

Email Security

What the hacker knows that you don't



E-Mail Security Challenges

This document outlines how hackers are exploiting vulnerabilities in e-mail systems, and describes the widely available hacking methods and tools they use. The main body of the paper attempts to describe the problem and its solution as non-technically as possible. The Appendix has the real technical detail for those who are interested.

E-mail systems such as Microsoft Exchange, Lotus Notes and GroupWise were constructed with a single purpose in mind, to accept and send the maximum amount of mail, and route that mail as efficiently as possible. Without question this has succeeded; e-mail is the most commonly utilised business communication tool on the planet, and its use is projected to continue to rise. In fact, the current volume of e-mail sent worldwide is now more than 50 billion messages per day, with that number expected to double by 2008.

E-mail's continually burgeoning popularity makes it an increasingly attractive target for individuals seeking to do harm, either for their own misguided personal satisfaction, or more likely, for financial gain. The first e-mail hackers found simple vulnerabilities in the operating systems and protocols of e-mail systems, and exploited these known weaknesses. Now, however, hackers and virus writers have become specialists, constantly developing new and innovative methods of overcoming the improvements made in today's security systems. The game of cat-and-mouse is unlikely to end any time soon, if ever. With every improvement in defensive techniques, hackers and virus writers modify their tactics in an attempt to circumvent these defences and wreak havoc on corporate networks.



Vulnerabilities of E-Mail Systems

But first some jargon busting

An IP address is simply the numerical address that the Internet uses to find out where it must deliver the data. Computers only understand numbers, so www.encrypted.co.uk might have an IP address of 12.3.456.7

TCP/IP Protocol is a standard used to ensure a common format in which data can be sent and received

Pinging is a method of sending data to a specific address (IP) to see if it is “alive”, i.e. it responds.

The Risks

Along with the many conveniences and efficiencies that e-mail use brings to an organisation, there are some inherent risks and vulnerabilities:

Internet communications were designed to enable seamless communication among multiple machines. As a result hackers seek to exploit the open nature of these designs (protocols) to attack organisations. The TCP/IP protocol was designed before there was much experience with the wide-scale hacking that is seen today and as a result, there are a number of general security flaws.

The first level of attack can discover services which exist on the target network. This involves a number of possible techniques to gather data on the remote network, known as sweeping and scanning. (See Appendix One)

Hackers are free to forge and change IP data with impunity

There are a range of attacks that take advantage of the ability to forge (or “spoof”) an IP address. While a source address is sent along with every IP packet, this source address isn’t actually used for routing to the destination. As such, the attacker can forge a source IP address, allowing the attacker to exploit the remote server while pretending to be someone else.

IP spoofing is used frequently as part of other attacks such as SMURFing, in which the source address of a broadcast ping is forged so that a huge number of machines that are pinged respond back to the victim indicated by the address, overloading it (or its link).

Many organisations have inbound e-mail gateways which validate the legitimacy of the inbound e-mail recipients. If the inbound e-mail address is valid, the e-mail is forwarded on to the addressee. However, if the e-mail address is non-existent, a response is dispatched to the sender notifying them of the invalid e-mail address. Hackers exploit this inherent “politeness” of the e-mail systems to gain access to valid addresses. They then unleash Directory Harvest Attacks (DHA), whereby a program guesses at possible e-mail addresses within a domain and attempts to send a message to that domain. In a situation such as this, the e-mail gateway rejects those addresses that are invalid. By process of elimination, addresses that are not rejected are deemed valid by the hacker, spammer, or virus writer and added to their database of legitimate addresses.

Servers can be instructed not to reject bad addresses; however, this can result in a never-ending increase in mail volume which must be processed by the organisation.

Social engineering

Unfortunately, the trusting nature of most people makes them vulnerable to social engineering from a hacker. In these attacks, a hacker may use a tool as simple as an Internet search to find legitimate e-mail addresses within an organisation. The hacker will then send an e-mail to the known valid address in order to elicit a response. If a response is received, the hacker will examine the headers in order to determine the path followed by valid mail within the organisation. Additionally, this information can be used to set up attacks at the machine level, or over the phone using more social engineering techniques, to glean login/password information.

Misguided belief in the firewall as adequate protection

A common misunderstanding is that firewalls recognise e-mail-borne attacks and block them.

Firewalls simply control network-based connectivity and usually perform no scrutiny upon traffic coming through on the standard e-mail port (port 25). The firewall administrator adds rules that allow specific types of network level traffic to go through the firewall. For example, a typical corporate firewall allows mail traffic to pass through unimpeded, thus the firewall assumes that any traffic being passed on port 25 is indeed e-mail. This assumption is faulty, since an attacker may also use port 25 to deliver an attack, thus bypassing any protection the firewall might provide.

How Hackers Attack

Multiple different mail servers are used in today's enterprises; chosen for performance, price, name recognition or any of a number of other reasons, servers such as Lotus Notes and Microsoft Exchange dominate the corporate e-mail landscape. Once a company has chosen a mail server, it is essentially beholden to that brand, as the primary server platforms are not interoperable. Each different mail server has its own set of known vulnerabilities, giving resourceful hackers ample opportunity to search for weaknesses. Once these weaknesses are identified, a single hacker can take down an entire rack of mail servers in the blink of an eye.

For the more technically minded see Appendix One to explain how they do it.

To execute a Denial-of-Service (DOS) attack, a hacker uses Trojans to take control over a potentially unlimited number of zombie computers, which then take aim at a single target and flood it with traffic in an attempt to overwhelm the server.

System Configuration Holes

Weaknesses in enterprise system configuration can be classified as follows:

- Default configurations – Most systems are shipped to customers with default, easy-to-use configurations. Unfortunately, “easy-to-use” can mean “easy-to-break-into” as well. Almost any UNIX or WinNT machine shipped can be exploited rather easily.
- Empty/Default root passwords – A surprising number of machines are configured with empty or default root/administrator passwords. One of the first things an intruder will do on a network is to scan all machines for empty passwords.
- Hole creation – Virtually all programs can be configured to run in a non-secure mode which can leave unnecessary holes on the system. Additionally, sometimes administrators will inadvertently open a hole on a machine. Most administration guides will suggest that administrators turn off everything that doesn't absolutely need to run on a machine in order to avoid accidental holes. Unfortunately this is easier said than done, since many administrators aren't familiar with disabling many common services.

Exploiting the Human Factor

Education of e-mail users by organisations regarding how hackers seek to exploit them has improved to the point that a large majority of e-mail users now have at least a rudimentary understanding of fundamental security. The basic message regarding not opening certain malicious attachment types, particularly .exe files, from unknown senders is widely known. This means the hackers are being forced to redouble their efforts in order to counteract the education that e-mail users are receiving.

Self-Propagation: The New Mission of Attacks

Hackers are becoming increasingly sophisticated and are no longer content with simply gaining access to networks to cause mischief and disrupt service. Whereas hackers first spread viruses through individual networks simply because they could, we now are seeing more and more attacks that involve the use of Trojans designed to spread a virus to as many computers as possible, with the intent of taking control of these machines for nefarious purposes.

Trojans

Trojans enter the victim's computer undetected, usually disguised as a legitimate e-mail attachment. Once the unsuspecting recipient opens the Trojan, the attacker is granted unrestricted access to the data stored on the computer. Trojans can either be hidden programs running on a computer, or hidden within a legitimate program, meaning a program that the user trusts will have functions they are not aware of. The following chart outlines some of the most common types of Trojans used by hackers:

Type	Purpose
Remote Access	Designed to give hacker access to the victim's machine. Traditionally, Trojans would listen for a connection on a port that had to be available to the hacker. Now Trojans will call out to hackers giving access to the hacker to machines that are behind a firewall. Some Trojans can communicate through IRC commands, meaning a real TCP/IP connection is never made.
Data Sending	Sends information back to the hacker. Tactics include key logging, searching for password files and other private information.
Destructive	Destroys and deletes files.
Denial-of-Service	Gives a remote hacker the power to start Distributed DoS (DDoS) attacks using multiple "Zombie" computers.
Proxy	Designed to turn the victim's computer into a proxy server available to the hacker. Used for anonymous TelNet, ICQ, IRC, etc. to make purchases with stolen credit cards, etc. Gives the hacker complete anonymity as trail leads back to infected computer.

Spreading Viruses via Trojans

Hybrid attacks that combine the use of Trojans and traditional viruses have become increasingly popular. An example of this is the notorious Nimba virus that used multiple methods to spread itself and managed to get past anti-virus software by using a behaviour not typically associated with viruses. Nimda managed to infect 8.3 million computers worldwide.

Typical Hacking Scenario

While not all hacker attacks are alike, the following steps outline what could be referred to as a “typical” attack scenario. Keep in mind that an attack on your organisation may look completely different from the one outlined below, as the methods used in attacks are constantly changing to adapt to improved security techniques.

Step 1: Outside Reconnaissance

The intruders will attempt to find out as much information as possible without actually exposing themselves. They will do this by finding public information or appearing as a normal user. In this stage, you really can't detect them. The intruders will do a 'whois' lookup to find as much information as possible about your network as registered along with your Domain Name. They will browse other public information, such as your public Web sites and anonymous FTP sites, and might search news articles and press releases about your company.

Additionally, many attackers will resort to social engineering steps in an effort to perform their outside reconnaissance. For example, an attacker might call an employee on the phone posing as a member of the Information Technology department. The attacker might then request personal information from the vulnerable employee such as username or password information. Unfortunately many unsuspecting employees when presented with a supposed “authority figure” will give any information at their disposal, thus putting the organisation at significant risk.

Step 2: Inside Reconnaissance

Here, intruders use more technically invasive techniques to scan for information, but still don't do anything physically harmful. They might do a “ping” sweep in order to see which machines are active. They might do a scan on target machines in order to see what services are available. They'll run programs in order to see what information is available. Hackers also will send e-mail to invalid users to receive error response so that they can determine information such as how many hops are involved in the mail system, where in the infrastructure the company does recipient checking on inbound e-mails, and other information that can be gleaned from the data captured in e-mail headers. At this point, the intruders have engaged only in “normal” activity on the network and have not done anything that can be classified as an intrusion.

Step 3: Exploit

At this point, the intruders cross the line and start exploiting possible holes in the target machines. The intruders might attempt to exploit well-known buffer overflow holes by sending large amounts of data, or may start checking for login accounts with easily guessable (or empty) passwords. The hackers may go through several stages of exploits. For example, if the hackers were able to access a user account, they will now attempt further exploits in order to get root/admin access.

Step 4: Foot Hold

At this stage, the hackers have successfully gained a foothold into your network by hacking into a machine. The intruders' main goal is to hide evidence of the attacks (doctoring the audit trail and log files) and make sure they can get back in again. They may install "toolkits" that give them access, replace existing services with their own Trojan horses that have backdoor passwords, or create their own user accounts. The hackers will then use the system as a stepping-stone to other systems, since most networks have fewer defences from inside attacks.

Step 5: Profit

This is where it can get really ugly for an organisation. The intruders now can take advantage of their status to steal confidential data, misuse system resources (i.e. stage attacks at other sites from your site), or deface Web pages, often receiving monetary rewards from behind-the-scenes benefactors.

Another scenario starts differently. Rather than attack a specific site, intruders might simply scan random Internet addresses looking for a specific hole.

Protect Your Organisation

As businesses place increasing reliance on e-mail systems, they must address the growing security concerns from both e-mail borne attacks and attacks against vulnerable e-mail systems. When enterprise e-mail systems are left exposed by insecure devices, hackers can enter the organisation and compromise the company's corporate backbone, rendering investments in information technology security useless.

The implications from a security breach can impact the company's reputation, intellectual property and ability to comply with government regulations. The only way for organisations to fortify their e-mail systems is to use a comprehensive e-mail security gateway to lock down the e-mail systems. This approach includes:

1. Locking down the e-mail system at the perimeter – Perimeter control for the e-mail systems starts with deploying an e-mail gateway. The e-mail gateway should be purpose-built with a hardened operating system, and intrusion detection capabilities to prevent the gateway from being compromised.
2. Securing access from outside systems – The e-mail security gateway must be responsible for handling traffic from all external systems, and must ensure that traffic passed through is legitimate. By securing access from outside, applications like Web mail are prevented from being used to gain access to internal systems.
3. Real-time monitoring of e-mail traffic – Real-time monitoring of e-mail traffic is critical to preventing hackers from utilizing e-mail to gain access to internal systems.

An e-mail security gateway should provide the following benefits:

Simplify Administrator Work

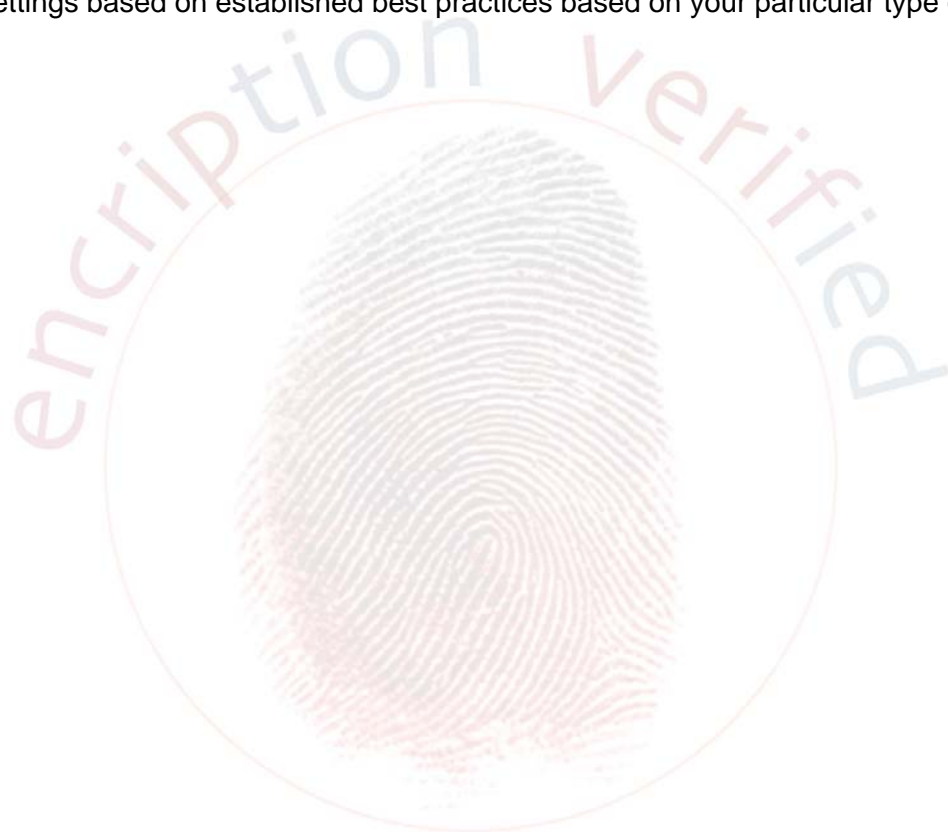
Rather than having multiple appliances from different vendors provide piecemeal protection for different areas of your e-mail network, the e-mail security solution that protects your organisation should be capable of protecting the entire e-mail system on its own. Comprehensive security must be purpose-built into the e-mail security appliance, not added as an afterthought.

Easy Integration

Integrating an intrusion detection/prevention system can be complicated, depending on your requirements. However, these systems must not complicate a network, and they should not require the administrator to spend additional time managing them.

Easy Configuration

Many intrusion detection systems are difficult to navigate and configure. A purpose-built e-mail security system containing intrusion detection and prevention should be easy to configure and manage, with settings based on established best practices based on your particular type of business.



Appendix One

The following sections are for the more technically minded and outline some of the vulnerabilities widely known within hacking circles. They also explain how hackers are able to take advantage of these security holes.

Scanning and Sweeping

- Ping Sweeps – Pings a range of IP addresses to find which machines are active. Sophisticated scanners will use other protocols (such as an SNMP sweep) to do the same thing.
- TCP Scans – Probes for open (listening) TCP ports, searching for services the intruder can exploit. Scans can use normal TCP connections or stealth scans that use half-open connections (to prevent them from being logged) or FIN scans (never opens a port, but tests if someone's listening).
- UDP Scans – Sends a garbage UDP packet to the desired port. Most machines will respond with an ICMP "destination port unreachable" message, indicating that no service is listening at that port. These scans are a little bit more difficult because UDP is a connectionless protocol.
- OS Identification – Identifies the operating system and applications by sending TCP packets. Each operating system's unique responses to inputs forms a signature that hackers can use to figure out what the target machine is and what may be running on it.

IMAP & POP Vulnerabilities

Hackers have found a number of issues in both IMAP & POP servers that are exploited. Items such as dictionary attacks can expose sensitive e-mail which is stored on an IMAP or POP server. There are countless tools available for performing these attacks and the graphical nature of many of these tools make it simple for even a novice to perform them. Additionally, weak passwords are common vulnerabilities in these protocols. Many organisations do not have adequate controls for password strength, thus end users will use passwords which can easily be broken. Lastly, there may be concerns about defects or bugs in various IMAP and POP services which can leave them susceptible to other types of exploits such as buffer overflows.

Denial-of-Service (DoS) Attacks

- Ping of death – Sends an invalid fragment, which starts before the end of packet, but extends past the end of the packet.
- Syn Flood – Sends TCP SYN packet (which starts connections) very rapidly, leaving the attacked machine waiting to complete a huge number of connections, and causing it to run out of resources and start dropping legitimate connections. A new defence against this is "SYN cookies." Each side of a connection has its own sequence number. In response to a SYN, the attacked machine creates a special sequence number that is a "cookie" of the connection, then "forgets" everything it knows about the connection. It can then recreate the forgotten information about the connection when the next packets come in from a legitimate connection.
- Loop – Sends a forged SYN packet with identical source/destination address/port so that the system goes into an infinite loop trying to complete the TCP connection.

Exploiting Software Issues

Software bugs can be exploited in the server daemons, the client applications, the operating system, and the network stack. Software bugs can be classified in the following manner:

- **Buffer Overflows** – Almost all the security holes you read about in the press are due to this problem. A typical example is a programmer who will set aside a specific number of characters to hold a login username. Hackers will look for these types of vulnerabilities, often sending longer strings than specified, including code that will be executed by the server. Hackers find these bugs in several ways. First, the source code for a lot of services is available on the net. Hackers routinely look through this code searching for programs that have buffer limitations. Hackers will also examine every place the program accepts input and try to overflow it with random data. If the program crashes, there is a good chance that carefully constructed input will allow the hacker to break into the system.
- **Unexpected Combinations** Programs usually are constructed using many layers of code, including the underlying operating system as the bottom-most layer. Intruders can often send input that is meaningless to one layer, but meaningful to another when constructed properly.
- **Unhandled Input** – Most programs are written to handle valid input. Most programmers do not consider what happens when somebody enters input that doesn't match the specification.

Examples of hackers using sophisticated means to get users to open e-mail attachments include the following:

- **Double Extension** – The Netsky, lovegate, and Klez viruses took advantage of this vulnerability. Malicious files are given double extension such as “filename.txt.exe” to trick the user into running the executable. NetSky actually would place 100 spaces between the extensions so the victim would not see the second extension. NetSky would also put the DOS command “COM” at the end of a string that appeared to be a Web address ending in .COM.
- **Password-Protected Zip File** – Virus writers encrypt the virus in a password protected zip and send the file to users with the password in the message body. Since the encrypted file skips virus scanning, the end user gets what they think is legitimate e-mail. Unfortunately, in most cases this message has a look of urgency and the unsuspecting user will many times go the extra mile to open the malicious attachment.
- **Plain Trickery** – Hackers harvest e-mail addresses from LDAP servers and spoofing the “from” field with names the victim would recognize so they open the e-mail and attachments, and by trying to trick the victim into accessing a Web site. Common tactics include sending e-mails with headings with “re:” or “Re: re: re:” included to make the victim believe it is a chain e-mail. Another common header tactic is including technical terms that make the victim believe that e-mail system error was encountered; MyDoom used this tactic effectively.