Task Force on Spam

Companion Document to *Best Practices for Internet Service Providers and Other Network Operators*

By John Levine

May 2005

Executive Summary

Electronic mail has remained the Internet's most useful application for the past 20 years, even as spam and other email abuse have increased. However, the serious threat posed by spam led the Government of Canada to establish a task force to implement *An Anti-Spam Action Plan for Canada.* Because Internet service providers (ISPs) and other network operators can prevent many kinds of email abuse and detect others, a working group was set up to look at technology and network management. The working group developed a set of technical best practices designed to increase the accountability of mail senders while minimizing disruption to mail users.

Considering the technical nature of the best practices document, this companion guide has been developed to explain the concepts involved and how they will serve to address the spam problem.

An Overview of Internet Email

Internet email was designed in an era when computers and networks were far slower and less reliable than they are now, but when the Internet and its predecessor, ArpaNet, had no abuse problems. For this reason, the standard used to deliver mail, known as Simple Mail Transfer Protocol (SMTP), provides robust facilities for delivering mail, but very weak facilities for tracing and sender accountability. SMTP is designed as a store and forward system, so that a sender's computer may deliver mail directly to a recipient's computer. However, more likely, the mail will be passed from computer to computer in several stages — such as from a user's PC to their ISP's outgoing mail server, then to the recipient ISP's incoming mail server, and then to the recipient ISP's internal delivery mail server for final delivery.

In Internet parlance, a computer attached to the Internet is called a "host." In any transaction between two computers, one is considered to be a "server," and the other a "client." In SMTP email, the sending computer is the client and the receiving computer is the server, so a mail server is a host that receives mail for its users. It is possible for the same host to assume different roles, so that a mail server that has received a message may then act as a mail client to send the message to another server.

Each host has an Internet protocol (IP) address, a number analogous to a phone number, which other hosts use to contact it. Since people find names easier to remember than numbers, host and mail systems also have names (e.g. **example.ca** and **ic.gc.ca**). The Internet's DNS keeps track of which names correspond to what IP addresses. Domain names are also used in email addresses; in an address like **pm@pm.gc.ca**, **pm.gc.ca** is the email address's domain.

The Best Practices

1. All Canadian registrants and hosts of domain names should publish Sender Policy Framework (SPF) information in their respective domain name server zone files as soon as possible.

SPF is a system intended to detect forged email. It allows the manager of a mail domain to publish a list of the IP addresses of hosts allowed to send mail using email addressed in that domain. If all of a domain's mail is sent from a single place, as is often the case for bulk mailers and small businesses, SPF can be a useful tool with which to detect forgery when mail claiming to be from that domain arrives from elsewhere.

Since much of the spam sent uses forged domains to hide its origins, the ability to identify forgery is key in tracing spam, and thereby preventing it from entering networks.

2. ISPs and other network operators should limit, by default, the use of port 25 by end-users. If necessary, the ability to send or receive mail over port 25 should be restricted to hosts on the provider's network. Use of port 25 by end-users should be permitted on an as-needed basis, or as set out in the provider's end-user agreement / terms of service.

Although it is technically possible for ISP customers to send email directly from their own PCs to the message recipient's computer, normally the user sends the message from their PC to the ISP's mail server, which then sends the mail to the recipient. The original motivation for this technique was to make email delivery more reliable, since the ISP's mail server can retry later if the message isn't immediately deliverable. Other advantages will be described later.

In order to avoid detection by the user's ISP, spammers and viruses often try to send mail directly from user PCs to recipients without using the ISP's mail server. Originally, spammers signed up accounts of their own, but now they use user PCs infected with worms or viruses that let the computers be remotely controlled by spammers (these controlled computers are known as "zombies").

If the ISP limits user access to port 25, which is the logical channel used for Internet email, so that a user can only send email directly to the ISP's mail server, the ISP will be aware of all mail a user's PC sends, and it can take action if it detects abusive behavior.

In a few cases, user PCs have legitimate reasons for contacting mail servers other than their ISPs' (e.g. when a telecommuter sends mail through his or her employer's mail system). ISPs, therefore, need ways to provide exceptions to the port 25 block; however,

such users make up a small enough fraction of the total user population that they can be handled individually.

Port 25 has been widely abused by spammers running zombie networks (or "botnets"). By monitoring and limiting the use of port 25, ISPs and other network operators can close off a major avenue for spamming. Canadian ISPs that have already implemented port 25 blocking have seen very significant declines in the amounts of spam originating on their networks.

3. ISPs and other network operators should block email file attachments with specific extensions known to carry infections, or should filter email file attachments based on content properties.

An Internet email message can consist of a combination of message text and data files attached to the message. While attachments are extremely useful (e.g. for sending documents or presentation files to a coworker), they are also heavily used by viruses, which create email messages and attach copies of the virus itself to them in order to infect the message recipients' computers when the recipients open the message. The types of attached files that viruses use have few legitimate uses, so ISPs can and should block them. (An "extension" is the part of a file name following a dot, and identifies the file type.)

Attachment blocking is useful both in incoming mail, to keep viruses from entering the network, and in outgoing mail, to detect virus-infected customers.

These efforts will minimize the risk that a spam message carrying a harmful program can enter a network and deposit viruses and worms used to create and run botnets. A reduction in botnets means a reduction in spam.

4. ISPs and other network operators should actively monitor the volume of inbound and outbound email traffic to determine unusual network activity and the source of such activity, and should respond appropriately.

An Internet user's mail behavior rarely changes radically from one day to the next. If a user has been sending out three messages a day and suddenly starts sending 10 000 messages in one day, it is much more likely that a virus has commandeered their PC to send spam than that they have made a lot of new friends. Conversely, if a user who normally receives a dozen messages a day starts receiving a thousand messages a day, it's almost certainly evidence of a problem.

If an ISP notes unusual activity and acts on it promptly, it is often possible to stop a spam mailing run or other abusive activity while it is in progress, rather than wait for the victims to notice and send in complaints afterward.

5. ISPs and other network operators should establish and consistently maintain effective and timely processes to allow compromised network elements to be managed and eliminated as sources of spam.

The largest single source of spam and other abusive network activity is now zombies — users' computers infected with worms or viruses that allow spammers and other criminals to control the computers remotely. Since the computers still work for normal purposes, users are rarely aware that their computers have been taken over.

When an ISP notices abuse coming from a user's PC, the ISP needs to first suspend the user's network access in order to stop the abuse, and then help the user remove the worm(s) or virus(es). The removal process is not simple, and generally needs a combination of anti-virus and anti-spyware programs, along with system updates from the user's software vendor to avoid reinfection. Few users can de-worm their own computers without help, and the ISP is usually a user's only resource with the knowledge to help them.

ISPs and other network operators that implement this best practice will be able to quickly and effectively stop spam originating on their networks.

6. ISPs and other network operators should establish appropriate intercompany processes for reacting to other network operators' incident reports.

ISPs and network operators frequently send each other reports of network abuse. If an ISP's user is sending out spam, the first time the ISP may hear about it is in reports from other ISPs, when the email hits recipients' ISP spam filters.

The more quickly and efficiently ISP's can notify each other about problems, the more quickly the ISP with the problem user can address the issue, and, thus, the less spam that will be sent, and received, across multiple networks.

7. ISPs, other network operators and enterprise email providers should communicate their security policies and procedures to their subscribers.

The effects of ISP security policies are often not immediately apparent to users. For example, if a virus filter rejects a particular type of message, users who might want to receive that type of message would only find out if the correspondent who sent such a message noticed it wasn't delivered and then notified the user.

Every ISP security policy involves a trade-off between prohibiting often-abused activities and preventing legitimate use of those activities. While the vast majority of activities that security policies affect have little or no legitimate uses, some policies do prevent small but significant amounts of legitimate use. If ISPs notify their users about their policies, the users who are affected can adjust the way they work, so that their methods are

compatible with the policies. For example, when an ISP applies port 25 blocks, users who need to contact mail servers on other networks can usually adjust their mail programs to use a different port that is not subject to the abuse problems of port 25.

A particular area of concern is spam filtering, since misconfigured spam filters can reject or delete significant amounts of desired mail. ISPs can make available descriptions of the filtering techniques they use, what happens to email that is identified as spam, and what options users have if they believe that valid mail has been mischaracterized as spam or vice versa.

Users are key in the fight against spam. Increasing their awareness of the issue, and of the measures in place to fight spam, will improve their understanding and ability to protect themselves and, by extension, the networks on which they are hosted.

8. ISPs and other network operators should implement email validation on all their Simple Mail Transfer Protocol (SMTP) servers (inbound, outbound and relay).

Traditionally, Internet email servers accepted mail from any source, intended for any destination. Unfortunately, spammers in the 1990s started abusing mail servers, and ISPs had to reconfigure the servers so that only ISPs' own users could use them to send email. ISPs can use a variety of techniques to recognize mail from their own users. The preferred technique is to use SMTP authentication (defined in Request for Comments [RFC] 2554), in which message senders authenticate themselves to a mail server, usually using the same user name and password they use to pick up incoming mail. This ties each message to a particular user.

Being able to identify the source of abusive outgoing mail will make it is possible to block that mail if it is spam.

9. Non-delivery Notices (NDNs) should only be sent for legitimate emails.

It is, technically, no harder to put a false return address on an email message than it is on a paper letter. As a result, nearly all abusive email now has forged return addresses. In the common case that a spam message or virus is sent to an invalid address, a resulting NDN would be sent back not to the spammer, but to the innocent party whose address was forged. Widely forged domains can get vast amounts of forged-address NDNs per day, at great inconvenience to the users.

One possible solution would be to abandon NDNs altogether, but that would be undesirable, since it would prevent legitimate senders from knowing if they had sent mail to a mistyped or obsolete address. Rather, network operators can often arrange for their mail servers to reject undeliverable mail at the time the sending host tries to deliver it. If the sending host is a legitimate mail server, the server will create an NDN or otherwise notify the sender. If it's a spamming program, the mail server will ignore the failure

without creating an NDN. Network operators can also apply spam and virus filters to undeliverable mail, and only send NDNs for mail not categorized as spam or a virus.

This best practice does not specifically prevent spam, but can minimize the inconvenience to users who have become the victims of spam by having had their return addresses forged.

10. ISPs and other network operators should ensure that all domain names, Domain Name System (DNS) records, and applicable Internet protocol (IP) address registration records (e.g. WHOIS, Shared WHOIS Project [SWI] or referral WHOIS [RWHOIS]) are responsibly maintained with correct, complete, and current information. This information should include points of contact for roles responsible for resolving abuse issues including, but not limited to, postal address, phone number and email address.

Every domain used on the Internet (e.g. **gc.ca** or **example.com**) is registered with one of a group of registries. At the time of registration, domain owners provide basic contact information, including their name, address, phone number and email address. These registries provide access to some or all of the registrant's information through a service called WHOIS. Once a domain is registered, the domain owner creates DNS records to publish the locations of mail servers, Web servers and other network servers available to Internet users.

Also, when a Canadian network requires IP addresses for its own use or that of its customers, it obtains those addresses from a regional registry known as the American Registry for Internet Numbers (ARIN). Again, at registration time the network provides its contact information. If the network then suballocates some of its assigned IP addresses to customer networks, it can use SWIP to add contact information to ARIN for the suballocations. ARIN also runs a WHOIS service that provides the registrant information for allocated IP addresses. Network operators can also create special DNS entries called "reverse DNS" or "rDNS" to document the domain name(s) assigned to each of their IP addresses.

DNS and WHOIS data are crucial tools for tracking Internet abuse, including spam. Since the domain names in email are so easily forged, the only reliable identification information in an incoming spam email or virus is the IP address of the host that sent it. WHOIS and rDNS allow a recipient to identify the party responsible for the IP address, so long as the WHOIS and rDNS data are accurate. It is, therefore, important for network operators to both ensure that they provide accurate initial WHOIS information and update it whenever the contact information changes.

Implementing this best practice could help identify spammers, and could also provide a way for legitimate emailers to identify themselves.

11. ISPs and other network operators should ensure that all their publicly routable and Internet-visible IP addresses have appropriate and up-to-date forward and reverse DNS records, WHOIS and SWIP entries. All local area network (LAN) operators should be compliant with Request for Comments (RFC) 1918 — "Address Allocation for Private Internets." In particular, LANs should not use IP space globally registered to someone else, or IP space not registered to anyone, as private IP space.

Each network has ranges of IP addresses assigned to the computers on that network. In most cases, network operators go through considerable effort to prevent duplication and to assign each allocated IP address to one specific computer. This makes it possible for any other computer anywhere on the Internet to contact that computer using that address. These universally accessible addresses are known as "publicly routable," or "visible," addresses.

It is also possible to create private networks, using IP addresses not accessible from the rest of the Internet, either for security purposes or when there are more computers on the network than there are available publicly routable IP addresses. For example, a home DSL or cable modem user is usually only assigned a single public address by their ISP. If the user has several computers, it is possible to connect all of the computers into a private home network, then use a device called a "router" to connect the private network to the Internet. The router makes the entire private network appear, to the rest of the Internet, to be a single computer with a single public IP address. Businesses frequently use this technique on a larger scale to create secure private networks with controlled access to the public Internet.

Since the computers on a private network are not connected directly to the Internet, the IP addresses on the private network don't need to be allocated by ARIN. In 1996, RFC 1918 defined the sets of IP addresses to be used on private networks; since then, all properly designed private networks have used RFC 1918 addresses.

Unfortunately, some private networks use non-RFC 1918 addresses — typically, addresses the system manager thought would never be assigned to anyone else. Non-standard addresses on private networks cause trouble for those tracking down network abuse, since it is most often impossible when looking at message or network trace data to tell whether the non-standard address is being used on a private network. Worse, in the case that the address has been assigned to someone else, using non-standard addresses can shift the blame to the legitimate users of the IP addresses.

This best practice is related to Recommendation #10, and its implementation could help in identifying spammers. It could also provide a way for legitimate emailers to identify themselves.

12. ISPs and other network operators should prohibit the sending of email that contains deceptive or forged headers. Header-tracing information should be correct and compliant with relevant RFCs, including RFC 822 and RFC 2822, and reference domains and IP addresses should have up-to-date, accurate registration information.

Each time an email message is handled by a mail server, the server adds a trace header to the top of the message. The set of headers on a message is intended to document the message's path through the Internet's mail system. Since each server merely adds its header to what's already present, it's not hard for viruses or spamming software to insert forged headers to disguise the actual source of the email. In addition, some legitimate email software doesn't use correct header formats and produces headers that, while not deliberately fraudulent, can be misleading.

Network operators can use their knowledge of the structure of their network to refuse email with many different kinds of forged headers. For example, if a message was sent from a host that has a particular IP address on the network, any message headers that make the message appear to have originated elsewhere can be presumed to be forged.

This best practice is related to Recommendation #10, and its implementation could help in identifying spammers. It could also provide a way for legitimate emailers to identify themselves.

References

All RFCs are available online (**www.rfc-editor.org/rfc.html**).

D. Crocker. RFC 822 — Standard for the format of ARPA Internet text messages. August 1982.

D. Karrenberg et al. RFC 1918 — Address Allocation for Private Internets. February 1996.

J. Klensin, ed. RFC 2821 — Simple Mail Transfer Protocol. April 2001. [Updates RFC 821]

J. Myers. RFC2554 — SMTP Service Extension for Authentication. March 1999.

J. Postel. RFC 821 — Simple Mail Transfer Protocol. August 1982.

P. Resnick, ed. RFC 2822 —Internet Message Format. April 2001. [Updates RFC 821]