

You've got too much mail!

Rarely welcome, too often offensive, sometimes ludicrous to the point of comical, spam is suddenly no laughing matter. Businesses worldwide receive some 7 billion spam messages each day and researchers expect that number to triple over the next three years*. For many companies these relentless barrages represent over half of their incoming e-mail traffic. All that junk mail not only costs corporations dearly in precious network resources and employee productivity but also carries with it serious legal liability as well as network security risks.

This paper analyzes architectural and technology alternatives that enterprises should take to fight spam, and recommends the most effective approaches.

* Based on a June 2003 study conducted by The Radicati Group

While managers count the costs of handling and storing unsolicited messages, e-mail users waste countless hours sorting through inboxes chock full of pitches for everything from cheap pornography to cheaper prescription cures for impotency. And they're calling for a stop to it.

Invariably, the place they call—or e-mail—is the IT department. The trouble is, most IT professionals find themselves at a loss to explain the lack of a silver bullet to fight phony e-mail.

“Kill the messenger” easier said than done

The reasons why spammers have maintained the upper hand are many. First, purveyors of e-mail-based promotions are highly motivated—because, remarkably, spam works. According to *Computerworld*, some one-quarter of one-percent or more of the people hit with a typical spam-marketing scheme actually buy an advertised product or service. Since it costs spammers almost nothing to send out millions of e-mails, that fraction of a percent response rate can quickly add up to huge profits.

Second, unlike many of their offerings, spammers are not stupid. They represent cagey adversaries adept at ever-changing ways of disguising commercial pitches as legitimate e-mail. If pressured by legislation, they move their operations to offshore ISPs. Or they hijack innocent computers from which they launch clandestine spam attacks.

Thirdly, spammers enjoy the advantage of having the initiative. It's almost impossible to predict, and thus implement countermeasures against, the next spam wave. Who could have foreseen the U.S. military would create a deck of cards containing its most-wanted list of Iraqis, and who might have imagined people would want to buy such a collection?

Expensive Junk

In lost productivity alone, industry research indicates that spam costs corporations \$874 per employee and, adding in IT costs, it will drain a total of more than \$10 billion from corporate coffers this year.

Source: Ferris Research, Nucleus Research

Fourth, spam is often in the eye of the beholder. Most organizations would, quite logically, categorize an e-mail containing “XXX” as pornographic spam, but not companies in the insurance industry, where XXX is used by a national credit rating agency and in the names of several state insurance regulations. E-mail newsletters, especially the commercially sponsored ones, represent another case where legitimate correspondence can set off false alarm bells. These e-mails share many of the same characteristics as spam, often including the identity of the sender. Many of the companies that distribute legitimate e-mail newsletters on behalf of various enterprises also blast out huge volumes of more blatantly commercial e-mail for spammers.

Finally, business use of e-mail is growing as fast as spam. No longer just an adjunct to phone and fax, e-mail has become the method of choice for business communications and thus a mission-critical resource. So as pressure to answer the spam problem comes up from users and down from all levels of the organization, including from the corner office, IT staff know they need to find a solution that eradicates spam without erroneously blocking or slowing crucial business communications.

Despite the inherent risks, IT managers know it’s time to take a stand. However, myriad vendors have jumped into this relatively immature market, which makes choosing the best defense difficult. This paper examines two key aspects of selecting an enterprise anti-spam solution:

- Where in your infrastructure to pitch the battle, and
- Which technologies offer the most effective weapons

Identifying the costs and risks

To identify the best enterprise anti-spam solution, it’s important to understand both the costs and the risks that spam poses to an organization. Spam negatively impacts organizations in five basic ways:

Reduced Employee Productivity – organizations lose productivity to the extent employees spend time viewing and deleting spam messages rather than doing productive tasks. Cost estimates vary, but one large organization with 52,000 employees—using a rule of thumb that a person takes approximately 12 seconds to scan and delete one spam message—estimated that each employee would spend 5½ hours per year deleting spam at a cost of \$14 million.

Increased Network Resource Costs – as the volume of spam e-mail increases so does the cost of system resources to support it. If 50% of e-mail coming into an organization is spam, then half of an organization’s e-mail servers (as well as related LAN bandwidth and storage backups) are dedicated solely to processing and storing junk mail. In addition, all this e-mail traffic requires additional bandwidth, network hardware, and archival storage capacity.

Increased IT Administration Costs – with more IT infrastructure and end-user problems comes the need for additional IT resources. This includes additional network and e-mail administrators, as well as additional help desk/technical support resources to assist end users.

Increased Legal Liability Risk – a significant proportion of spam includes pornography, offensive, or hate-based content that enterprises cannot allow into their organizations. Otherwise, they risk creating a “hostile workplace.” For example, HR organizations at many large enterprises fear their companies may become the targets of lawsuits unless they quickly take proactive steps to stem the tide of pornographic spam.

Reduced Security and Control – IT security staff worry that spammers will take advantage of the biggest weakness of their security architecture—people. Tricking end users into opening messages or attachments containing malicious applications or viruses, or tricking the user into divulging sensitive company information represents a serious and increasing security exposure. In addition, some of the cures are worse than the disease: desktop solutions often allow end-users to ignore established security practices and define their own spam policies, while outsourced solutions require organizations to hand over control of their confidential, mission-critical e-mail stream to a third party. Add to this the personal cost of various spam-based scams to employees, who may be tricked into spending money on unwanted items, divulging personal data that the spammers can then use for identity theft, or falling prey to fraudulent get-rich-quick schemes.

Understanding these costs and risk is important in evaluating the appropriateness of the alternative anti-spam solutions examined in this paper.

Where to Mount a Spam Defense

The first step in determining how to stop spam is deciding where to stage your defenses. The good news is, spammers have only one entryway into the enterprise. Whereas viruses can penetrate a network from multiple points (web access, web-mail, floppy disks, SSL, etc.), spam enters from only one place—the Internet gateway. Along this relatively narrow path of attack, IT departments can choose where best to lay their defenses: at the e-mail client, on the enterprise e-mail servers, at the perimeter of the network, or outside the network (where outsourced solutions filter e-mail before it hits the corporate firewall). As organizations gain more experience in fighting spam, the pros and cons associated with fighting spam at each of these points are becoming more clearly understood.

“It is an advantage to choose the time and place for battle.”

Sun Tzu, *The Art of War*

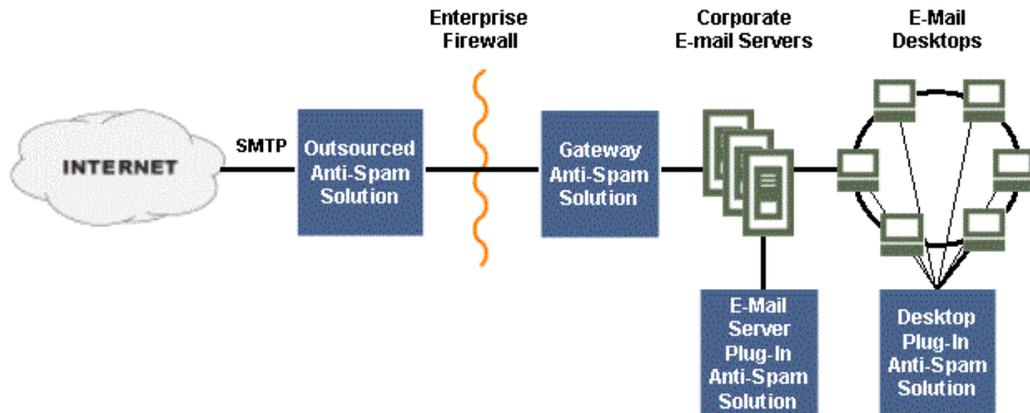


Figure 1: IT professionals can choose a number of places to implement anti-spam measures.

Desktop Anti-spam Solutions

A number of client-side anti-spam solutions available today work reasonably well, although they are primarily consumer-focused. Most client-side products allow end users to fine-tune spam blocking capabilities to suit their individual needs.

Unfortunately, less sophisticated users sometimes find it difficult to get the most out of these applications. Also, the fact that end users spend time managing an additional desktop application incurs a measurable productivity cost, particularly if the software requires training before people can use it effectively. Supporting any new desktop application also puts demands on already scarce IT resources. The cost of updating the software, deploying patches, and help desk support count among the key concerns. In addition, most desktop anti-spam applications lack enterprise-class features such as centralized management, control and reporting. For example, this can make it troublesome to implement corporate-wide anti-spam policies. And, quite understandably, most organizations would rather not leave it to their employees to decide whether or not pornography or online gambling are appropriate in the workplace.

In general, desktop solutions are most appropriate for small businesses and individuals, but most enterprises are shunning desktop anti-spam solutions in favor of more centralized approaches that IT departments can easily manage.

E-mail Server-based Anti-spam Solutions

Another class of anti-spam solutions sits at the corporate e-mail server, deployed as integrated plug-ins for Microsoft Exchange, Lotus Notes, or Novell Groupware. Blocking spam at this point in the network saves a significant amount of employee productivity: end users do not waste time sorting through spam because they never see it, nor do they spend valuable time managing an anti-spam application. Since the organization is taking positive action to ensure that end users do not view pornographic or hate-based spam, this approach provides good risk reduction in terms of legal liability. In terms of security, these solutions tend to perform fairly

well, as they inherit the security features of the e-mail servers on which they reside, and they significantly reduce end-user exposure to spam scams.

Nonetheless, this approach has only minor impact on network resource costs: spam still traverses the corporate network, so most of the costs of network bandwidth and mail server storage/capacity are still incurred. In addition, e-mail server plug-ins, especially ones that do complex processing such as anti-spam filtering, tend to degrade the performance and uptime of mail servers. Another area where this approach falls short is in IT management costs for larger organizations (generally over 500 mailboxes), which tend to have multiple e-mail servers. Not all of these anti-spam products are architected for centralized management and administration, meaning that IT staff may need to replicate management tasks for the anti-spam solution at every mail server.

“By cutting off spam at the entry point, a gateway-based approach can reduce overall e-mail traffic by up to 45%.”

Server-based solutions are cost effective for smaller organizations that have only a single e-mail server for the entire organization (usually organizations with less than 500 mailboxes). However, this is not a cost-effective approach for organizations with multiple e-mail servers.

Gateway-based Anti-spam Solutions

Gateway-based anti-spam solutions reside at the perimeter of the corporate network where the MTA (Mail Transfer Agent), or e-mail relay, routes e-mail to and from the Internet. These products sit between the organization’s firewall and its corporate e-mail servers (typically in the DMZ), and generally take one of two forms:

E-mail Relay Plug-in – an add-on application integrated with the SMTP relay that processes inbound and outbound e-mail. It adds a spam-filtering step to inbound message processing.

E-mail Firewall – a standalone application that provides a broad set of e-mail hygiene functions. Capabilities generally include:

- Anti-spam
- Anti-virus
- Outbound content filtering for policy and regulatory compliance
- E-mail relay services

Some products also incorporate encrypted e-mail capabilities and intelligent e-mail archiving support. The core technologies of an e-mail firewall are an SMTP relay for routing messages and a content filter/policy engine for analyzing e-mail based on identity and content.

Gateway-based anti-spam solutions generally do a good job of improving employee productivity, as they eliminate most spam before it reaches the end-user. These solutions also provide good network resource benefits because spam is blocked before it enters the corporate network. By cutting off spam at the entry point, a gateway-based approach can reduce overall e-mail traffic by up to 45%. This reduces the need for corporate mail servers, as well as hardware for routing and archiving e-mail traffic. Centralizing anti-spam efforts at the gateway further eliminates the need to manage individual spam applications installed on every e-mail server, or on myriad desktops and laptops typically scattered throughout an enterprise. In addition, centralization streamlines administration of anti-spam measures, such as updating enterprise-wide spam policies, and managing quarantine queues.

Gateway solutions are generally robust. Because e-mail firewall solutions are often used as an organization's primary e-mail relay, they have been designed for high availability. While e-mail relay plug-ins also are generally robust, integrating products from two different vendors always entails some performance and reliability risk. In addition, as many organizations today must cope with a patchwork of different e-mail systems, a gateway anti-spam solution operating at the perimeter offers the advantage of working outside fragile heterogeneous mail-server operating environments. In terms of security, these types of solutions ensure that end users are not exposed to injurious e-mail, and generally provide strong perimeter security capabilities (either as part of the native e-mail firewall application, or as capabilities inherited from the host e-mail relay).

Besides these advantages, gateway-based solutions may benefit from access to the SMTP protocol-level and network-level information available at the gateway and thus deliver more effective spam-blocking. Certain types of data, such as inbound IP addresses and SMTP envelope information, are typically not passed on to mail servers and desktops. This makes it difficult for these systems to use this crucial information to help block spam. For instance, a gateway solution can often detect if the number of envelope recipients on a message exceeds a pre-established threshold—a key indicator that the message is probably spam. Such heuristic analysis is not possible at the e-mail server or at the desktop because that information is not available.

Running a gateway-based solution requires some amount of time and resources, including software, network resources and IT administration. For smaller organizations that do not currently have an e-mail relay (and expose their corporate mail server directly to the Internet), these costs may be significant. For larger organizations that already have an e-mail relay, the hardware and software to host these gateway protection systems may already be in place.

Gateway-based solutions are the best approach for larger organizations (with more than 500 mailboxes). Smaller organizations that leverage an e-mail relay may also benefit.

Outsourced Anti-spam Solutions

Outsourced anti-spam solutions offer organizations a “managed service” that offloads the problem of spam fighting to an external third party. These solutions work by redirecting the organization’s inbound e-mail stream to the third-party outsourcer. The outsourcer filters spam out of the e-mail stream, and then passes the remaining e-mail on to the organization for standard delivery through their e-mail relay, corporate mail servers, and finally down to desktops.

Outsourced solutions, as with other centrally managed solutions, are effective in reducing spam volume and thereby improving end-user productivity. Note that a number of the outsourced anti-spam solutions are used by major ISPs such as MSN and AOL. Their customers benefit because the outsourcers have developed a strong understanding of consumer-targeted spam. The downside is that the public nature of these services allows spammers to freely experiment against the spam filters. For this reason, some outsourced services have lower capture rates.

Like gateway-based solutions, they eliminate all network resource costs associated with spam, because they filter spam before it enters the corporate network. And because all of the system administration work is done externally, IT resource requirements are relatively low.

The tradeoffs for an organization of choosing an outsourced solution include reduced control over the reliability and security of its own systems. Outsourcing is usually attractive to organizations that do not want to spend time fighting spam and are willing to have a third party decide what constitutes spam and how to handle it. However, giving up this control is often a difficult challenge for larger organizations with more mission-critical security and uptime requirements. Additionally, because all of the organization’s e-mail is routed through a third party, outsourced anti-spam solutions can present a significant problem in a number of industries that have e-mail security issues, such as financial services organizations that handle sensitive customer financial information and healthcare providers and payers who must comply with HIPAA privacy and security regulations. Organizations are also exposed to some risk in terms of the unknown reliability of the outsourcer’s systems—if the outsourcer’s system goes down, the customer’s e-mail stops in its tracks. And while these solutions do not require additional hardware, software or administration, it is industry practice to price these costs into the service fees. While the price for one year of service might appear attractive over the short term, over a three-year payback period these costs often exceed those of hosting anti-spam solutions in-house.

Outsourced anti-spam solutions are cost effective for smaller organizations that don’t have the expertise or resources to manage the spam problem internally. Larger organizations that rely on e-mail as a business-critical resource should weigh the risks of this approach against the longer-term value of owning and controlling their own anti-spam resources.

Best Practices Recommendation – Where to Mount a Spam Defense

In summary, a couple of clear choices stand out. For smaller organizations with less than 500 mailboxes that may have limited IT resources, the best solutions to evaluate are e-mail server based plug-ins, or outsourced anti-spam services. Larger organizations that need reliability, security and control, should investigate gateway-based solutions.

	DESKTOP	E-MAIL SERVER	GATEWAY – RELAY PLUG-IN	GATEWAY – E-MAIL FIREWALL	OUTSOURCED
End User Productivity					
Network Resources					
IT Productivity					
Legal Liability Risks					
Security & Control					
TOTAL					

Strongest Weakest

Best Methods for Identifying Spam

Combating spam continues to be a complex, evolving process. The current state of the art involves two distinct steps:

- Identifying spam
- Disposing of suspected spam messages

“He who looks below the surface of things, wins with ease.”

Sun Tzu, *The Art of War*

To date, the marketing of anti-spam products has largely focused on newer and sexier ways of identifying spam. These technologies have been developed by the plethora of startups founded to address the problem. Nonetheless, in evaluating an anti-spam solution, flexibility in handling the disposition of suspected spam messages should rank with equal importance, particularly for organizations concerned about losing business messages.

Identifying Spam

In terms of identifying spam, there are two key criteria by which the effectiveness of a solution should be measured: capture rate, and false positive rate.

Capture Rate

Many technologies exist for fighting spam. When comparing these different approaches, one key attribute to look at is “capture rate”, which measures effectiveness in blocking unwanted e-mail. Capture rate is defined as the percentage of messages identified as spam, divided by the actual number of spam messages received. This number can be hard to nail down in a real-world setting, because the actual number of spam messages received is usually only possible to ascertain through the cooperation of end users.

Despite some claims to the contrary, no enterprise anti-spam technology can deliver a 100% capture rate on real-world spam. The key in evaluating these products is to find a solution that provides a high capture rate (80%-90%) on a continuing basis, even as spam evolves and changes over time. Given the spammer’s proven ability to evolve clever new ways to cloak spam in the disguise of legitimate e-mail, achieving consistently high capture rates over time requires that someone maintain and update the anti-spam solution. While first-generation anti-spam products enabled this to be done manually by an IT administrator, next-generation products are automating this process by providing an “update” service similar to anti-virus engine updates. A viable update service must be supported by some type of anti-spam lab staffed by personnel who continuously monitor the environment for new spam tricks and techniques, and develop timely updates or entirely new types of countermeasures. These second-generation products have the potential to significantly reduce the cost of the fighting spam.

False Positives

The other key measure of effectiveness is the “false positive” rate, which measures the percentage of valid messages that were incorrectly identified as spam (in other words, legitimate business messages that were blocked). This metric is particularly important where e-mail is mission critical, and where organizations cannot afford to lose business messages. A zero false positive rate remains the Holy Grail in the anti-spam industry. Towards this end, a number of vendors provide a “quarantine queue” mechanism that holds suspected spam for manual review by an IT administrator. This approach, while effective, can become somewhat expensive as IT administrators spend much of their time reviewing and releasing e-mail. In addition, a few vendors are coming out with mechanisms to push these “quarantined” messages down to end users to help reduce IT costs without losing business messages. Organizations are mixed about whether this is a good approach to reduce costs, or whether it simply ends up re-burdening end users with productivity and liability issues. Given these trade-offs, most in the industry feel that less than one tenth of one percent (0.1%) represents a more realistic goal.

Five Layers of Defense

A comprehensive anti-spam solution involves multiple layers of defense, incorporating different but complementary technologies. The following framework identifies five basic layers of spam defense that an organization should look for in the anti-spam solutions it evaluates.

One: Anti-spam Engine

All anti-spam solutions include a core engine that analyzes and identifies suspected spam messages. There are a number of different spam engine technologies on the market today, offering varying levels of effectiveness. The table below provides a comparison of the leading anti-spam engine technologies and approaches, including their relative strengths and weaknesses.

TECHNOLOGY	PRO	CON
Lexical Analysis Applies content filtering to each e-mail message to identify suspected spam, based on word and phrase lists. Note that spam solutions vary in the depth to which they apply lexical analysis – solutions should analyze message subject, body, attachments, and HTML tags, and support some way to identify “disguised” text	<ul style="list-style-type: none"> Proactively identify spam based on common phrases/words. Word weightings allow for tuning of spam capture based on word frequency. Regular expression and pattern matching technology coupled with lexical analysis allows for more precise analysis (for example, “G.a.p.p.y” text can be identified). Allows an enterprise to tune or customize its definition of spam. 	<ul style="list-style-type: none"> Some difficulty in detecting bizarre spellings of words (e.g. ‘V.I.4.G.R.A’). Not effective on the increasing amount of HTML-based spam that contains no text, just URLs and images. Is time consuming and error prone if not combined with a dynamic update service, maintenance and testing of word lists. No understanding of the context of words used in a message (e.g. the word “breast” is an appropriate term in the Healthcare industry).
Heuristics-based Analysis Uses a set of rules to analyze an e-mail message to determine the likelihood that it is spam. Can be combined with lexical analysis to provide improved spam identification.	<ul style="list-style-type: none"> Proactively identify spam based on multidimensional attributes. Equally effective on HTML, graphics-based, and text-based spam. Examples include: <ul style="list-style-type: none"> E-mail with two or more embedded images has a higher likelihood of being spam Backdated or future-dated e-mail has a higher likelihood of being spam 	<ul style="list-style-type: none"> Heuristic rules tend to be complex and difficult to build, test, and maintain. If not combined with a dynamic update service, maintenance can be time consuming for the IT administrator.
Signature-based Analysis Maintains a database of ‘hashes’ of previously identified spam messages, and compares each incoming e-mail to this database. Messages that match are identified as spam.	<ul style="list-style-type: none"> Fast, effective identification of known spam messages. Equally effective on HTML, graphics-based, and text-based spam. 	<ul style="list-style-type: none"> Catch rates of these methods have degraded as spammers learn to deliberately introduce random variations into their messages. Must be used in conjunction with a dynamic update service. Updates must be virtually continuous for maximum spam blocking effectiveness. Dependent on vendor to maintain spam database.
Bayesian Analysis Algorithm that can be trained to automatically differentiate aggregate textual attributes of spam and non-spam messages. Applied to identify probability that any given message is spam.	<ul style="list-style-type: none"> Potential approach to intelligent, learning algorithm for blocking spam. Works well at the desktop, as the user can train the algorithm according to his/her needs. 	<ul style="list-style-type: none"> Not effective on the increasing amount of HTML-based spam that contains no text, just URLs and images. Not proven at enterprise scale – Bayesian learning algorithm difficult to apply automatically at the gateway.

TECHNOLOGY	PRO	CON
Natural Language Processing Artificial intelligence technology that combines morphemic, syntactic, and pragmatic analysis to correlate text with categories of meanings.	<ul style="list-style-type: none"> • Can be effective in identifying subtle, text-rich spam messages. • Can filter messages based on multi-word concepts, rather than individual keywords. 	<ul style="list-style-type: none"> • Not effective on the increasing amount of HTML-based spam that contains no text, just URLs and images. • Not designed to handle content that is created with the explicit goal of avoiding automated analysis using lexical or grammatical obfuscation. • Performance has typically been an issue.
Challenge/Response Requires senders to verify their authenticity before the e-mail is received by the recipient. The sender of the message will receive a challenge e-mail in response to their original message.	<ul style="list-style-type: none"> • Authenticates the sender (although this is not foolproof). • Can work for very occasional consumer e-mail users. 	<ul style="list-style-type: none"> • Sending out a challenge for every spam message received significantly increases outbound e-mail volume by 50% or more. • Fails on all automated e-mails, even if they are valid. • Fails on all mass mailings, even if they are valid. • Can be irritating to valid senders. • Spammers are figuring out ways to automate responses to challenges. • Delays message delivery if sender is not checking e-mail.
Collaborative Filtering End users vote on which messages constitute spam.	<ul style="list-style-type: none"> • OK for consumers applying spam blocking at the desktop. 	<ul style="list-style-type: none"> • Not tailored for enterprise messaging traffic. • Can lead to high incidence of false positives. • If implemented at the gateway, performance can be an issue. • Enterprises usually prefer to control these types of policies for liability and security reasons.
Cocktail Solution that combines multiple analysis methodologies to produce the highest degree of spam identification accuracy.	<ul style="list-style-type: none"> • If implemented well, will block a broader range of spam variations with the least number of false positives. • The most effective approach for enterprise e-mail traffic that can vary widely from company to company and industry to industry. 	<ul style="list-style-type: none"> • If done poorly, has the potential to equal the sum of false positives of each individual method. • If the underlying architecture is not sound and scalable the cocktail approach can introduce performance issues that could lead to e-mail delivery problems.

Overall, the most effective solutions tend to incorporate multiple, overlapping anti-spam technologies, and provide a mechanism for dynamically updating the anti-spam engine to maintain effectiveness against evolving spam tactics (similar to how anti-virus engines are updated).

Two: Address and Hacker Protection

In addition to anti-spam engines, an enterprise's first line of defense involves protecting its networks and e-mail infrastructure from attacks and address harvesting.

In particular, spammers often try to harvest e-mail addresses from enterprise e-mail domains by sending an exhaustive set of fabricated, likely e-mail addresses (often tens or hundreds of thousands of e-mail messages). In this so-called Directory Harvest attack, spammers attempt to identify valid e-mail addresses through a process of elimination based on bounce-backs or protocol responses. They then send spam to these addresses, and sell the lists to other spammers over and over again. To make matters worse, entry barriers are extremely low, as novice spammers can easily and cheaply purchase the harvesting tools used to execute these attacks.

Technologies used to defend against these types of attacks often include the e-mail relay, and possibly e-mail policy engines found in e-mail firewall products. In addition to Directory Harvest attacks, these technologies provide protection against Denial-of-Service (DoS) attacks, open relay hijacking, and address re-writing to obscure internal domains. In general, e-mail firewall products provide the strongest defenses in this area since they incorporate a native e-mail relay.

	DESKTOP	E-MAIL SERVER	GATEWAY – RELAY PLUG-IN	GATEWAY – E-MAIL FIREWALL	OUTSOURCED
Directory Harvest Attack Protection					
Denial-of-Service Attack Protection					
Open Relay Hijack Protection					
Address Re-writing					
TOTAL					

 **Strongest**  **Weakest**

Three: Proactive Blocking

Although sometimes heavy-handed, one of the more effective anti-spam defenses involves proactively blocking known spam sources at the Internet gateway. This technique deflects spam before it hits the anti-spam engine, thus eliminating the need for analysis and processing. Proactive blocking can provide a significant performance improvement in spam processing/message throughput. The key technology that enables these capabilities is the SMTP relay, with possible help from a rules engine/policy engine. The technologies and techniques used in this type of spam fighting include RBLs (real-time black hole lists), RDNS (reverse DNS lookup), and enterprise and end-user defined blacklists and white lists. Notably, RBLs are starting to wane in popularity because they tend to be so aggressive as to often block non-spam sources. In fact, many organizations are finding their business partners, customers, and sometime themselves turning up on these lists because they have been unknowingly used as an open relay to send spam on to others. Note that these capabilities are generally integrated as part of an e-mail firewall, but for e-mail relay plug-in anti-spam solutions, this must be managed separately as part of the native e-mail relay capabilities.

	DESKTOP	E-MAIL SERVER	GATEWAY – RELAY PLUG-IN	GATEWAY – E-MAIL FIREWALL	OUTSOURCED
RBLs					
Reverse DNS					
Custom Blacklists					
TOTAL					

 **Strongest**
 **Weakest**

Four: Identity-based Spam Filtering

One of the reasons that the spam crisis exists today is because it is extremely difficult to accurately identify the sender. The SMTP protocol and the Internet allow people (spammers in particular) to remain anonymous. The net effect is to negate all of the well-meaning legislation being pushed today. These legislative measures are useless because it is virtually impossible to identify a spammer who wants to hide.

However, if a way can be found to accurately authenticate the identity of senders on the Internet, both legal remedies to spam and other identity-based spam fighting would become effective. A number of identity validation/authentication technologies exist today, and this area is just starting to emerge as an effective partial solution to the problem of spam. By allowing known or trusted e-mail from partners, customers, and other employees to automatically bypass the spam filtering engine and enter the network, IT staff can create an environment that minimizes loss of business messages and maximizes throughput. E-mail from unknown or non-trusted sources still must pass through the standard spam-filtering process.

The technologies that make it possible to “know” or “trust” the sender of e-mail primarily revolve around digital signatures and encrypted e-mail. For instance, if you can validate the identity of an e-mail sender by looking up his/her digital certificate, the probability decreases that the message is spam. And if it does turn out to be spam, you have now identified a party against which to pursue legal recourse. Technologies in this area include:

- **S/MIME** - an Internet standard for e-mail encryption for years, supports digital signatures and thus strong authentication. While hard to deploy at the desktop, S/MIME is a viable server-to-server solution and is used today by business partners in finance, healthcare, and government to create ‘trusted networks’.
- **TLS-based messages** - an emerging standard that allows e-mail servers to set up an encrypted channel between them—to the extent the other server can be trusted, the message can be allowed in.
- **Directory integration** (either LDAP or some other internal directory format) - works by confirming that a recipient is valid before passing an e-mail message into the network.

- **Outbound e-mail recipient caching** - makes it possible to identify current or recent correspondents based on outbound message traffic, which is assumed to be valid.
- **Harvesting of digital certificates from inbound e-mail** - used to "remember" trusted senders.

	DESKTOP	E-MAIL SERVER	GATEWAY – RELAY PLUG-IN	GATEWAY – E-MAIL FIREWALL	OUTSOURCED
S/MIME					
TLS					
Directory Integration					
Outbound E-mail Cache					
Inbound Certificate Harvesting					
TOTAL					

 **Strongest**
 **Weakest**

Five: Customized Spam Definition/Tuning

Since every organization will have its own definition of “spam”, the ability to tune and tailor this definition becomes crucial. Some types of technologies lend themselves to customization more than others. Signature-based and heuristic technologies tend to be maintained by an external vendor lab, so some form of override capability must be present in this type of anti-spam engine to support tuning. On the other hand, technologies such as lexical analysis are quite easy to modify.

The best approach for handling this tuning is to support some form of override or exception processing technology as part of the anti-spam engine. These include:

- **White Lists** — allow administrators or end users to identify specific messages or message types that should be allowed to enter the organization, whether or not the message was identified as suspected spam. This is often used to allow in e-mail newsletters for example, which can be hard to distinguish from spam in many cases.
- **Policy Engine-based Exception Lists** — identify specific message attributes that should be granted override access to the network. This type of mechanism is often used to ensure that industry-specific messages/content are not blocked by the anti-spam engine - for example, allowing a message containing words such as “penis” or “breast” to reach its intended recipient in a healthcare organization.
- **Directory-based Anti-spam Policies** — allow IT departments to apply different policies to different individuals/groups/departments in the organization. For example, this approach enables an

administrator to allow the marketing department to receive advertising that would normally be blocked as spam.

	DESKTOP	E-MAIL SERVER	GATEWAY – RELAY PLUG-IN	GATEWAY – E-MAIL FIREWALL	OUTSOURCED
Custom Whitelists					
Policy Engine					
User/Group Policies					
TOTAL					

 **Strongest**  **Weakest**

Best Practices Summary – Methods for Identifying Spam

In summary, a number of clear recommendations stand out for enterprises looking for an anti-spam solution. One of the most important criteria is to get a high ongoing capture rate (80-90%). In addition, it is important to ensure that the solution provides a mechanism for dynamically updating the anti-spam engine. A number of products on the market today provide Internet-based update services for this purpose, similar to how anti-virus engines are updated. This type of automated update service also lowers the total cost of ownership (TCO).

Equally important is finding a solution that validates the capture rate in your environment or via customer references, rather than relying on vendor statistics. In terms of false positives, the solution should ensure a very low rate (less than 0.1%). Note that zero false positives can be achieved by investing personnel costs in an IT administrator to monitor quarantine queues. Also, the solution should provide a mechanism to ensure that end users can retrieve business messages caught as false positives (e.g. personal quarantine queues, digest of quarantined messages, etc.).

Overall, the most effective enterprise solutions include all five layers of defense against spam, and specifically incorporate multiple, overlapping anti-spam engine technologies in a cocktail approach.

	DESKTOP	E-MAIL SERVER	GATEWAY – RELAY PLUG-IN	GATEWAY – E-MAIL FIREWALL	OUTSOURCED
Anti-spam Engine					
Address and Hacker Protection					
Proactive Blocking					
Identity-based Filtering					
Customized Spam Tuning					
TOTAL					

 **Strongest**  **Weakest**

Conclusion

Fighting spam presents the basic challenge of fighting a fast mutating virus—how to eliminate the invading, morphing organism without harming the patient. IT managers are faced with the delicate balancing act of trying to knock out huge volumes of bogus e-mail without blocking a single business-critical communication. And they must do this with limited human and financial resources. By using the spammer’s Achilles heel—a single point of entry into the network—it is possible to mount a cost-effective defense at the perimeter of the network using an e-mail firewall. Outsourced solutions, while good solutions for small organizations are generally too risky for the enterprise; client-side applications don’t scale; and anti-spam applications installed on myriad e-mail servers multiply administration costs and headaches.

Key to the e-mail firewall’s effectiveness is a cocktail of technologies that combine to maximize spam blocking while minimizing false positives. In addition, owing to its strategic placement at the network gateway, an e-mail firewall offers a number of advantages—not the least of which is access to the corporate-wide information that allows sophisticated identity- and policy-based tuning of the anti-spam engine. It further integrates many anti-spam functions, which results in significant savings on administration costs and TCO.

FOR MORE INFORMATION, PLEASE CALL 650.216.2121

Tumbleweed Communications Corp
 700 Saginaw Drive
 Redwood City, CA 94063
 Phone 650.216.2000
 Fax 650.216.2001
www.tumbleweed.com
info@tumbleweed.com

Copyright © 2003 Tumbleweed Communications Corp. All rights reserved. Tumbleweed is a registered trademark of Tumbleweed Communications Corp. All other brand names are trademarks of their respective holders.

ACEASWP0703