
Sender ID Framework

Protecting Brands and Enhancing Detection of
Spam, Phishing, and Zero-Day Exploits

A White Paper

Published: April 2007

Microsoft®

The information contained in this document represents the current view of Microsoft Corp. on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKE NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT. Complying with all applicable copyright laws is the responsibility of the user. Microsoft grants you the right to reproduce this white paper, in whole or in part, specifically and solely for the purpose of personal education.

Microsoft may have patents, patent applications, trademarks, copyrights or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights or other intellectual property.

Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place or event is intended or should be inferred.

Microsoft, MSN, Hotmail, Windows, Windows Live Hotmail, SmartScreen and Outlook are either registered trademarks or trademarks of Microsoft Corp. in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

© 2007 Microsoft Corp. All rights reserved

TABLE OF CONTENTS

Table of Contents	3
Executive Summary	4
Assessing the E-mail Threat	5
Sender ID, a Critical Component of Your Defense	6
Why Implementing Sender ID is Good for Business	12
Benefits of Reducing Unwanted E-Mail	13
Assuring That Your Domain Is Not Spoofed	14
How Windows Live Hotmail Uses Sender ID to Reduce Spam, Improve Deliverability & Help Protect Users	16
Case Study Overview	16
Situation	17
Solution	17
Benefits	20
Conclusion	22
For More Information.....	22
Acknowledgements	22

3.8 Billion

The average number of messages sent to Hotmail that are estimated to be spam *each day!*

20 Million

The average number of forged messages that fail Sender-ID and are not delivered to Windows Live Hotmail users *each day.*

98%

The percentage of Phishing messages that are spoofed and can be caught by Sender ID

Up to 98%

The percentage of messages spoofed from established brands that have been caught by Sender ID

90%

The percentage of e-mail marketers who have implemented Sender ID

8%

Percent of additional spam detected by Sender ID

30 minutes

The amount of time to read this paper and understand Sender ID benefits

\$0

The cost to authenticate your outbound e-mail via Sender ID.

EXECUTIVE SUMMARY

Unsolicited e-mail has become an expensive and at times dangerous problem for users, organizations and companies who depend on e-mail to communicate with friends, colleagues, and customers. Unfortunately, increasing volumes and the malicious nature of spam continues to undermine productivity, trust, and confidence in e-mail.

Sender ID offers a relatively simple and cost effective approach to fighting spam and phishing by detecting e-mail forgery or sender spoofing. Sender ID authenticates inbound e-mail to verify it is from the domain it says it is from. Messages that have been authenticated by Sender ID are less likely to be spam and those that fail Sender ID are more likely to be spam. Anti-spam solutions consider the Sender ID authentication result when determining if a message gets delivered, helping to improve deliverability of legitimate e-mail.

Organizations that have implemented Sender ID have realized the following benefits¹:

- **User protection** – Over 90 percent of phishing and virus laden e-mail forge the “from line” or sender’s address to deceive users into thinking the e-mail is from a legitimate source. Sender ID detects spoofed e-mail from authenticated domains and protects users from these attacks which often contain dangerous payloads and try to trick users into divulging personal or corporate information.
- **Brand protection** – Sender ID identifies and validates messages that claim to be from your organization, helping to protect your brand from spammers that claim to be you. Windows Live Hotmail uses Sender ID to help block over 20 million forged e-mails each day, helping to protect those brands and their customers.
- **Enhanced deliverability of legitimate e-mail** – Organizations and marketers who have adopted Sender ID, and have a positive reputation, realize increased deliverability with up to 85 percent fewer messages mistakenly marked as spam.
- **Return on Investment** – Sender ID was designed as a no or low-cost solution, easy to implement and manage with no measureable impact to system performance. Inbound authentication is available from over a dozen commercial and open source solutions. Together, user and brand protection along with enhanced e-mail deliverability provides a significant ROI.

¹ Data based on Windows Live Hotmail internal analysis completed April 2nd, 2007

ASSESSING THE E-MAIL THREAT

Situation

The SMTP protocol provides no formal way for a receiving e-mail system to validate that a message was sent from the apparent user or domain. This makes it very easy for an unscrupulous sender to spoof a legitimate domain or brand identity and deceive a recipient into opening a message. It also compromises brand identity among spoofed organizations and reduces the confidence of e-mail senders and receivers alike in the value of e-mail as a means of communication.

Solutions

The Sender ID Framework provides a way for receiving e-mail servers to authenticate an incoming e-mail message and validate its source.

Benefits

Sender authentication using Sender ID decreases the likelihood that your outbound e-mail will be identified as spam and reduces the number of unwanted messages allowed into your organization.

The purpose of this paper is to demonstrate the business and technical value of the Sender ID Framework (SIDF) and its critical importance in aiding in the detection of spam and phishing, helping to protect against unwanted and potentially malicious e-mail, and helping to protect domains and brands from spoofing, while increasing the deliverability of legitimate e-mail.

Using SIDF as part of your anti-spam solution adds business and technical value by increasing user and customer satisfaction while decreasing levels of deceptive and malicious e-mail that enters your network. SIDF provides the following primary benefits:

- Enhance anti-spam solutions in detecting and identifying unwanted e-mail.
- Help protect brands from being spoofed or used for malicious purposes.
- Increase the deliverability of legitimate e-mail.
- Improve user trust and confidence in e-mail.

Implementing SIDF, for most organizations, is straightforward and relatively inexpensive:

- Inbound authentication is technology *agnostic* as it relies on DNS.
- SIDF functionality is supported by many anti-spam solutions.
- Brand and domain protection with SIDF is done by simply adding a record to DNS.
- No additional client software is needed.
- SIDF authentication is a simple DNS query and doesn't affect overall performance.

Spammers send approximately 55 *billion*² messages a day; add to that 1.3 billion mailboxes worldwide³, and the result means roughly 37 spam messages per day for each mailbox — nearly double the estimate from a year ago and a cost to companies in the billions of dollars. Unscrupulous individuals and companies send spam from everywhere on earth. Unchecked, those messages end up in front of your customers and employees, enticing them to open virus-laden attachments or to reveal personal or corporate information that criminals can use against that person or your company. SIDF, along with the reputation of the sending domain and respective IP addresses, can help protect your users and your brand from the threat and inconvenience of unwanted e-mail.

² IronPort Systems, Inc. (28 June 2006). http://www.ironport.com/company/ironport_pr_2006-06-28.html

³ Radicati Report: MICROSOFT EXCHANGE AND OUTLOOK, MARKET ANALYSIS, 2005-2009

Glossary of Anti-Spam Technologies

IP Connection Filtering

The most rudimentary method of connection-level spam protection involves manually managing lists of SMTP hosts using SMTP connections that you choose not to accept.

Block List Filtering

A more dynamic means of providing connection-level protection, block lists are lists of IP addresses that are either known sources of spam, open relays, or part of an IP scope that should not include an SMTP host. Anti-spam solutions query the block list provider when evaluating the inbound SMTP connection.

Recipient and Sender Blocking

Another way to reduce spam manually is to define individual senders or domains from which you do not want to accept messages.

Content Filtering

Content Filtering distinguishes between characteristics of legitimate e-mail and spam. Content Filters assess the probability that an e-mail message is either a legitimate message or spam.

Anti-Phishing

Phishing is a type of deception designed to steal personal and private information. In phishing exploits, the phisher attempts to get the user to disclose valuable personal data, such as credit card numbers, passwords, account data, or other information, by convincing the user to provide it under false pretenses.

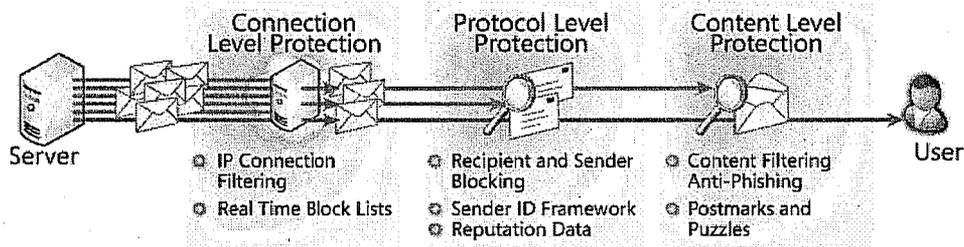
Postmarks and Puzzles

E-mail clients can create a message-specific puzzle and solution, known as a *postmark*, which is attached to each outgoing message. The postmark requires a number of CPU cycles to create and decipher. Spammers generally don't have the time or computational resources to attach complex individual puzzles and solutions to thousands of outgoing messages, so they don't use them.

SENDER ID, A CRITICAL COMPONENT OF YOUR DEFENSE

Today's anti-spam solutions are composed of several methods or layers for detecting spam. During each stage of inbound SMTP delivery, different technologies are used in different ways. Ideally, unwanted e-mail is blocked at the network perimeter, reducing exposure to inbound threats while having the least impact on computing resources.

However, at the edge or connections level it can be difficult to determine if a message is spam until the sender's identity, reputation, and content can be evaluated. This section shows how Sender ID integrates into a typical anti-spam solution.



Levels of Protection

The first level of protection against unwanted e-mail is at the initial connection between the sending SMTP server and the receiving SMTP server. Connection-level protection is among the most beneficial layers of defense against spam, because the spam never enters the organization and has the least impact on computing resources. Two common types of connection-level protection are *IP Connection Filtering* and *Real Time Block Lists*, explained in detail in the left column.

After the inbound SMTP message has advanced beyond connection-level protection, the next layer of defense is at the SMTP protocol level. The SMTP dialog between the sending SMTP host and the receiving SMTP host is analyzed to verify the permission of the sender and recipients, and to determine validity of the sender's SMTP domain name for the SMTP host that sent the message. Along with the Sender ID Framework and reputation data, explained in detail below, *Recipient and Sender Blocking* is considered protocol level protection.

After the spam filter applies connection-level and protocol-level filtering technologies, the next line of defense is to analyze message content for characteristics that may indicate spam and malicious e-mail. *Content Filtering*, *Anti-Phishing*, and *Postmarks and Puzzles* are all content-level filtering methods that establish a rating that determines if the inbound e-mail is legitimate, malicious, or un-solicited spam.

Tools for creating SPF records

<http://www.microsoft.com/senderid/wizard>

How Sender ID Works

The Sender ID Framework⁴ verifies that messages originate from the indicated domain as seen by the user in their e-mail client or Web interface. As an example, when susans@nwtraders.com sends a message, Sender ID attempts to verify that the message comes from an SMTP host belonging to or authorized to send mail on behalf of nwtraders.com. After authentication, reputation data can be used to rate or score the sender, nwtraders.com. The combination of SIDF authentication and IP reputation data, in addition to the other anti-spam counter-measures, helps to determine if the message is legitimate and trustworthy, spam, or malicious.

Walking through figure 1 to the right, for each inbound message, Sender ID obtains the sending host's IP address from the incoming TCP connection. Sender ID acquires the sender's domain name from the SMTP "MAIL FROM" or EHLO command or from a message header field determined by the Purported Responsible Address (PRA)⁵ – an algorithm that determines the address responsible for sending the message^(step A). Sender ID then queries the DNS for a text record, known as a Sender Policy Framework (SPF) record, that lists all authorized SMTP hosts^(step B). Using the SPF record, Sender ID attempts to authenticate the message and assigns a weight to the message depending on the verdict result^(step C). The authentication is authoritative because the DNS zone that hosts the queried SPF record is controlled and secured by the sending organization.

Figure 1 shows Fabrikam Marketing sending a message from nwtraders.com. This message may or may not be spam. Fabrikam Marketing might be contracted by Northwind Traders to send their newsletters or to respond to support requests. It is the SPF record that will authorize who is authorized to send e-mail on behalf of nwtraders.com.

Once the Sender ID process is complete, the results and scoring are combined with scoring of other anti-spam technologies to determine the message's trustworthiness.

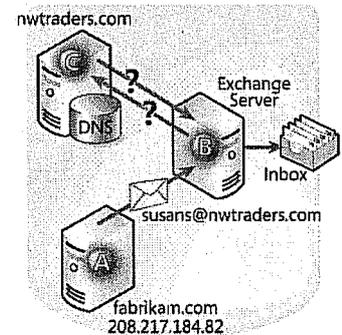


Figure 1

⁴ RFC 4406 <http://www.ietf.org/rfc/rfc4406.txt>

⁵ RFC 4407 <http://www.ietf.org/rfc/rfc4407.txt>

The effort necessary to create and manage an SPF record in DNS is technically simple and straightforward. Most organizations with only a few dozen outbound e-mail services can create an SPF record by using one of the many tools or wizards available on the Internet⁶. The most challenging effort is not technical, but one of “inventorying” third parties and their respective IP addresses, which groups within an organization may have delegated or authorized to send mail on their behalf. In our example above, Northwind Traders hires Fabrikam Marketing to create and send a newsletter to Northwind Traders customers. Fabrikam Marketing will send the newsletters, but they will be addressed from nwtraders.com.

In this scenario, if IP addresses are not included in their SPF record, receiving networks might reject the message or mark it as forged. Other examples of third-party suppliers that may need to be added to an SPF record include outsourced customer or product support, public relations, and shipping confirmations.

Sender ID and Reputation Data Work Together

The relationship between SIDF and reputation data is similar to that between a driver’s license and a driving record. Just as SIDF vouches only for the origin of the message, not its content, a driver’s license validates only that a driver has passed a test and says nothing about driving ability or recent infractions.

Driving records document the driver’s reputation. The better the reputation, the more privilege the driver enjoys, such as lower insurance rates and so on. Similarly, the better a sender’s reputation, the fewer restrictions may be placed on their e-mail, such as restrictions on attachments or attachment types, the number of messages they’re allowed to send into the organization, and so on. While SIDF and reputation data offer independent benefits, their combined value is greater, providing quantifiable business value to both the sender and the receiving network.

Reputation data is available from multiple sources. The majority of Internet Service Providers (ISPs) have developed their own reputation systems based on spam complaints and user feedback loops. Factors that influence a sender’s reputation include how many

⁶ <http://www.microsoft.com/senderid/wizard>

messages a particular domain has sent in the past, the frequency of such e-mailings, and how many messages users report as spam. Other sources of reputation data include trap accounts (a.k.a. honey pots), which are used exclusively to detect address book attacks⁷, or e-mail harvesting campaigns by spammers. If a message is sent to a trap account, it is likely a result of an e-mail harvesting campaign, and the sending domain's reputation is impacted accordingly.

Senders with no reputation may receive a neutral or negative score and be faced with limitations to their e-mail. These limitations may include blocking attachments and disabling images, inbound throttling (limiting the amount of e-mail from a given IP address daily), and other steps to mitigate potential user risk while the sender's reputations can be established and validated.

Marketers and high-volume legitimate e-mailers who send on behalf of an organization can enjoy the benefits of their customer's positive reputation. By adding their sending IP addresses to their customer's existing SPF record, the sender can realize the benefit of the existing reputation of all other IPs within the same record. This functionality of SIDF and usage of SPF records provides added flexibility to marketers and high-volume legitimate mailers.

The Importance of Change Management for SIDF

As with all security processes, to maintain the SPF record it is necessary to employ a change management process. That process will define how new sources of SMTP messages from your domain are communicated and approved for addition to the SPF record. In the example above, the marketing department would follow the change management process for adding the outbound SMTP hosts for fabrikam.com to the nwtraders.com SPF record. The change management group at Northwind Traders would approve the SPF change and request the updated SPF records with new or replacement IP address be published in the DNS records for contoso.com . It is recommended that such changes to the SPF records should be made 48 hours in advance of any major e-mail campaign, to allow for DNS replication and SPF record cache deployments to be updated.

⁷ An address book attack sends messages to random domain recipient names. Messages that do not bounce as *unknown recipient* are identified by the spammer as legitimate e-mail addresses and are added to e-mail lists that are sold.

Sender ID – Value to Online Banking

Today, a growing number of financial institutions throughout the world have successfully deployed SIDF and published SPF records. A real-world example is a large international bank with more than 16 million online users and over 4,800 U.S. branches. This North America-based bank was an early adopter of SIDF and was highly motivated to protect its customers and associates from phishing attacks and deceptive e-mail. Prior to implementing SPF, the bank had to perform several tasks to ensure that all mail legitimately purporting to be from one of the bank's several domains had an associated SPF record.

The first step was to obtain executive support for the project in order to gain management's backing and apply the appropriate priority to activities necessary to achieve the goal of protecting the bank's brand and its customers. Understanding the risk to online banking and user trust and confidence, e-mail authentication received approval and was placed on an implementation fast track.

The second task was to create a comprehensive inventory of all entities sending mail using the bank's domains. This included internal mail systems and applications as well as outsourced vendor systems. During the inventory process, the bank involved several key internal constituents whose involvement was vital to the project's success.

These groups included:

- Advertising/Marketing
- Helpdesk/customer service
- Event marketing
- Corporate e-mail systems administration
- Transactional e-mail alerts
- Public relations
- Investor relations
- Branch Banking
- Human resources
- E-product delivery

Once inventory, communication, and administrative processes were in place, the next step was deployment. A clearly written policy was distributed stating the purpose of the project and the inventory. *The policy emphasized the importance of the requirement to keep the inventory updated:* it would result in less e-mail destined for customers that may be identified as spam. This policy included:

-
- A domain/host policy that defines a data collection and update methodology, risk analysis, and key contacts.
 - A communication plan that includes DNS administration, international/domestic business units, third-party senders, and supply chain management.
 - Providing a process for lines-of-business and vendors to update their inventory information. This is critical to keeping SPF records updated – again emphasizing that outdated SPF records may result in legitimate e-mail being treated as spam by receiving domains.

The project was moved through planning, testing, staging, development, and production cycles to ensure predictable behavior. Today, the bank enjoys the benefits of SIDF and the knowledge that it has taken the steps necessary to protect its brand and customers.

Helping in the Fight Against Viruses

As mentioned in the introduction, Sender ID verifies that messages originate from the domains they claim to be from. Unfortunately, messages from spoofed domains can be the most dangerous because they often contain viruses and/or zero-day exploits.

Sender ID can provide your infrastructure and client PCs with an added layer of protection against these threats. These exploits take advantage of the lag between the time when a security vulnerability is discovered and the time when antivirus signature updates are made available. During this lag, Sender ID and reputation data protection may be one of the only ways to prevent such exploits from entering your organization.

"Sender ID is providing real business value to the entire ecosystem, helping to protect consumers and millions of businesses from online exploits. With adoption now over 42% of all legitimate e-mail worldwide, Sender ID is dramatically helping ISPs and receiving networks to improve the deliverability of and trust of legitimate e-mail."

*John Scarrow
General Manager Anti-Spam Technologies
Microsoft Corporation*

WHY IMPLEMENTING SENDER ID IS GOOD FOR BUSINESS

Whether your organization depends on e-commerce, e-marketing, or e-mail to function properly, Sender ID is good for all stakeholders, including your business, customers, stockholders, and employees. Any organization with an online presence, who depends on e-mail to conduct business, or uses the Internet to sell goods and services benefits from implementing Sender ID because it affords an increased level of protection from phishing exploits and other spoofed e-mail messages that can create doubt and suspicion. Companies that use Sender ID help protect users and network resources from unwanted e-mail and the dangerous payloads it often carries.

Prime Example

GoDaddy.com is the world's largest domain name registrar, with approximately 20 million domains under its management. GoDaddy.com also provides online Web-based e-mail services with more than 4 million mailbox customers receiving over 30 million e-mails a day. To protect those customers, in 2004 GoDaddy.com implemented Sender ID to work with other anti-spam technologies to identify both wanted and unwanted e-mail. The results of this effort have included the following:

- Over 50,000 SPF records for Go.Daddy.com customers have been published.
- Approximately 30 percent of all legitimate, "non-spam" inbound e-mail messages received by GoDaddy.com have an associated SPF record.
- 88 percent of those e-mail messages pass Sender ID.
- 8.5 percent of those e-mail messages fail Sender ID the remainder are neutral.

By implementing Sender ID, GoDaddy.com not only has reduced false-positives, but also reduced the amount of spam that reaches its customers' Inboxes and helps protect customers from dangerous e-mail. For a minimum investment in time and resources, GoDaddy.com and its users have realized a reduction in delivered spam and increased level of customer satisfaction.

Benefits of Reducing Unwanted E-Mail

In addition to reducing spam and improving user productivity, organizations can realize additional benefits from implementing SIDF. These benefits come as a result of keeping unwanted e-mail out of your organization and include:

Protect Corporate Assets

Spam can propagate viruses and keystroke loggers that destroy, corrupt, or steal corporate assets. Spoofed or forged e-mail provides these criminals with a direct path to your user desktops.

Protect Users and Their Identity

Users in most organizations have become accustomed to receiving e-mail via their corporate account that may be personal or otherwise non-work-related. Therefore, if a user gets a message from their bank asking them to reset their password or reveal their account number, it is not that far-fetched to imagine the user responding. Phishing schemes like this are a common threat to all e-mail users.

Sender ID can play an important role in reducing the number of phishing messages allowed into your organization, thereby reducing the threat to your users.

Decrease Resource Consumption

When a message enters an organization and is delivered to the recipient, it consumes network and storage resources. For example, GoDaddy.com receives 30 million messages a day, more than 5 million of which have an SPF record. Of the e-mail messages with an SPF record, more than 400,000 fail the Sender-ID test. If GoDaddy.com did not identify those messages as spam when they entered the organization, the messages would consume storage resources and in some cases require archiving and backup for long periods of time. Stopping unwanted e-mail before it enters an organization frees up network and storage resources for more streamlined business operations.

Inexpensive Implementation Costs

Sender ID has two components: an SPF record inserted in the DNS zone file, which identifies the SMTP servers that are allowed to send e-mail from your domain; and the inbound server protocol. Sender ID is a royalty-free standard, compliant with IETF, RFC and the open source communities, and is now available in the majority of e-mail MTA solutions.

In many cases, implementing inbound e-mail authentication via Sender ID can be as simple as selecting a check box in the e-mail administrator's management console, as is done in

several solutions including Alt-N, Barracuda, Postfix, Port25, Sonicwall, Sendmail, SecureComputing and Microsoft® Exchange Server 2003/2007.⁸

Assuring That Your Domain Is Not Spoofed

When users receive e-mail that is falsely addressed from your domain, it can open up several risks and vulnerabilities – for example, it can affect your brand image and overall confidence in your organization, as the user is unable to distinguish between legitimate and deceptive e-mail from your domain. Today, more than 27 percent of Fortune 500 companies have implemented Sender ID, while 37 percent of financial institutions⁹ and over 90 percent of marketing-related domains have implemented Sender ID. As of April 2007, an estimated 9 million domains have adopted Sender ID worldwide, a threefold increase over the last 12 months. More than 43 percent of legitimate e-mail today is SIDF-compliant.

In addition to creating and sustaining customer confidence in your brand, SIDF-authenticated e-mail also helps to protect your customers from ID theft and loss of privacy. According to the Anti-Phishing Working Group, more than 95 percent of all phishing messages are forged, attempting to deceive the recipient into trusting and opening the e-mail and divulging personal information.

How eBay and PayPal Use SIDF to Help Maintain Customer Confidence

As much as any other organization, eBay and PayPal have a major stake in upholding online user trust and confidence and preventing spammers from spoofing their domains. As two of the most common targets of phishing exploits, both eBay and PayPal rely on Sender ID¹⁰ to help assure their customers that e-mail messages sent to them from the companies' domains are legitimate, and just as importantly, to help receiving networks block deceptive e-mail purporting to have come from the companies' domains.

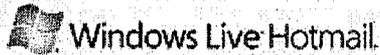
As with any Internet brand, if customer confidence in eBay and PayPal diminishes, their business will suffer and customers may be less likely to use their services and trust or open legitimate e-mail sent by the companies.

⁸ For a listing of SIDF-enabled solutions, visit <http://www.microsoft.com/senderid>

⁹ Source: April 2007 study by the Authentication & Online Trust Alliance www.aotalliance.org

¹⁰ http://ebaydeveloper.typepad.com/dev/2005/03/preparing_for_1.html

Taking a leadership position, eBay has been working closely with several anti-spam vendors, industry working groups, and ISPs to campaign for the implementation of e-mail authentication protocols including Sender ID. Based on their confidence in SIDF and on escalating threats to users, eBay and PayPal have now begun asking receiving networks to more definitively block and delete any e-mail from eBay.com that fails the Sender ID test.



How Windows Live Hotmail Uses Sender ID to Reduce Spam, Improve Deliverability & Help Protect Users

Case Study Overview

Profile:

Windows Live Hotmail, the next generation of MSN® Hotmail, is a free e-mail service provided by Microsoft. Windows Live Hotmail is one of the largest online e-mail service providers with over 280 million customers in 220 countries.

Business Situation:

Due to the popularity of the service and number of active users, Windows Live Hotmail receives an enormous amount of unwanted e-mail daily. Without efficient and accurate spam protection, user trust, satisfaction, and usability would be tarnished.

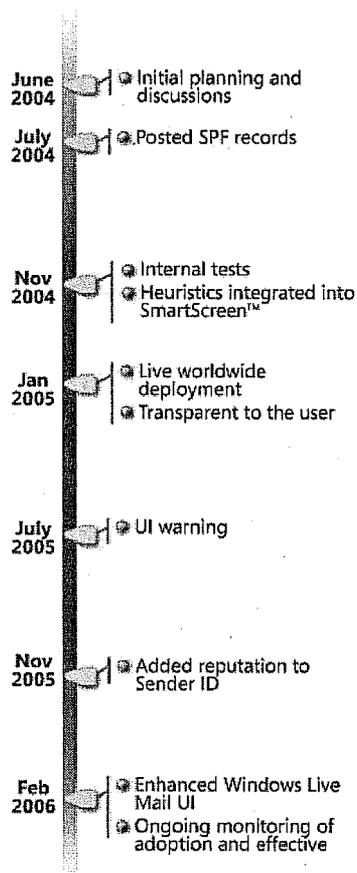
Solution:

The Windows Live Hotmail team worked with Microsoft research in the development of the Sender ID Framework to reduce the amount of unwanted e-mail sent to Hotmail customers.

Benefits:

- Sender ID yields an 8% improvement in spam detection.
- 87-percent reduction in false positives (for senders with good IP reputation data) compared to those who have not adopted Sender ID.
- A reduction in the amount of Phishing messages that customers are exposed to.

Windows Live Hotmail (Hotmail) is a leading provider of free Internet based communications services. Hotmail provides e-mail to more than 280 million e-mail accounts in over 220 countries, and processes over 4.5 billion e-mail messages per day. Of those messages, an estimated 85–90 percent are classified as spam. For years, Hotmail, together with other ISPs and receiving networks, tried to control the growing problem of unsolicited e-mail by installing a variety of anti-spam solutions. One such solution that has had a significant impact in fighting spam is the integration of Sender ID into Hotmail's overall anti-spam infrastructure. Sender ID results and its integration with sender reputation increased the amount of identified spoofed e-mail and spam and significantly reduced false-positives for legitimate senders. The result is that legitimate e-mail from domains that support Sender ID receive enhanced deliverability, while unwanted spoofed or phishing e-mail from spammers doesn't. User confidence, satisfaction, and productivity increases as users spend less time managing unwanted e-mail.



Situation

Windows Live Hotmail is a free e-mail service that provides users with up to 2 gigabytes of mail item storage as well as contact and calendar management capabilities. Due to Hotmail's popularity and large user base, it has become a prime target for spam. The Hotmail network receives upwards of 4.5 billion messages a day. Of these, more than 85% is classified as spam, unsolicited, and not relevant to the intended recipient.

Faced with an escalating level of deceptive tactics by spammers, and without the protection provided by innovative anti-spam countermeasures, Hotmail users would be exposed to an unacceptable level of Spam-In-The-Inbox (SITI), negatively affecting user experience as well as Hotmail's quality of service and system infrastructure.

Solution

Recognizing the inherent limitations to traditional content-filtering mechanisms, Hotmail recognized the importance of building upon the success of Microsoft patented SmartScreen™ technology and unique customer feedback loop, with investments in reputation and authentication technology. In early 2003, Microsoft researchers began to investigate the emerging threats and increasing levels of e-mail being forged worldwide, resulting in the development of a proposed specification known as Microsoft Caller ID for e-mail. In early to mid-2004, Microsoft merged Caller ID with SPF to create the Sender ID Framework, providing greater interoperability and flexibility for receiving networks. In late 2004, Hotmail began production testing of the Sender ID Framework (SIDF), with a system-wide rollout in January 2005. This implementation integrated the SIDF result into the combined spam confidence level (SCL) score generated by Microsoft SmartScreen™ anti-spam technology, and provided user warnings on e-mail messages confirmed to be forged. These forged messages were placed in the user's Junk E-mail folder, with any links and images disabled.¹¹

Microsoft's implementation of Sender ID uses the sender's domain in the Purported Responsible Address (PRA) as the domain to test because this is the domain spoofed by

¹¹ Microsoft SmartScreen is patented anti-spam technology used by Windows Live Hotmail and other Microsoft products that includes direct feedback from Hotmail users plus hundreds of thousands of machine learning algorithms.

4.5 Billion

The approximate number of e-mail messages sent to Hotmail each day

3.8 Billion

The average number of daily e-mail messages sent to Hotmail that are estimated to be spam each day

3 Billion

The average number of daily e-mail messages received without an SPF record to authenticate the sender

300 Million

The average number of daily e-mail messages received with an SPF record to authenticate the sender

275 Million

The average number of daily e-mail messages that pass the Sender ID, resulting in improved deliverability

25 Million

The average number of daily e-mail messages that fail the Sender ID and are junked or deleted, offering enhanced protection for Hotmail users

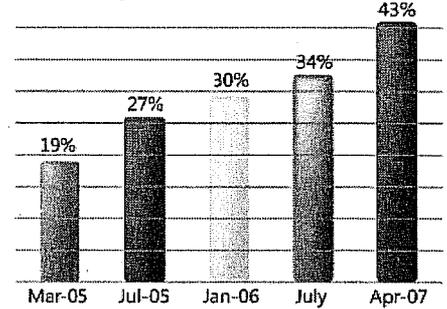
unscrupulous spammers. The address given in the SMTP *Mail From* command can also be checked, but it is not the address seen by the end user and is not the address that is used if the e-mail is replied to.

Today, more than 43 percent of all legitimate e-mail sent into Hotmail is SIDF-compliant, meaning that the sending domains have a valid SPF record published for the server sending the mail into Hotmail. This rapid adoption is a strong achievement because it represents the majority of large-volume senders. For example, 85 percent of all e-mail marketers and advertisers have

adopted Sender ID for their outbound marketing e-mail¹²; when coupled with sender reputation data, this helps their e-mail reach the intended recipients, and keeps spammers from spoofing their domain.

The table below shows some of the top domains and the amount of e-mail that Sender ID identifies as spam or legitimate e-mail.

SIDF Adoption - Volume of "Good" mail



Domain	Pass delivery rate	Non-Pass delivery rate	Spoof rate
ebay.com	100.00%	0.00%	47.90%
freelotto.com	18.26%	0.00%	2.63%
gmx.net	100.00%	4.00%	58.10%
hotmail.com	96.99%	14.83%	39.87%
msn.com	100.00%	6.20%	54.50%

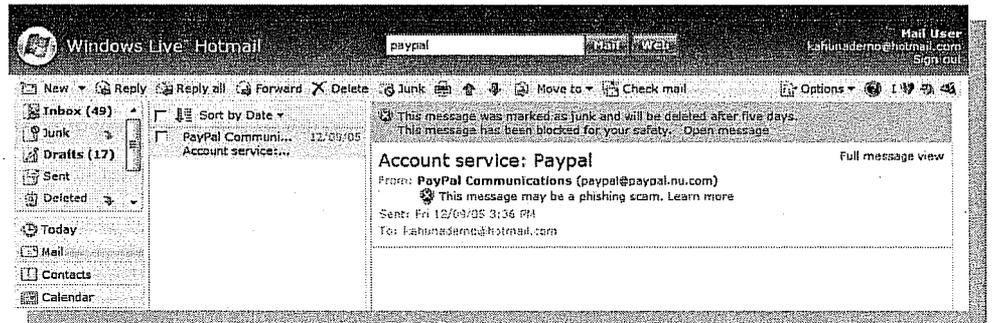
Domain	Pass delivery rate	Non-pass delivery rate	Spoof rate
paypal.com	100.00%	0.00%	32.40%
prodigy.net.mx	100.00%	4.80%	50.40%
telefonica.net	100.00%	2.00%	75.00%
verizon.com	100.00%	0.80%	66.80%
yahoo.co.jp	73.40%	0.00%	98.40%

As shown by the data above, if a message passes Sender ID, it is very likely to be delivered to the user's mailbox for these high-volume domains. Conversely, a message from these domains that does not pass the Sender ID test is very unlikely to be delivered and is likely

¹² April 2nd, 2006 DNS query and ESPC survey

to be spoofed e-mail. Sender ID not only helps protect these domains' brands, but also reduces the number of unwanted e-mail messages that Hotmail users receive.

Windows Live Hotmail applies the results of the Sender ID test to the results of other anti-spam technology to determine the overall reputation of the sender. For example, if contoso.com has a positive reputation, and a message is received from contoso.com and passes the Sender ID test, the message is not likely to be spam. Conversely, if nwtraders.com has a positive reputation, and a message is received from nwtraders.com but fails the Sender ID test, the message is likely to be spam. Messages like this may be deleted, blocked, or delivered to the Hotmail user's Junk E-mail folder, based on the SPF record syntax as defined by the domain holder.



When a user opens an e-mail message in the Junk E-mail folder, a message is displayed that warns the user that the sender's identity could not be verified. The Hotmail user is then given the opportunity to report the e-mail as "junk" the message.

When a user chooses to report the message as "junk" the senders IP address is entered into the Hotmail Junk e-Mail Reporting system (JMR), and subsequent messages classified from other Hotmail users are also tracked in JMR, forming one portion of the senders overall reputation.

Windows Live Hotmail uses Sender ID combined with SmartScreen and user based reputation data to determine if a message is spam. As the table above shows, just because a sender's domain passes the SIDF test does not mean that the message is not spam, nor is it a guarantee that the e-mail will be delivered into the Inbox. It is the combination of Sender ID, reputation data, and the other anti-spam technologies that provides Hotmail with its comprehensive anti-spam and anti-phishing solution. The following table provides

a *hypothetical* example how SIDF can add to results from content filtering and other reputation data, to produce the combined score.

E-mail Scoring with SIDF, Reputation & Content Filtering Data		
Test	Test result	Weight
Sender ID	Pass	+0.5
Reputation	Favorable	+3.0
Content Filter	Probable spam	-1.0
Net e-mail score		+2.5

In this example, you can see that passing the Sender ID test has more weight than having e-mail content identified as potential spam by the content filter. This is important for organizations that send content that is similar to content often used in spam. For example, financial institutions may send newsletters with legitimate information that the recipients want or need. Content filtering alone may mark this type of e-mail as spam. However, since it came from a source that passed Sender ID and has a good reputation; the message is delivered to the recipient. Applying the positive reputation to such mailers can reduce the likelihood of e-mail being marked as spam or even deleted.

Benefits

SIDF deployment at Windows Live Hotmail has demonstrated that e-mail authentication plays a significant role in helping protect customers from spam while also helping legitimate e-mail get delivered.

As the number of compliant domains increases, the number of spoofed messages delivered to Hotmail inboxes decreases, allowing Hotmail to place added weight on SIDF compliant e-mail.

High-volume mailers that implement and maintain an SPF and send legitimate e-mail have a positive reputation and enjoy an 87 percent reduction in false-positives and 85 percent fewer false negatives as compared to a random sampling of non-SIDF-compliant mail.

Sender ID and Reputation Data Working Together

While Sender ID authentication is crucial, reputation is no less important. The combination of authentication and reputation offers significant business value by reducing the volume of messages that the majority of users consider spam. Receiving networks like yours can employ similar reputation approaches with your current anti-spam solution or with off-the-shelf anti-spam solutions that include SIDF working with reputation data.

Different Domains, Different Reputations

Another lesson learned from the Windows Live Hotmail deployment is that senders should segment their outbound mail to maximize deliverability. For instance, separate subdomains and IP addresses can apply to different categories, such as transactional mail (statements, shopping cart and shipping confirmations, and so on), marketing mail, and corporate communication.

The rationale for segmentation is simple: each of these domains can earn a different reputation. A user might want a monthly statement from his or her bank, but not bulk mailings from the bank's marketing department. As a result, the monthly statements could receive a high reputation, while the bulk mailings score low.

A variation on this lesson involves third-party providers working on behalf of another organization. If an organization relies on an external firm to send messages in its name — for example, its weekly electronic sales announcements — the organization from which the message is sent should list its vendor's outgoing mail servers in the SPF it publishes to its DNS zone file. If it does not, the SIDF check will return a "fail" verdict because it cannot find the vendor's servers in the purported organization's SPF file.

Focusing the Benefits

As the first worldwide implementation of Sender ID, Windows Live Hotmail has realized improved spam detection in conjunction with other anti-spam technologies and reputation data. Measurable results have included a reduction in "Spam In The Inbox" (SITI) as well as improved deliverability, reliability, and confidence in legitimate e-mail. Businesses who have adopted outbound authentication have also benefited from enhanced customer communications and brand protection, thereby realizing a competitive advantage.

For More Information

Microsoft Sender ID Information
www.microsoft.com/senderid

Implementation Tips for the Sender ID Framework—Creating Your SPF Record
<http://www.microsoft.com/downloads/details.aspx?familyid=B7CE1CAC-D884-4216-82FE-379F875663FF&displaylang=en>

Sender ID SPF Record Wizard
www.microsoft.com/senderid/wizard

Anti-Phishing Working Group (APWG)
www.antiphishing.org/

Authentication & Online Trust Alliance (AOTA) www.aotalliance.org/

Direct Marketing Association (DMA)
www.the-dma.org

E-mail Sender Provider Coalition (ESPC)
www.espcoalition.org/

Messaging Anti-Abuse Working Group (MAAWG) www.maawg.org

For more information
senderid@microsoft.com

Acknowledgements

Thanks to the following people for their substantial contribution to the writing and development of this paper:

Contributors

Brian Holdsworth

Craig Spiegle

Harry Katz

John Scarrow

Kelly Sieben

Peter Ollodart

Ryan Colvin

Steve Young

Special thanks to the teams at Wadeware, GoDaddy.com, eBay and PayPal for their input.

CONCLUSION

There is no silver bullet or single solution to stop or combat spam, phishing and online deception — it takes a combination of innovative technologies, user education, effective and strong enforcement, and collaboration with industry, business, and governments. The Sender ID Framework is an example of a simple yet innovative, cost effective and easy-to-deploy solution, developed in collaboration with organizations throughout the world.

SIDF has two parts: a DNS record that identifies SMTP servers authorized to send e-mail, and an authentication mechanism that uses that DNS record to verify that inbound e-mail is from an authorized server. Together with reputation data, SIDF plays an important role in the fight against spam by authenticating the sender and applying reputation data. This enables valid messages that might otherwise be identified as bad to be delivered to the Inbox, and conversely keeps messages that are spoofed and do not pass authentication out of the Inbox. In doing so, SIDF helps protect users from unwanted e-mail, delivers the e-mail that users want, and helps keep company brands protected from bad messages that may hurt their reputation and expose their customers to risk.

Sophisticated spammers recognize that domains that have implemented SIDF are highly resistant to spoofing and phishing attacks and are not worth their time. As adoption of SIDF has increased, we have witnessed spammers moving to softer targets, providing early adopters of SIDF a competitive advantage. Large international banks, online retailers such as eBay and PayPal, and online service providers such as GoDaddy.com and Windows Live Hotmail have all implemented Sender-ID and have benefited from the protection it provides their brand and their customers.