

Anti-Spam Technical Alliance Technology and Policy Proposal

Version 1.0

Anti-Spam Technical Alliance (ASTA)

11 June 2004

Table of Contents

Executive Summary	3
Background	3
Scope	3
Summary of Best Practices	4
Summary of E-mail Forgery	4
Introduction	6
Considerations.....	8
Curbing Spam through Best Practices	9
Recommendations for ISPs.....	9
Close All Open Relays.....	9
Monitor formmail.pl and Other CGI applications	9
Configure Proxies for Internal Network Use Only	10
Detect Compromised Computers.....	10
Implement Authenticated E-mail Submission.....	11
Remove Remote Access to CPE	12
Implement Rate Limits on Outbound E-mail Traffic	12
Control Automated Registration of Accounts	13
Close Web-based Redirector Services Susceptible to Abuse.....	13
Develop Complaint Reporting Systems and Subscribe to Existing Systems	14
Recommendations for Bulk E-mail Senders.....	14
Recommendations for Consumers	15
Curbing E-mail Forgery	17
Anti-Forgery Technologies: IP Address Approaches.....	17
Methods of Implementation.....	17
Anti-Forgery Technologies: Content Signing Approaches.....	18
Summary	19

Executive Summary

Background

The Anti-Spam Technical Alliance (ASTA) is a collaborative effort between six leading Mailbox Providers and the Internet community to establish technical and non-technical solutions for handling unwanted and unsolicited e-mail (spam). ASTA founding members include America Online, British Telecom, Comcast, Earthlink, Microsoft, and Yahoo!. We came together because we share a common set of experiences and problems associated with spam. It is our intent to promote an inclusive process that embraces a broad range of ideas while ensuring that all proposals and recommendations address economic, technical, user, and resource impacts. To that end, we have involved representatives of state, federal, and international governments, consumer advocacy groups, marketers, large commercial e-mail senders, enterprises of all sizes, technology providers, industry standards groups, and trade groups representing other interests.

Scope

This Statement of Intent (SOI) document presents best practices and technologies that ASTA members are implementing to help secure the e-mail infrastructure and bring increased accountability. We fully recognize that this document does not provide all solutions to the spam problem and intend to update and enhance this SOI over time. However, we feel that our recommendations, if implemented on a large scale, can be successful in improving e-mail messaging and Internet communications in general. Because we represent a large percentage of mailboxes on the Internet, we hope to provide the necessary leadership to ensure large-scale adoption within a relatively short period of time.

We expressly solicit feedback from interested parties on the proposed solutions. We also recognize that there is not one easy solution to the spam problem; solutions for handling spam are technically challenging and may take considerable time and effort to implement. This is why we have chosen to pursue multiple approaches: those that can be implemented quickly as well as those that may take more effort and time but have a greater impact. We are hopeful that implementing these approaches will provide significant benefits while the Internet community as a whole continues to develop the infrastructure required to completely solve the spam problem.

We recognize that our recommendations may be applicable to other areas of the Internet and encourage the adoption of these technologies in other protocols as appropriate, for example, instant messaging (IM), Short Message Service (SMS),

Hypertext Transport Protocol (HTTP), Network News Transport Protocol (NNTP), and so on.

Summary of Best Practices

Much of the spam today is generated by spammers exploiting a security flaw in the configuration and application software used to support Internet communications. This includes proxy software, mail server software, and CGI scripts that act as open relays or proxies upon initial installation. More recently, criminals have been creating viruses and worms specifically designed to infect and compromise personal computers around the world. These computers, sometimes called *zombies*, can then be controlled by spammers and used to send out spam on their behalf. Security is paramount in order to combat these issues.

This SOI describes best practices that all Mailbox Providers should follow. These include basic policies for testing and securing open relays as well as unintentional open proxies and CGI scripts. There are a number of recommendations that squarely put responsibility on the sending organization for the transmission of abusive e-mail. This frees the receiving organization from the responsibility of having to continually create new methods for filtering bad content. The primary goal is to stop spam origination points as close to the true source as possible.

The best practices section for bulk e-mail senders and marketers provides high-level recommendations for organizations engaged in legitimate broadcast mailings. This area does not attempt to settle the opt-in versus opt-out debate. We do not recommend specific enhancements as a number of third-party organizations are working to increase the ability of legitimate marketing organizations to self-manage their systems.

The best practices section for consumers encourages end-users to seek out information from several public web sites and their own Mailbox Providers about the available tools to help combat the spam epidemic. We also ask Mailbox Providers to take a stronger role in communicating with their customers and engaging them in a comprehensive awareness program.

Summary of E-mail Forgery

Another major issue with spam is that most spam e-mail falsely represents who actually sent the e-mail. Historically, sender forgery has been used to trick spam filters, but the trend lately has been to use forgery to trick the recipient of the e-mail. We and many industry groups recognize this as a critical problem.

We have begun testing several technologies to provide a more secure e-mail identity. These technologies secure the domain portion of the e-mail address, which is located to the right of the @ (at) symbol. The technologies differ in how they provide this security but have similar end goals.

We believe the use of these technologies can provide a basis for restoring trust in the sender of an e-mail by reducing forgery, eliminating erroneous bounce messaging, and preventing identity theft. We understand that these technologies alone will not solve the e-mail identity problem. The combination of technologies to secure the domain portion of the e-mail address and implementation of several best practices involving e-mail submission and SMTP (Simple Mail Transfer Protocol) authentication must be implemented in parallel to help secure the sender identity (left side of the @ symbol).

Introduction

As of late 2003, well over half of the e-mail messages received by many organizations, ISPs¹, and Mailbox Providers on a typical day were spam—unwanted and unsolicited e-mail sent indiscriminately to users. While spam has been an annoyance to Internet users for many years, the spam nuisance today is significantly worse, both in the quantity and the nature of the material received. There is no “silver bullet” that addresses all aspects of the spam problem; however, this document proposes technology and policy approaches to begin to address this complex problem. Implementing the described approaches will provide a foundation to control the abuse of the Internet e-mail system.

There are a number of methods in use to manage the volume and nature of spam. Many organizations employ filtering technology. Others use publicly available information about potential sources of spam. Still others construct elaborate rules that determine which senders are allowed to connect or deliver mail to their networks and which are blocked. These policy and technology measures can be effective under certain conditions, but over time, their effectiveness degrades due to increasingly innovative spammer tactics. The burden to keep up with these new tactics falls to ISPs, Mailbox Providers, enterprises, and consumers. However, these approaches fail to address the root problems: first, sending junk e-mail is a profitable business for spammers; and second, e-mail messages today do not contain enough reliable information to enable recipients to consistently decide if messages are legitimate or forged.

In this document, we suggest a framework for new mechanisms that enable an e-mail sender to provide proof that an e-mail message is legitimate and not from a spammer. With this framework in place, more effective spam control mechanisms can be built to reduce both the amount of spam delivered and the amount of legitimate e-mail that is blocked in error. Although the effectiveness of the proposed changes depends upon widespread deployment of the new mechanisms, we believe that this goal is achievable. Deployment of the proposed changes by even a small number of Mailbox Providers can lead to the success of this initiative. We are committed to deploying these mechanisms and working to improve them over time.

The initial deployment will be a series of live, real-world tests that demonstrate the feasibility of the proposals and allow for feedback from the community to refine them into suitable solutions.

¹ These include web hosting providers, telecommunications providers, educational institutions, and others who manage networks or provide access to the Internet.

Comments, questions, suggestions, and other feedback regarding this proposal are critical for this process to work. Contact information is provided at the end of this document. We believe that through coordinated action in the Internet community, significant progress in deterring spam can be made.

Considerations

There are several important, non-technical considerations in regard to this effort. First, the introduction of programs to better establish sender identity and encourage the development of trust or reputation systems may involve privacy and legal issues. Notably, the potential impact to free speech is a serious consideration. We remain sensitive to these issues and seek to emphasize technologies and approaches that allow for freedoms of expression to continue.

Second, it is important to note that some of the ideas or approaches presented in this document have been put forth and recommended by others in the community. It is through the collective work and contribution of many technical and business professionals that these recommendations are made possible. In fact, there are several Internet Engineering Task Force (IETF) Request for Comments (RFCs) that address some of the goals set forth in this document. We believe that some of these ideas have not been implemented due to two primary reasons:

1. *The approaches require critical mass of Mailbox Providers to drive standards adoption.* The founding members of ASTA represent a large percentage of the mailboxes on the Internet and are thus in a strong position to address this issue. The support of the biggest commercial e-mail senders will help provide the necessary critical mass for these solutions.
2. *There is no single, simple solution that will serve the requirements of the global Internet population.* We are cognizant of the geographic, cultural, technical, and behavioral differences between the millions of individuals and international organizations using the Internet. We also recognize that organizations other than ISPs, such as those providing mail forwarding services and those hosting multiple domains on a single mail server, have unique concerns about how our proposed solutions impact their operations. It is our intent to respect these differences and allow for flexibility in as many of the proposals as possible.

Our goal is to foster new technologies that can reduce spam while maintaining the flexibility that has already made e-mail such a pervasive and indispensable communications tool. We support approaches that are inexpensive to implement and do not unfairly burden smaller organizations and networks. We seek representation of all viewpoints and desire to be as inclusive as possible in working towards change. We cannot, however, guarantee that new technologies will be implemented without requiring changes to the Internet's e-mail infrastructure. We believe that the challenge facing the community is so significant as to require new ways of thinking about e-mail delivery, security, and verification.

Curbing Spam through Best Practices

Today, almost all spam contains at least one major forgery and is sent using fraudulent means, usually through a security exploit on an ISP network and/or a customer's personal computer. Better awareness of spammer behaviors, closer to the source of the security breach, can significantly reduce spam levels using current technology and protocols.

Important Note: We believe in the "Good Neighbor" policy. Simply put, all abusive traffic emanating from an ISP on port 25 is the responsibility of that ISP to control. If the ISP does not reasonably control abusive traffic, it is at risk of being blocked by other ISPs. This policy applies equally to network and backbone providers and their downstream customers.

Recommendations for ISPs

We recommend a number of best practices that organizations that provide e-mail services and network connectivity (generically referred to as ISPs) should implement as applicable to their environment. The best practices described in this section have already been adopted by most responsible organizations, but we encourage broader global adoption. The combined effect of implementing these approaches can serve to minimize opportunities for spammers.

Close All Open Relays

Issue: Mail servers that allow third parties (unrelated to the owner of the server) to relay mail through them without any formal authentication are considered *open relays*. Open relays let spammers remain anonymous.

Recommendation: Open relays should be reconfigured as secure relays. ISPs should build systems to test all remote mail servers that connect to that ISP to ensure that they are not configured as open relays. ISPs that can not afford to build these systems should subscribe to third party open relay lists available from many anti-spam organizations.

Monitor formmail.pl and Other CGI applications

Issue: Mail-generating Common Gateway Interface (CGI) scripts, formmail.pl, CGI E-mail and similar programs are a specialized set of mail software. For example, formmail.pl is commonly used in feedback forms for web site visitors. When a web site visitor submits their name and address in a form provided by formmail.pl, an e-mail message is generated and sent to the web site owner. However, formmail.pl can be installed insecurely, allowing a spammer to

generate and send e-mail to anyone. When this happens, formmail.pl turns the web server into an open relay.

Recommendation: ISPs, especially web hosting organizations, should regularly scan for misconfigured or outdated programs that can be used to create e-mail. Note that while formmail.pl and other CGI scripts have been exploited most recently, there is the potential for other programs to be targeted by spammers on large scales. Rate limiting on an ISP's outbound systems can help prevent the amount of damage a single insecure script can cause (see the Implementing Rate Limits on Outbound E-mail Traffic section).

Configure Proxies for Internal Network Use Only

Issue: Open proxies, like open relays, allow third parties to anonymously send e-mail through them to any address. However, open proxies are more dangerous in that they allow communications on ports other than the standard e-mail communications port. For example, open proxies can be used in mounting Denial of Service (DoS) attacks, hosting web sites on commandeered personal computers without the owner's knowledge, using commandeered personal computers to redirect users to a hacker's web site, registering e-mail accounts anonymously, and for other malevolent purposes. Since proxies do not normally have logging configured, it is also more difficult to track down open proxy abusers than those using open relays and CGI scripts.

Recommendation: Off-the-shelf proxy software like WinGate, WinProxy, and others should be configured to only allow users on the internal network to use the proxy. Details can be provided by the appropriate technical support organization for the software. ISPs should test their customer's proxies to determine if any are misconfigured and could allow for third party abuse.

Detect and Quarantine Compromised Computers

Issue: Using viruses, worms, and malicious software (malware like BAGEL, MyDoom, SoBig, and so on), hackers and spammers have intentionally deposited millions of "back door" open relays and proxies on the personal computers of unsuspecting users. The spammer community uses this network of compromised computers to generate billions of unsolicited e-mail messages. In addition, hackers have used this network of computers to mount Distributed Denial of Service (DDoS) attacks on web sites, register fraudulent accounts, and lay the groundwork for future anonymous hacking activities.

Recommendation: All ISPs should develop methods for discovering compromised computers. These techniques include the monitoring of abuse complaint feeds from other ISPs (see the Develop Complaint Reporting Systems and Subscribe to Existing Systems section), volumetric tracking of how much

mail is being sent by an internal customer's computer (see the Implement Rate Limits on Outbound E-mail Traffic section), and internal network polling and testing to determine when a computer is open to third party abuse (such as open relay or proxy testing). Computers that show signs of infection should be removed from the network or quarantined until the virus, worm, and/or malware can be removed.

Important Note: As discussed earlier, the Good Neighbor policy requires that ISPs and network providers be responsible for all traffic emanating from their systems on port 25. This is especially important in the case of traffic from a compromised computer since it may include viruses and/or worms that threaten other ISP networks.

For many consumer-oriented ISPs, the simplest solution to stop e-mail worms and spam from their network is to block outbound port 25 traffic. However, blocking port 25 can be problematic for customers who need to run their own mail server or communicate with a mail server on a remote network to submit e-mail (such as a web hosting company or a hosted domain's mail server). In the first case, an ISP should develop a capability to identify customers who have a legitimate need to run a mail server, and then not block port 25 connectivity for these customers. In the second case, ISP customers should use the standard Mail Submission Port, port 587, and ISPs should avoid blocking this port. (For more information, see the Implement Authenticated E-mail Submission section.) A BCP document that discusses this recommendation is also available at:

<http://www.ietf.org/internet-drafts/draft-hutzler-spamops-00.txt>

Implement Authenticated E-mail Submission

Issue: Spammers take advantage of ISPs that do not require senders to authenticate themselves before sending e-mail.

Recommendations: ISPs should implement authenticated e-mail solutions that require valid credentials from users before they can send mail. This can be accomplished in many ways, but it is usually implemented by requiring some form of a username and password to validate the user account on the system.

Note: The importance of an authenticated e-mail infrastructure is directly related to the central issue of verifying the identity of an e-mail's sender. Once the identity of the sender is verified, the recipient can trust the validity of an e-mail. If an abuse issue arises, an ISP can trace the origination of the e-mail to the offender to stop further abuse.

We recommend implementing the strongest possible authentication methods available. For example, although the SMTP AUTH protocol is a good form of

authentication, it sends the password of the user in clear text format across the public Internet. A stronger form of authentication is the deployment of SMTP AUTH with the STARTTLS protocol. This protocol encrypts the password, analogous to the way that HTTPS secures passwords on a web site.

We further recommend that SMTP authentication be implemented on the standard Mail Submission Port, port 587, and that ISPs encourage their customers to switch their mail client software (for example, MS Outlook, Eudora, and so on) to this port. Using this port will provide seamless connectivity that does not depend on if a network allows port 25 traffic.

Details on the Mail Submission Port and authenticated SMTP can be found at the following:

<http://www.ietf.org/rfc/rfc2476.txt>

<http://www.ietf.org/rfc/rfc2554.txt>

<http://www.ietf.org/internet-drafts/draft-hutzler-spamops-00.txt>

Remove Remote Access to CPE

Issue: Hackers and spammers exploit Customer Premises Equipment (CPE), such as routers and broadband/SOHO routers, as relay points or proxies to forward spam. Normally, the hacker or spammer must know the CPE password to program the CPE to perform this function.

Recommendation: All network providers and consumers should ensure that remote access to CPE is turned off, or at least that the CPE does not respond to a known default password, for example, a blank password for the admin user.

Implement Rate Limits on Outbound E-mail Traffic

Issue: Over the last year, hackers have begun to conspire with spammers, resulting in new e-mail viruses and worms that commandeer personal computers for use by spammers. Viruses such as mydoom can compromise millions of computers in the span of several days. These computers can then start generating high volumes of spam. The situation has been widespread at ISPs that do not require e-mail authentication. But ISPs that do employ account authentication have also seen an increase in the hijacking of accounts via other techniques such as password phishing and trojans with keystroke loggers.

Recommendation: ISPs should implement rate limits on outbound e-mail traffic through the ISP's primary SMTP gateway hosts. These limits should be based on To/Cc/Bcc recipient counts per unit of time from an end user account or server IP address (using e-mail messages per unit time or recipients per message

methods will not be successful). Typically these rate limits vary widely, but the goal is to prevent a compromised account from sending spam to millions of recipients in a short timeframe. A suggested set of rate limits for a consumer-oriented ISP is a maximum of 150 recipients in 1 hour and 500 recipients in 1 day combined with some limit of spam complaints. As discussed earlier, we also recommend that ISPs implement a secure SMTP authentication technique so that only registered members are able to use the ISP's mail servers.

Control Automated Registration of Accounts

Issue: Spammers and hackers have found methods for automatically registering millions of accounts with ISPs, especially free e-mail account providers. These accounts can be used for many purposes including sending spam and mounting DoS attacks.

Recommendation: ISPs should develop and implement methods for blocking the automated generation of accounts. For pay-for-service ISPs, this may involve requiring and preauthorizing a payment method before allowing access to an account and/or e-mail privileges. For free ISPs, this may involve tests, for example, a Turing² test, to ensure that a registration request is not being generated by an automated script.³

Close Web-based Redirector Services Susceptible to Abuse

Issue: Web based redirection services are used by many organizations to count click-throughs in the ad-serving areas of their organizations. Unfortunately, many of these redirection services are open to use by anyone, not just the clients of those organizations who are legitimately serving advertisements. This lets spammers create their own advertised sites and products that appear to be endorsed by an ISP that is running an insecure redirection service. Not only is it difficult for ISPs and filtering programs to block spam URLs when a redirector is used, but also consumers can be tricked into thinking a spam URL is legitimate.

Recommendation: We recommend securing all web-based redirectors that can be used by third parties without permission.

² Using current Turing technology may have an adverse impact on users who have visual disabilities.

³ Rate limiting registrations based on IP addresses is not an effective solution due to the vast number of open proxies available on the Internet.

Develop Complaint Reporting Systems and Subscribe to Existing Systems

Issue: Detecting spam is especially difficult without knowing exactly what e-mail recipients believe is and is not spam. Because of this, it is critical for ISPs to provide methods for customers to provide feedback about what they consider spam. Using this feedback, ISPs can analyze the data to determine the type and level of spam they receive and the methods spammers are using to circumvent filtering and blocking systems.

Recommendation: All ISPs should develop a system for customers and external parties to report spam. The system should be simple to use and keep the content of the original e-mail intact so that it can be used to improve filtering and potentially trace spammers for litigation purposes.

Larger ISPs with sufficient volumes of complaints and resources may develop more sophisticated systems in which complaints can be automatically uploaded into a database. This database can be used to trace spam from internal sources not caught by rate limiting and to provide feedback to other ISPs about spam originating from their networks.

Additionally, as required by the e-mail RFCs, all organizations should have abuse and postmaster mailboxes that are monitored at least daily for abuse reports coming from external sources.

Several ASTA members are developing or have deployed capabilities to share complaints with other organizations. ISPs interested in learning more about spam emanating from their networks are encouraged to visit ASTA members' individual web sites for more information.

Note: We are working with the Anti-Spam Research Group (ASRG) to establish proposed abuse sharing formats for the IETF standards process. These formats will permit sharing abuse complaints among service providers. More information about this format should be available shortly.

Recommendations for Bulk E-mail Senders

In general, all organizations, commercial or otherwise, should follow generally accepted industry practices for sending e-mail. The following is a partial list of these practices:

- Do not harvest e-mail addresses (defined as collecting e-mail addresses, usually by automated means) through SMTP or other methods without the owners' affirmative consent.

- Always provide clear instructions to customers about how to unsubscribe or opt-out of receiving e-mail. Promptly respond to these requests.
- Do not use or send e-mail that contains invalid or forged headers.
- Do not use or send e-mail that contains invalid or non-existent domain names in the From or Reply-To headers.
- Do not employ any technique to hide or obscure any information that identifies the true origin or the transmission path of bulk e-mail.
- Do not use a third party's Internet domain name or be relayed from or through a third party's equipment without permission.
- Do not send e-mail that contains false or misleading information in the subject line or in its content.
- Monitor SMTP responses from recipient's mail servers. Promptly remove all e-mail addresses for which the receiving mail server responds with a 55x SMTP error code (for example, *user doesn't exist*).
- Consider working with an e-mail accreditation provider to demonstrate to ISPs and Mailbox Providers that your company's mail meets the highest industry standards.
- Use different IP addresses based on the type of mail or customer. For example, use one IP address for order confirmations, another for customer service, and another for customer newsletters. This can help avoid having a problem on one type of mailing affect all communications or customers.

Other practices have been defined by various groups, including the Mail Abuse Prevention System (MAPS), the Internet Architecture Board (IAB), Network Associates, Inc. (NAI), and the Direct Marketing Association (DMA). In addition, some organizations and state and federal governments may have stricter requirements and restrictions for communicating with customers or constituents. For more information, see the following:

<http://mail-abuse.com>

<http://www.iab.org>

<http://www.nai.com>

<http://www.the-dma.org>

Recommendations for Consumers

Consumers share some of the responsibility in curbing spam. Many users today have not activated available options in their computer software that can protect

them from spammers and hackers. Consumers should learn about these features and ensure that they are using the right level of protection for their needs.

Consumers should install firewalls on their PCs and use up-to-date anti-virus software, along with other screening tools, to detect incoming viruses, malware, and harmful or suspicious code.

We recommend that ISPs and Mailbox Providers proactively raise awareness of tools available to customers to fight spam. We encourage ISPs to emphasize this topic when customers sign up for service and to send reminder communications to new and current customers. In addition, ISPs should devote a comprehensive part of their web sites to informing customers about how they can fight spam.

In addition to your ISP's email support and documentation pages, the Federal Trade Commission (FTC) and FirstGov for Consumers also provide information for consumers. For more information, see the following:

<http://www.ftc.gov>

<http://www.consumer.gov>

Curbing E-mail Forgery

When the SMTP protocol was developed, the Internet was small, and participants could be trusted. At that time, the fact that the SMTP protocol provided no mechanism for sender authentication at the user and server levels and no guarantee that any of the e-mail message headers were accurate were not major concerns. Today, however, one of the key challenges facing the Internet community and anti-spam technologists is that messages do not contain enough reliable and verifiable information to let recipients (users and ISPs) decide if messages are legitimate or spam. In particular, it is not possible to reliably identify the true sender of an e-mail message by looking at the message headers and content. Spammers take advantage of this fact and disguise the origin and true sender of their messages by forging the sender address and domain name. This is called *domain spoofing*. Virus writers also use this technique to hide their location and to encourage recipients to open and install e-mail viruses and worms (typically by forging the address of a company or person that is well known to the recipient). Likewise, identity thieves use spam to distribute phishing scams that attempt to entice users to provide their personal information.

Although the problem of identifying the sender or origin of e-mail is complex, there are two promising new methods that organizations can implement to lay a foundation for future advances: IP address approaches and content signing approaches.

Anti-Forgery Technologies: IP Address Approaches

Currently, the only trustworthy attribute in an e-mail message header is the Internet Protocol (IP) address of the server that is transmitting the e-mail.⁴ Given this, IP addresses can be used by e-mail receivers to verify other attributes in the message header, such as the sending domain, and thus help reduce the common forms of phishing and forgery that are rampant today. This verification loop can be done using the existing DNS infrastructure combined with fairly simple changes to the receiver's mail systems.

Methods of Implementation

One way to implement an IP-based sender authentication system is for domain owners to publicly publish the IP addresses of mail servers authorized to send

⁴ It is possible to forge the IP address; however, we recognize that this is difficult to achieve on a large scale for session-based traffic. In addition, if hackers gain the ability to forge IP addresses on a large scale, this presents a problem that is inherently more dangerous than spam to the Internet's infrastructure and community.

mail on behalf of their domains. This allows the receiver of the mail to validate the domain information in the headers of an incoming message. This would be appropriate for any organization wishing to defend itself against domain spoofing, and especially valuable to organizations whose brands are compromised by spammers and phishing schemes such as financial services and e-commerce companies. The existing DNS system can be used to publish these mail server IP addresses. If the sending mail server is not listed as legitimate for the domain, the receiving system can consider the message a forgery attempt and apply a local policy, such as quarantining the message, putting it in a special folder for the recipient, or blocking it altogether. For cases when the sending server is verified as legitimate, the receiving server may also mark the e-mail in some way to notify the recipient that the message has a higher level of trust.

One side benefit to an IP-based approach is that it provides information about outbound mail delivery servers. The publishing of outbound IP address information in DNS can be used to help build and maintain Internet mail whitelists, sometimes referred to as a *web of trust*.

We are working closely with the Internet community to test IP based authentication technologies. Throughout the implementation process, we will provide implementation feedback that, along with other industry-wide feedback, will enable subsequent improvements to the specification. The goal is to provide the best long-term, industry-wide IP-based authentication solution.

Anti-Forgery Technologies: Content Signing Approaches

Another approach to sender authentication uses a technology called Content Signing (CS). CS systems use public/private key pairs to generate the signatures that are used for sender verification. The public keys may be made broadly available through a variety of key exchange mechanisms or via publication in a directory or in DNS. The private keys are stored securely on the domain's mail servers. When a user sends an e-mail, the mail server uses the stored private key to automatically generate a digital signature for the message. When the recipient's mail server receives the e-mail, it retrieves the sender's public key and uses it to verify the digital signature in the message. This verifies both the sender's identity and the integrity of the message body (that the e-mail content was not modified during delivery).

As with IP-based sender authentication, we believe that content signing technologies are an important component of a long-term industry solution. We are working closely with the Internet community to test content signing technologies. Throughout the implementation process, we will provide implementation feedback that, along with other industry-wide feedback, will enable subsequent improvements to the specification. The goal is to provide the best long-term, industry-wide content signing solution.

Summary

Implementing the recommendations in this document can help reduce many of the worst types of spam, forgery, and spoofing that occur in e-mail. These technologies will not stop spam entirely, but implementing these recommendations will significantly enhance the Internet community's ability to trace the source of spam and hold senders accountable for their actions. This ability will provide the foundation for building future solutions.

We look forward to the community response to these recommendations. We invite participation from all segments of the community to assess the validity and impact of these proposed solutions and their accompanying technical specifications.

To participate in the testing of the identity technologies, readers are encouraged to join the IETF working group MARID (MTA Authorization Records in DNS) at:

<http://www.ietf.org/html.charters/marid-charter.html>