WHITE PAPER

March 1998

Internet Solutions
Business Unit

Compaq Computer
Corporation

CONTENTS

Intrusion Detection in the Enterprise Network: Managing Hacker-Related Risk

This White Paper addresses the use of network-based intrusion detection technology to mitigate hacker-related risk in Microsoft Windows NT-based networks. It is intended for information technology managers and users who have a need to understand the information protection concepts and technologies appropriately applied when minimizing hacker-related risks.

.

COMPAQ

NOTICE

Intrusion Detection in the Enterprise Network: Managing Hacker-Related Risk
First Edition (March 1998)
ECG041/0298

I. INTRODUCTION

Effective enterprise security is all about managing risk. An effective enterprise security program is based on a multi-layered security policy that recognizes and addresses the complexity of the operational risks associated with any information system. Recently, the internetworking of previously isolated private networks with value added networks; public networks, such as the Internet; and other enterprises' private networks, has become a trend greatly complicating risk management.

Protecting an information system is, in large measure, based on providing essential services while isolating the information system from those risks that would damage the system or the information it stores, processes, or transmits. Should the isolation measures fail, and an event occur that could damage the system, a prudent security policy provides the means to detect the failure and report the event to an assigned person for resolution.

II. THE THREAT

A key challenge presented in protecting interconnected networks, intranets and extranets, is to maintain isolation from individuals with no legitimate need to access the system: "outsiders". Internetworking, particularly with the Internet, has grossly increased the number of outsiders with potential to access internal networks. This potential has been interpreted as a challenge by a significant number of those individuals, motivating them to develop techniques for penetrating system security. Penetration techniques are conveniently grouped under the concept of "hacking" with the practitioners of hacking being "hackers."

Unfortunately, insiders as well as outsiders can engage in hacking activities. In fact, the insider is much more likely to be successful in hacking activities, because the insider is not hampered by the protection measures implemented at the network perimeter (e.g. firewalls) to keep the network isolated from the outsider.

Vulnerabilities

Hackers exploit security holes, or vulnerabilities, in a network to gain access. Vulnerabilities can be introduced in a number of ways, the two most common are intentional vulnerabilities and human error.

- ***Intentional Vulnerabilities:*** It is often necessary to accept the need for a vulnerability in order to perform a task of high value to the enterprise. For example, use of some network protocols and services, including connectionless and anonymous protocols, have significant associated vulnerabilities. If a threat can exploit those vulnerabilities without detection, enterprise security is immediately compromised.
- ***Human Error:*** Firewalls provide the primary technology to isolate the information system from outsiders, including hackers. Other important isolation technologies used to supplement a firewall include operating system security mechanisms, such as authentication and access control privileges. All of these mechanisms require configuration and administration. During either configuration or administration, errors can be made that create vulnerabilities, which are easily exploited by a minimally competent hacker. Vulnerabilities in a network, created by human error, are likely to exist because of the dynamic nature of a network. New components are constantly being introduced and old components taken away. Enterprises experiencing rapid growth, reorganizations, down sizing, merger activity, and other such changes, are prime candidates for this type of problem. Equally troublesome is the lack of properly trained and experienced personnel. It is often necessary to assign inadequately trained personnel system administration duties. The result is a high likelihood of human error in the administration of some network components. When those components include firewalls, web servers, and other perimeter devices, hackers can be expected to easily attack and exploit any vulnerability created through administrative error.

The Internet Related Risk

Today, many organizations are connecting their internal networks to the Internet to meet important business objectives that include:

- ***Increased employee productivity:*** Allowing access to the information and services on the Internet can increase employee productivity.
- ***Enhanced communication with customers, suppliers, and partners:*** Organizations need to make corporate information available to users in the outside world, including customers, suppliers, and business partners.

- ***Increased revenue/reaching more customers:*** One of the major attractions of the Internet is that it enables organizations to reach customers in much greater numbers and in far more locations than with conventional commerce vehicles.
- ***More efficient use of technology:*** Using the Internet as a Wide Area Network (WAN) backbone (intranet), or to interface other companies' networks (extranets), provides an economical medium for connecting local and metropolitan area networks to WANs.

While connecting to the Internet offers numerous advantages, it also exposes internal networks to millions of outsiders. The result is that the Internet creates a high bandwidth channel into the heart of a network. Normally that channel will be guarded by a firewall. However, if a threat, a hacker, can exploit vulnerability, either in the firewall or elsewhere, a successful attack against a network can be launched. The likelihood of such an event happening is referred to as "risk."

Many network managers feel that their networks are relatively secure and that it is unlikely that a hacker could gain unauthorized access. This often results in a false sense of security. One of the major failings of a large percentage of security policies is that they do not provide a means of detecting and reporting security events as they happen. A hacker could be active, but undetected, in the network. Therefore, in reality, the manager has no way of knowing if a hacker is active in the network or not. In virtually all environments today, the security measures in place are not designed to alert management when an active attack is in progress. Therefore, the only time the network manager knows a hacker is active is if the hacker is careless or destructive.

Addressing the Issue

If the operational risk associated with a network is to be managed, the ability to recognize and respond to an attack, while it is happening, is required. Otherwise, the potential of the attacker to inflict damage cannot be minimized. An attack recognition and response capability integrated into the security mechanisms of an environment is required to address this issue.

Typically implemented in software, an attack recognition and response system continually monitors network traffic, looking for known patterns of attack. When it detects an unauthorized activity, the software responds automatically with predetermined actions. It may report the attack, log the event, or terminate the unauthorized connection. Attack recognition and response software operates in concert with other security mechanisms to address the risks associated with hacking.

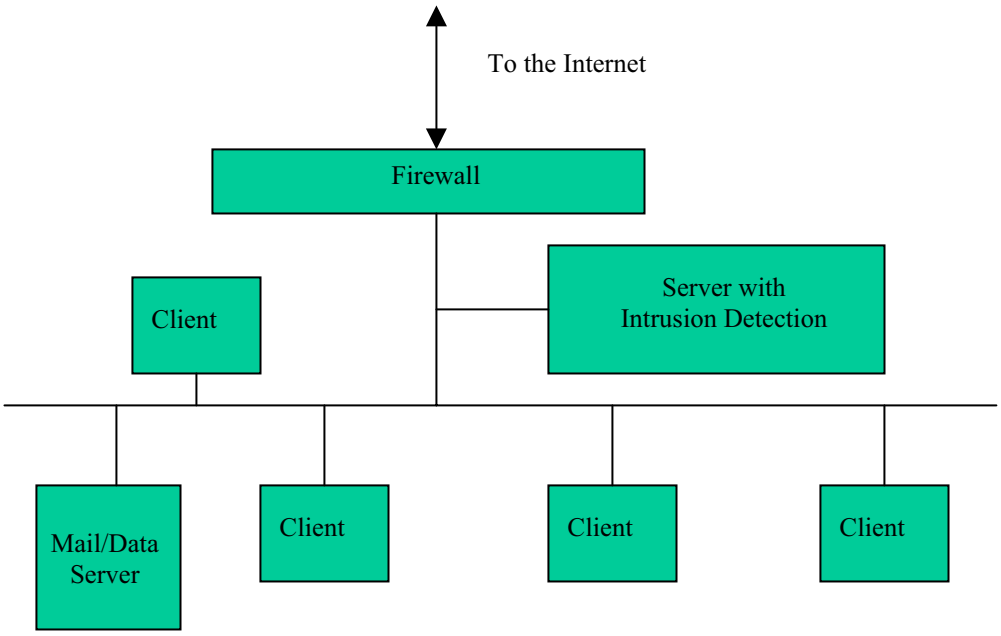

Figure 1: Server with an Intrusion Detection Tool in a Simple Client/Server Architecture

In client/server architecture, the intrusion detection application typically runs on a server under the Microsoft Windows NT Server operating system (see Figure 1).

III. INTRUSION (ATTACK) DETECTION AND RESPONSE

Attack Detection

Attack recognition and response software typically detects attacks using either rule-based or statistical anomaly approaches:

- ***Rule-based:*** This approach draws from a library of known attack patterns or unauthorized activity, and watches for those specific types of attacks. This is similar to the technique used in virus detection. The attack pattern library is updated continually as new types of attacks are discovered. Database updates are provided either by downloading new copies of the database automatically, or in new software releases.
- ***Statistical Anomaly:*** This approach operates on the assumption that users and networks always exhibit a predictable pattern of behavior, and do not depart from this pattern over short periods of time. A deviation is considered an attack.

Attack Response

Attack recognition and response software can be configured to react automatically to an attack in a variety of ways, including:

- Log the event along with associated information.
- Alert appropriate personnel through console messages, e-mail, or pagers.
- Terminate the offending connection.
- Call a user-defined script or program.
- Perform a combination of these actions.

Requirements for Effective Attack Recognition and Response

Attack recognition and response software must meet a number of requirements to provide truly effective protection against attacks. The major requirements include:

- ***Real-time Reporting:*** The attack recognition and response software must be capable of detecting, reporting, and reacting to suspicious activity in real-time. Software that merely logs events is ineffective. After-the-fact detection is like a burglar alarm that goes off long after the burglar has fled. In addition, many attackers erase logs during the break-in, so the intrusion cannot be detected by merely scanning an event log.
- ***Update Capability:*** Just as there is a progression of new computer viruses, hackers continually find new ways to break into computer systems. Therefore, attack recognition and response software must be capable of continually adding to its knowledge base of known break-in patterns and unauthorized activity. The update function should be performed frequently and relatively easily. If the attack signatures are integral to the intrusion detection system, software releases should be issued when appropriate, and automatically downloaded. If the signatures are contained in a table or database, the file should be automatically downloaded when appropriate.
- ***Run on and Support Popular Network Operating Systems:*** The software must support existing network infrastructure, including network operating systems, such as UNIX and Windows NT.

- ***Ease of Configuration.*** Configuration should be easy, without sacrificing effectiveness. The attack recognition and response software should provide a default configuration so that administrators can deploy it quickly and optimize it over time as information accumulates. In addition, the software should provide sample configurations to guide administrators in setting up the system.
- ***Flexibility:*** To provide maximum flexibility, all control options should be configurable.
- ***Manageability:*** Rapidly rising network management costs present a significant problem for organizations. Attack recognition and response software must be easy to manage so that it does not contribute to this problem. Management of the software over the network, from a central location, is essential. In addition, the software should be easy to integrate with the existing network management infrastructure. This requires compliance with network management standards such as SNMP.
- ***Adaptability:*** Today's business environment is dynamic. Organizations are continually changing, driven by many factors including reorganizations, mergers, and acquisitions. Therefore, security policies are also in flux. To remain effective, attack recognition and response software should be easy to adapt to changing security policies. This ensures that these policies can be implemented in fact, as well as on paper.
- ***Non-obtrusiveness:*** The software should operate in a non-obtrusive way. That is, it should not degrade network performance. It should be transparent to authorized users so that it does not hamper productivity. In addition, it should not alert the intruder to its presence.
- ***Reporting Features:*** Reporting features should be easy to use and configure. A graphical user interface (GUI) should be available to support reporting.
- ***Robust:*** The product should detect all known attacks.
- ***Secure****:* Intrusion detection technology runs as an application on an operating system. The vendor should provide information related to the secure configuration of the operating environment to protect the application. The product should also provide integrity controls to ensure that it cannot be easily modified or infected with malicious code (e.g., viruses).
- ***Performance.*** Because the application is analyzing packet traffic, it must be able to process large numbers of small packets at high speed. The vendor should address throughput constraints and performance degradation based on network traffic loads.

Intrusion Detection in a Security Policy

Compaq believes that intrusion detection technology is an affordable and prudent measure that will significantly improve an enterprise's ability to manage operational risk related to hacker activity and conduct safe computing in an interconnected network environment. It is recommended that intrusion detection technology be required by enterprises with network connectivity that spans enterprise borders (intranets and extranets); for example, public network connectivity or connections to other enterprise's networks. Intrusion detection technology should be used in conjunction with perimeter defenses (e.g., network firewalls), and other appropriate technology, to provide robust and durable layered protection for the enterprise's information assets.

A network-based intrusion detection system is subject to potential misuse. The system is, at its root, a network sniffer. Therefore, it could be misused to gather inappropriate information. Compaq recommends that access to the application be mediated based on appropriate entries in the Access Control Lists associated with the application's directory and files, and limited to those individuals with direct responsibility for network security. This issue should be addressed in the Security Policy.

All protection technologies have some associated vulnerability. Compaq is concerned about potential vulnerabilities associated with this technology. For example, we believe that caution

should be exercised in relation to the automatic disconnection of hosts associated with an attack, because it may be possible to masquerade in such as way as to create a denial-of-service. In this scenario, the attacker's intent is to create a denial-of-service for a specific host. By assuming the host's IP address and directing an attack, it may be possible to spoof the intrusion detection system into believing the host is the attacker. Automatic disconnection would complete the denial-of-service attack on the host. This scenario can be defeated by investigating attacks before hosts are disconnected. An environment-specific security policy should address this issue.

Compaq has not performed, or reviewed, a comprehensive vulnerability assessment of intrusion detection technology. Therefore, we believe it is circumspect to assume that vulnerabilities may exist that have yet to be uncovered. These issues are prudently addressed in the Network Security Policy through layering and best practices.

IV. THE REALSECURE SOLUTION

RealSecure, from Internet Security Systems (http://www.iss.net), is a real-time monitoring attack recognition and response system. It monitors packet flow over a network in real-time and analyzes packets for known attack patterns and unauthorized activity using a rule-based approach. When RealSecure detects an attack, it reacts automatically according to its configuration. RealSecure can react in four ways when it detects an attack:

- It can alert appropriate personnel through administrator console messages, e-mail, or pager alerts.
- It can log the event, along with associated information.
- It can kill the event by terminating its connection.
- It can initiate a user-supplied script.

Organizations can strengthen security with RealSecure in three primary ways:

- As an effective, second line of defense behind firewalls; intercepting and disposing of attacks that get through the firewalls, or that originate inside the firewalls.
- As a means of measuring the effectiveness of the current network security mechanisms. For example, RealSecure can detect when a Telnet session is being established from the Internet, even if the firewall has been configured to disallow Telnet sessions from the Internet.
- By providing feedback on vulnerabilities, which have been accepted by the organization to provide access to a required service, identified by a vulnerability assessment tool. RealSecure can determine the number of attempts to exploit the vulnerability. If the number of attempted break-ins is unacceptable, the organization may choose to review the decision to make the service available.

RealSecure provides useful reports on its findings to help organizations assess and tighten their security.

Advanced Architecture

RealSecure consists of three integrated components running on a central server:

- ***Recognition Engine:*** The Recognition Engine monitors the network in real-time, detecting and reporting attacks. It reports events to the Administrator's Module.
- ***Response Engine:*** The Response Engine reacts automatically to attack events as they are recognized, triggering pre-specified actions ranging from logging attacks and alerting the administrator to terminating offending connections.
- ***Administrator's Module:*** The Administrator's Module provides management of the Recognition and Response Engines, managing all Recognition and Response Engines from a single GUI, simplifying network management.

Recognition Engine

The Recognition Engine leverages Internet Security System's in-depth knowledge of how systems are attacked. The Engine can detect low-level IP attacks that can bypass packet filter firewalls, such as those using IP spoofing, whereby an attacker uses the identity of a trusted correspondent. The Engine can also detect high-level attacks, such as from those Web, FTP, NFS, NIS, or e-mail sessions.

The administrator can configure the Recognition Engine to implement specific security policies. The Engine can react to (or ignore) connections based on specific packet types, source, destination IP addresses or address ranges, port numbers, or particular types of attack patterns. This flexibility

enables the administrator to set up custom monitoring for individual hosts and networks. Because the Recognition Engine is a passive monitor, operating like a sniffer with built-in security knowledge, it doesn't degrade network performance.

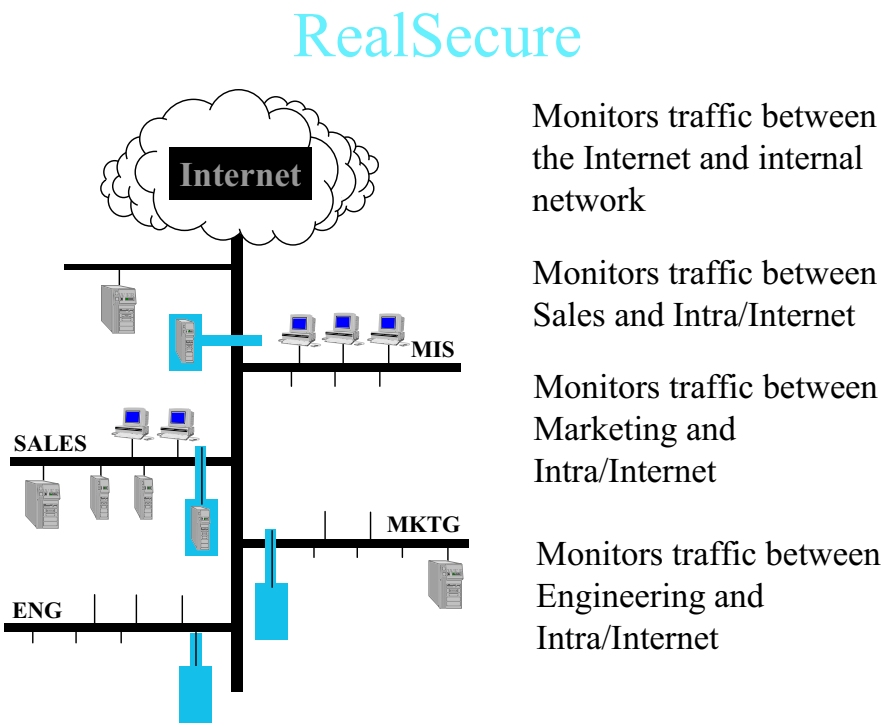

Figure 2: The RealSecure Real-time Attack Recognition and Response System

As Figure 2 shows, an organization can place multiple Recognition Engines in strategic locations in its network topology. An organization can also use multiple Recognition Engines in parallel at a single location to accommodate high bandwidth connections, such as T3 access to the Internet. Such usage is configurable.

Response Engine

The Response Engine can be configured to react automatically to detected attacks in a variety of ways:

- ***Alert Appropriate People:*** The Response Engine can alert the administrator through the Administrator Module, send e-mail messages, and activate pagers.
- ***Log Attack:*** The Response Engine can be configured to log the attack. It can log the event only or the entire attack session, including all the hacker's keystrokes. The administrator can play back the session later, in its entirety, through an easy-to-use GUI control panel.
- ***Terminate Session:*** The Response Engine can be configured to reset the connection on both the attacker's machine and the target machine when it detects an attack. It terminates the session by sending a Reset (RST) packet to the attacker's machine. It sends an RST packet to the target machine, spoofing the attacker's IP address.

- ***Initiate User-supplied Scripts:*** The response can also be customized with user-supplied scripts that are activated when an attack is detected. In this way, reaction can be tailored to the specific needs of the organization.

The Response Engine's ability to react automatically provides proactive security, without requiring administrator intervention.

Administrator's Module

The Administrator's Module provides a single point of management and control for all Recognition and Response Engines in the network. The administrator can configure the Recognition and Response Engines from the Administrator's Module. In addition, the Module collects and presents events reported by the Recognition Engines in an easy-to-read GUI display in an NT window.

When a Recognition Engine detects an attack, it reports it to the Administrator's Module. The Module displays the event in real-time, as it is happening, optionally including the hacker's keystrokes and a copy of the hacker's screen. To allow for easy interaction, attack events are classified and displayed by high, medium, or low priority.

Manual Response

In addition to automatic response, the administrator can respond manually to reported attacks in a variety of ways using the Administrator's Module GUI:

- ***Request Additional Information:*** The administrator can request that the Engine provide more detailed information on the reported attack. This information can include the packet source and packet data, such as e-mail headers.
- ***Instruct RealSecure to Log the Event:*** If automatic event logging has not been configured for this event, the administrator can request that RealSecure log the reported event.
- ***Instruct RealSecure to Kill the Event:*** If automatic session termination has not been configured for this event, the administrator can request that RealSecure terminate the session.

The administrator can combine automatic and manual responses to maintain the exact level of control required.

Easy Configuration

The administrator configures RealSecure through the easy-to-use Administrator's Module GUI. The configuration specifies the types of checks to be performed by the Recognition Engine, and the response to be initiated by the Response Engine for each type of attack detected. Through the GUI, the administrator can custom-tailor a security model of the network to match the organization's security policy.

RealSecure includes a default configuration to allow an organization to get up and running quickly. The default configuration is biased toward tight security to ensure the protection of the monitored network. RealSecure also includes a variety of sample configurations, at different levels, to provide starting configurations for various types of security policies.

Meaningful Reports

RealSecure can generate meaningful reports from its event log files. These reports can include such information as the amount of data processed by a web server each day, or the number of connections that were killed each day and by whom. The Administrator's Module can display these reports in graphical form, such as bar or pie charts, for easy review and analysis.

Administrators can use the reports to optimize security. The suggested method of implementation is to start with extremely tight filtering, using the default configuration. The administrator can use the information in the reports to tune filtering over time, as an understanding is gained of normal network activity. Repeated tuning reduces the number of alerts without relaxing security.

V. CONCLUSIONS

In order for organizations to compete effectively in today's business environment it is essential that they increase their use of networks, including:

- Expanding the reach of their internal networks by interconnecting LANs into WANs
- Opening their internal networks to outside organizations to gain higher levels of interaction with their business partners
- Taking full advantage of the power of the Internet

However, the increasing use of networks brings with it increased vulnerability to network break-ins. Traditional network security solutions such as firewalls, operating system security mechanisms, and encryption protect internal networks to a large degree, but hackers still manage to penetrate. Break-ins—whether internal or external—can be costly, and expose sensitive and confidential information to the outside world.

RealSecure augments existing security mechanisms and helps reduce the gap between an organization's security policy and its actual security practice. It enables organizations to measure the effectiveness of their current network security mechanisms in implementing security policy. It also provides an effective second line of defense behind existing mechanisms.

RealSecure running on Compaq servers provides a robust solution for Security Policy requirements related to minimization of the hacker-related threat. In this way, enterprises can maintain their competitive edge while keeping their operational risk within acceptable limits.