

# In The Boxing Ring



## IN THIS ISSUE

### 2. EMAIL - SMTP, POP3 AND IMAP4

A presentation on Internet eMail, its core standards and the SMTP, POP3 and IMAP4 protocols. In particular, information is given as to why Network Box (as well as other gateway appliances) can quarantine mail in some eMail protocols, but not others.

### 3. ANTI-SPAM AND WHITELISTING / BLACKLISTING

A summary of the Network Box Anti-Spam system, and how whitelisting (and blacklisting) can be effectively used to tune the system.

### 4. NOVEMBER 2009 FEATURES

As usual, we will be deploying our on-going enhancements and improvements as well as maintenance features.

### 4. NOVEMBER 2009 HINT

Tips on how to cleanly shut down, start up, or restart your Network Box appliance.

## Network Box Technical News from Mark Webb-Johnson, CTO Network Box

### Welcome

Welcome to the November 2009 edition of 'In the Boxing Ring'. In this edition, we focus on eMail.

Turn to page 2 for a presentation on Internet eMail, its core standards with the SMTP, POP3 and IMAP4 protocols (as implemented by Network Box). This presentation is intended to give a good grounding on these core protocols, and the benefits and drawbacks to each one. In particular, information is given as to why Network Box (as well as other gateway appliances) can quarantine mail in some eMail protocols, but not others. The difference between 'from' and 'envelope sender' is also discussed.

On page 3, we give a summary of the Network Box Anti-Spam system, and how whitelisting (and blacklisting) can be effectively used to tune the system.

Page 4 details the usual monthly features summary and hints.

As usual, if you have any feedback, or comments, they are always appreciated. You can contact us here at HQ via eMail ([nbhq@network-box.com](mailto:nbhq@network-box.com)). Or, drop by our office next time you are in town.

You can also keep in touch by following our Network Box Security Response twitter feed at:

[twitter.com/networkboxhq](https://twitter.com/networkboxhq)

Mark Webb-Johnson  
CTO, Network Box Corporation  
November 2009





## Internet eMail SMTP, POP3 and IMAP4 Protocols

Internet eMail, as we know it, dates back to the early 1980s with the publication of RFC 821 and the SMTP protocol. This protocol, still in use (relatively unchanged) today, is the core Internet protocol responsible for passing eMail messages from one mail server to another. While other proprietary solutions exist (such as Microsoft Exchange inter-exchange and Lotus Notes protocols), SMTP is by far the most popular mail exchange protocol in use today.

### The SMTP Protocol

SMTP (Simple Mail Transfer Protocol) provides for differentiation between an eMail envelope (comparable to a physical mail envelope with recipient listed on the front as well as a sender return address) and eMail message (including headers and body - comparable to the physical paper letter inside a mail envelope). The protocol provides for three stages of transmission - first the servers identify each other and indicate their capabilities, then a message envelope is transmitted, and finally the message itself is transmitted.

While normally the same, there is no requirement for the sender and recipient listed in the eMail message headers to be the same as those listed in the envelope. Importantly, there is, as standard, no validation of the header or envelope addresses - it is trivial to forge SMTP eMail (although many optional extension mechanisms such as SPF and digital signature schemes attempt to address this). The other primary difference between physical mail and SMTP eMail is that SMTP eMail can be addressed (in the envelope) to multiple recipients - kind of like listing all the recipients on the outside of an envelope and expecting the post office to make photocopies of the mail for each delivery.

Once a message has been delivered to the final SMTP mail server mailbox responsible for end user delivery, the envelope is usually discarded.

SMTP servers can choose to either accept an eMail, temporarily defer

acceptance, or permanently reject - and this can be done at either the envelope or message stage of the transmission. If a message is undeliverable (DNS error or the next-hop SMTP server indicates a permanent reject), the message is bounced - ie; a Non Delivery Receipt (NDR) is generated back to the envelope sender. In the case where there is no envelope sender, the message is discarded.

SMTP is a message transmission protocol. While an optional extension is available to TURN around a transmission (i.e.; the destination server would become the transmitter and send back any queued messages), this is not commonly deployed nowadays. In normal use, SMTP is a PUSH protocol. Servers with messages queued for delivery will connect to the destination (using SMTP) deliver the messages and then remove them from the queue. Messages which have been in the queue for too long (typically several days) will typically be bounced (using the above NDR mechanism) as undeliverable.

SMTP is an extremely flexible protocol. The receiving SMTP server can accept a message (then chose to forward it on, quarantine, or drop it), defer or reject it. It can also change the envelope, in order to redirect the message elsewhere. Network Box makes good use of all these facilities.

So, if SMTP is a transmission protocol, how does a mail client (such as Outlook, Thunderbird, Eudora or Apple Mail) receive its eMail messages? While proprietary solutions exist (such as Microsoft Exchange and Lotus Notes) two common Internet standards exist - POP3 and IMAP4.

### The POP3 Protocol

POP3 is the Post Office Protocol version 3 and was originally documented in RFC 1939 in the mid 1990s. It is a relatively simple protocol designed to allow a mail client to poll a mail server, authenticate the user to a mailbox, get a list of messages in the mailbox, transfer those messages to the client, and then optionally delete the messages from the mailbox on the server (presumably after

adding them to a local copy of the mailbox on the client).

The POP3 protocol is, however, very limited. Once a client requests a message, there is no mechanism for the server not to deliver it (other than to terminate the connection with an error). For this reason, advanced functionality such as message redirection and quarantining is not possible when using the basic POP3 protocol.

### The IMAP4 Protocol

IMAP4 is the Internet Message Access Protocol version 4 and was originally documented in RFC 2060 around the same time as POP3. When compared to POP3, it provides a much more rich environment (including facilities such as support for folders within the mailbox, retrieval of individual parts of a message, and sophisticated search mechanisms). While POP3 was designed to allow the mail client to keep a local copy of the mail box (where the advanced functionality is performed), IMAP4 was designed to be client-server (so that the mail client does not need a local copy and all the messages can be maintained on the mail server).

The IMAP4 protocol does, however, suffer from much the same limitations as POP3. There is no mechanism for a server to refuse a request for a message (without a nasty error message), and hence advanced functionality such as message redirection and quarantining is not possible when using the basic IMAP4 protocol.

### Conclusion

The SMTP protocol is used to send email messages from mail server to mail server. Highly reliable, and tightly integrated to the DNS system, it is the core backbone for eMail on the Internet.

The POP3 and IMAP4 protocols are used to retrieve messages from a mail server, for display by a mail client.

Most Internet mail clients will have configuration screens to define the SMTP server used to send mail to and POP3/IMAP4 server to receive from.



## Anti-Spam and Whitelisting / Blacklisting

The fight against spam involves being given an eMail message, and making a determination as to whether that message is spam or not.

Network Box uses a sophisticated Anti-Spam system deploying a comprehensive collection of technologies to maximize the rate of successful detection of spam, while minimizing the false-positives. Such technologies include: co-operative spam checksums, signatures and spam scoring, white lists and black lists, heuristics, real-time IP and URL blacklists (aka reputation), URL to IP mapping and blacklists, URL categorization, domain age, bayesian filtering, challenge/response systems, digital signatures, OCR, fuzzy signatures, and relationship based databases.

The output of the anti-spam system is a determination as to whether the message is definitely spam/ham (blacklist or whitelist) or the likelihood that it is spam (expressed as a score, with messages scoring higher than a defined threshold being more likely to be spam).

Messages which are either blacklisted or score sufficiently high (and thus have a high likelihood of being spam) are fed back into the self-learning system to train the system what spam looks like (to improve future spam detection performance). Messages which are whitelisted are also fed back into the self-learning system to train the system what ham looks like (to reduce the future false positive rate).

Self-learning systems include such technologies as bayesian statistical analysis and relationship databases.

It is thus vitally important that the whitelisting and blacklisting systems are understood and used effectively - in order to avoid mis-training the self-learning systems.

Whitelisting and Blacklisting are normally performed using the sender address of an eMail. The problem is that, as previously discussed, it is trivially easy to forge the sender.

The Network Box relationship system takes advantage of the fact that in the vast majority of cases the spammer knows who you are, but not who you correspond with. When whitelisting and blacklisting, despite the forging problem, you can still take advantage of this.

The key is to avoid whitelisting or blacklisting domains that the spammer would know you correspond with. Such domains include your own (approximately 1% - 3% of current spam is forged to appear to come from your own domain) and popular common domains (such as yahoo, hotmail, gmail and popular ISPs).

For unwanted bulk spam, blacklisting is normally ineffective (as the spammer changes sender eMail address with each message he sends out). For such cases, the best is to forward (as an attachment) missed spam to [spam@network-box.com](mailto:spam@network-box.com) and let the Network Box experts handle it.

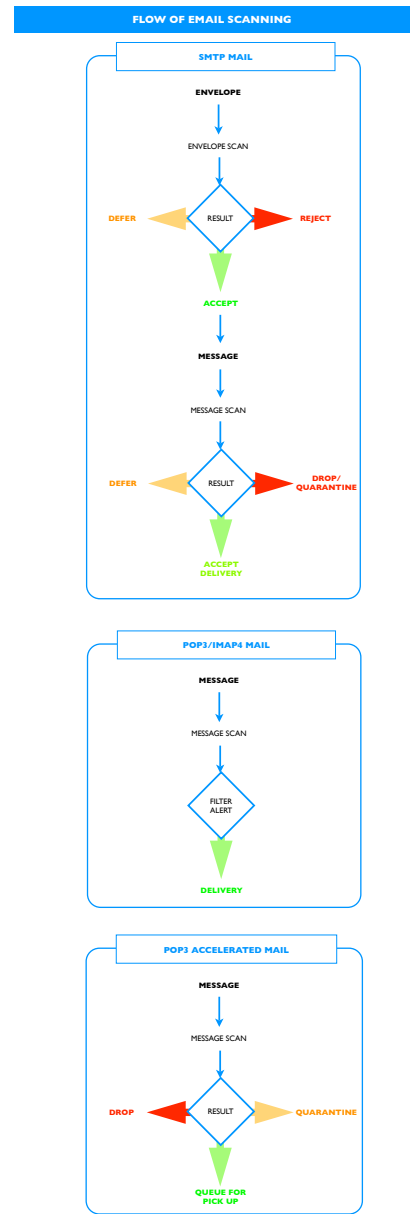
For unwanted mailing list style eMails, blacklisting is very effective.


For wanted (but incorrectly marked as spam) eMails, whitelisting the sender's eMail address is very effective - so long as it is not on your own domain.

For senders on your own domain, the best approach is to arrange for all such eMails to be sent outbound (either from the LAN, via VPN, or via an authenticated SMTP connection).

Such outbound eMails are not scanned for spam and can never be blocked as such.

If you can ensure that all mail from your domain leaves outbound via defined IP address ranges, then you should also implement SPF. Doing so would immediately stop 1% to 3% of spam forged as sender you, as well as stop spammers from forging you to other people checking SPF records. If everyone implemented SPF, there would be no more forgery of sender eMail addresses.





### November 2009 Features

On Tuesday, 3<sup>rd</sup> November 2009, Network Box will release our patch Tuesday set of enhancements and fixes. The regional NOCs will be conducting the rollouts of the new functionality in a phased manner over the next 7 days. This month, these include:

- Some minor enhancements and fixes to my.network-box.com and mail portal interfaces (mostly involving quarantine release, searching, and the administrative display of challenged eMails for customers using the challenge/response system).
- A new option to allow the suppression of the malware table in the Mail Portal eMail report.
- Relaxing of the default restriction to allow whitelisting of individual eMail addresses on common domains (but to maintain the default restriction against whitelisting entire common domains).
- Support for NBIDPS in the weekly reporting system. If you have deployed this system, you will now see IPS alerts reported in the weekly report.
- Enhanced support for syslog and eMail alerts for SSL VPN connections.

In most cases, the above changes should not impact running services or require a device restart. However, in some cases (depending on configuration), a device restart may be required. Your local NOC will contact you to arrange this if necessary.

Should you need any further information on any of the above, please contact your local NOC. They will be arranging deployment and liaison.

### November 2009 Hint

As it is important to cleanly shut down / reboot Network Boxes, we offer various mechanisms to help you. These include:

- For models with a front-panel display and keypad, you can select SHUTDOWN or REBOOT from the display.
- For models connected via serial console, the VPANEL facility will give you the same functionality as a front-panel display and supports both SHUTDOWN and REBOOT.
- For models with a power switch, depressing the switch once will initiate a clean shutdown (and the front-panel display will show progress).
- For all models, the my.network-box.com administrative interface offers “Reboot Box” and “Shutdown Box” options under the BOX / CONTROL screen.

Cleanly shutting down or rebooting your Network Box will ensure that (a) all in-progress updates are cleanly aborted, (b) all disks are cleanly unmounted, (c) all services are cleanly stopped, and (d) the database and unfinished transactions are written to disk.

Failing to cleanly shutdown or reboot your Network Box may cause damage to hardware and/or system files. We recommend that the power plug should only be pulled in unavoidable circumstances.

Mark Webb-Johnson  
 CTO, Network Box Corporation  
 November 2009

### OCTOBER 2009 NUMBERS

Key Metric)	#	% difference (since last month)
PUSH Updates	1,361	-12.2
Signatures Released	220,552	-2.8
Firewall Blocks (/box)	613,146	-3.9
IDP Blocks (/box)	190,456	+6.0
Spams (/box)	65,247	+9.8
Malware (/box)	4,060	+19.3
URL Blocks (/box)	106,173	-14.2
URL Visits (/box)	2,947,567	-11.2

### NEWSLETTER STAFF

**Mark Webb-Johnson**  
 Editor

**Michael Gazeley**

**Jason Law**

**Nick Jones**

Production Support

**Network Box Australia**

**Network Box Hong Kong**

**Network Box UK**

Contributors

### SUBSCRIPTION

Network Box Corporation  
[nbhq@network-box.com](mailto:nbhq@network-box.com)  
 or via mail at:

**Network Box Corporation**

16th Floor, Metro Loft,  
 38 Kwai Hei Street,  
 Kwai Chung, Hong Kong

Tel: +852 2736-2078

Fax: +852 2736-2778

[www.network-box.com](http://www.network-box.com)