E-Mail Authentication Is Here, but Has It Arrived Yet?

George Lawton

The explosion in spam and phishing attacks has cost victims time and money. Companies have had to buy and implement protective technology, and individuals either have had to waste time going through in-boxes packed with unsolicited messages or have lost money to phishing scams originating with bogus e-mail.

The attacks have also dented public confidence in e-mail. A recent survey by Forrester Research, a market analysis firm, found that 20 percent of respondents refuse to open e-mail or attachments, even those that look legitimate, because they lack confidence in the system, explained company analyst Jonathan Penn.

Attempts to solve the problem by the multimillion-dollar mail-filtering industry or via legislation have failed. In light of this, companies such as Cisco Systems, Microsoft, and Yahoo are pursuing e-mail authentication, designed to determine whether mail has come from a reliable source and is thus legitimate.

Vendors have coalesced around two primary authentication proposals: the Sender ID Framework (SIDF) and Domain Keys Identified Mail (DKIM). Some businesses are already adopting one or both approaches in trials or full deployments in their e-mail systems. However, both technologies have limitations.

Industry observers say that authentication by itself is not likely to seriously curtail spam and phishing.

Thus, industry experts say, authentication will serve as only one tool to combat spam and phishing.

ADDRESSING E-MAIL PROBLEMS

Since January 2004, e-mail security vendor MX Logic has found that each month, between 60 to 93 percent of all mail has been spam. There has also been a high volume of phishing. In these attacks, phishers send e-mail messages that link to counterfeit bank, credit card, or e-commerce Web sites and trick victims into entering Social Security and account numbers, passwords, and other sensitive data, which the attackers can then exploit.

Phishers can also send mail with attachments or Web site links that ultimately infect visitors with viruses that turn their computers into zombies that hackers can use for large-scale attacks.

Companies have used several approaches to identify spam, including the most popular, content filtering, which examines messages for keywords commonly found in unsolicited mail.

Other antispam and antiphishing approaches include creating blacklists of known spammers; making senders retype a keyword, requiring a human response that hampers spammers' ability to automatically send bulk mail; and using reputation services that analyze e-mail senders' behavior and history to determine the possibility that their messages are spam or phishing related.

However, these antispam and antiphishing approaches don't combat e-mail address spoofing and other problems related to the Simple Mail Transfer Protocol's inability to adequately authenticate a sender's identity.

Spammers and phishers attempt to generate reputable-looking return addresses using methods such as hacking into unprotected e-mail servers or falsifying names and addresses in the mail sender field.

By detecting spoofed addresses, authentication could enable companies to identify spam or phishing attempts. The organizations could then set policies for their e-mail systems that either inform users about the messages and let them decide what to do, block the transmissions altogether, or even bounce them back to the senders. The latter policy could overload or at least reduce the capacity of attackers' systems and thus encourage them to stop sending spam to the organizations.

SENDER ID FRAMEWORK

SIDF combines Microsoft's Caller ID technology with the Sender Policy Framework (SPF)—formerly called Sender Permitted Form—developed by Meng Weng Wong, founder of e-mail service provider Pobox.com, and backed by companies such as America Online, EarthLink, and Google.

With SIDF, Internet service providers (ISPs) and businesses register their mail server's IP addresses with the Internet Corporation for Assigned Names and Numbers, which manages and coordinates the Domain Name System. The addresses are stored in DNS databases.

Software on a message recipient's client or e-mail server reads the sender's return IP address. It then queries the DNS databases and compares the return address to those the sender has listed. If the address is one of those listed, the system assumes the return address hasn't been spoofed and the message is legitimate. Otherwise, the system labels the message as being possibly related to spam or phishing.

With SPF, messages have the return address in the header. Caller ID e-mail has the IP address in the message's body. Because SIDF combines SPF and Caller ID, it checks both the header and message body.

The Internet Engineering Task Force (IETF) was considering an SIDF standard last year but terminated the process after some participants expressed concern that Microsoft was trying to patent parts of the technology, said John Levine, chair of the Internet Research Task Force's Anti-Spam Research Group.

This shouldn't have been a concern because even if Microsoft received such patents, it would let any company use SIDF royalty-free, said Samantha McManus, business strategy manager for Microsoft's Technology Care and Safety Group.

Concerns

About 57.4 million domains are currently registered, based on statistics compiled by domain-registration utility vendor Name Intelligence. However, only about 2 million domains currently publish IP-address information for SIDF systems, noted McManus.

SIDF has not been used more widely because some ISPs are concerned the technology will break their systems' e-mail-forwarding capabilities, explained Forrester's Penn. When a system forwards mail, the message looks fraudulent to SIDF applications because it didn't come from the original sender's IP address.

According to critics, users should be careful with SIDF because it is not an accepted standard, has not been tested thoroughly for security, and has shortcomings. For example, with complex systems, such as those with multiple mail servers in different locations or those that are outsourced, SIDF can be too complicated to scale well or function efficiently, noted Ray Everett-Church, cofounder of the Coalition against Unsolicited Commercial E-mail.

A challenge for many organizations—especially large ones with scattered facilities or small ones with little expertise or staff—is conducting thorough inventories of their IP addresses for subsequent registration and then adequately maintaining their address lists, noted Penn. The failure to do so can cause SIDF systems to function improperly by, for example, yielding false spam findings.

Authentication tries to determine whether e-mail has come from a reliable source.

An MX Logic study of 15.8 million e-mail messages in August 2005 found that 8.7 percent had SIDF records, and of those, 83 percent were from spammers' domains, noted Scott Chasin, the company's chief technology officer.

In these cases, spammers register the domains they use to send unsolicited mail, just like any other SIDF participant. This is a concern because if spammers have SIDF records, their messages look legitimate. This is why spammers are among the technology's biggest adopters, Chasin explained.

DKIM

DKIM is the leading cryptographic-based e-mail authentication approach. It combines two digital-signature-based technologies: the Yahoo-supported DomainKeys and the Cisco-supported Identified Internet Mail.

Cisco Distinguished Engineer Jim Fenton said his company and Yahoo started working on their separate approaches about two years ago and joined forces early this year to develop a viable proposal for submission to the IETF. The companies are now forming an IETF working group to study the technology for possible standardization.

How DKIM works

As Figure 1 shows, DKIM attaches encrypted digital signatures to outbound mail headers so that recipients' servers can verify the incoming mail's origin.

The technology uses public-key cryptography to create a digital signature. In a public-key system, a single algorithm creates a public and private key pair. Users send their public key to individuals who might want to send encrypted messages. Only the user's corresponding private key can decrypt them.

If DKIM users receive a message with a signature encrypted by a public key that their private key won't decrypt, the system identifies it as being related to spam or phishing. Otherwise, the system considers senders to be who they say they are and judges the message to be legitimate.

To make DKIM work, both senders and receivers must implement e-mail server software upgrades. Unlike SIDF, DKIM does not break e-mail forwarding.

Concerns

Some mail programs, such as services that send multiple customized versions of mail to recipients, can modify message headers and garble the information that DKIM uses to authenticate senders, said Cisco's Fenton.

Also, DKIM adds about 15 percent overhead in cryptographic processing and requires major upgrades of software, as well as perhaps hardware to handle the increased processing load, noted Forrester's Penn.

AUTHENTIC CHALLENGES

E-mail authentication won't eliminate spam by itself, and this has disappointed some industry observers, said MX Logic's Chasin. Proponents say they already knew this but still value the contributions the technology can make.

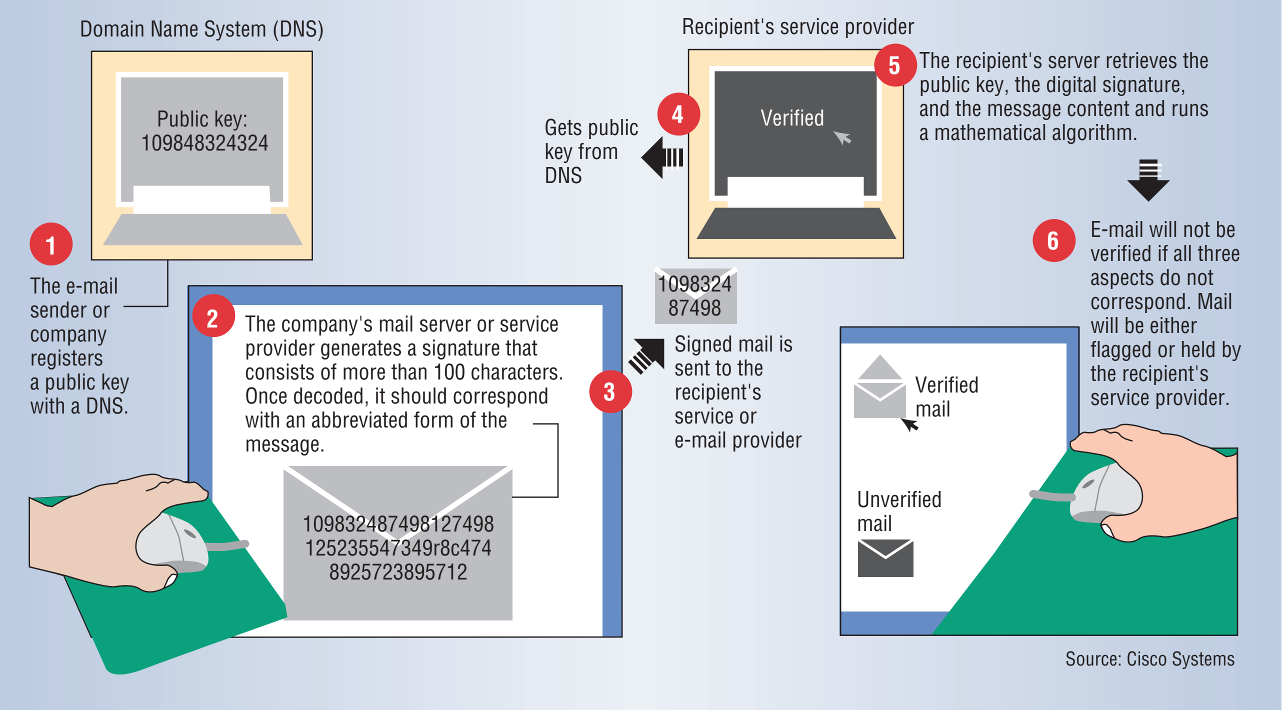
Domain Name System (DNS)
Public key: 109848324324
1
The e-mail sender or company registers a public key with a DNS.
2
The company's mail server or service provider generates a signature that consists of more than 100 characters. Once decoded, it should correspond with an abbreviated form of the message.
109832487498127498125235547349r8c4748925723895712
3
Signed mail is sent to the recipient's service or e-mail provider
109832487498
Recipient's service provider
4
Gets public key from DNS
Verified
5
The recipient's server retrieves the public key, the digital signature, and the message content and runs a mathematical algorithm.
6
E-mail will not be verified if all three aspects do not correspond. Mail will be either flagged or held by the recipient's service provider.
Verified mail
Unverified mail

Source: Cisco Systems

Figure 1. Domain Keys Identified Mail is one of two major e-mail authentication approaches.

According to Chasin, political infighting among vendors and other technology companies—such as the controversy over Microsoft's attempted SIDF-related patents—has delayed progress in adopting standards and implementing e-mail authentication.

Companies face challenges in determining how to incorporate authentication into their current infrastructures. For example, with SIDF and DKIM, users must develop policies for handling messages identified as spam- or phishing-related.

Neither technology keeps spammers from taking over computers via viruses and using them as zombies to deliver mail. By using other computers and their legitimate IP addresses to send their mail, spammers don't have to spoof return addresses, and their messages pass authentication tests. According to Chasin, a recent study found that 48 percent of spam came from zombies. Thus, e-mail authentication won't eliminate a large portion of spam.

In addition, the ability of spammers and phishers to use authentication technology to make their own messages appear to be legitimate limits the approach's effectiveness.

According to Forrester's Penn, SIDF is simpler and will be widely deployed before DKIM, which could require five years to gain traction in the marketplace.

However, he and many other industry observers say that because SIDF and DKIM take different approaches and each has shortcomings, they will end up complementing, rather than competing with, each other.

Meanwhile, both will work with existing technologies such as filtering, noted Eric Allman, founder and chief technology officer at Sendmail, an e-mail security and infrastructure vendor.

At some point, Pobox.com's Wong stated, researchers must enhance reputation systems. He said they also must work on accreditation management systems, which let e-mail users post a bond or pay a fee to be recognized as reliable senders as long as they address subsequent complaints against them.

The ePrivacy Group has proposed the Trusted E-mail Open Standard, which has not gained as much support as SIDF and DKIM. Everett-Church said TEOS goes further than current authentication schemes by building in layers of signatures through which senders can make assertions about the types of messages they're sending and recipients can decide more accurately how to handle them.

However, he added, authentication will help control spam only when authorities use it to enforce compliance with antispam laws. "We have to hold bad actors accountable," he explained, "and enforce existing consumer protection laws." ■

George Lawton is a freelance technology writer based in Brisbane, California. Contact him at glawton@glawton.com.

Editor: Lee Garber, *Computer*, l.garber@computer.org