



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Slippery Slope or Terra Firma? Current and Future Anti-Spam Measures

A short time ago, it would have been unthinkable that electronic mail or "email" would have grown as popular as it has over the last few years, but many corporations now consider it to be their primary means of communication. Consumers feel the same way with the advent of email enabled PDA's, phones, and pagers as well as Internet mail, which is accessible just about anywhere from the home, office, library, or even Internet cafes. With this growth in popularity, marketers have realized the potential of reaching million...

Copyright SANS Institute
Author Retains Full Rights



AD

**Slippery Slope or Terra Firma?
Current and Future Anti-Spam Measures**

**Charlene LeBlanc
GIAC Security Essentials Certification
Practical Assignment
Version 1.4b Option 1
June 20, 2003**

© SANS Institute 2003. Author retains full rights

Abstract

A short time ago, it would have been unthinkable that electronic mail or “email” would have grown as popular as it has over the last few years, but many corporations now consider it to be their primary means of communication. Consumers feel the same way with the advent of email enabled PDA’s, phones, and pagers as well as Internet mail, which is accessible just about anywhere from the home, office, library, or even Internet cafés. With this growth in popularity, marketers have realized the potential of reaching millions of consumers at a fraction of the cost for printed material, radio, or television. Thus, this new kind of marketing has been given the moniker of the famous canned ham SPAM® for its being readily available all over the world and considered to be made up of just about every part of a pig. Since most of “spam” is unwanted mail, many technology companies, network administrators, and politicians have taken up the fight against its widespread use and abuse.

This paper will help to explain current methods that are used to send out spam, combat spam, and legislation in place to hold abusers accountable and will also take a look at what might be on the technology horizon with more robust filtering methods and perhaps a better SMTP standard.

Introduction

Spam mail or unsolicited commercial email (UCE) has been growing at an exponential rate over the last couple of years. It takes next to nothing to compose an email and send it out to any address that you can add into the message. Consumers have now grown tired of receiving email for products, services, and even pornographic material that they did not ask for, much like printed advertising reaches them through the United States Postal Service®. Email administrators grow weary of spam as well because of the amount that comes into their mail servers unnecessarily expanding the size of their databases, degrading performance, and in some cases containing computer viruses. If it takes only seconds to throw out the printed advertising when it reaches your home or office or delete email permanently off of a server, why then are consumers and administrators so upset? The answer is simple. Most people and corporations consider email addresses to be private and spam to be an invasion of privacy. They also see time wasted in spending time every day in deleting spam.

How can spam’s growth be thwarted? The answer is not always an easy one to answer. One person’s junk is another’s treasure. Methods to combat spam wind up being a slippery slope because the problem is never completely resolved. Spam is sent out by a variety of methods and there are many that fight its spread. A few states have enacted laws but, this problem might land on terra firma if better controls are put in place.

Slippery Slope

Current Assault

“The history of the advertising industry shows that the lower the cost of the direct marketing technique, the greater the risk of abuse” according to a study by the European Union. (<http://www.netinsites.com/article3.cfm?ArticleID=97>)

The abuse has grown to epic proportions. Hotmail users receive more than one billion pieces of spam mail each day and Brightmail, an anti-spam filtering company, reported over 4.8 million spam attacks over the course of one month. Most of those same attacks may have contained thousands of the same email messages. (<http://www.netinsites.com/article3.cfm?ArticleID=97>) The average amount of spam messages one email account receives during a year is 2,200 according to Jupiter Research and Brightmail estimates that almost forty percent of all Internet mail traffic is unwanted. (Fordahl)

Due to the spam mail onslaught, politicians are now scrambling to bring bills before state and federal legislatures in order to appear to constituents as if they are taking a hard line against this burgeoning issue. Even with legislation and stiff penalties, can these new laws really be enforced when spammers, those that send the unsolicited email, can hide behind spoofed IP addresses, anonymous mailers, and relay thousands of messages off of open relay mail servers?

Computer programmers realized the need for a solution to this issue long before politicians. They began to develop different methods of stopping spam from hitting customers. To date, methods designed have run the gamut from scanning email messages for specific keywords to looking at the number of skin-tone pixels in graphic attachments to determine if the image is pornographic or not. While these methods provide some sort of defense, they don't catch everything.

RFC Standards

One of the real reasons why spam mail has taken off as much as it has could be traced back to a now twenty year old standard. The first statement of the SMTP standard, RFC 821 and written in 1982 and adhered to by many email server applications, states that, “The objective of Simple Mail Transfer Protocol (SMTP) is to transfer mail reliably and efficiently.” (<http://www.faqs.org/rfcs/rfc821.html>) But, it does not discuss any security measures for authentication or identifying the sender. There is a VRFY or verify command that can be used on the part of the sender to identify if the recipient's address does exist on the server that it is sending to but, this does not actually stop the message from being received as this command is usually used via a Telnet session to the receiving server via port 25 and is not in the normal email transaction that takes place.

Below is an example of a normal SMTP transmission (S: denotes Sender R: denotes Recipient):

Receiving Server: 220 mail.someotherdomain.com ready

S: HELO mail.somedomain.com
R: 250 OK

S: MAIL FROM:<Red@somedomain.com>
R: 250 OK

S: RCPT TO:<Blue@someotherdomain.com>
R: 250 OK

S: DATA
R: 354 Start mail input; end with <CRLF>.<CRLF>

S: This is a test.
S: <CRLF>.<CRLF>

R: 250 OK
S: QUIT
R: 221 Closing connection

If a user did not exist at that domain, the receiving server responds back with a standard message that the user is not located in that domain or is not known. Depending on the settings of the receiving server, it may also forward the email to another server within a similar or sub-domain or may forward the message to a postmaster account so that an administrator of that domain can possibly ensure it's delivery to the intended recipient.

Relaying an email messages is similar to forwarding but, the messages are not sent to a sub-domain but to a completely different domain altogether. When sending a message to an email server that will relay messages, the message is accepted and instead of the original sender becoming the source sender of the email, it is replaced with the relaying servers' information. The process continues until the recipient address is found residing on the destination server and the message is delivered. If the domain is correct but, the recipient address is incorrect, the message is delivered to a postmaster account on that server.

Here is the process for relaying as defined by RFC 821:

Conceptually the elements of the forward-path are moved to the reverse-path as the message is relayed from one server-SMTP to another. The reverse-path is a reverse source route, (i.e., a source route from the current location of the message to the originator of the message). When a server-SMTP deletes its identifier from the forward-path and inserts it into the reverse-path, it must use the name it is known by in the environment it is sending into, not the environment the mail came from, in case the server-SMTP is known by different names in

different environments.

Using source routing the receiver-SMTP receives mail to be relayed to another server-SMTP. The receiver-SMTP may accept or reject the task of relaying the mail in the same way it accepts or rejects mail for a local user. The receiver-SMTP transforms the command arguments by moving its own identifier from the forward-path to the beginning of the reverse-path. The receiver-SMTP then becomes a sender-SMTP, establishes a transmission channel to the next SMTP in the forward-path, and sends it the mail.

The first host in the reverse-path should be the host sending the SMTP commands, and the first host in the forward-path should be the host receiving the SMTP commands.

(<http://www.faqs.org/rfcs/rfc821.html>)

There are some measures in place to prevent unauthorized relaying with some mail server applications. These measures involve specifying an IP address of the domains that are unauthorized to relay and their messages rejected. Many ISPs are non-restrictive in that they will allow just about any transaction to be relayed because they may host many domains. These servers become hapless victims and are turned into workhorses by a spammer. Email administrators noticing the increased email traffic from the larger ISPs are hesitant to block any mail from these domains because of the possibility of blocking legitimate mail from being delivered to their customers.

Overall, there are no measures in the standard that state that a mail message has to only come from a particular entity but, implies that mail can be received by just about anyone given that the transmission follows the basic transmitting rules and syntax. In April, 2001, RFC 2821 came out suggesting that the use of the VRFY command not be implemented if so desired and if a receiving mail server was given a VRFY command, it could echo back either 500 Command not recognized or 502 Command not implemented. (<http://www.faqs.org/rfcs/rfc2821.html>) This would suggest that this command was being exploited by spammers hoping to harvest email addresses.

Spammer Tactics

Because the standards for email transport and delivery are somewhat wide open, there are a variety of methods that spammers will use to get their spam out one way or another and make it hard for anti-spamming methods to be affective. There are also many methods of harvesting email addresses so that spam is more targeted. The following are brief descriptions of some of these methods.

IP Address Spoofing

Spammers will “spoof” or forge the IP address if their IP address is part of known offender listings. By spoofing the IP address, messages have a better chance of getting past a filtering server that users RBL’s or “Real-time Black hole Listings. The spoofed

address will often be that of a well-known ISP or will be a non-routable Internet address such as 10.10.10.1, etc. Another reason for spoofing the IP address is due to a recipient server having reverse DNS enabled. If the spoofed address is a well-known address that has reverse DNS entries known on name servers, then the mail will be received by the recipient server despite the email address being forged.

False Display Address or From Field Entries

Messages will appear as if they are from a reputable company because the From field is filled out as say support@Microsoft.com. A closer look at the email header may reveal that the message actually came from some lesser known company or ISP instead of really originating from Microsoft or another company. Also, some from fields are filled out with single names so that spammers hope to entice their victims by possibly using a name that the recipient may know and then they will open up the email message.

Phony Subject Lines

Just as hackers will change the subject lines of messages in order to entice the recipient to open up their message and then click on their attached file that is a virus, spammers will do the same thing. In the past, spammers would use common words in their subject line such as “free gift” but, are now resorting to variations like fr33 g1ft to get their mail past spam filters looking for specific keywords.

Image Only Email

Image only email refers to spammers that will send just an image in their messages that often contains pornographic material but, will then have messages superimposed on the image with a URL address to a porn site, etc. This type of email, when opened, can cause a serious threat to a child’s welfare if opened by a child.

Anonymous Mailers

As the term applies, anonymous mailers will allow an email sender to appear anonymous. The process was originally a good one if users posting to news groups did not want their identity revealed. It is now an exploited method because spammers found that this could be a great place to send their spam. These mailers work in such a way that once the message is received by the intended recipient, the header information is virtually untraceable. To send a message through these mailers, you compose a message and address it to the mailer instead of to your intended recipient, such as to anonymous_remailer@mail.com. Next, special characters are added to the message body usually in the format of two colons (: :). The next line is usually where text is added that will request the message be sent to a different address: Request-Remailing-To: somebody@anotherdomain.com. A blank line is next followed by the body of your message. The anonymous remailer companies then have special software to search for the “Request-Remailing-To:” field and will then change the recipient or To: field to the intended recipient rather than to the anonymous_remailer@mail.com address. (<http://www.strassmann.com/pubs/anon-remail.html>)

Open Relays

RFC 821 describes relaying as being able to allow a mail server to send mail to one or more mail servers before eventually reaching the desired recipient. If a spammer uses a relay server, there is possible means with which to trace back to the original sender but, a professional spammer would not use just any relay server for this reason. He will use an “open” relay server as these types of servers can hide their identity. The spammer has basically hijacked the server into becoming a workhorse by means of masquerading as a legitimate client of the ISP to which the open relay server belongs. The hijacking occurs when a user sends his user name and password via either the POP3 or SMTP commands to log into their account. The spammer then steals this information and thus the masquerade begins.

Email Address Harvesting

Spiders

Spiders are robots (programs) that go out and scan all web pages of a particular site for email addresses. Since almost all web pages are composed by HTML or Hypertext Markup Language, the common search criteria is to look for a “mailto” tag. Once found, the entry after the tag is then added to a mailing list. These same robots will also search web pages for guestbook signings and message boards in which a person may post their email address and then it is posted on the web page. Popular email spider programs are EmailSiphon and Earthonline Nitro. The following link has more information on email harvesting programs: <http://www.geocities.com/bulksolutions05/>. (http://www.internet-tips.net/Email/SPAM_spiders.htm)

Spam Replies

Spam replies refer to the option in an email message to “opt-out” or unsubscribe from a message by clicking on a link or replying to the email with a subject line of “Remove.” Instead of being removed by the spammer, you are often considered “a live one” in that the spammer has tricked you into confirming that your email address is active and that your message was viewed. The email address is then put on a mailing list and the list may be sold to other marketers over and over again. (Karp, Jack. “Protecting Yourself From Spam.” Techtv.com.

<http://www.techtv.com/cybercrime/privacy/story/0,23008,3339325,00.html>)

Signing Up For Free Stuff

If you notice something on the Internet that you would like to subscribe to because it is free, such as a newsletter, you probably would subscribe to it. Why not it’s free, right? Wrong. Your email address would now be put on a list with the company that you subscribed to. In turn, they will share with other partner companies or sell the listings to make money. (http://www.internet-tips.net/Email/SPAM_how.htm)

Internet Yellow and White Pages

Internet Yellow and White Pages are listings for Internet users to share specific information about themselves, such as business or personal information. Spammers

will search these pages for listings that contain email addresses and will add any that are found to their address lists.

Email Servers

Some email servers are still allowing the VRFY command and also the SMTP standard EXPN or Expand command to expand the membership of the list which will reveal all of the members' email addresses. This is a well known tactic and spammers will use this tactic to try and reap addresses by pounding away at known servers with this command enabled and will check it frequently for any changes. (http://www.internet-tips.net/Email/SPAM_how.htm)

AOL Profiles

Users on the Internet that use AOL or America Online, one of the largest ISPs in the world, will often unknowingly have their email addresses in their user profile. These profiles are able to be searched for email addresses. (http://www.internet-tips.net/Email/SPAM_how.htm)

Hacking

"Hacking" is a term used for computer programmers that write code to do malicious activities on a computer for their own gain. In terms of email hacking, this involves similar types of people using any type of malicious code necessary to gain access to database files from various companies so that they can gain thousands of email addresses.

Newsgroups

Many people frequent newsgroups, which are public forums, for placing and finding information to just about anything. Email addresses are posted within these forums in order to allow users to communicate outside of the forum. Spammers will search these listings religiously looking for email addresses or will employ email spiders to do the work for them.

Guessing

If spammers don't have access to many validated email address listings, they may guess at an email address. For example, some email addresses are made up of a person's name, especially if they work for corporations. Spam mail may then be addressed to several different permutations of "John Smith," such as johnsmith@somecompany.com or johnathansmith@somecompany.com. When an email address is not valid, most email server applications will echo back a failure message via a mail daemon running on the server. Spammers will either allow the return message to be received, at which point they will sometimes remove the failed address from their homemade lists, or will disallow any return mail.

DNS Registries

Domain name registries require complete contact information when setting up a domain. Email addresses are often placed in these registries and the lists are freely shared on the Internet and can be found via WHOIS databases.

Spam-Fighting Tactics

Since spammers use many methods, how can consumers and email administrators combat spam? A cottage industry has sprouted up from small companies developing programs at either the client level or at the server level. These anti-spam products were few and far between just a short time ago but, now there are many choices in the marketplace. Because what is one man's treasure is another man's junk means that a human element has come in to play to determine what legitimate email is and what isn't.

With this in mind, the technology companies' products vary widely. Below are a few descriptions of how the most common anti-spam products work.

Keyword Filtering

Some anti-spam filtering programs will filter out mail depending on keyword lists that can either be a set group of words or can be one's that the end-user or administrator can devise. These have some sort of effectiveness but, if certain words, such as those that are anatomically correct, can block legitimate mail to health organizations. Most keyword filtering software can be at the client, server, or DMZ level.

Subject Line Filtering

The subject line of an email message is scrutinized with this type of product and was somewhat effective at one time until spammers grew wise of this practice and started to develop ways in which to thwart a filter. They accomplished this by changing the spelling of words, such as substituting the letter O with a zero. Wildcards can be used but, it is an administrative nightmare to be all inclusive with permutations of properly spelled words. The message bodies are not scanned nor are attachments. Although ineffective, this method is still widely used.

White List Filtering

"White list" filtering is one in which a server maintains a list of authorized senders, usually by IP address, and will reject anything that is not contained on this list. This is a very restrictive type of filtering and could reject legitimate mail from entering an ISP or corporation. Because of its restrictions, this filtering is used more on the consumer level. Outside of the legitimate email being rejected, maintaining the list becomes a headache for an end-user or administrator with web sites and ISP's using more than one email server and perhaps sub-domain. For example, Bellsouth in the Miami region has over 15 email servers that will resolve to mail.mia.bellsouth.net as shown below via an nslookup command:

```
> server mail.mia.bellsouth.net
Default Server: mail.bellsouth.net
Addresses: 205.152.58.4, 205.152.58.5, 205.152.58.6, 205.152.58.7, 205.152.58.8,
205.152.58.129, 205.152.58.130, 205.152.58.131, 205.152.58.132, 205.152.58.133,
205.152.58.134, 205.152.58.135, 205.152.58.136, 205.15 2.58.1, 205.152.58.2,
205.152.58.3
Aliases: mail.mia.bellsouth.net
```

Black List Filtering

“Black list” filtering is somewhat the opposite of “white list” filtering but, with a catch. Black lists are made up by organizations that oppose spammers or sites with open relays and will place known offenders on the list. This method is very effective but, may also be very damaging. If an ISP that does not do business with spammers just happens to handle a huge load of email traffic for a known offender, then this ISP may get put on the list even if they aren’t the direct entity sending the messages. Having this ISP on a black list or what is also known as an RBL (Real-time Black-hole Listing) can block mail directly from this entity or any mail that is relayed to this entity because as stated earlier in RFC 821, the source address is replaced with the relaying server’s information. Many popular anti-spam appliances, such as McAfee’s WebShield series products, use black lists. Some well known organizations that maintain such lists are “MAPS” or Mail Abuse Prevention System (<http://mail-abuse.org/>) and ORDB or Open Relay Data Base (<http://www.ordb.org>).

Legislation

How can offenders be held accountable? Recent anti-spam legislation has only gone into affect in a few states. With the Internet being a global institution, this legislation will only put a bandage on the problem but will not allow it to fully heal. Some state bills that have been passed into law will cite an offender a certain dollar amount per incident only if the abuse was found to be sent from the state in which the law resides or if it was sent to a user within that state. Since spammers make it a point to hide their identity, anti-spam legislation will not do much to resolve the issue as it is very hard to trace spam back to its original source. Spammers may also move from state to state. If a spammer resides in a state with an anti-spam law and indicted by another state, there might be conflicting interpretations and penalties as well as jurisdiction.

Other state laws have only gone so far as to make it a requirement for spam to be labeled as either ADV: for advertisement or ADV:ADULT or ADV:ADLT to warn recipients that the material sent contains pornographic material that is not suited for minors but does not state in the law that the recipient had to have asked for the email to be sent. The lack of specifics only makes the problem worse because spam is still received unsolicited.

Below are summaries of some current state level laws that were passed to fight unsolicited mail in any form, including electronic mail:

Indiana Code

Title 24. Trade Regulations; Consumer Sales and Credit

Article 5. Consumer Sales

Chapter 22. Deceptive Commercial Electronic Mail

Added by House Bill 1083, approved April 17, 2003

A third party’s domain name cannot be used without permission, the subject line must not be false or misleading, and cannot falsify its origination and routing information.

Messages that are unsolicited require a label of ADV: or ADV:ADLT as a prefix in the subject line and have clear opt-out instructions. If the sender knows that the intended recipient resides in Indiana or if address information of the recipient is obtained via a request from the sender to the domain registrant, this law will apply. Damages can be awarded as either total damages incurred provided a plaintiff can prove the actual dollar amount suffered or may be awarded \$500 for each message that violated the code. (<http://www.spamlaws.com/state/summary.html>)

State of California

BUSINESS AND PROFESSIONS CODE SECTION 17530-17539.6

Passed in 1998, UCE requires fully honored opt-out methods and contact information and must have subject-line labeling prefixes of ADV: and ADV:ADLT. An email provider can sue the sender of the UCE if the provider's policies have been violated provided the sender has a notice of the policies. The law is applicable to email sent to a California resident by way of the provider's facilities located in California. A per incident fine of \$500 shall be awarded for damages. (<http://www.spamlaws.com/state/summary.html>)

State of Washington

CHAPTER 19.190 RCW

COMMERCIAL ELECTRONIC MAIL

Washington enacted its first anti-spam law in March 1998 but, quickly amended it in May 1999. The state sees commercial email sent via a third party's domain without consent, a message contains false or missing routing information, and contains a deceptive or false subject line to be illegal. The law only applies if the sender openly knows a recipient is in Washington or if the person or entity that registered the domain name of the recipient's address can confirm by request that the recipient does reside in Washington. Damages are awarded to a recipient of a message in the amount of \$500 or actual damages, whichever is greater. An additional award of damages to an interactive computer that received the violating email can be granted to a plaintiff in the amount of \$1,000 or actual damages. The greater amount is awarded. (<http://www.spamlaws.com/state/summary.html>)

These state laws are very similar in nature regarding what constitutes unsolicited email and on damages and penalties. This would suggest that lawmakers in these states used the Federal Junk Fax Law or the Telephone Consumer Protection Act of 1991 as a guideline in drafting their laws. This federal law also carries a per incident fine of \$500.

No federal law has passed, but many bills have been brought before Congress. To date, there have been 26 bills presented before the Senate and House of Representatives. None of them have passed into federal law. Some of these bills have been around since 1999 – most notably the “CAN-SPAM” act of each year – and have had several revisions because Congress is not in agreement on what should or shouldn't be considered spam or whether or not the spam should be solicited or not.

As of May of 2003, there are seven bills before the 108th Congress of the United States and yes, the CAN-SPAM Act is there as well with its current revisions. Perhaps most of these federal bills have not passed because of the consensus of “free speech” as stated in line 1 under Section 2 (a) of the CAN-SPAM Act as stated below:

SEC. 2. CONGRESSIONAL FINDINGS AND POLICY.

- (a) FINDINGS- The Congress finds the following:
(1) There is a right of free speech on the Internet.

<http://www.spamlaws.com/federal/108s877.html>

The majority of Congress wants to allow the sending of unsolicited email advertisements as long as there are opt-out abilities for the recipient so that they may be removed from an email address listing for a specific product or service. If this type of spamming is permitted, then email server administrators will have to allow this type of email to go through their servers. This puts a tremendous burden on email servers because of the increased workload and inflation of the databases. Left unchecked, email servers' performance could be severely degraded. In addition, a database may rise to the maximum amount of disk space available because of a sudden rise in spam mail. This could render a server useless until an administrator has freed up space or recovered the database(s). The onus is then on the consumer to decide whether or not he wants to receive the email message. The human element will prevail.

New Technology

With tools in place and more states enacting laws against spam, isn't it enough? The tools that consumers and email administrators have to use only go so far. Spammers work out ways to get around these tools on a consistent basis. Legislation is not rising to its potential because of the ease at which spammers can hide. To process a case once it goes to trial will be difficult to prove because of the burden of proof in showing that a spammer willingly sent a message to someone in a state with an anti-spam law. How can a spammer possibly know that someuser@somedomain.com actually resides in that state with the anti-spam law when most people do not post their home address along with their email address? Both current technology and legislation are band-aids on an open wound. Hope is not lost, however, because newer technology and a proposed SMTP standard are on the horizon and could possibly lead to a solution to the problem of spam and adhere to applicable law.

Bayesian Filtering

A new twist on an old theorem shows some real promise. This twist is based on the Bayes Theorem, which is named after Thomas Bayes. According to the site, <http://whatis.techtarget.com>, the theorem has the following principle:

“Based on probability theory, the theorem defines a rule for refining a hypothesis by factoring in additional evidence and background information, and

leads to a number representing the degree of probability that the hypothesis is true.”

The theorem was then adapted for email purposes to look at the probabilities that a message was indeed spam and states the probability for something to indeed be true must add up to a value of 1. A false value must equal 0. A neutral value, .5. Values are then assigned within this framework to specific phrases or words that would not be commonly found in a regular email message. Most words within an advertisement email have certain combinations such as “buy now” or “free shipping.” A higher value of .85 might be given to the phrase of “free shipping” and a lesser value given to the phrase of “buy now.” The reason for the lower value for “buy now” is that there is a probability that it could be a phrase that is used in a regular email in which the sender instructs a stock broker to buy shares of a stock. Other phrases that are common for spam mail may appear in more parts of the message and thus create a possible value of 1 or true. Based on the settings of the software, the message is either quarantined or deleted. Some proponents of the Bayesian method have called it an adaptive process because of its perceived ability to learn over time. (<http://members.aol.com/johnp71/bayes.html>)

White List Verification Filters

As discussed earlier, a white list is a list of authorized senders that a server will allow mail to pass through. A verification process is added on to this idea and is then invoked when the server receives a message from a non-listed sender. Another name for this method is called TDMA or Tagged Message Delivery Agent and is OSI certified. The software process usually resides at the MTA (Message Transfer Agent) and automatically generates and sends a challenge/response message back to the original address. A simple phrase or question is put into the body of the message and requires human intervention in order to act upon the phrase or question. For example, the simple message might be a question of “What is this picture of?” and contains a picture of a puppy. The correct answer would be DOG or PUPPY and if the server receives a reply and the message contains the correct answer, the sender’s address is automatically added to the white list to allow future mail to be accepted. If, after a pre-determined time frame, the challenge/response message is not answered, the sender’s address is not added to the list and email is rejected. (<http://www.nwfusion.com/columnists/2002/0812netbuzz.html>)

There are a few drawbacks to this method despite its promise. One, if a user is supposed to receive a newsletter that they signed up for and the sender has an automated process with which to send out the newsletters, the challenge/response message sent back for verification may be overlooked as it requires someone to act upon the message. Another drawback is that some legitimate marketers that use email as an effective tool may not allow messages to return back from the sent address as the email message is generated by an automated process. Marketing email may also contain a separate address for recipients to reply to such as “for more information, please send an email to feedback@somedomain.com,” which is different from the automated email address of perhaps customersupport@somedomain.com. If a

spammer does act upon the challenge/response message and is now authorized to send mail to that server, then the purpose of denying spam is useless. For administrators, this could be a very useful tool for two reasons. First, it is an automated process of adding users to the white list and second, maintaining a list of authorized senders takes up less management time than a list of blocked senders. (<http://www-106.ibm.com/developerworks/linux/library/l-spamf.html>)

Distributive Adaptive Blacklists

This method allows an MTA to call upon a list of known spammers to determine whether or not a message is spam. It is very similar to regular blacklists but, uses a statistical method for spam determination. Those that are on the list of known spammers do not get there merely by someone reporting the addresses to blacklist companies. Some people who maintain and develop these lists will create “honey pots” or mailbox addresses that will attract spammers in order to tweak their statistical methods and make them more affective. There are at times higher rates of “false positives” meaning legitimate mail may be blocked, but it is recommended that other tools be used in concert with this approach. (<http://www-106.ibm.com/developerworks/linux/library/l-spamf.html>)

Rule-Based Rankings

If a message contains regular expressions that this type of spam filtering software looks for, each expression is given a specific value. Other expressions may subtract from the overall value when found. Exceed the set value and the message is filtered. Software of this type can be customized over time to look for newer phrases that spammers will use. As with other spam filters, there is a level of false positives but, having the ability to customize the filter is an attractive option. (<http://www-106.ibm.com/developerworks/linux/library/l-spamf.html>)

New SMTP Standard

A new standard is being touted by several advocacy groups and corporations, most notably Microsoft and the ePrivacy group. The Trusted Email Open Standard (TEOS) was introduced by these groups on April 30, 2003. As the current SMTP protocol suggests “Simple Mail Transport Protocol” email is quite simple and results in great reliability and speed. With no authentication processes or validation of content, this protocol is now vastly outdated and abused. Simplicity is not what many users and administrators are willing to give up. Re-designing a new protocol from the ground up would be time consuming for both development and global acceptance. The authors of the TEOS protocol recognized this fact and have designed the framework around the original standards. (<http://www.eprivacygroup.net/teos/>)

The following section is a detailed overview of the TEOS standard as developed by the ePrivacy Group and its sponsors of the proposed standard.

The TEOS framework is based on three principals:

1. Best practices
2. Enabling technology
3. Oversight

Best Practices

In order to allow for greater acceptance of this proposed standard, the proposal calls for a breakdown of principle-based best practices. These best practices are broken down as follows:

A	A Minimum Standard for Accountability	Sender Identity (Level I) Optional Assertion of Message Type Minimal cost
B	Bulk Sender Trusted Email Certification Programs	Sender Identity (Level II) Required Assertions - Message Type - Relationship/Permission - Standardized Opt-out Optional Assertions (program dependent)
C	Consumer Oriented Trusted Email Certification Programs	Sender Identity (Level III) Required Assertions - Message Type - Relationship/Permission - Visible Assertions - Secure Seal with One-click Verify - Trusted Opt-out - Privacy Policy Link Dispute Resolution Process Trust Authority Oversight Optional Assertions (program dependent)

<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>, p. 11

Since SMTP does not require senders to fully disclose their origination, email headers are often spoofed – a monumental problem in order to enforce any kind of state or federal legislation against spam mail. Given this, the Basic Standard has different levels to ensure sender identity based on the most common types or originators – Private, Bulk Sender, and Consumer Oriented. Each level has different criteria that specific senders would adhere to. Level A has the least cost factor, both in cost and computation, and still guarantees a sender’s identity by incorporating a level 1 digital certificate that is different from what commercial Certificate Authorities issue. This level will also allow for optional assertions for message type. This is considered to be a Trusted Email Domain Identity (TEDI). The optional assertions or Trusted Email Type Assertions (TETA) would be a three character prefix in an e-mail’s subject line such as FAF for Friends and Family. Some proposed assertions are listed in the graphic below:

Trusted Email Type Assertions	
1. Unsolicited advertisements	(ADV) ¹⁵
2. Adult	(ADT) ¹⁶
3. Permission-based advertisements, offers	(CRM)
4. Invoices, statements, notices and customer service correspondence	(CSC)
5. Subscriptions	(SUB)
6. Official government email	(GOV)
7. Business to business or employee	(BIZ)
8. Personal, friends and family	(FAF)
9. Non-profit, charitable	(NPE)

<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>, p. 11

These assertions, while not a requirement at Level A, are still considered a benefit to both the sender and the receiver as it will denote the nature of the message on the part of the sender and alert the receiver as to its nature. By adding these designations, it would also allow for a more standardization of subject lines and would eliminate deceptive practices. More affective filtering capabilities would also be achieved if these assertions are implemented and those that are legitimate would have a higher rate of success in reaching the intended recipient.

Level B for bulk senders would require not only the TETA assertions, but would add a relationship and/or permission assertion and a standardized opt-out clause with a verifiable link. The bulk sender would have to have a next higher level of certificate but not at a much greater cost. A higher level certificate will account for stricter means of accountability with Level C adhering to the strictest level. Level C or Consumer Oriented Trusted Email Certification Programs would require all of the assertions previously mentioned and goes a step further by requiring a visible assertion or logo from a Certificate Authority, a Trusted Opt-Out clause, and links to verify the validity of the certificate. A link for consumers to find out how privacy information is handled must also be included. This level would also be for corporations that have dispute resolution methods in place for a consumer to make a claim against a product or service.

The certificates and assertions will require some slight modifications to a providers' server as well as to an email gateway and spam filtering programs. Trusted Email Send and Receive Engines (TESE and TERE) would be developed to allow an email server to verify both the identity and the assertion placed on a message. The diagram below shows where these engines should be placed along with spam filters:




Figure 2: Detailed diagram of the email chain, showing role of list providers, and location of Trusted Email Send and Receive engines (TESE and TERE).

<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>, p. 14

By adding these engines and assertions, trust and accountability would be achieved. These aspects are sorely missing because of spam mail. To further guarantee these aspects, members of Level C are subject to a Trusted Email Oversight Board oversight. This board would periodically review the assertions being made and the validity and currency of the digital certificate of the senders.

Enabling Technology

In order for adherence of “best practices” and “oversight” to be implemented, technology must comply with some technological requirements.

Level A participants must have the basic level of sender identity placed into their email messages. A Level 1 certificate will ensure that the SMTP source of the message is verifiable and the TERE engines on a receiving server could identify any fraudulent Trusted Email messages. The engine would also flag messages from Trusted Email compliant domains that are not compliant by this minimum level.

In addition, other mechanisms will allow for a revocation of identity if there is an identity failure, sender-generated declarative statements to transmit data, non-replayable trust elements so that false trust elements cannot be obtained by other messages, and non-repudiation of trust elements.

Bulk Sender Trusted Email Certification Programs

Senders from this category will adhere to the methods mentioned above, but will add further requirements to process data. The requirements are:

- Securely verifiable 3rd party trusted identity. Messages can have many 3rd party trusted identity statements and receivers can verify the trusted identity statements.
- Revocation of 3rd party trusted identity. Allows certifiers to revoke certification statements if their criteria are not met. Each certifier could have varying criteria.
- Securely verifiable 3rd party trusted declarative statements. This will encompass the trusted declarative statements by certifying 3rd parties.
- Special case: Self-certification. Institutions, such as the government and special corporations, will not be required to have a 3rd party certification as they are able to self-certify themselves. They must still adhere to the other requirements of accountability.

Consumer Oriented Trusted Email Certification Programs

There are specific communication requirements that must be made when sending to single recipients and to gateway receivers. These include:

- Display of trust mark to recipient. The mark will certify that all of the methods of verification and processing have been completed by a 3rd party's set trusted identity and declarative statements.
- Special case: Self-certification. Those that are self-certified will not have to comply with any special technical requirements but still must maintain the minimum requirements to send their messages via this mechanism.

Standardization Overview

In order to comply with the standards outlined, the following will explain the technical side of meeting these standards.

Basic Components

Cryptographic Medium

Protecting data through asymmetric cryptography can satisfy many of the requirements as outlined below:

Key header:

Base64(<publicExponent>):

Base64(<modulus>):

Base64 (signed_digest(Base64(<publicExponent>):

Base64(<modulus>)))

Signature header:

Base64(signed_digest(<senderAddress><rcptAddress>

<Assertions><message_specific_data>

<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>, p. 24

Cryptographic Algorithms

RSA Public Key Cryptography
SHA-1 Hashing

The generation of messages must have an SHA-1 hash process and an RSA signature process per message. The verification of a message must have at least the same processes but with cached pre-verification of the signing keys.

Declarative Statement Language

This area is designed with logically equivalent XML-based and human-readable and compact machine-readable markers. A reliable and federated standard can be achieved by having a name-based standard structure and a scalable architecture as outlined in the example below:

Human and Machine Readable Form
Namespace declarations

```
<assertion>  
  <parameter namespace> message_type </parameter>  
  <operator namespace> is_equal_to </operator>  
  <value namespace> customer_service </value>  
</assertion>
```

Bytecode Form
AC040001AA4828C1

<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>, p. 25

How the Standard Works

There are four technical components of the Trusted Email Open Standard: distributed verification, identity confirmation, program participation/standing, and exception handling. All of these work in concert to resolve spam.

Distributed Verification and TERE

Software running as a TERE or Trusted Email Receive engine will allow translation and source verification of machine-readable code within messages. Within this code will be the source identity and assertions, if any. This software can be distributed freely via ISP's for implementation. By using this software, it can authenticate the sender's identity via a certificate, read the message assertions, and validate a sender's currency

in real-time via the email trust program. Conversely, an ISP will achieve verification of the message to applicable laws, determination of sender compliance, and can enforce additional parameters that an ISP may decide to use.

Consumer Visible Seals

In order to achieve identity confirmation and program participation/standing, the ePrivacy Group stresses that organizations create a consumer-oriented email trust program that is governed by an independent oversight panel. A participating entity in this program will be allowed to implement its own program and could install an appliance with the required software within the entity's network and the Internet. This appliance would automatically generate an individually encoded Trust Stamp or seal for every message sent out by their email server. This stamp or designation would be a visible assertion that a company is participating in this program and is complying with its standards. This technology has already been developed so further development is not necessary.

Exception Handling

The last part of how the standard works is with exception handling. Despite the best intentions of this proposed standard, there will be those that find this technology confusing. Because email volumes can be quite large, there must be a means to deal with these issues that are deemed as "exceptions" and are automated.

The type of exception can be determined through an automated process from the input of a message recipient. Once the type is verified, the exception can be resolved or escalated. Technical and dispute are two types of exception categories. When a serious attempt to abuse a system has been attempted, an exception is escalated. If there are varying differences between the sender and receiver, the dispute is escalated. These would require human responsiveness by a trusted email program operator in the event of a technical escalation and a dispute resolution operator for a dispute resolution.

Oversight

The possibility of this proposed standard coming to fruition can only be based upon the development of a governing body consisting of many bodies that are independent of the entities that they oversee. In order for this body to succeed, there must be a great acceptance from all those involved such as ISP's, consumers, and corporations and governments. Each body within the oversight group must have adequate input to ensure that trust, independence, and authority are achieved. Difficulty will arise if this does not happen. This new oversight board will be seen as the Trusted Email Oversight Board.

Responsibility for creating standards will be a primary role for the oversight board but it will also govern and rule over all disputes. The disputes are not the same as consumer

disputes. Disputes presented before the board will be in the form of standardization changes, certificate changes, etc.

The oversight board will authorize third party organizations the ability to design a Consumer Oriented Trusted Email Certifications program that can adhere to privacy concerns of consumers and email practices on behalf of an organization. These same third parties can then develop, operate, and govern these programs. Dispute resolution would also be handled by these third parties with intervention from the Trusted Email Oversight Board when extraordinary issues arise.

Corporations may want to develop different levels of the minimum standard as described under Levels A and B of the “Best Practices.” The oversight board must approve or disapprove of any requests of this nature are brought before it.

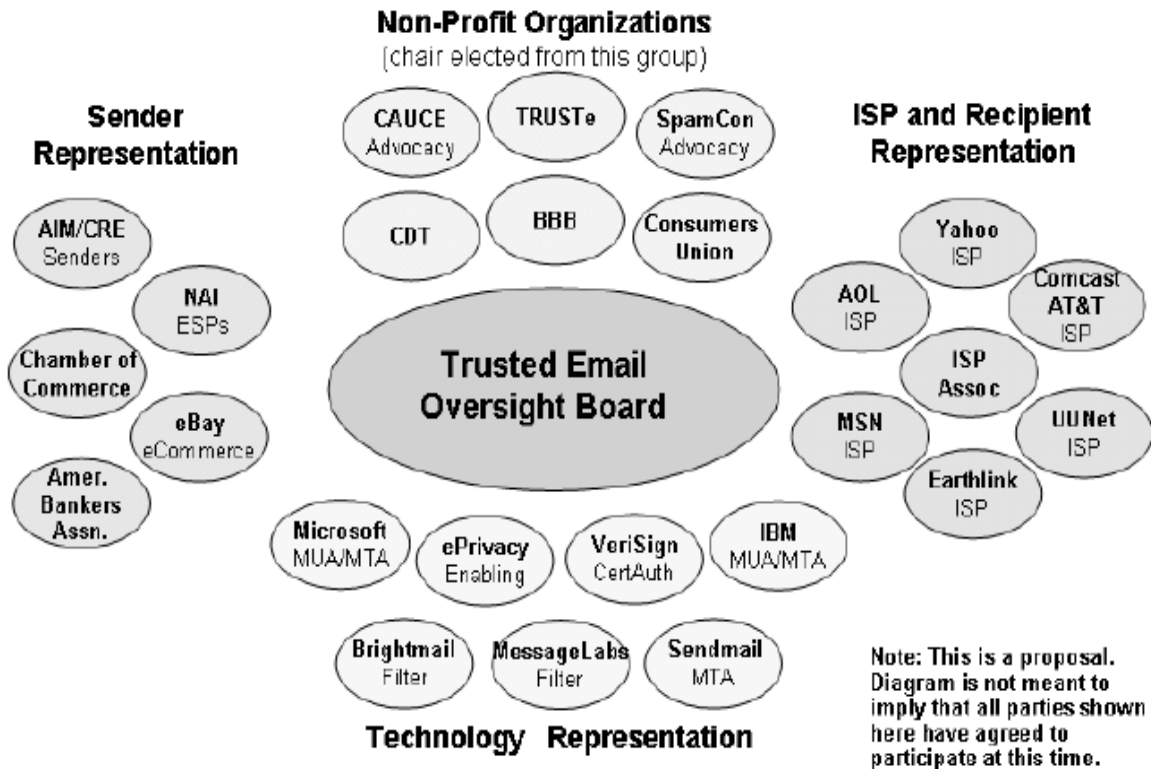
Framework for the Oversight Board

The Oversight Board will govern the Trusted Email Open Standard with emphasis placed on credibility, equal interests, scalability, and international expansion. Membership to a board-based Trusted Email Oversight Board should include consumer advocates, privacy lobbyists and representatives in the email industry. The architecture of this group will be flexible, scalable, and internationally adaptable.

The board may also sponsor companies that issue certificates and participate in activities that support the standard. The chain of command would be from the board first and then to the companies and organizations. This board would also serve in an arbitration capacity if disputes between entities arise.

The following diagram is the proposed architecture of the Trusted Email Oversight Board.

© SANS Institute 2003. All rights reserved. This document is the property of SANS Institute. No part of this document may be reproduced without the written permission of SANS Institute.



<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>, p. 33

At each level of this proposed standard, accountability is ensured. If this proposed standard is adopted by the Internet community, the assault on spam would be significantly reduced and consumer trust would be restored. The technical changes that this standard would implement would be a benefit to all entities that handle email because it not only has a minimal cost, but would practically remove any liability for damages in a court case for anti-spam law violations. This is achieved because of the additional processing engines that can validate identity of the sender as well as the currency of the certificate in real time. In addition, having an oversight board will allow approved changes to the practice of handling email and will give the global community the right to voice their positions. Overall, this standard is a definite step in the right direction.

Conclusion

It's been twenty years since the first SMTP standard entered into the Internet realm, but was not abused until a relatively recent time ago. When the Internet started to grow rapidly, email went right along with it. It was never enhanced much to allow for accountability and sender identity mechanisms to be developed because it worked very well as is. Soon companies and agencies discovered how wide an audience could be reached by getting their very own web site. To communicate with customers, they turned to email. With SMTP's reliability and simplicity, it continues to be a viable part of many businesses and organizations. The simplicity, however, has allowed for its steady

erosion of trust – not in an inability to send mail, but in the inadequacy to identify senders. Technology companies and programmers have not re-designed the SMTP protocol because the task is perceived as being a huge task when considering how long the protocol has been in existence. Instead of revision, they turned to developing applications that would scan messages to reduce the advent of spam mail. These applications are becoming impractical. Politicians now scramble to enact legislation that will be fair both to legitimate marketers and corporations, but will also take a strong stand on violators. Because of personal opinions with what is or isn't spam mail tends to prevail, legislation may stall or may be scaled back. Those bills that have passed don't necessarily deter an offender from continuing to abuse the system.

A tighter mixture of technology, standardization, and legislation, but not at the expense of revamping the SMTP protocol, may be the answer in resolving the email spam issue.

References

Newspaper

Fordahl, Matthew. "Trying to can the spam." Sun-Sentinel Newspaper. 23 March 2003. 3G

Internet

Altunergil, Oktay. "Bayesian Filtering with bogofilter and Sylpheed Claws" 30 January, 2003. URL: <http://linux.oreillynet.com/lpt/a/3167>

Cobb, Stephen. "The Trusted Email Open Standard – In Ten Bullet Points or Less" URL: <http://cobb.com/spam/teos.html>

Garden, Alex. "Newsletter # 70 The Spam Onslaught" 8 August 2002." URL: <http://www.netinsites.com/article3.cfm?ArticleID=97>

Graham, Paul. "A Plan for Spam" August 2002. URL: <http://www.paulgraham.com/spam.html>

Klensin, J. "RFC 2821 – Simple Mail Transfer Protocol" April 2001. URL: <http://www.faqs.org/rfcs/rfc2821.html>

Lowe, Richard and Claudia Arevalo. "Email Spiders" URL: http://www.internet-tips.net/Email/SPAM_spiders.htm, http://www.internet-tips.net/Email/SPAM_how.htm

McNamara, Paul. "Spammers Can't Spell "Cat" Network World Fusion 12 August 2002. URL: <http://www.nwfusion.com/columnists/2002/0812netbuzz.html>

Mertz, David. "Six approaches to eliminating unwanted e-mail" 1 September 2002. URL: <http://www-106.ibm.com/developerworks/linux/library/l-spamf.html>

Mosher, Sue. "News from the antispam front" 13 May 2003 URL:
<http://www.exchangeadmin.com/Articles/Index.cfm?ArticleID=39014&pg=1&show=854>

Pezzullo, John C. "Bayes' Theorem Calculator" URL:
<http://members.aol.com/johnp71/bayes.html>

Postel, Jonathan B. "RFC 821 – SIMPLE MAIL TRANSPORT PROTOCOL" Information Sciences Institute. August 1982 URL: <http://www.faqs.org/rfcs/rfc821.html>

Robichaux, Paul. "New Hope in the Spam Battle?" 2 May 2003 URL:
<http://www.exchangeadmin.com/Articles/Index.cfm?ArticleID=38919&pg=1&show=756>

Schiavone, Vincent; Brussin, David; Koenig, James; Cobb, Stephen; Everett-Church, Ray. "Trusted Email Open Standard – A Comprehensive Policy and Technology Proposal for Email Reform" An ePrivacy White Paper May 2003 URL:
<http://www.eprivacygroup.net/teos/TEOSwhitepaper1.pdf>

Strassman, Paul A. "Risk-Free Access Into the Global Information Infrastructure via Anonymous Re-Mailers" 30, January 1996. URL:
<http://www.strassmann.com/pubs/anon-remail.html>

Tschabitscher, Heinz. "What You Need to Know about Bayesian Spam Filtering" URL: http://email.about.com/cs/bayesianfilters/a/bayesian_filter.htm,
http://email.about.com/cs/bayesianfilters/a/bayesian_filter_2.htm

<http://www.eprivacygroup.net/teos/>

<http://www.ftc.gov/bcp/online/pubs/online/inbox.htm> April 2002

<http://www.junkbusters.com/ftc.html#2.16>

<http://www.leginfo.ca.gov/cgi-bin/displaycode?section=bpc&group=17001-18000&file=17530-17539.6>

<http://mail-abuse.org/>

<http://www.ordb.org/>

<http://www.spamlaws.com/state/summary.html>



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Dubai 2018	Dubai, AE	Jan 27, 2018 - Feb 01, 2018	Live Event
SANS Las Vegas 2018	Las Vegas, NVUS	Jan 28, 2018 - Feb 02, 2018	Live Event
Cyber Threat Intelligence Summit & Training 2018	Bethesda, MDUS	Jan 29, 2018 - Feb 05, 2018	Live Event
SANS Miami 2018	Miami, FLUS	Jan 29, 2018 - Feb 03, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZUS	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS London February 2018	London, GB	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Southern California- Anaheim 2018	Anaheim, CAUS	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Secure India 2018	Bangalore, IN	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Dallas 2018	Dallas, TXUS	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Brussels February 2018	Brussels, BE	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Secure Japan 2018	Tokyo, JP	Feb 19, 2018 - Mar 03, 2018	Live Event
Cloud Security Summit & Training 2018	San Diego, CAUS	Feb 19, 2018 - Feb 26, 2018	Live Event
SANS New York City Winter 2018	New York, NYUS	Feb 26, 2018 - Mar 03, 2018	Live Event
CyberThreat Summit 2018	London, GB	Feb 27, 2018 - Feb 28, 2018	Live Event
SANS London March 2018	London, GB	Mar 05, 2018 - Mar 10, 2018	Live Event
SANS San Francisco Spring 2018	San Francisco, CAUS	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS Secure Osaka 2018	Osaka, JP	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS Secure Singapore 2018	Singapore, SG	Mar 12, 2018 - Mar 24, 2018	Live Event
SANS Paris March 2018	Paris, FR	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS SEC460: Enterprise Threat Beta	OnlineCAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced