5 Essential Steps to Safer Email

How to Protect Small and Medium Businesses from Email Threats.

BY TOM GILLIS

WITH A FOREWORD BY
MICHAEL OSTERMAN

A MESSAGING NEWS PRESS PUBLICATION

DISCLAIMER: The law in this area changes rapidly and is subject to differing interpretations. It is up to the reader to review the current state of the law with a qualified attorney and other professionals before relying on it. Neither the authors nor IronPort make any guarantees or warranties regarding the outcome of the uses to which this material is put. This paper is provided with the understanding that the authors and IronPort are not engaged in rendering legal or professional services to the reader.

Table of Contents

> “The combination of effective technology and a focus on best practices can help messaging managers to maintain email’s role as the corporate world’s most critical application.”
>
> **MICHAEL OSTERMAN**
> President and Founder, Osterman Research, Inc.

Forward

by Michael Osterman

It almost goes without saying that email is the most critical application in use by organizations large and small. Email is used in virtually every organization by at least some, if not all, employees—and its use is growing at a rapid pace. Organizations increasingly use email as the primary method for communicating with employees, managers, customers and prospects. Four out of five organizations use email for critical activities like transmitting and accepting proposals, finalizing agreements and transmitting business-critical records of all sorts. Email has become the de facto file transport mechanism for almost all organizations and the best way for employees to communicate while at home, traveling, and at their desks.

The Growing Problem with Inbound Email

The dominance of email for corporate communication has been driven in large part by its extremely low variable cost, its ease of use and the fact that the SMTP standard has made email interoperable worldwide. However, these factors have also made email one of the most vulnerable infrastructure elements currently running on corporate networks and the avenue through which an enormous number of threats have entered these networks. For example, email is the primary avenue by which viruses, worms and Trojan horses enter corporate networks, causing problems that range from irritating pop-ups to the complete destruction of corporate data. Email has become dominated by spam with the result that three out of every five email messages received by the typical email user is an unwanted message. More insidiously, email is also the vehicle used by criminals to fraudulently obtain sensitive personal information like credit card or bank account numbers through what are known as phishing attacks.

Problems Start Inside the Organization, As Well

The problem for those charged with maintaining the integrity of their corporate email systems, as well as those who use those systems, does not stop there. In addition to viruses, spam and phishing attacks, organizations are increasingly vulnerable to information that is sent not only to their users, but also by them. Audits of email content sent from corporate networks reveal that email users often—and typically inadvertently—send messages that contain sensitive corporate data like passwords, credit card numbers, intellectual property, and financial information. Further, many employees will say things in email that can have a serious impact on corporate reputations, often with embarrassing results for their employers when this information is leaked to third parties or is brought out during a legal action. Further complicating the issue is the growing array of regulations, such as Sarbanes-Oxley and the Health Insurance Portability and Accountability Act (HIPAA), that focus on the security and preservation of email content.

Email is Relied on More and Trusted Less

Serious problems are being caused by this growing array of email threats. For example, while email is an excellent method for legitimate marketers to inform prospective customers about their offerings, spam has caused recipients to be very distrustful of any sort of marketing message received via email. Users increasingly employ email for sending file attachments, but viruses and other threats carried in attachments have forced organizations to increasingly block email attachments. The net result is that email is becoming increasingly important as a critical business tool, and trusted less by the people who need it.

Changes in the Email Landscape

To combat the growing threats posed by viruses, worms, Trojan horses, spam, phishing, spyware, and other threats introduced to the organization through email; and to protect organizations from employees who often inadvertently send sensitive content out of the organization; people who manage email systems for their organizations must do more simply to maintain email's usability and utility. A number of new protocols, techniques and best practices are emerging for protecting organizations from the growing variety of external and internal email threats, including domain authentication, traffic shaping, development of better email policies, user education, and other techniques and practices. New offerings from a growing array of vendors promise to combat email threats more effectively—while reducing the quite serious problem of false positives (tagging messages as threats when they are, in fact, valid messages), the bane of email threat management systems.

The Bottom Line for Messaging Managers

IT staff and others charged with maintaining the integrity of their corporate email systems must continually do more with resources that typically do not grow as quickly as the threats that face them. Consequently, they must become more effective with the tools and techniques they have available. The good news is that vendors, like IronPort Systems, are responding to this challenge by introducing increasingly sophisticated systems that more effectively prevent threats from entering or leaving networks and that allow those who manage email systems to handle these threats more efficiently. In addition to deploying more capable systems, however, those who manage email systems must become more proactive by educating users about the dangers of email, establishing corporate policies about email use, and ensuring that users are familiar with and comply with these policies.

In short, the combination of effective technology and a focus on best practices can help messaging managers to maintain email's role as the corporate world's most critical application. The booklet that you're about to read will help you understand the key issues involved in protecting your email system—helping your users to get the most out of email and helping those who manage email for your organization to do so efficiently and effectively.

President and Founder, Osterman Research, Inc.

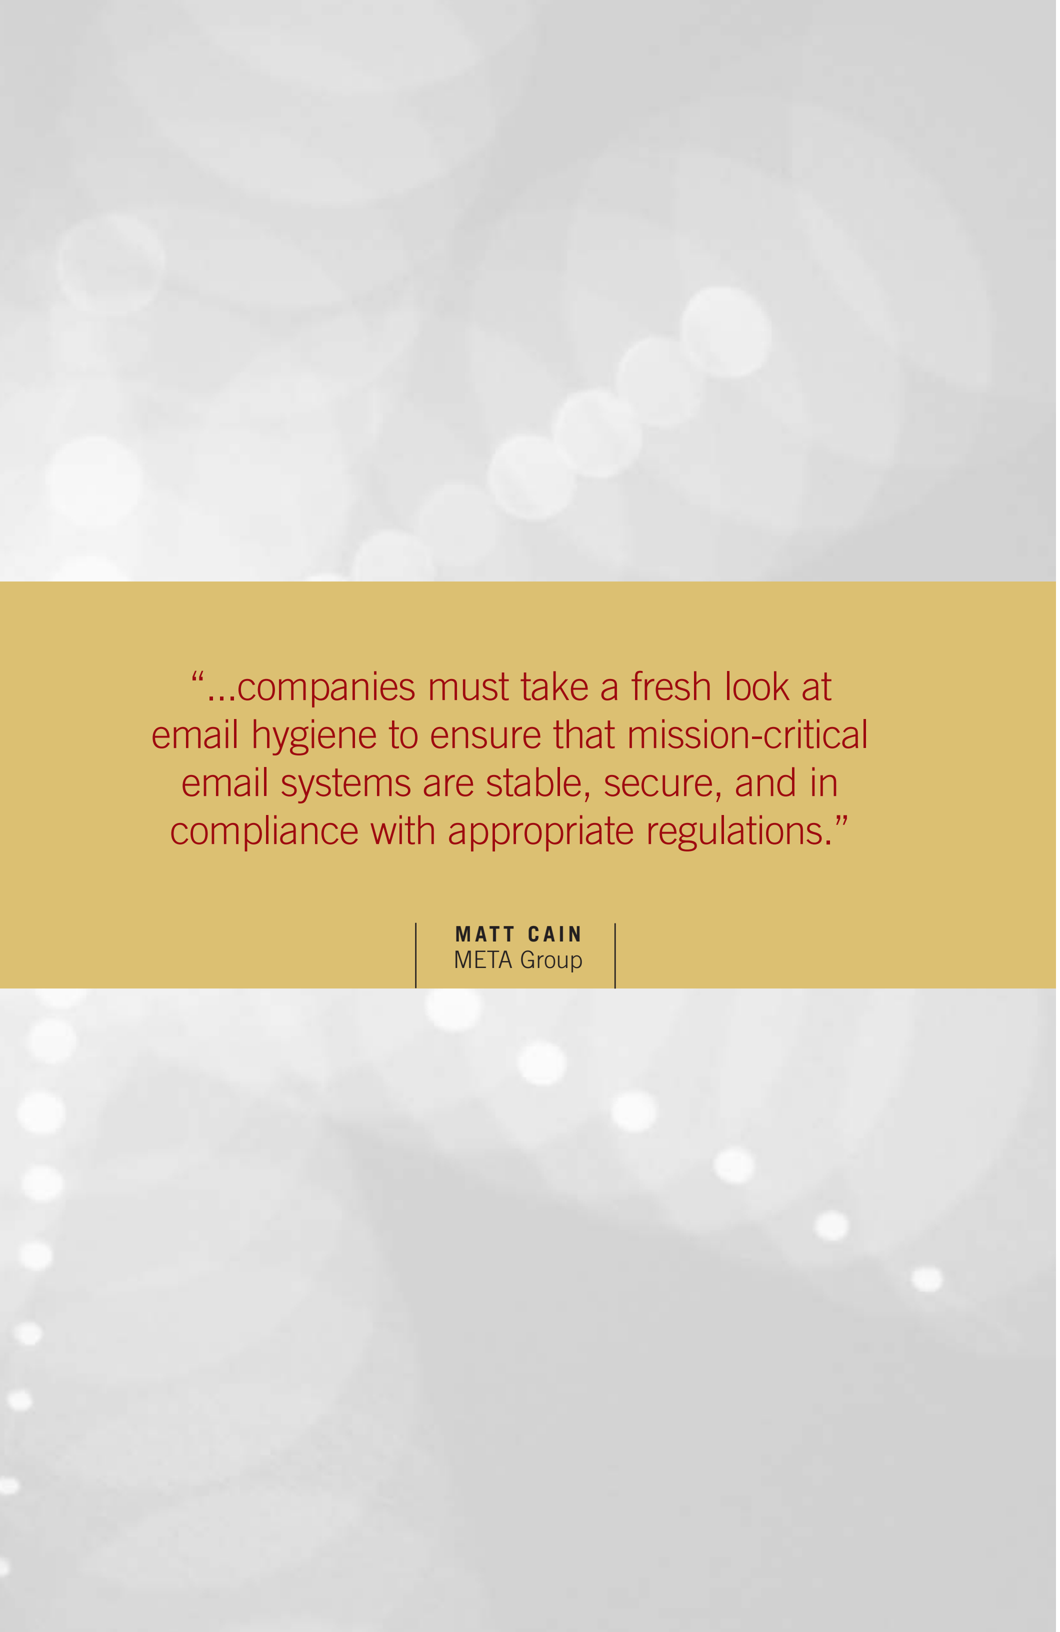

> “...companies must take a fresh look at email hygiene to ensure that mission-critical email systems are stable, secure, and in compliance with appropriate regulations.”
>
> **MATT CAIN**
> META Group

Introduction

Email has become the world’s most important form of business communication. The low cost, high efficiency, and ubiquity of email makes us wonder what life was like before its widespread adoption. Today the question is no longer “do you have an email address?,” but rather “what’s your email address?”.

But email is a victim of its own success. The very attributes that make it so compelling for business communication, have also made it attractive to those who use it for illicit and illegal forms of marketing. Today’s business email systems must contend with an ever growing volume of spam, viruses, fraudulent or “phishing” email, and (the latest scourge) email borne spyware. In addition to these inbound threats, companies are growing increasingly aware of the need to stop outbound threats— intellectual property leaving the company by email or outbound email subject to regulatory requirements.

FAST FACT

Today’s users are often notified by an email recipient that a virus was detected in an email they sent, however, the user never sent the email.

Sagging under the weight of these unending threats, the infrastructure used to send and receive mail is entering a period of rapid change. New authentication protocols are being developed to attack the spam and virus problems at their core. Also, new techniques and standards are being developed for the handling of bounce messages, a huge headache for the entire Internet community.

This booklet will attempt to cover the basics needed for a modern email security solution:

1. Stopping Spam
2. Stopping Viruses
3. Protecting Your Identity
4. Outbound Scanning
5. Fixing Email

1

STOPPING SPAM

> “Within the next few years, U.S. businesses will be spending almost $2 billion annually trying to keep spam from their doorstep.”
>
> IDG

STEP 1

Stopping Spam

The first generation of email security solutions used a simple approach to stopping spam—keyword analysis. These early filters would look for words typically found in spam (words like “free”, “Viagra”, or other more spicy language). The filters would typically use a scoring algorithm—if the word “free” occurs next to “Viagra” than it’s probably spam. The problem with this approach was twofold. The first issue being that it would frequently trap legitimate messages—Viagra is actually a product used in business, and the word free is almost unavoidable in the business lexicon. The other drawback to keyword filtering is it is relatively easy for spammers to defeat by using a zero instead of the letter o (I L0ve Y0u) or adding blocks of text that would fool the filters.

Nearly all modern spam systems have moved to a two-layer defense. The outer layer is known as a reputation filter. A reputation filter asks the simple question, “who is sending this email?” before accepting it. By examining the reputation or sending history of a given sender, the vast majority of spam can be eliminated before it even enters the network.

At the heart of any reputation system is a database that identifies “good guys” from “bad guys”. As might be expected, the quality and accuracy of a reputation filter is directly tied to the quality and accuracy of the underlying database. Ironport Systems invented the concept of reputation filtering in 2003. The underlying database behind IronPort’s solution is *SenderBase*®—the world’s first, largest, and most accurate reputation database. *Senderbase* collects data from more than 100,000 networks that make up over 25 percent of the entire world’s email traffic. This massive data footprint means that *SenderBase* can detect the sending patterns of literally every mail server on the Internet, in real time.

IronPort is the only company in the industry that shares this valuable data with entities outside of its customer base. IronPort has made *SenderBase* available to select ISPs and open source programs. IronPort also makes *SenderBase* data available via a web portal at www.senderbase.org. This open policy has lead to widespread adoption of *SenderBase* as the default reputation service—and the more entities that use the *SenderBase* data, the better the quality of data. Note that *SenderBase* is not licensed to other commercial solutions—it is embedded into the IronPort appliances.

FAST FACT

***SenderBase* is the first and largest traffic monitoring network.**

SenderBase measures objective parameters about a given mail server. In total, more than 150 different parameters are measured, such as how much mail does an IP send, do they accept mail in return, how long have they been sending mail from a given IP address, and what is their country of origin. This data is then rolled up into a score, using a statistical algorithm. The score is made available to all IronPort appliances. The appliances then use the score to determine how much, if any, mail to accept from a given sender. Coupling the reputation score with the ability to rate limit a given sender is another IronPort innovation. The system has the capacity to "push back" and slow down senders that appear suspicious, but not necessarily block them outright. This capability allows the IronPort appliance to deal with questionable senders—senders that appear to be spamming, but not conclusively. By rate limiting these senders the most hostile mail can be kept out of the system, without introducing the false positive problems associated with first generation systems.

The process of reputation filtering is very similar to consumer credit systems. Every sender has a reputation, and the email transactions of that sender are tracked, just like the transaction history of any individual is tracked by a credit bureau. These transactions are rolled into a score, and the score is made available to merchants in the credit example, and to receiving mail servers in the email example. The merchant then makes a determination on how much credit (if any) to extend, just like the receiving IronPort appliance makes a determination of how much (if any)mail to accept. This simple but powerful concept is very effective—blocking as much as 80 percent of incoming spam at the connection level, before it even enters the network. Network level blocking has the added benefit of saving bandwidth and system resources.

IronPort Reputation Filters™ take known good, trusted senders and will route them directly to the anti-virus scanners. Known bad senders are (typically) blocked. Senders in the middle are rate limited and sent to a second stage of filtering, known as context analysis. IronPort has developed a *Context Adaptive Scanning Engine* (*CASE*). *CASE* technology does a second examination of messages, asking 4 basic questions—who sent it, what does it contain, where does it direct the user, and how was it constructed. It's almost a rule set of common sense: examine the who, what, where and how of a message. For example, if the *CASE* is analyzing a message that contains multiple references to "Viagra" (the what?), this message is considered suspicious. But if the message in question is coming from a known pharmaceutical company (who?) and doesn't contain any links to an external online pharmacy (where?), then the *CASE* will determine the message is valid. If this message had been examined without the benefit of a full context analysis it would likely have been marked as spam.

FAST FACT

Most leading spam filters have copied IronPort's reputation filtering.

The combination of reputation filtering to sort out the obvious good from bad and then the more careful context filtering to evaluate a message in its full context makes for a "1-2 punch" against spam. This architecture has been deployed at more than 25 percent of the world's largest enterprises and has proven to be very effective at stopping spam, without suffering from false positives.

Many small and medium businesses have had experience with first generation spam filters that rely on simple key words to identify spam. As discussed earlier, these filters have proven to be fairly inaccurate, letting spam in and occasionally blocking legitimate messages. To account for this, first generation solutions include complex end user controls that allow end users to whitelist or blacklist certain senders. End users also need to check a quarantine to review spam messages and make sure none are legitimate. These tools are really a work-around for an inaccurate spam filter. While the IronPort solution supports much of this end user facing functionality, most IronPort users choose not to enable it. Therefore, all the end users know is that email works again—and their IT team is genius.

FAST FACT

A modern spam filter is accurate enough to not require an end user quarantine or end user control.

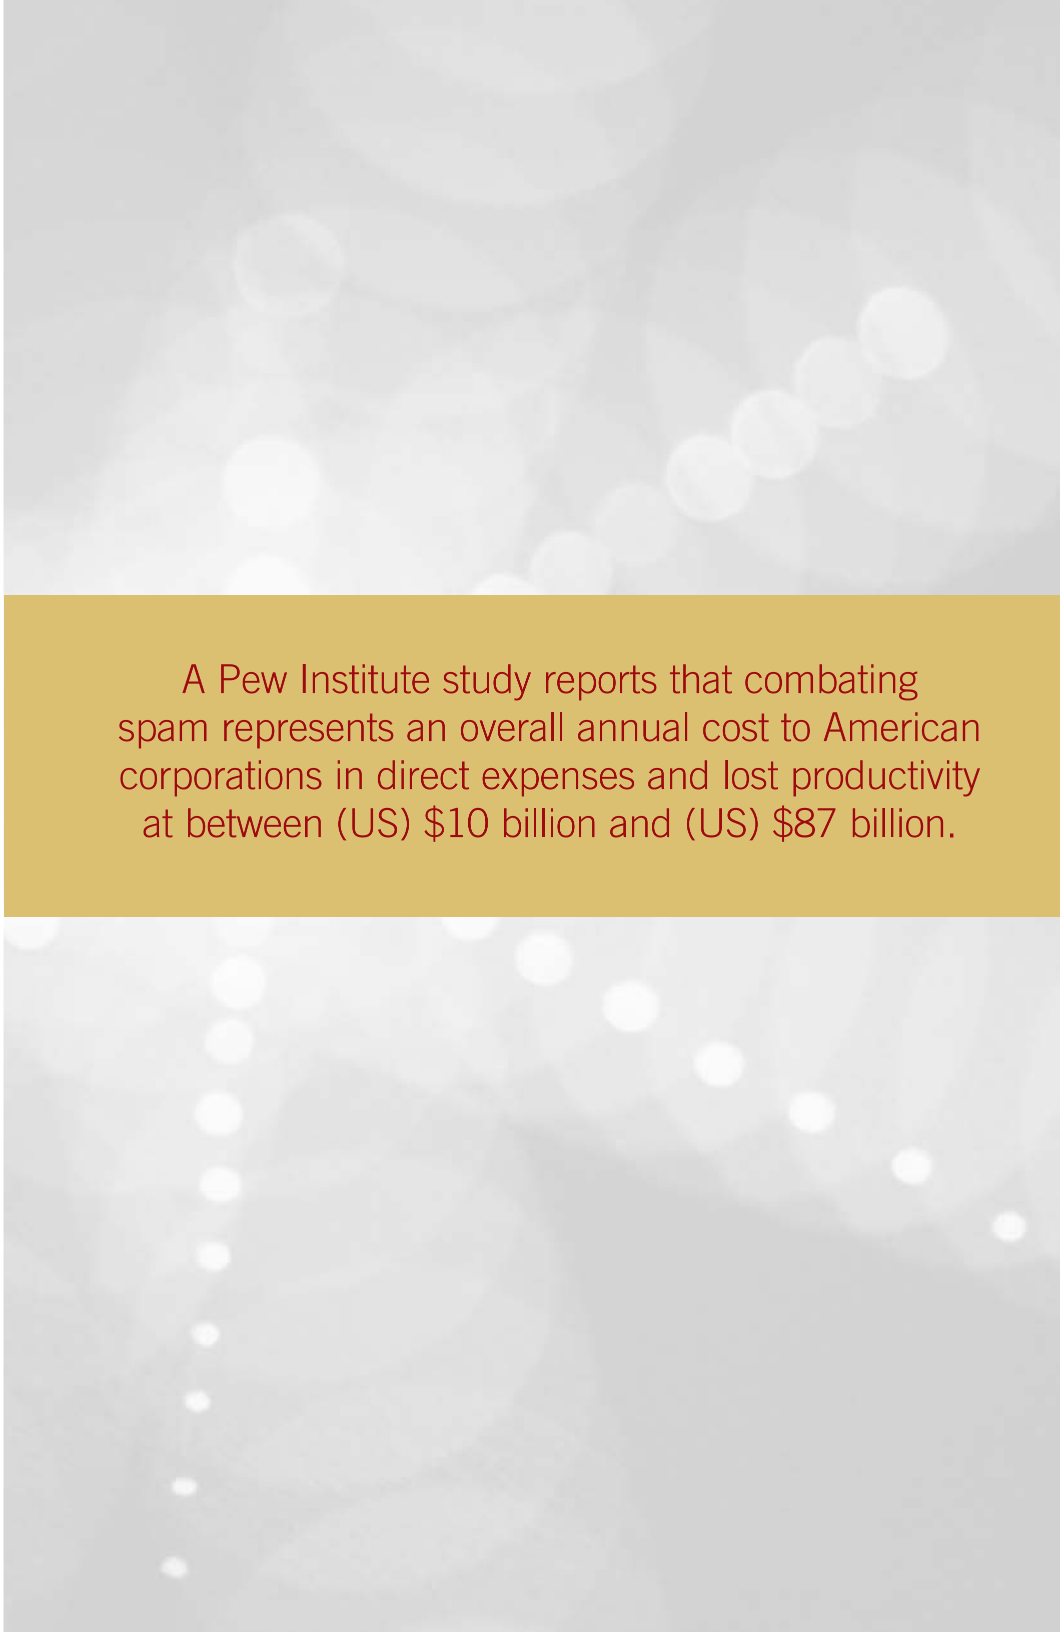

A Pew Institute study reports that combating spam represents an overall annual cost to American corporations in direct expenses and lost productivity at between (US) $10 billion and (US) $87 billion.

2

STOPPING VIRUSES

> "With the really good viruses, people don't even know they're being attacked."

RICHI JENNINGS
Lead Analyst, Ferris Research's Spam and Boundary Services Practice

STEP 2

Stopping Viruses

It may not be common knowledge, but spam and viruses are originated by the same people. 90 percent of the viruses in the past year have been designed to leave behind a small SMTP engine that is used to hijack an unsuspecting consumer PC and send out spam. So it's ironic that the biggest sources of spam on the Internet might be the PC your mom or dad have connected to a cable modem, spewing out spam unbeknownst to them.

These "zombie" PCs have proven to be very effective tools to help spammers fool less sophisticated spam filters. So in order to keep their army of zombie PCs alive and growing, spammers need to create new viruses to infect unsuspecting PCs.

The traditional defense against viruses rely on a "signature" or a series of bits that identify malicious attachments. While signatures remain a critical component of any virus defense system, they have an inherent weakness. No matter how good the anti-virus signature vendor, it takes a finite amount of time—usually about 13 hours—to detect, isolate, characterize, and create a signature for a new virus outbreak. So the bad guys simply design new virus variants every few weeks and get them to spread rapidly in the window when signatures are being developed. This is why, despite the widespread use of signatures, email-borne viruses continue to be a major problem for IT teams.

FAST FACT

More than 90 percent of viruses spread by email.

To contain the rise in rapid outbreaks, a two layer filtering system is needed, similar in nature to the two layer spam filtering systems found in leading edge solutions. The inner layer is a signature based anti-virus filter. The outer layer is a preventive anti-virus solution. With IronPort's solution, the outer layer of virus defense is known as *IronPort Virus Outbreak Filters*.™ The concept is to identify a new outbreak based on a traffic anomaly and then quarantine or "pause" suspicious mail until the traditional anti-virus signatures have been developed.

IronPort has a unique asset that is used to battle these outbreaks—the *SenderBase Network*. Because *SenderBase* measures such a large population of email—more than 25% of the world's email traffic—IronPort can detect the propagation of a new virus the instant it begins. These rapid outbreaks are designed to spread around the world in a matter of hours, in an attempt to beat the signatures. There is no form of human communication that mirrors this type of massive propagation.

IronPort has created a *Threat Operations Center* (*TOC*) to monitor *SenderBase* and look for anomalies in global email traffic that indicate a new outbreak. For example, a good outbreak indicator might be a sudden increase in password protected zip files, corresponding with an increase in mail coming from IP addresses that have never sent mail before, because these are really infected PCs spreading the virus. The *TOC* is staffed with technicians and statisticians that create algorithms to automatically detect these anomalies, and also to provide manual oversight to an automated system to ensure it isn't being manipulated by the engineers that created and propagated the virus. Having IronPort's highly trained analysts in the loop means that the IT team at every corporation doesn't need to be reviewing and reacting to every new outbreak—IronPort takes care of that effort.

FAST FACT

The technicians in the IronPort *TOC* speak 36 languages.

When the *TOC* team issues a new outbreak alert, it automatically pushes a rule out to the IronPort appliances. The appliances have a unique dynamic quarantine that scans and re-scans mail as more fined grained information becomes available. For example, the moment an outbreak occurs, the *TOC* may issue a very course rule such as "quarantine all .zip files". This rule is automatically sent to the IronPort appliances in the field and, at any time of the day or night, initial virus defenses are activated—.zip files are put in a special quarantine on the appliance. Within five minutes the *TOC* technicians may determine that the outbreak is associated with .zip files that are sized between 50 and 55 KB. A new rule is created and pushed out to the appliances. The dynamic quarantine then re-scans all quarantined messages, releasing anything that isn't a .zip between 50 and 55 KB. As the outbreak rages around the world, the *TOC* team may create another (more fine grained) rule—the outbreak is .zip between 50 and 55KB and contains the word price in the file name. This rule is pushed out and the dynamic quarantine re-scans, narrowing in on the outbreak. The concept is illustrated in Figure 1.

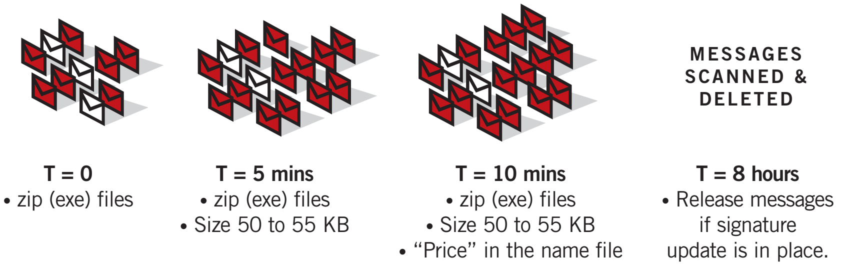

Figure 1: Dynamic Quarantine in Action

The *TOC* team also coordinates with IronPort's anti-virus signature partner, Sophos. Both teams share information on the outbreak and coordinate the release of the updated signature. When the signature is known to have been updated on the remote IronPort appliance, the *TOC* will issue a rule that says "scan quarantine with Sophos" and the Sophos engine will scrub all messages in the quarantine, deleting or stripping all messages that match the signatures.

This two layer filtering system with a preventive outer layer and a content-based reactive inner layer represents comprehensive virus protection. *IronPort Virus Outbreak Filters* have been in production for more than a year and have stopped over 150 outbreaks an average of 13 hours ahead of signature availability, providing state of the art protection for corporate networks and admin free defense for the IT team. An outline of IronPort's response to recent outbreaks is provided in Table A.

VIRUS NAME	IRONPORT'S EARLY DETECTION ADVANTAGE
Multiple "Bagle" Variants	DETECTED **41:43 hours** BEFORE ANY OTHER TECHNOLOGY
"MyTob.EC"	DETECTED **15:22 hours** BEFORE ANY OTHER TECHNOLOGY
"Sober.J"	DETECTED **10:23 hours** BEFORE ANY OTHER TECHNOLOGY
"Zotob-C"	DETECTED **2:51 hours** BEFORE ANY OTHER TECHNOLOGY

Table A: IronPort's Virus Outbreak Filters Lead the Industry

No system can provide perfect security against spam or viruses (although some vendors make claims and guarantees to this effect). *IronPort Virus Outbreak Filter* technology, combined with Sophos anti-virus signatures, yields the most effective virus defense system on the market—in production at more than 20 percent of the world's largest enterprises. The same engine that protects companies like Cisco, Juniper, Network Appliance, Dell and Intel is powering *IronPort C10* email security appliance designed for small and medium sized businesses. Plug it in, and spam and viruses just go away.

"*IronPort Virus Outbreak Filters* is a big winner for us. We know that our network is protected, even as we wait for for AV signature updates."

MARK DIAL
E-Messaging Team Manager, Tellabs, Inc.

3

OUTBOUND SANNING

“Enterprises today are struggling to deal with a complex regulatory environment full of costly, unfunded mandates, while still managing tight budgets.”

RICH MOGULL
Director, Gartner Research

STEP 3

Outbound Scanning

KEEPING GOOD PEOPLE FROM DOING BAD THINGS

There are two factors at work that are driving interest in outbound scanning—regulatory compliance and protection of intellectual property. Regulatory compliance can be put into three basic buckets—the Health Insurance Portability and Accountability Act (HIPAA), the Gramm-Leach-Bliley Act (GLB) and the Sarbanes-Oxley Act (SOX).

HIPAA requires that any entity dealing with personal healthcare information (PHI) put very specific safeguards in place to ensure this information is protected. The language around “safeguards” is open to interpretation, but in plain English it means any company that is transmitting patient health information such as doctor’s appointments, medical charts, etc., put in place encryption and access controls to make sure this information isn’t accidentally or intentionally exposed to unauthorized eyes. The intent of this act is to ensure that if an employee has a serious illness and leaves a job, they won’t be denied health coverage at their new employer.

For any company in the health care or insurance industry, this is a very significant regulation that requires a deep understanding of the act itself and a thorough review of enterprise wide workflow to ensure safeguards are in place, back to front. But, for small or medium sized business that are not specifically in the healthcare industry, there are still potential risks—especially since the HR team at a smaller enterprise may not have experience with the specifics of HIPAA. It’s not inconceivable that an HR generalist may email a patient’s history to an insurance provider—a clear HIPAA violation.

The IronPort appliances have built-in logic to provide basic safeguards against HIPAA violations. There is a pre-populated dictionary that looks for data that would appear on a patient healthcare record, such as treatment codes, medical names, and codes of drugs. These are very specific data types, so it is relatively easy for the system to identify healthcare information entering or leaving the firewall—and to flag it for encryption or review. Having this type of protection in place is excellent practice to ensure someone doesn't unknowingly violate HIPAA requirements and expose the enterprise to legal liability.

There are similar requirements for GLB. GLB applies to financial institutions (banks, mortgage companies, credit unions, etc.). It has specific guidelines for safeguarding personal financial information. The IronPort appliance has a set of pre-populated filter terms to look for GLB triggers—things like social security numbers or credit card numbers. It is unlikely that a small or medium sized enterprise outside of the financial industry would need GLB filters, but this is a must have item for any enterprise dealing with financial services.

FAST FACT

HIPAA and GLB have very specific guidelines for healthcare and financial services companies.

The third act is much more broad—Sarbanes Oxley. The spirit behind Sarbanes Oxley is intended to protect the market from inaccurate financial reporting in the wake of the Enron scandal. The majority of these regulations apply to internal finance and accounting procedures. However, there are some steps that can be taken with outbound email. Email coming to and from certain groups, such as legal or finance, should all be archived. This can be done with a few simple mouse clicks in *IronPort Email Security Manager*™. Beyond this, filters that look for financial indicators, such as social security numbers, can be employed. However, this is a coarse protection layer (and likely to snag legitimate mail) so it is not recommended outside of the financial regulations driven by GLB.

Many small and medium sized businesses are looking for the lowest cost and easiest way to comply with regulations. IronPort email security appliances offer capabilities that meet this need. For more sophisticated compliance capability, IronPort has partnered with industry specialists such as PostX and PGP corporation. Together with PostX and PGP, IronPort can yield a fully integrated solution which can identify mail that needs special handling, and automatically take the appropriate measures such as archiving or encrypting the messages. These more sophisticated solutions tend to apply to the specific requirements called out in the healthcare and financial services industries.

Beyond compliance, there is growing interest in protecting the intellectual property of the enterprise. There have been many high profile cases of accidentally releasing sensitive information, such as press releases or earnings announcements. There is an important element of common sense that needs to be applied here. If an employee is intent on stealing and distributing information, it is very difficult to stop them. Information can leave the enterprise on a USB drive or even on the back of a napkin. However, in many of the most publicized cases, the information leak was done accidentally. This is a problem that can be greatly reduced with some simple safeguards. The IronPort appliance has the ability to implement basic rules which look for sensitive information that should not be destined for the outside world. These rules can be tied to information about the employee sending the mail (e.g. are they from the engineering department?) and to the destination of the mail (e.g. is it destined to a competitor?). Having simple rules that look for mail coming from engineering going to a competitor, or a rule that looks for mail coming from the marketing team that contains the words "press release draft" can be very useful in detecting accidental dissemination of information. Think of it as a goalie for your private information.

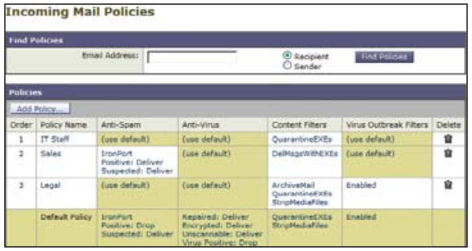
Incoming Mail Policies

Find Policies

Email Address: ● Recipient ○ Sender

Policies

Order	Policy Name	Anti-Spam	Anti-Virus	Content Filters	Virus Outbreak Filters	Delete
1	IT Staff	(use default)	(use default)	QuarantineEXEs	(use default)	🗑
2	Sales	IronPort Positive: Deliver Suspected: Deliver	(use default)	DelMsgsWithEXEs	(use default)	🗑
3	Legal	(use default)	(use default)	ArchiveMail QuarantineEXEs StripMediaFiles	Enabled	🗑
	Default Policy	IronPort Positive: Drop Suspected: Deliver	Repaired: Deliver Encrypted: Deliver Unscannable: Deliver Virus Positive: Drop	QuarantineEXEs StripMediaFiles	Enabled	

"Email Security Manager serves as a single, versatile dashboard to manage all the services on the appliance." — ***PC Magazine 2/22/05***

These rules are easy to implement using *IronPort Email Security Manager. Email Security Manager* provides a graphical representation of the various rules being implemented for various groups, making it easy and quick to implement policies that ensure good housekeeping and prevent good people from doing bad things.

4

PROTECTING YOUR IDENTITY

"Whether you employ one part-time worker or 100,000 full-time professionals, any time you allow employees access to your email system, you put your assets, future, and reputation at risk."

NANCY FLYNN
Executive Director, The ePolicy Institute

STEP 4

Protecting Your Identity

There are two major email pitfalls that every IT manager needs to be aware of—bounce handling and outbound commercial mail.

Bounce handling refers to how a mail gateway responds to incoming mail that has an invalid address. There are two modes of response—conversational bounces and delayed bounces. A conversational bounce occurs during the SMTP conversation. This means that before the receiving mail server has acknowledged receipt, it checks a directory (such as Microsoft Active Directory) to make sure the address is valid. If the address is valid the receiving mail server responds with "OK I have it" or if the address is not valid the receiving mail server says "Sorry I can't accept it". The advantage of this approach is that the bounce message that is being delivered directly to the sending mail server using the same connection or "conversation" that the message arrived in, so the bounce message cannot be redirected or spoofed. The disadvantage is that it effectively exposes the corporate directory to anyone. Spammers will routinely launch "dictionary attacks" where they guess at likely email addresses to see what gets through (e.g. bob@acme.com, charly@acme.com, etc.). Since a valid/invalid message is delivered in the conversation, in a matter of minutes a spammer can have a full list of valid email addresses at a corporation, which in turn can be sold on the internet for $50 or so, resulting in huge volumes of spam.

To protect their directories, most companies have chosen to issue delayed bounces. With a delayed bounce, the receiving mail server accepts all incoming mail. Then it checks for valid addresses. If the address is invalid, it will generate a separate email message back to the sender with a notification of why the message couldn't be delivered. This separate email coming back is much harder

FAST FACT

A "conversational" bounce cannot be redirected, but can expose the corporate directory.

for a spammer to use to automatically harvest a corporate directory, thus delayed bounces protect the corporate directory.

Since spammers send mail at large volumes, they don't want millions of delayed bounces coming back to them. So they will typically forge the return address of their spam. This is the root of the "misdirected bounce" problem. One common spammer tactic is to use the address of a known spam trap as the return address. Thus when the legitimate corporate mail server proceeds to respond to a spam with a delayed bounce message, that bounce message is sent to a spam trap operated by a blacklist—and the corporation finds itself blacklisted for being a spammer. This is a very common problem, and a huge source of frustration for the IT department. It can be very difficult to get "un-blacklisted" since many blacklists are run by volunteers and don't provide customer service. The other danger with misdirected bounces is they can be used to create distributed denial of service attacks. If a spammer sends out one million messages with a return address of postmaster@acme.com, acme.com is likely to get 750,000 delayed bounce messages from 750,000 different mail servers on the Internet. These misdirected bounce attacks have caused multi-day mail outages at major banks and ISPs, but also at small and medium businesses that did nothing to instigate it.

FAST FACT

More than 90 percent of spam has a forged return address, causing "misdirected bounce" problems.

IronPort appliances have a unique solution to this problem. IronPort has a "secure bounce" mode where it will issue a conversational bounce to a trusted sender, but not issue a bounce at all to a suspicious sender. It keeps track of the number of invalid address attempts from a given sender. When the sender exceeds the threshold, the IronPort appliance knows it is a directory harvest attack and continues to accept messages, but does not issue a response at all—fooling the attacker and protecting the directory. The number of invalid attempts allowed is tied to the reputation of the sender. A reputable sender like a Fortune 500 company is allowed many invalid address attempts; an unknown or disreputable sender is allowed only a few invalid address attempts. This advanced technology operates totally autonomously to make the headache of bounce handling just "go away", without exposing the corporate directory.

The other email pitfall facing most IT teams is the handling of commercial email. Most companies have some type of commercial email—newsletters, transactions confirmations, investor updates, etc. These emails are typically produced by a database and sent out automatically. When machines generate and send mail, mistakes can sometimes happen. A new operator might hit the send button 10 times, thinking it wasn't working. Or the marketing department might get a great new list of names they downloaded for $50 (see previous section on directory attacks). Either way, the receiving mail servers might view this incoming mail as spam—and begin dropping all mail from the corporate sender.

An excellent practice is to separate machine generated commercial mail from employee mail. With traditional equipment, this requires installing and maintaining two mail gateways, an unacceptable cost for many companies. But the IronPort appliances have a unique capability called *Ironport Virtual Gateways*™ which can segment different classes of mail and put them on different outbound IP addresses. So employee generated mail goes on one IP address, machine generated mail on another, and delayed bounces (if used) can be put on a third IP address. Think of this as powerful segmentation that will limit damage to only one IP address and not allow an accident to impact the reliable delivery of vital employee generated mail.

5

FIXING EMAIL

> "Email is mission-critical to my business. Technology provides great tools to fight the email security battle. But there is no substitute for a solid understanding of the many management issues surrounding email."

MARK FITZGERALD
Messaging Manager, KeyCorp

STEP 5

Fixing Email

Spam, viruses and fraudulent email have put a massive stress on email infrastructure. The root cause behind this scourge lies in the email protocol itself, SMTP. SMTP was developed in the late 1980's when the Internet was primarily a tool used for technical people, such as university professors, to collaborate and share information over unreliable data links. To facilitate this, SMTP has provisions that allow an email message to be forwarded from one machine to another, hopping its way to a final destination. At the time this was a trusted network, there was never reason to believe that a message wasn't actually being sent from the person it purports to be from. As a result, the protocol has no capability to validate a sender. So when a message arrives at a mail server at a company and says that it is from george.bush@whitehouse.gov, there is no way for that receiving mail server to know if it really is or isn't from whitehouse.gov. This core weakness is what allows spam to come from a seemingly legitimate sender, or viruses appear to come from someone an end user knows, or fraudulent email to appear from a trusted bank or trading site.

Plugging this hole in the email protocol SMTP will go a long way towards attacking spam and viruses at their core. But it turns out that adding authentication into the email protocol is a relatively complex undertaking, mostly because there are more than 20 million email servers active on the Internet. The approach that the Internet community has been taking is to create an overlay protocol that sits on top of SMTP. The two leading proposals are called "Sender ID" and "DomainKeys". These two proposals are very different and largely complimentary. But they are fundamentally changing the way email works.

Sender ID uses a "path based" approach, were the sender publishes a list of all IP addresses that are allowed to send mail on their behalf. This approach has the advantage of being light weight and easy to implement. At a bare minimum, a corporation should publish the IP addresses of its outbound mail servers as Sender ID records. If the corporation uses an email service bureau, the IP addresses of this entity should be included as well. Receiving mail servers will scan incoming messages and go back to the purported sender to see if the Sender ID record includes the IP address of the server that actually delivered the message. So for example, if whitehouse.gov published a Sender ID record of 1.2.3.4, the receiving mail server that just got a message from george.bush@whitehouse.gov can verify that the server that delivered the message was actually 1.2.3.4.

FAST FACT

Sender ID is backed by Microsoft and many ISPs.

The big challenge with Sender ID is known as "the forwarding problem". Many people maintain permanent email addresses at their universities or other institutions. So if George Bush was sending an email to joe@university.edu, but that message got forwarded to joe@acme.com, acme.com would see a message from george.bush@whitehouse.gov but it would not be delivered by IP address 1.2.3.4—the server identified in whitehouse.gov's Sender ID record. The forwarding problem prevents the receiving mail server from taking definitive action when incoming mail does not match a Sender ID record. However, as Sender ID is gaining in acceptance, an intelligent receiving mail server will view a positive Sender ID authentication as a very good thing and weight the message towards "not spam". Sender ID failure does not mean for sure a message is spam—but it will be a mark against the message as the receiving mail server looks at a variety of factors and scores the message as either spam or not.

The other emerging standard is known as DomainKeys or DK for short. DK uses a cryptographic stamp embedded in the message header. Invisible to the end user, this stamp allows the receiving mail server to definitively authenticate a message. The stamp is applied by the sending mail server which uses a "private key" to make the stamp. When a receiving mail server sees the stamp, it goes back to the purported sender and gets the public key. If the stamp decrypts properly, the message is known to be legitimate. If the stamp doesn't decrypt properly, the message is known to be fraudulent.

FAST FACT

DomainKeys is backed by Cisco and Yahoo!.

DK solves the forwarding problem because the message can be forwarded many times and the stamp travels along with it. The main challenge to DK is that it requires a fairly significant change to both the sending and receiving mail servers. IronPort appliances make it simple to start applying DK stamps to outgoing mail. Major ISPs such as Yahoo! are now looking for DK stamps. When they see a DK stamp that authenticates properly, they expose an icon to the end user stating "this sender is trusted". A user can see the DK stamp system in action if they have a Yahoo! mail account and get messages from Amazon.com. Amazon is using IronPort appliances to do DK stamping, and Yahoo! decrypts this messages and displays the trust icon.

The important thing for companies to realize is that spam, viruses, and fraud are forcing the Internet community to change the way email works. These new authentication technologies will take years to implement globally. However, as consumers and mail users begin to become aware of the trusted aspects of authenticated mail, the absence of authentication will become more and more suspicious. Consider this analogy: Email authentication is like having a driver's license. An individual does not *have* to have a license to get on an airplane. But the lack of a license and the unwillingness to authenticate makes an individual suspicious, and they may be subject to searches, delays and disruptions. Similarly, as the adoption of email authentication grows, the unwillingness of a corporate mail server to authenticate will make it increasingly suspicious and subject the sender to delays and disruptions in their mail flow. IronPort Systems has made it easy for busy IT staffs to implement authenticated email, and make sure their email infrastructure will keep working—today and tomorrow.

Conclusion

Email security is an ongoing endeavor. Because spam, viruses and fraud are a profitable business, the resources and tactics employed by those who generate this scourge are ever changing. As a result, your email security vendor needs to be committed to innovation. IronPort Systems leads the industry in technical innovation—with the largest research and development team in the industry and the world's most demanding networks as customers. IronPort has a full range of products that use its advanced email security engines, packaged in affordable and easy to use 1U appliances. These appliances allow administrators to "plug it in and make spam and viruses go away". Regardless of whether or not you select an IronPort appliance, there are five essential steps to safer email:

1. Use a leading edge spam filtering system that combines reputation and content analysis. A leading edge spam filter should be accurate enough to avoid the need for an end user quarantine or end user whitelist and blacklist controls. These end user facing features just create work for end users and tickets for the IT team.

2. Traditional signature based anti-virus systems are not sufficient. These systems are widely deployed and yet the world is still plagued by email viruses. The IT team should look for a solution that includes an outbreak control mechanism—it can pay for itself in one outbreak. *IronPort Virus Outbreak Filters* leads the industry in response time.

3. Scan outbound email. Healthcare and Financial Services companies have very specific email filtering requirements. All other industries have light requirements, but some safe guards need to be employed to stop good people from doing bad or dumb things.

4. Protect your identity and reputation. Conversational bounces expose the directory. Delayed bounces lead to blacklisting or DDoS attacks. IronPort has a unique "secure bounce" solution that mitigates this problem. Segment outbound mail. Put commercial mail on one outbound IP, employee mail on another, delayed bounces on a third. This practice will protect your reputation on the Internet.

IronPort C10 Email Security Appliance specifically designed for small and medium sized businesses.

5. Look to the future and stay ahead of the game. Set up a Sender ID record for outbound mail, and look for a solution that supports outbound DomainKeys (DK) signing. Lack of authentication will look increasingly suspicious in the coming 12 months and will lead to disruptions in outbound mail delivery. Look for a vendor that has the R&D resources to stay ahead of email threats. Spam, viruses and fraud email is "good" business and is fueling innovation. Look for a vendor that can out innovate the "bad guys" and keep your email system running trouble free.

TOM GILLIS is a recognized leader in the dynamically charged and high growth email security industry, with in-depth knowledge of the challenges surrounding secure network infrastructure. As an author, speaker and industry executive Gillis has made invaluable contributions to the email community. Gillis has held positions at iBEAM Broadcasting, SGI, and Boston Consulting Group (BCG). He is currently the Chief Marketing Officer for IronPort Systems.

MICHAEL D. OSTERMAN is the founder and principal of Osterman Research, Inc. Osterman has more than 20 years experience in the market research industry, conducting research for a wide variety of technology-based clients, including Microsoft, Lotus, Hewlett Packard, Sun Microsystems, Nokia, USinternetworking and Qwest, among many others. Mr. Osterman has written numerous articles for a variety of trade publications, and is currently author of a twice-weekly, online column on messaging issues published by Network World. He is a panelist and speaker at various industry and vendor-sponsored events.

MNP MESSAGING NEWS PRESS

www.messagingnews.com