proactive scanning to examine computer systems and network devices vulnerabilities in enterprise network environments. With iNETPATROL™ Information Technology network security professionals can test any TCP/IP-based device including workstations, servers, hubs, switches, and routers. iNETPATROL™ includes exclusive patent-pending scanning engine tests to provide thorough perimeter audits of firewalls and any other external visible devices. Report options include executive summaries, detail reports and corrective-action recommendations. Most importantly, Network Security Systems conducts continuous research to identify newly-discovered intrusion techniques and provides ongoing updates and enhancements to assure that the scanning engine vulnerability testing database is always current.

LANPATROL™ is a 'next generation' network appliance that provides the same level of computer and network security penetration testing as the online service and with the same browser-based interface. It is designed to be installed anywhere on the network to provide users with the ability to scan networks from inside the firewall and report on any vulnerability or exploit found. *Together, the two products provide a complete and comprehensive network security assessment solution across diverse internal and external network topologies.*

The following is a summary of key features that distinguish iNETPATROL™ and LANPATROL™:

- *Profiler*—determines what type of node and what OS that node is running.
- *Interrogator*—the scanning engine that performs auditing tests to assess the security vulnerabilities and exploits of computer and network systems.
- *Exploiter*—auditing tests to exploit vulnerabilities found on computer and network systems.
- *War Dialer*—automated dialing of a telephone exchange to determine what systems may have an answering modem attached.
- *Analyzer*—network traffic and protocol analysis.
- *Reporter*—automated reporting and viewing of the vulnerabilities found
- *Online Help*—automated online product documentation and troubleshooting
- *Security Checks*—vulnerability test definitions database search
- *Continuous updates*—from the NSS Labs, Network Security Systems' research division, updates are automatically added into the product scanning engines as they are detected.

Conclusion

Attacks against network assets continue to rise dramatically. Sophisticated tools developed by experienced hackers are now distributed freely across the Internet. These tools allow complicated attacks to be staged by relatively inexperienced "hobbyist hackers." At the same time, the e-Business revolution demands that companies continue connecting mission-critical systems to the Internet. As we have seen, first-generation network intrusion detection systems are simply not living up to their initial promise. Today's next-generation approach to penetration testing can be a vitally important tool in the security arsenal for diagnosing vulnerabilities and exposures on an organizations' computer systems and networks.

For further information on Network Security Systems, please visit our website at <http://www.netsecuritysys.com>.

