

THE NEED FOR AN IN-HOUSE SMTP HONEYPOT

Vinoo Thomas and Nitin Jyoti
McAfee Avert Labs, India

In the good old days spammers scanned the Internet aggressively for open relay servers to send spam. But open relays are out of vogue these days – so much so that the Open Relay Database has shut down as a result of changes in spammer tactics [1].

Today’s spammers, in collusion with malware authors, infect thousands of machines on the Internet, turning them into spam relay zombies. Many Internet users are ignorant of the fact that one can get infected just by visiting a malicious website – without any user intervention. All it takes is the click of a link and one is directed to sites that serve a cocktail of browser and application exploits that attempt the ‘drive-by’ installation of malware on one’s machine. To add to these woes, legitimate sites are also being compromised and abused to serve malicious content and infect unsuspecting users [2].

Once infected, a zombie machine connects to a command-and-control server run by the spammer. This server provides a constantly updated live feed of email addresses and content for spamming. The content could be anything from malicious e-cards or pump-and-dump stock scams, to advertisements for online retailers of pharmaceutical drugs or sexual aids. Each zombie machine is capable of sending hundreds of spam messages per minute. With the average personal computer having better bandwidth and processing power these days, commanding large numbers of zombies keeps virus authors and spammers in business.

On the corporate front, organizations have traditionally blocked outbound SMTP traffic on port 25 that originates from the local area network (LAN) and virtual private network (VPN) segments. This prevents an internal machine that has been infected with a mass-mailer from spamming the outside world. A spammed email can be traced back to its source using the IP information contained in the mail header – imagine if a security company had an internal infection and the originating IP of the spammed mail were to be traced back to that organization! It would be a public relations nightmare.

By blocking port 25 at the firewall, an organization can prevent a mass-mailer from spreading. However, by blindly blocking outgoing SMTP traffic, we also lose valuable data on threats that use port 25.

The way to hold on to that data is with sticky fingers: by redirecting all SMTP traffic originating from the LAN and

VPN segments to a honeypot, we can learn much more about real-time infections that occur within the internal network. This allows the network administrator to capture information about the following types of threat:

- Backdoors/password stealers: if a keylogger, password stealer, or backdoor were to send a notification mail back to the trojan’s author, the captured email content would reveal not only the IP address of the compromised machine on the network, but also the kind of sensitive data that was being relayed to the attacker.
- Mass-mailers: copies of the worm can be harvested from the spammed attachment and can be sent to an anti-virus researcher for analysis. Infected machines can be identified and isolated from the network.
- Spam bots: spammed emails can be captured and used to identify machines on the internal network that are being used as spam-relay zombies.

IMPLEMENTATION

At McAfee Avert Labs in Bangalore, India, we implemented a honeypot on an internal network on which the administrator was reporting frequent infections and had trouble isolating the rogue machines. To accomplish this, we used *MailPot* – an open-source SMTP/ESMTP honeypot from *iDefense* [3]. As is typical of any honeypot, *MailPot* sits in the demilitarized zone of the firewall [4], and using iptables we redirect any traffic passing through port 25 from the internal network to our SMTP honeypot.

Here’s an example of an iptable rule on the firewall:

```
iptables -t nat -A PREROUTING -i <interface id> -s <affected network address> -p tcp -dport 25 -j DNAT -to <ip address of honeypot>
```

Figure 1 shows a screenshot of the *MailPot* interface, illustrating some of the sample mails that we captured.

In the example above we see a variety of email captures, ranging from trojan victim online notifications, to mass-mailers and malicious postcard spam. From the

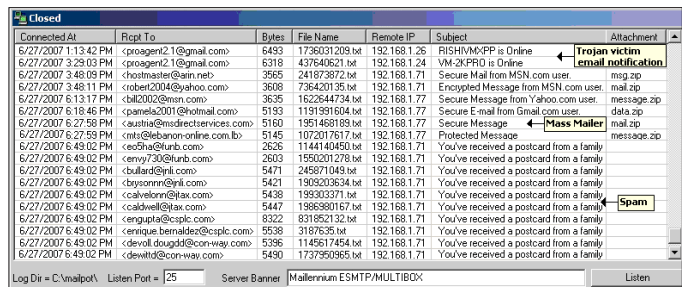


Figure 1: MailPot captures.

originating IP address that's displayed, the network administrator can pinpoint the source of the infection. Also, the captured mass-mailer samples can be collected and submitted quickly to an anti-virus vendor for analysis and inclusion in their signature files.

Note: Mass-mailers such as W32/Sober@MM [5] use the non-standard SMTP port 587 [6] in addition to port 25 to spread. Redirecting SMTP traffic on this port is also highly advisable.

A useful alternative is for the mail administrator to embed a special email address into the standard corporate ghost image for workstations, servers and laptops. If anything shows up in this special mailbox, it means that a mass-mailer or email-harvesting trojan has crawled the disk and found this special address. That's a sure way of knowing you've had an internal infection.

Every good idea has a downside, however. Such an email address can get polluted over time and would start receiving 'regular' spam and viruses. One countermeasure is to create unique bait email addresses for every new workstation image that is rolled out.

ANALYSING CAPTURED CONTENT

During a six-month period, our honeypot captured emails that originated from the internal network, and we were able to identify the offending machines from the originating source IP. Based on the content of the captured email messages, we classified them as shown in Figure 2.

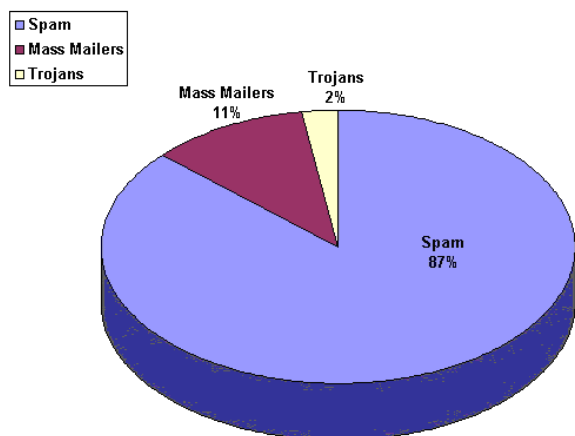


Figure 2: Breakdown of captured email content.

Only 11% of the captured email included executable attachments. Just 2% were mails containing infection notifications or captured cached passwords that were relayed to the trojan author. The other 87% was spam. A

high percentage of this content was image spam and PDF spam – techniques that spammers have used effectively to subvert traditional detection by anti-spam vendors.

As we scanned the captured emails for malicious code, we found the first half of 2007 belonged predominantly to W32/Stration [7] (a.k.a W32/Warezov), and W32/Nuwar [8] (a.k.a. Storm), variants. With both families improvising with every new variant, these mass-mailers have caused scores of mini-outbreaks that have overloaded legacy mail servers.

SPREAD THE WORD, NOT THE VIRUS!

We've highlighted the benefits of redirecting vs. blocking internal SMTP traffic in this article. Any mass-mailer or spam-like activity on the internal network can be detected against proactively via this setup, which goes a long way toward containing and isolating the source of infection or attack. This solution is scalable, cost effective, and relatively easy to implement.

With malware authors putting so much thought and creativity into keeping the spam juggernaut rolling, it will be interesting to see how well we can combat these threats. For every countermeasure there is a counter-countermeasure. We lose only if we stand still. And what would be the fun in that?

REFERENCES

- [1] http://www.theregister.co.uk/2006/12/22/ordb_shutdown/.
- [2] McAfee Virus Information Library. HTool-MPack: http://vil.nai.com/vil/content/v_142501.htm.
- [3] iDefense MailPot: http://labs.iddefense.com/software/malcode.php#more_malcode+analysis+pack.
- [4] Wikipedia DMZ: [http://en.wikipedia.org/wiki/Demilitarized_zone_\(computing\)](http://en.wikipedia.org/wiki/Demilitarized_zone_(computing)).
- [5] McAfee Virus Information Library. W32/Sober@MM: http://vil.nai.com/vil/content/v_137072.htm.
- [6] The Internet Engineering Task Force message submission: <http://www.ietf.org/rfc/rfc2476.txt>.
- [7] McAfee Virus Information Library. W32/Stration@MM: http://vil.nai.com/vil/content/v_140419.htm.
- [8] McAfee Virus Information Library. W32/Nuwar@MM: http://vil.nai.com/vil/content/v_140835.htm.