



February 2005

Understanding E-Mail Hygiene:

*A Primer on Threats to E-Mail Infrastructures and
How Best to Protect Messaging Systems*

A META Group White Paper

“The value of e-mail to an organization has never been greater. Corporate business can grind to a halt with system outages, as internal communication is severely compromised and links to business partners, suppliers, and customers are severed. In fact, a META Group survey of 500 business people found that 80% of those surveyed indicated that e-mail was more valuable than the telephone for meeting daily communication needs. Paradoxically, however, as the value of e-mail increases, the threats and risks to system stability are steadily rising.”



Contents

Executive Summary	2
Mail Hygiene Market Dynamics.....	3
<i>Software</i>	<i>3</i>
<i>Hosted Services</i>	<i>4</i>
<i>Appliances.....</i>	<i>4</i>
Convergence of E-Mail Systems and Hygiene	5
<i>Technology Strategies.....</i>	<i>7</i>
Emerging Hygiene Techniques.....	8
The Future of E-Mail Hygiene.....	9
Bottom Line	10

Executive Summary

The value of e-mail to an organization has never been greater. Corporate business can grind to a halt with system outages, as internal communication is severely compromised and links to business partners, suppliers, and customers are severed. In fact, a META Group survey of 500 business people found that 80% of those surveyed indicated that e-mail was more valuable than the telephone for meeting daily communication needs. Paradoxically, however, as the value of e-mail increases, the threats and risks to system stability are steadily rising.

E-mail remains the primary vector for virus transmission. Although most companies now have aggressive antivirus defenses, virus writers are rising to the challenge. The number of viruses circulating on the Internet has never been greater, and the payloads have never been more destructive. Viruses now rapidly propagate by exploiting a user's contact list and heighten the chances of further propagation by spoofing the sender's address (making outbound virus filtering mandatory). Viruses routinely create backdoor access to the PC, so that it can be exploited by spam-sending and denial-of-service campaigns. We believe that about 60% of organizations were impacted by e-mail bourn viruses in 2004.

Further assaulting the integrity of the e-mail system is the rising tide of spam. Spam not only creates a major distraction for end users, but it can also clog inbound message transfer agents (MTAs), overwhelm storage systems, and create human resource issues via salacious content. In 2004, we estimate that about 70% of inbound SMTP traffic was spam (up from 50% in 2003), with no respite in sight. We expect that number to climb to 80% in 2005. We do not expect government initiatives such as the CAN-SPAM Act to have any impact on spam, nor do we expect emerging sender authentication programs, such as Sender Permitted From and Sender ID to have significant impact on spam volume for the foreseeable future (though it will greatly reduce spoofing).

Adding to the woes of mail managers are sophisticated hacker attacks targeting port 25 (SMTP). Hackers can bring down e-mail relays by flooding the system with an overwhelming volume of e-mail, and they also are using buffer-overload attacks to the same effect. More benign yet possibly equally destabilizing are mail loops (e.g., ping-ponging out-of-office notifications). Port 25 has also become a common vector for unauthorized, hacker-initiated malicious intrusions. Spammers are getting into the act as well with so-called dictionary harvest attacks, where they bombard an MTA with common names and harvest those that do not return an error message. In this way, they build up a base of legitimate names for a particular domain and then launch a spam attack.

The cost of all this unsavory activity is high. Many of these activities compromise system uptime, and organizational productivity is sapped. Mail systems being down also creates a major public relations problem for the IT group, especially e-mail managers. In many cases, companies are forced to upgrade their messaging infrastructure to accommodate blights such as spam by adding additional mail relays or investing in greater storage capacity to handle rising volumes. Furthermore, without proper defenses, IT groups are constantly in reactive mode, which creates a sense of crisis that is detrimental to overall performance and morale.

The remainder of this white paper is divided into five sections to examine various facets of the mail hygiene discipline:

- The first section covers overall market dynamics, focusing on the three prevalent delivery models for mail hygiene: traditional software load, appliances, and hosted services. We also examine the suite versus best-of-breed debate.
- In the second section, we look at the need to build hygiene services directly into the mail system and the role security plays.
- The third section goes deep into the technology, discussing various approaches and best-practice architectures for maximum perimeter protection.
- The fourth section examines emerging technologies for more effective mail hygiene practices, including reputation filtering, advanced end-user controls, and next-generation policy management.
- The final section looks at emerging needs for mail hygiene services, such as regulatory compliance, content scanning, and outbound filtering.

Mail Hygiene Market Dynamics

Vast user need for sophisticated e-mail hygiene services has created a hypercompetitive market. Venture capitalists are pouring hundreds of millions of dollars into startup companies, more established vendors are broadening their product portfolios, and merger and acquisition activity is rampant. We estimate that there are more than 60 vendors in this space, creating a vast and confusing market for end-user organizations. In addition to this cacophony of vendors, there is a panoply of delivery options.

Software

The most popular delivery model is still the traditional approach of buying software and loading it onto a server. Typically, this approach provides the most flexibility by allowing a broad range of services on the server, and also allows hardware upgrades to accommodate growing capacity needs (e.g., disk space, memory,

processors). Yet problem diagnosis can be complicated, given that multiple vendors provide multiple system components (e.g., application software, storage, operating systems, server).

Hosted Services

The hosted model routes all message traffic through an off-site managed service provider network, where messages are examined for spam probability and viruses. The end-user organization changes its MX (Mail Exchange) record to point to the service provider, which then routes all clean mail to the recipient mail relay. These services offer fast time to value, no capital expenses, and outsourced operational duties. We see this model becoming increasingly popular, particularly with smaller organizations.

Appliances

There is an increasing demand for hygiene services bundled directly with hardware — a so-called appliance. Perhaps the biggest benefit of an appliance to an organization is time to value. The advantages start with procurement, where organizations buy a single SKU rather than making the separate purchases of application, OS, and hardware required for the traditional model. Furthermore, with an appliance, there is no software load time or patch loading, no compatibility concerns, and a shorter burn-in period. There is also increased security at the OS level, since most appliances are based on a stripped-down variant of Linux, where ports have been shut down and unnecessary daemons/services disabled.

Off-the-shelf Unix versions can be similarly hardened, but their configuration requires extra time and skill sets. In addition, since the OS is often a proprietary implementation, hackers have a tough time knowing where the vulnerabilities are. However, a hardened, proprietary OS is no absolute guarantee against corruption; we know of several instances when appliances have been brought down by malicious code. From a user perspective, it is critical that appliances be evaluated in terms of maturity and sturdiness:

- Has it been on the market for some time with a good record?
- Does it include redundant hardware?
- Is there a single image of the OS and the application?

Different appliance suppliers have different notions of security, and organizations must closely scrutinize the appliance security proposition. Organizations should also take a close look at the administration tools of appliances before purchasing. In general, these tools are attractive because they are device-specific, but occasionally, we see administration tools so stripped down that they tend toward being cryptic. We also like to see administrator access via a Web interface, not an

OS prompt, to configure the software for reduction of inadvertent and malicious activity. Although we have yet to see a formal study, anecdotal evidence suggests that appliances are more reliable compared with the traditional model, due to stripped-down code and parts. Yet the advantage may be mitigated by the inability of administrators to troubleshoot or repair the device on their own.

Appliances have some operational advantages as well. Known operator error potential is reduced, due to the integration and single image of the OS and application as well as the ability to “turn on” additional features as needed with the enablement of a license key.

The hosted and appliance delivery model will increasingly take share from the software model. The market will move from a 2004 estimate (new sales, by seat) of 20% appliance, 15% hosted, and 65% software to 40% appliance, 20% hosted, and 40% software by year-end 2006.

Because of the vast array of competitors, we generally recommend that organizations limit the number of vendors used for mail hygiene. Use of one vendor, for example, for antivirus, antispyware, and denial-of-service attack prevention creates procurement and operational efficiencies, allowing an administrator to configure and manage multiple disciplines from a common management console. Longer term, we believe a similar dynamic — focusing on a single vendor that meets multiple needs — will play out with multiple communication modalities, such as instant messaging, port 80 traffic, and VoIP.

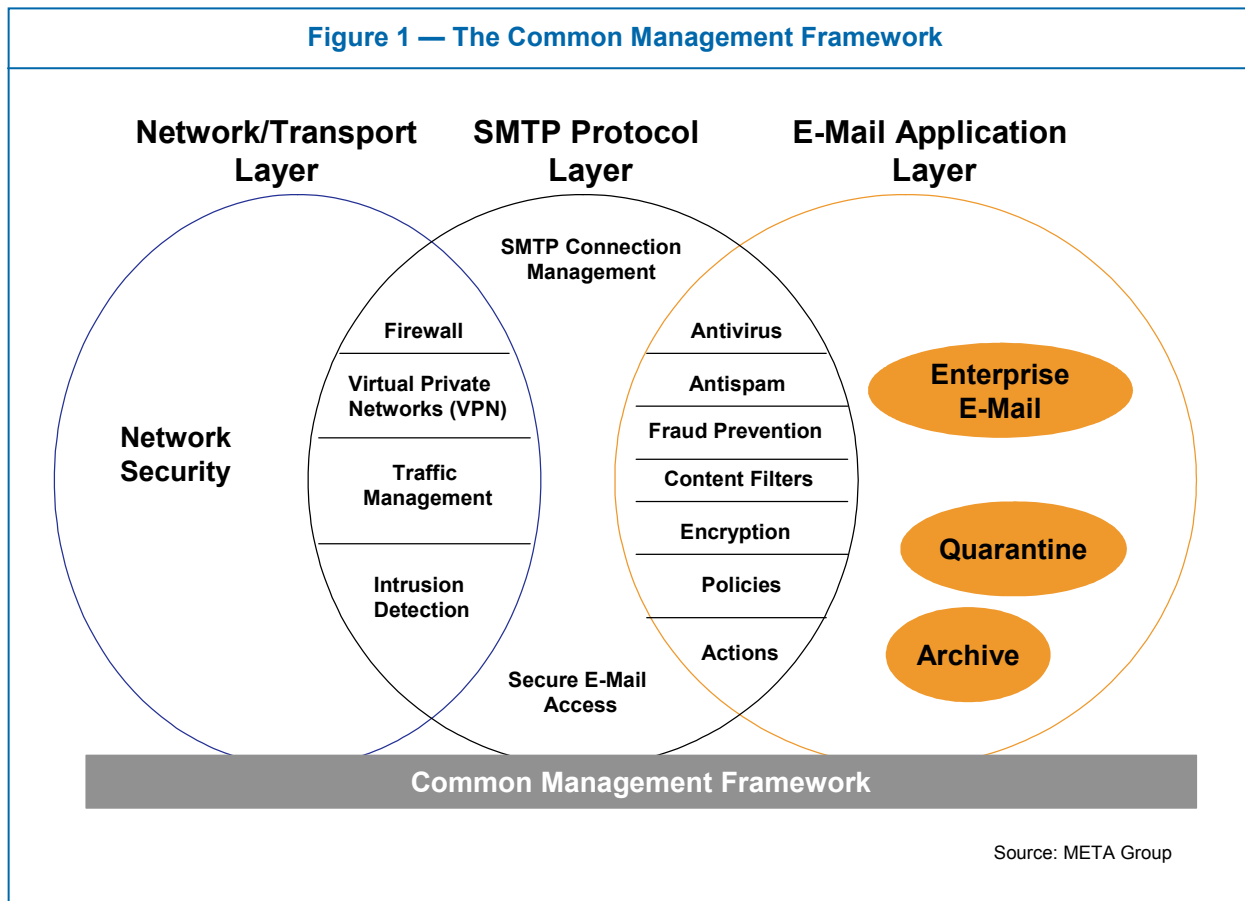
Convergence of E-Mail Systems and Hygiene

Although the traditional model for procuring hygiene services is to use a hygiene specialist vendor that is separate from the e-mail system supplier, there is growing evidence of a merging of the two disciplines. IBM, for example, has added support for real-time blackhole lists (RBLs — which are lists of domains thought to be known suppliers of spam) to its Domino mail system and is adding spam-blocking services to its Workplace Messaging system. Vendors such as IBM, Microsoft, Mirapoint, Ipswitch, and SendMail offer both hygiene solutions and mailbox servers. Microsoft, for example, is getting into the game with its proposed “Edge” server, a mail relay server (due mid-2006) that will offer native hygiene services (e.g., spam/virus blocking) as well as act as a platform for plugging in third-party hygiene services. Microsoft already offers a spam-blocking service for deployment in Exchange 2003 environments and will add additional capabilities in its Exchange Server 2003 Service Pack 2, due 2H05.

Understanding E-Mail Hygiene: A Primer on Threats to E-Mail Infrastructures and How Best to Protect Messaging Systems

The rationale behind a single-vendor approach to all e-mail components is that companies can experience procurement and operational efficiencies from sourcing multiple components for the message life cycle, including message routing, storage access, management, and services for securing all these steps en route. Yet modularity is also important, so that companies can break out or integrate components to address these individual messaging needs (e.g., breaking out gateway and routing functions for the enterprise DMZ).

Figure 1 — The Common Management Framework



Yet end-user organizations must be careful about relying exclusively on e-mail system vendors' hygiene services. In the case of both IBM and Microsoft, for companies to have maximum protection, we believe they must augment native services with third-party options. Therefore, due diligence on the part of the purchaser is warranted to ensure adequate depth and range of hygiene services prior to purchasing from the system vendor. However, the combination of single-source e-mail system and hygiene services can also lead to efficiencies via use of a common management console (see Figure 1) and the ability to turn system

functionality on and off as needed. We also believe it can ultimately help push hygiene services down into the internal mail operations, where there is a growing need for disciplines such as content filtering and regulatory compliance.

Technology Strategies

The history of mail hygiene services is one of constant battle between mail hygiene vendors and spammers, virus writers, and hackers. Just two years ago, for example, effective spam defenses typically were made up of no more than support for RBLs and some simple content-filtering rules. Spammers ultimately found workarounds to beat the defenses — for example, using spoofing to beat RBLs and URLs and deceptive content to beat content-filtering services. Since that time, best-of-breed spam defenses have expanded to include header analysis, heuristics, URL scanning, signatures, whitelists/blacklists, connection analysis, and Bayesian filters. New standards and techniques for blocking spam continue to emerge, such as sender authentication services and blocking services at the TCP/IP and SMTP connection levels.

A common technique used by spammers is impersonation of legitimate sending domains to help bypass filters and encourage users to open messages. An emerging technique, generically called sender authentication, will help minimize spoofing by registering the sending MTA IP addresses in the DNS, enabling the recipient MTA to query the DNS and validate that the sending IP address is registered to the sending domain.

There are two ways to identify the sender: either by “mail from” being contained in the message envelope, or by the purported responsible address (PRA) being in the message header, around which Microsoft has some patent claims. Microsoft supports the PRA model combined with “mail from.” This combination is known as Sender ID. The open source community supports only the “mail from” approach, which is known as Sender Policy Framework (SPF). We believe vendors will implement both approaches, with senders indicating via a “scope” parameter whether the information they publish in DNS pertains to the PRA, to “mail from,” or to both. Receivers would check PRA, “mail from,” or both, depending on their software. It is critical that enterprises ensure their hygiene vendor supports sender authentication and that they also register their IP addresses to support sender authentication.

Also of interest is a growing push to block spam at the network edge — the idea is to block obvious spam before it gets into the network, thereby diminishing the load placed on deep message interrogation by dedicated mail hygiene servers. It also results in increased security and lower operational burdens for mail managers. Vendors, for example, have determined that spammers often display certain behavior during the SMTP HELO conversation string — when the recipient and sender MTA

first establish a connection. Vendors will not disclose the specific behavior for fear of tipping off spammers, but when this common behavior is identified, vendors issue a 550 SMTP error message indicating that access is denied.

Other vendors are starting to use another spam-blocking method called “traffic shaping” or “IP throttling.” In the case of “grey listing,” vendors again correlate message flow and type with a particular IP address. But instead of dropping the connection, they slow down delivery rates and issue an SMTP 451 error message indicating that the connection is temporarily unavailable. SMTP relays of legitimate sending organizations will retry later to get the message through, but a spammer — which is typically paid on volume of messages sent — will quickly lose patience and move on to another recipient MTA.

Emerging Hygiene Techniques

Because of the constant warfare between spammers and spam-blocking services, there is constant innovation on both sides. Among the more promising emerging techniques for spam blocking are use of reputation filters and granular end-user controls over blocking services.

Some spam-blocking companies filter millions of messages per day. From this large volume, they are able to glean intelligence about the sending patterns of a particular IP address. If there is a high correlation between a particular IP address and an unusual volume of mail or certain types of mail coming from the same address, connections from that address will be refused, at least for a period of time. Some companies issue a 550 SMTP error message to the sender indicating that access has been denied. In this scenario, some vendors have a technician examine the mail flow, determine whether the messages are spam, and then act accordingly.

Companies also do in-depth log analysis to determine the validity of sending IP addresses. In addition, vendors have automated the process for validation of reputation. In cases of real-time mail flood attacks, for example, the system can shut down connections, though only after human oversight. With these IP-based, reputation-filter approaches, vendors estimate that they stop between 5% and 8% of all spam flowing through a network. Although there is a common belief that IP addresses are spoofable, SMTP connections require that a confirmation packet from the recipient MTA be sent to the sending IP address. Although this is not technically impossible, in practicality it is extremely difficult to spoof an IP address.

Because spam is defined differently by individual users (e.g., some users may enjoy trade magazine updates, while others consider them spam), it is important that users be able to assert some control over spam filters. Therefore, we recommend

implementation of end-user controlled whitelists and blacklists. Whitelists enable the user to allow any messages from a particular sender or domain to come through the spam filter unmolested. In this way, if a spam filter routinely blocks spam messages from a domain, the user has the ability to override the filter and have those messages delivered to the inbox. Likewise, a blacklist allows users to block all messages coming from a particular sender or domain. Thus, if the filter determines that messages coming from a domain are not spam, the user still has the option to block those messages, resulting in a cleaner inbox.

The Future of E-Mail Hygiene

There is no doubt that e-mail hygiene needs will grow exponentially during the next several years. We forecast growing needs for regulatory compliance, content filtering, and outbound filtering services as well as increased need for overall policy-based control of mail hygiene activities.

During the past few years, we have seen a growing movement toward regulating e-mail correspondence through regulations such as the following:

- **SEC 17a-3/4:** Specifies archival and supervision of broker/dealer communication
- **HIPAA:** Specifies that healthcare data be sent securely
- **GLB:** Protects consumer information
- **Other:** Other regulations in the pharmaceutical, insurance, and government sectors that now apply to e-mail traffic

Organizations must understand which regulations apply to them and how best to comply (working closely with the legal department, hopefully). In addition to having a heightened regulatory atmosphere, many organizations now apply internal records management guidelines to e-mail, creating the need, for example, for the archiving of messages relating to financial transactions or HR disputes for several years.

Part of the solution for regulatory compliance is granular and accurate content-filtering services, which, for example, allow an organization to develop a customized list of words and phrases that would be most likely to indicate the need for regulatory oversight for the particular message. Yet content filtering provides other valuable services for an organization. Salacious content — pornography, hate mail, and offensive jokes — creates a hazard for an organization, raising the possibility of hostile workplace violations. The human resources department is particularly interested in this interaction and should be the leader in establishing policies.



Understanding E-Mail Hygiene: A Primer on Threats to E-Mail Infrastructures and How Best to Protect Messaging Systems

Content filtering is also valuable for the outbound message stream. Many companies now are concerned about proprietary data leaving the company without proper authorization, including content such as draft press releases, product plans, internal broadcast memos, and other intellectual property. Outbound content filtering can halt messages carrying this type of information, or at least quarantine the message until it is reviewed by the relevant authorities. However, outbound filtering should not be limited to content. Messages should always be scanned for viruses as well as salacious content. We believe companies must police outbound traffic more aggressively as part an overall mail hygiene program.

Since e-mail hygiene spans so many disciplines, its management can be quite difficult. We recommend a directory-driven, policy-based approach, where application of diverse hygiene services can be applied in a granular fashion to groups of users. Companies may decide that certain groups in the company are covered by various regulations that require supervision or archival services. By being able to apply policy to a group of users (e.g., the finance department), companies do not have to apply the cumbersome compliance duties to the entire enterprise. Companies may also decide to apply different lexicons for content-filtering services of different groups of users. Thus, it is imperative that companies examine the vendors' capabilities to implement policy-based management services, not only for content-filtering services, but also for management of all the mail hygiene disciplines mentioned above.

Bottom Line

At most organizations, e-mail hygiene has largely been limited to virus protection, and more recently, spam management. Given the criticality of the messaging infrastructure, however, it is incumbent on all companies to develop a much broader approach to message protection, encompassing denial-of-service attacks, intrusion detection, content filtering, and regulatory compliance. Companies must adjust budgets accordingly by adding additional dollars for protection procurement and additional headcount for operational duties.

Organizations must strive for maximum uptime of the mail system, including perimeter components such as the mail relay, and create a hygienic mailbox environment for end users. E-mail is the communication vector of choice in the 21st century. Therefore, enterprises must ensure that proper protection is in place to preserve the integrity of this valuable asset.

Matt Cain is a senior vice president with Content & Collaboration Strategies, a META Group advisory service. For additional information on this topic or other META Group offerings, contact info@metagroup.com.



About META Group

Return On IntelligenceSM

META Group is a leading provider of information technology research, advisory services, and strategic consulting. Delivering objective and actionable guidance, META Group's experienced analysts and consultants are trusted advisors to IT and business executives around the world. Our unique collaborative models and dedicated customer service help clients be more efficient, effective, and timely in their use of IT to achieve their business goals. Visit metagroup.com for more details on our high-value approach.



About Mirapoint

Mirapoint was founded in 1997 with the vision of delivering the email application in a reliable, high-performance, purpose-built appliance, much like what Cisco had done years earlier with routing for networks, which became increasingly ubiquitous and critical to businesses. Mirapoint saw message networks becoming similarly important for business and personal communications, especially with the convergence of data and IP networks with telephony and wireless networks. In 1998, Mirapoint shipped its first products and the industry's first dedicated email server and security appliances. Today, Mirapoint has over 800 customers and 60 million mailboxes worldwide representing some of the most demanding enterprise, service provider and education environments. Key customers include Ford, Cisco, RSA Security, British Telecom, CB Richard Ellis, Oxford University, University of Georgia, and others.

As the only vendor in the industry with an end-to-end offering for building message networks, Mirapoint addresses the complete lifecycle of an Internet message from routing, storage, access, management and security. Addressing email-borne threats like spam, viruses and inappropriate content, Mirapoint offers the powerful, industry-recognized RazorGate email security appliance. A significant differentiator of Mirapoint's offering is that in addition to providing an effective email hygiene appliance with RazorGate, Mirapoint also offers the Message Server appliance for customers that demand complete security from the network edge to the core mail server and the user's inbox. Mirapoint's Message Server works with Outlook clients and offers a rich set of enterprise-class features including group calendaring and address book services, as well as flexible access from web and wireless clients. For more information about Mirapoint, visit <http://www.mirapoint.com>.