Global Information Assurance Certification Paper

Email Security Threats

GIAC Security Essentials Certification (GSEC)
Practical Assignment - Version 1.4b
Option 1

By
Pam Cocca

September 20, 2004

TABLE OF CONTENTS

ABSTRACT

Email security has become a hot topic in Information Technology circles as new exploits and vulnerabilities affecting the most popular email clients and operating systems continue to make headline news on a regular basis. When you consider that a recent META group survey found 80% of survey respondents said they consider e-mail more valuable than phone for business communications,[1] it is no wonder that email security is a priority concern for many organizations.

In this paper I will outline the various threats to email security, focusing on those that are of particular concern. I will then review some of the most recent advancements in the industry that are aimed at solving some of these issues.

WHAT DOES EMAIL SECURITY INVOLVE?

The three main principles of Information Security involve maintaining the confidentiality, integrity, and availability of information resources. These three principles can be directly applied to the area of email security as well. Confidentiality of email involves making sure it is protected from unauthorized access. Integrity of email involves a guarantee that it has not be modified or destroyed by an unauthorized individual. Availability of email involves ensuring that mail servers remain online and able to service the user community. A weakness in any one of these three key areas will undermine the security posture of an email system and open the door to exploitation.

WHAT ARE THE THREATS TO EMAIL SECURITY?

Viruses

Email security is threatened by a range of issues. One of the most publicized and high risk of all the issues is viruses. Viruses are so dangerous because they often deliver extremely destructive payloads, destroying data, and bringing down entire mail systems. As a result they are a major drain on corporate IT departments and users.

According to an ICSA Labs 2003 Virus Prevalence Survey, in 2003 nine of the Top 10 reported viruses were mass mailers. Also, all of the viruses that were responsible for actual disasters during that time were either Internet worms or mass mailer viruses. To make matters worse, both of these virus types tend to stay around longer than other types, even after anti-virus products have included protection against them in their products.[2]

The following table (taken from the ICSA survey) shows virus encounters per month for the period January – December 2003. Note that of the Top 10 reported viruses, nine were either mailers or mass mailer viruses. The exception to this was Blaster, which was a worm that exploited a DCOM RPC vulnerability but did NOT contain any mass-mailing functionality.[3]

2003 Rank	Virus Name	Encounters
1	W32/Yaha	32
2	W32/Klez	29
3	W32/Mimail	22
4	W32/BugBear	18
5	W32/SirCam	12
6	W32/Sobig	7
7	W32/Dumaru	6
8	W32/Swen	5

9	W32/Lovgate	4
10	W32/Blaster	2

Taken from the ICSA Labs 2003 Virus Prevalence Survey, Table 2: Top Viruses for 2003

In the same ICSA survey, it was identified that email as the source of virus infection has been steadily increasing:

Virus Source	**1996**	**1997**	**1998**	**1999**	**2000**	**2001**	**2002**	**2003**
Email Attachment	9	26	32	56	87	83	86	88
Internet Downloads	10	16	9	11	1	13	11	16
Web Browsing	0	5	2	3	0	7	4	4
Don't Know	15	7	5	9	2	1	1	3
Other Vector	0	5	1	1	1	2	3	11
Software Distribution	0	3	3	0	1	2	0	0
Diskette	71	84	64	27	7	1	0	0

Taken from the ICSA Labs 2003 Virus Prevalence Survey, Table 10: Sources of infection, 1996-2003

The impact of viruses on organizations is huge. The impact goes far beyond money, resources, and effort required to recover from such incidents. It also includes loss of productivity, corrupt and/or lost data, and loss of user confidence.

SPAM

Another major threat to email security today is SPAM, often cited by organizations as being their *number one* concern. Otherwise known as junk email, SPAM is considered a security threat not only because the volume of it can affect system availability, but also because it can carry viruses, malicious code, and fraudulent solicitations for private information.

It is an ever-growing problem that is of particular concern to information security professionals. Analysts IDC reported in a study earlier this year that spam represented 32 per cent of all email sent on an average day in North America in 2003, doubling from 2001. That figure is less than the 50 per cent or more junk mail statistic commonly cited by email-filtering firms like MessageLabs and Brightmail but it still represents a serious problem.[4]

Businesses lose money when SPAM overloads network and server resources. Even with spam filtering mechanisms in place employees inevitably end up

spending inordinate amounts of time sorting through messages trying to distinguish legitimate emails from SPAM.

In a survey conducted by Information Security/SearchSecurity.com on the business implications of spam, lost productivity (92 percent) and clogged email servers (62 percent) were cited as the most painful consequences of spam. However, a growing concern is the threat of virus propagation via spam. Many security professionals fear that virus writers and spammers could get together and collaborate on more invasive ways of compromising networks and circumventing filters.[5] That will make the line between what is a virus and what is spam more fuzzy than it has already become, when you consider the consequences of spam on resource utilization.

Well-intended SPAM?

Described by research firm Gartner as 'Friendly Fire', the volume of email being sent by well-meaning friends and family to employees is on the increase.[6] Although the statistics available on this particular issue vary greatly, SurfControl Inc. cited in their whitepaper 'Fighting the New Face of Spam' that friendly junk email could cost a company with 500 employees nearly $750,000 each year.[7] Although email from family and friends may pose less of an overall security risk the *volume* of it can certainly affect *availability*. And there is increased risk also if you consider that many home users sending these 'friendly' emails are sending them from less secure systems than we find on a corporate network where often virus definitions are out of date and systems are unpatched. This makes it even more important for organizations to ensure they have systems in place to protect against not only the obvious, but even the seemingly well-intentioned.

Phishing

Phishing, also known as identify theft, is a newer threat to email security that was relatively unheard of one year ago. Phishing is the process whereby identity thieves target customers of financial institutions and high-profile online retailers, using common spamming techniques to generate large numbers of emails with the intent of luring customers to spoofed web sites and tricking them into giving up personal information such as passwords and credit card numbers.

It is a problem that has literally exploded over the last year. A study released by Gartner Research in May estimates that 76 percent of all known phishing attacks had occurred since last December. The Anti-Phishing Working Group (www.antiphishing.org), an industry association of more than 200 organizations, reported 1,125 unique phishing attacks in April, up from 402 in March and nearly seven times the number reported in January.[8] It is expected that these numbers will continue to climb drastically as security professionals struggle to find an effective solution to the problem.

Phishing has the potential to be highly lucrative for the 'Phisher', the individual or organization staging the attack. For the most part, a phishing attack is easy and cheap to engineer, is extremely hard to trace, and even if only a small percentage of recipients respond to requests for personal information – the return on investment can be very high. Although early phishing attacks were marked by misspellings, improper grammar, and less than perfect imitations of corporate logos and websites, Phishers are becoming more sophisticated in both the quality of their scams and the techniques they are using making this a growing security risk.

Gartner estimates the direct cost to companies of phishing attacks was $1.2 billion in 2003.[9] Given the sharp rise in the number of phishing attacks so far reported in 2004, its obvious losses in 2004 will exceed last year's numbers. The impact of phishing attacks against organizations doesn't stop with direct losses. Companies are also faced with downtime during an attack, having to issue new credentials to customers who have compromised their personal information, potential liability, and damage to their corporate image. And if phishing can't be brought under reasonable control consumers are going to become extremely reluctant to do business online (therefore loss of consumer confidence).

WHAT CAN WE DO?

There is a variety of mail security products on the market today, aimed at addressing the various threats to email security. They come in the form of special software that you can load on an existing mail server or on a dedicated mail gateway platform, or in the form of a hardware appliance that acts as an email gateway. Another option for companies is to outsource mail security to an outsourced service provider. All of these scenarios typically offer a similar feature set, although there are definite differences among competing products in terms of what they have to offer.

Some of the common features in mail security products today include content filtering services such as antivirus, antispam, HTML tag removal, script removal, block of attachments by file type, scanning of inappropriate content, confidentiality checks, and disclaimer enforcement. Antispam methods supported by most products include real-time blackhole lists (RBL), heuristics, confirmation process, Bayesian filtering, open relay protection, size and bandwidth control, and encryption.

Despite all the advancements in email security products, we continue to see an increase in the number of security related issues. Virus writers are continuously looking to exploit vulnerabilities in systems and software, and make every attempt possible to cover their tracks. Spammers are constantly changing the

appearance of spam and masking its source to avoid it being blocked before it reaches it's target. It is evident in both of these scenarios that one of the biggest challenges in solving the virus and spam problem is in identifying the origin of email messages. As a result the industry is crying out for radical changes to the email infrastructure that will bring these problems under control. Some of the major initiatives over the last year intended to address these ongoing issues involve *Sender Authentication*. They include the Sender Policy Framework, Caller ID for Email, the Sender ID Framework, DomainKeys, and Accreditation and Reputation Services.

Sender Policy Framework

One of the first technologies developed to authenticate the sender of an email message was the Sender Policy Framework (SPF). It is a technology created by Meng Wong (founder of email service firm pobox.com) that aims to identify the origin of email messages.

How does it work?

Currently, all domains already publish MX records to name servers (DNS) on the internet to let everyone know what machines *receive* mail for their domain. This is done so that mail servers know where to send mail destined for those domains.

SPF functions by publishing a type of reverse MX record (the SPF record) as well, which specifies what machines *send* mail from their domain. So, when a message is received from a domain, the recipient of the message can look up the SPF record of the sending domain and compare this to what is contained in the MAIL FROM: field within the message to verify that it did in fact come from where it should be coming from. So, if you were to receive a message from not_a_real_address@anydomain.com, the receiving mail server would go out and lookup the SPF record for anydomain.com. That ip address would then be compared to what is contained in the email header for the sending machine's ip address. If they match, the email would pass and you could be fairly sure the sender is who they say they are. If the message would fail the SPF test, it would indicate the message was a forgery and most likely that the sender is a spammer.

It is important to realize though, that SPF was designed to protect the *envelope sender*, in other words the return-path that shows up in "MAIL FROM", not the header "FROM". Most implementations of SPF today only use the return-path as the subject of authentication due to technical challenges. For those wishing to protect the "From" header, proponents of SPF recommend using a cryptographic technology such as S/MIME, or PGP.[10]

Caller ID

Another technology aimed at Sender Authentication, developed by Microsoft, is called Caller ID for Email. Similar in many ways to SPF, Caller ID specifies what is called a Purported Responsible Address (PRA) record, instead of an SPF record. The difference between the two is basically the algorithm used to determine the address that is checked for authenticity. SPF uses the visible email address of the sender, while PRA checks the record against the most recent sender of the email message. So, PRA indicates where the email came from most recently, SPF indicates from where the email initially came.

After Microsoft announced it's plan earlier this year to pursue the standardization of it's Caller ID technology,[11] it ended up proposing a hybrid specification to the Internet Engineering Task Force (IETF) combining it's Caller ID technology with SPF. The hybrid solution is known as the **Sender ID Framework**, and also comprises a third specification called Submitter Optimization.

The Sender ID Framework

- companies would publish their SPF records in DNS
- receiving mail server would look up the SPF record in the Sender's DNS record
- receiving mail server would determine the PRA, the compare the PRA to legitimate IP addresses in the SPF record
- a match would indicate a pass (the origin of the email has been authenticated)

The PRA (Purported Responsible Address) is the email address of the entity most recently responsible for injecting a message into the email system. It would be different from the initial author/sender if the message has traveled multiple hops. It is derived from the message headers (Resent-Sender, Resent-From, Sender, From). [12]

Sender Optimization is an optional extension to the SMTP MAIL command that would allow the receiver to check for spoofing BEFORE the message is sent across the internet. It allows the sender to declare the PRA within the SMTP protocol.

If implemented, a SUBMITTER= parameter would be specified on the MAIL FROM: command, if the PRA is different from the MAIL FROM. This would be a necessary requirement for mailing list servers and mail forwarders where the MAIL FROM will almost never match the PRA.

Some implementation examples of Sender Optimization where the submitter parameter would be used (taken from http://www.microsoft.com/mscorp/twc/privacy/spam_senderid.mspx):

NORMAL MAIL SUBMISSION
S: 220 alumni.almamater.edu ESMTP server ready
C: EHLO example.com
S: 250-alumni.almamater.edu
S: 250-DSN
S: 250-AUTH
S: 250-SUBMITTER ***(SUBMITTER extension advertised in EHLO response)***
S: 250 SIZE
C: MAIL FROM:<alice@example.com> SUBMITTER=alice@example.com ***(SUBMITTER parameter added to MAIL command)***
S: 250 <alice@example.com> sender ok
C: RCPT TO:<bob@alumni.almamater.edu>
S: 250 <bob@alumni.almamater.edu> recipient ok
C: DATA
S: 354 okay, send message
C: From: alice@example.com
C: (message body goes here)
C: .
S: 250 message accepted
C: QUIT
S: 221 goodbye

MAILING LIST
S: 220 example.com ESMTP server ready
C: EHLO listexample.com
S: 250-example.com
S: 250-SUBMITTER ***(SUBMITTER extension advertised in EHLO response)***
S: 250 SIZE
C: MAIL FROM:<owner-list1@listexample.com>
SUBMITTER=owner-list1@listexample.com ***(SUBMITTER parameter added to MAIL command)***
S: 250 <owner-list1@listexample.com> sender ok
C: RCPT TO:<alice@example.com>
S: 250 <alice@example.com> recipient ok
C: DATA
S: 354 okay, send message
C: Received By: ...
C: From: bob@woodgrove.com
C: Sender: owner-list1@listexample.com ***(Sender header added to message)***
C: To: list1@listexample.com
C: (message body goes here)
C: .
S: 250 message accepted
C: QUIT
S: 221 goodbye

MAIL FORWARDING
S: 220 woodgrove.example ESMTP server ready
C: EHLO alumni.almamater.edu
S: 250-woodgrove.example
S: 250-DSN
S: 250-AUTH
S: 250-SUBMITTER ***(SUBMITTER extension advertised in EHLO response)***
S: 250 SIZE
C: MAIL FROM:<alice@example.com>

SUBMITTER=bob@alumni.almamater.edu ***(SUBMITTER parameter added to MAIL command)***
S: 250 <alice@example.com> sender ok
C: RCPT TO:<bob@woodgrove.example>
S: 250 <bob@woodgrove.example> recipient ok
C: DATA
S: 354 okay, send message
C: Resent-From: bob@alumni.almamater.edu ***(Resent-From header added to message)***
C: Received By: ...
C: (message body goes here)
C: .
S: 250 message accepted
C: QUIT
S: 221 goodbye

The Sender ID Framework has been debated by the IETF for the past several months. Among several issues with the proposal is Microsoft's attempt to patent technology used for the Caller ID component, which may end up meaning Sender ID would require users to sign a license agreement. This has angered many in the open source world, and has somewhat soured the support of some that had previously backed the technology. The Sender ID proposal was most recently dealt a setback on September 11th when the IETF reached consensus that Microsoft's patent claims should not be ignored and their insistence on keeping the technology secret was unacceptable.[13]

After the results of the IETF vote, Microsoft indicated it will continue with its plans to develop its own proposal for Caller ID. They stated however, that they will use the Purported Responsible Address (PRA) to authenticate the source of email messages although they will continue to publish both SPF and PRA records (they will only check the PRA).

In the meantime, the proposal for the Sender ID Framework is not necessarily dead. The IETF ruling allows for negotiation, if Microsoft considers removing licensing restrictions. Given that some of the biggest email providers in the world (AOL, Microsoft, Yahoo, Comcast, Earthlink, and BT) have been promoting Sender ID, and products like Sendmail are adding support to their mail transfer agents, Sender ID is likely still to be further debated at the IETF.

DomainKeys

Domain Keys is a technology proposal developed by Yahoo, that provides a mechanism for verifying both the domain of each email sender and the integrity of the messages sent using DNS and an RSA public/private key method to digitally sign messages.

Overview of how the technology works:

- The owner of a domain generates a public/private key pair to use for signing all outgoing messages. The public key is then published in DNS,

and the private key is made available to their DomainKey-enabled outbound email servers.
- When an email is sent by an authorized user within the domain, the DomainKey-enabled mail system automatically uses the stored private key to generate a digital signature of the message. The signature is then pre-pended as a header to the email and the email is sent on to the target recipient's mail server.
- At the receiving end, the DomainKeys-enabled email system extracts the signature and the claimed FROM: domain from the email header and retrieves the appropriate public key from DNS for the domain.
- The public key is then used to verify that the signature was generated by the matching private key. This would prove that the email was sent by (and with permission of) the claimed sending FROM: domain and that its headers and content weren't altered during transfer.
- The receiving email system could then apply local policies based on the results of the signature test.[14]

Yahoo's proposal for how DomainKeys would be implemented suggests that the receiving email servers would be doing the verification. However, they suggest that end-user mail clients could also be modified to verify signatures and take actions based on the results. The benefits of having the receiving mail servers do the verification would be a reduction in the number of MTAs that have to be changed to support an implementation of DomainKeys, a reduction in the number of MTAs involved in transmitting the email between a signing system and a verifying system (thus reducing the number of places that can make accidental changes to the contents), and removing the need to implement DomainKeys within an internal email network.[15]

The use of DomainKeys has the potential to add a fair bit of processing load to outbound mail servers. However, the validity of the signatures and keys used in the technology and therefore the validity of the content in the messages and message headers would be almost guaranteed. Even in the situation where the keys might be compromised, new keys could be regenerated quickly.

It has been proposed that a technology like DomainKeys be used as a compliment to a technology such as SPF. SPF could be used to initially validate a legitimate message and then once it would pass the SPF pre-screen, it could go on to stand up against the DomainKeys test for further validation. This would avoid all the extra processing DomainKeys would require unless it was necessary.

Yahoo submitted their proposal for DomainKeys to the IETF in August of 2004 for consideration. It is too early to know if it will be approved, however it has already gained support from the likes of Sendmail, Port25 Solutions Inc., and CipherTrust.

Accreditation and Reputation Services

In December of 2003, the Aspen Policy Institute held a policy conference to discuss a framework for what they called "The Accountable Net". A key element of the framework called for a rethink of the email infrastructure, involving implementation of sender authentication technologies, and establishment of accreditation and reputation services.[16] They suggested that reputation and accreditation will make it possible to confidently distinguish between good senders and bad, bringing credibility to email.

The concept of accreditation and reputation services begins with authentication technologies. Those technologies provide a mechanism for being able to sort domains into 'good' and 'bad' by verifying the origin of email. Once identified a database of known good domains as well as known bad domains could be built and in effect establish a simple reputation system. A persistent reputation profile could be established for each sending domain that could then be tied into anti-spam policy systems and shared between service providers.

However, an important consideration in this model is the domains that don't make it onto either the good list or the bad list for one reason or another (i.e. recently registered domains). This is where accreditation comes in.

A domain that hasn't made it onto the good list yet, would be able to sign up with a third-party organization that offers accreditation services and that would publicly vouch for them as senders. Once the domain has been around long enough that it would make it onto a well-recognized 'good' list, they would no longer require endorsement by such a third party.

SUMMARY

Recent technological advancements and theories about how to solve the SPAM problem are no doubt promising. It is an important security threat to address, especially if our fears are realized and virus writers increasingly take advantage of spamming techniques for the purposes of virus propagation.

However, developers of the various sender authentication schemes have long maintained that their technologies alone will not stop SPAM. This is even more evident after a recent CipherTrust study indicated that spammers are supporting SPF faster than legitimate email senders, with 38 percent more spam messages registering SPF records than legitimate email.[17]

Reputational analysis based on Accreditation and Reputation services promises to pick up where sender authentication leaves off. The success of reputational

analysis however, is going to depend not only on it's widespread adoption by anti-spam vendors but also on end users' reporting senders as spammers.

SPAM strikes at the heart of the three main principles of Information Security. Given that and the challenges associated with the ever-changing face of SPAM, it is expected to continue to be the number one security concern facing organizations for some time to come.

---

[1] META Group Inc. "Spam, Viruses, and Content Compliance: An Opportunity to Strategically Respond to Immediate Tactical Concerns." August 2003.

[2] Bridwell, Larry. "ICSA Labs 9th Annual Computer Virus Prevalence Survey." 2004. URL: https://www.trusecure.com/cgi-bin/download.cgi?file=wp_vps2003_report.pdf&ESCD=w0169 (August 30, 2004)

[3] Bridwell, Larry. "ICSA Labs 9th Annual Computer Virus Prevalence Survey." 2004. URL: https://www.trusecure.com/cgi-bin/download.cgi?file=wp_vps2003_report.pdf&ESCD=w0169 (August 30, 2004)

[4] Leyden, John. "One third of email is now spam." April 2004. URL: http://www.theregister.co.uk/2004/04/20/idc_spam_survey/ (August 30, 2004)

[5] Plante, Amber. "Stuffing Spam." May 2004. URL: http://infosecuritymag.techtarget.com/ss/0,295796,sid6_iss386_art765,00.html (August 15, 2004)

[6] Trudeau, Paris. "Fighting the New Face of Spam." URL: http://www.surfcontrol.com/general/assets/whitepapers/New_Face_of_Spam.pdf

[7] Trudeau, Paris. "Fighting the New Face of Spam." URL: http://www.surfcontrol.com/general/assets/whitepapers/New_Face_of_Spam.pdf

[8] http://infosecuritymag.techtarget.com/ss/0,295796,sid6_iss426_art874,00.html

[9] http://infosecuritymag.techtarget.com/ss/0,295796,sid6_iss426_art874,00.html

[10] http://spf.pobox.com/faq.html

[11] http://www.microsoft.com/presspass/features/2004/Feb04/02-24CallerID.asp

[12] http://www.microsoft.com/mscorp/twc/privacy/spam_senderid.mspx

[13] http://news.com.com/2100-1032-5364075.html

[14] http://antispam.yahoo.com/domainkeys

[15] http://www.ietf.org/internet-drafts/draft-delany-domainkeys-base-01.txt; section 1.4

[16] Jahnke, Art. "Can White Lists Defeat Spam." May, 2004. URL: http://comment.cio.com/soundoff/050604.html

[17] Callaghan, Dennis. "Spam spotlight on reputation." eWeek 6 Sept. 2004: 12.