ROBERT HASKINS

ISPadmin: anti-spam roundup

Robert Haskins has been a UNIX system administrator since graduating from the University of Maine with a B.A. in computer science. Robert is employed by Shentel, a fast-growing network services provider based in Edinburg, Virginia. He is lead author of *Slamming Spam: A Guide for System Administrators* (Addison-Wesley, 2004).

haskins@usenix.org

IN THIS EDITION OF ISPADMIN, I TAKE a look at a few relatively current events in the area of anti-spam. Although not directly service-provider related, spam is certainly an area that is near and dear to the network service provider's operations.

The areas that I look at in this column include:

- Reputation
- Image spam
- SpamHINTS
- Sophos "top 12"
- Dlink SECURESPOT

Reputation

The area of reputation as it applies to anti-spam continues to mature. Although "reputation" can mean many different things to different people, I use "reputation" here to mean the likelihood a particular IP has emitted spam in the past. A simple reputation service would be a DNS blacklist service such as Spamhaus SBL [1] or DSBL [2]. However, these systems are simple "on" or "off" reputation systems; the systems described below assign a likelihood-of-spam probability in much the same way a personal credit scoring system assigns probabilities.

OPEN SOURCE SOFTWARE

Unfortunately, there has been little progress in the open source arena on anti-spam reputation-based systems. GOSSiP [3] appears to be at a standstill. I guess this means that a viable reputation service needs to have a commercial entity behind it (Trend Micro, Symantec, etc.). However, similar open source–like data-sharing anti-spam schemes currently exist in the form of the Distributed Checksum Clearinghouse (DCC) [4] and Vipul's Razor (written by Vipul Prakash, a founder of Cloudmark). Perhaps some would consider DCC commercial in the sense that it is backed by the commercial entity Rhyolite Software, but as far as I am aware, it doesn't generate revenue directly from running the DCC service.

COMMERCIAL SOLUTIONS

In the commercial arena, the IP reputation-based solutions keep on coming. One of the newer ones is from Simplicita Software [5]. The Simplicita Reputation Knowledge Server (RKS) can take IP feeds from many different sources (MTA logs, firewalls, DNS-based blacklists, etc.) and allow network operators to block inbound messages from IPs determined to be bad. The RKS solution allows easy manipulation of these lists, adding and expiring IPs using many different schemes.

Another approach to the IP reputation is the Symantec 8100 series appliance [6], formerly known as Turntide. (Disclaimer: My employer is a Symantec customer and user of the 8100 and SBAS products mentioned here.) In the most common implementations, you simply put the device “in line” in front of your mail transfer agents (MTAs), behind either switch(es) or router(s). This solution works at the TCP level to block SMTP connections from IPs that have sent spam messages to your email infrastructure. The most egregious spamming IPs are “ratcheted down” in the rate at which they can send messages to the inbound or outbound MTAs sitting behind the 8100 appliance.

The good news about the 8100 device is that deploying them can significantly reduce the need for additional MTAs in your email architecture and they can reduce the number of spam messages hitting your MTA. However, the appliance isn’t a perfect solution, as it does not eliminate the need for additional filtering capability, requiring an MTA-level anti-spam filter such as Symantec Brightmail AntiSpam [7].

The other bit of bad news is that the appliance often blocks legitimate email from hosted domain forwarders, resulting in complaints from customers who have email forwarded to local accounts that come from such domain email forwarders. Of course, the 8100 device allows the administrator to whitelist IP address space, but the problem is where to draw the line. If you whitelist too many senders, then there isn’t much point in having the device to begin with.

The 8100 device does allow placement of IPs into a number of “classifications,” ranging from essentially unlimited access to a very slow rate of accepting SMTP connections. One solution is to “lock” the domain-hosting IP into a specific classification, not allowing unrestricted access, but slowing it down somewhat. It’s not a perfect solution, but it’s better than the alternatives.

One resource I have found to be particularly useful in managing the 8100 devices is Ironport’s Senderbase [8]. When one needs to whitelist an email service provider and the provider is unwilling or unable to tell you what their outbound servers are, the Senderbase data can give you a good idea of what IPs to start with. Although not perfect, it gives one a place to start.

Image Spam

One of the newer spamming techniques is the use of embedded images in email messages. This is a particularly difficult method of spam to deal with. (I personally have been receiving image spam for at least two years now.) Some commercial anti-spam vendors recommend not allowing images as a solution to this problem. This isn’t very practical, as this would only work for those that don’t regularly receive images. I would guess that isn’t very many users. Barracuda Networks’ [9] solution is to perform optical character recognition on image attachments [10]. This sounds interesting, but I wonder how well it works. Pretty soon, we’ll be seeing watermarks in spam images to throw off the automated detection of spam via OCR!

SpamHINTS

SpamHINTS [11] is a research project by Richard Clayton at the University of Cambridge. His approach is to look at network traffic patterns and glean spamming IPs from the changes in patterns over time. This strategy will be very interesting if it works. The problem I see with this approach is that if the spammers can make their traffic look just like legitimate traffic, then there won't be anything to find. I suspect that a traffic-based approach will work for the high-volume spammers who send their junk via a relatively small number of IP addresses, but detecting small volumes spread out over widely disparate networks will prove difficult to identify using this method.

Sophos "Top 12"

I cringe every time I see a lot of press coverage of the "biggest spammer" lists such as that recently published by Sophos [12]. It's not that I have anything against such lists; but I have reservations about how accurate the data really is and what it really means. But first, does it really matter that North America allegedly originates 23.1% of spam? I think that the SpamCop stats [13], which show spammers by net blocks, are much more useful. With per-network spam information, we (Internet users) can put pressure on the egregious spammers with the SpamCop information. Back on the geographic lists, do we really care what country (or continent, for that matter) originates the most spam? What are we going to do, complain to George Bush because the United States originates the most spam?

Regarding the accuracy of the data, I suspect that the geolocation data has gotten better over time, but I still wonder how accurate it really is, with VPNs, inexpensive bandwidth, and all the other little details that make geolocation difficult.

Dlink SECURESPOT

The Dlink SECURESPOT [14] takes the security appliance to the extreme, for the consumer (SOHO) market. In a box half the size of a deck of cards, the device performs a whole host of security-related functions, including:

- Parental control
- Pop-up blocking
- Virus protection
- Spam blocking
- Spyware protection
- Identity protection
- Firewall protection
- Network reporting

SECURESPOT was designed by Bsecure Technologies [15], who handle the "service" side of the product, updating firmware and maintaining the lists of "bad guys" to block as part of the solution. Like any new product, it will take time for the support issues to be straightened out. I would imagine that this product would take some tweaking to work within the existing user's PC environment, given software firewalls, anti-spam, anti-virus, and other security-related software already running on the PC.

Of course, it remains to be seen how well this device works when compared to software equivalents such as Symantec's Norton Internet Security

[16]. The best design would be to integrate this functionality into the existing SOHO firewall/router and not have an extra box, but I suppose you have to start somewhere. I suspect that if this product takes off, Dlink will probably integrate the SECURESPOT functionality into some of their other products.

REFERENCES

[1] Spamhaus SBL: http://www.spamhaus.org/sbl/index.lasso.

[2] DSBL: http://dsbl.org/main.

[3] GOSSiP: http://gossip-project.sourceforge.net/.

[4] DCC: http://www.rhyolite.com/anti-spam/dcc/.

[5] Simplicita: http://www.simplicita.com/.

[6] Symantec Mail Security 8100: http://www.symantec.com/Products/enterprise?c=prodinfo&refId=852.

[7] Symantec Brightmail AntiSpam: http://www.symantec.com/Products/enterprise?c=prodinfo&refId=835&cid=1008.

[8] Ironport's database of sending IPs: http://www.senderbase.org/.

[9] Barracuda Networks: http://www.barracudanetworks.com.

[10] http://www.networkworld.com/news/2006/071906-barracuda.html?fsrc=rss-security.

[11] SpamHINTS: http://www.spamhints.org/.

[12] Sophos top 12: http://www.sophos.com/pressoffice/news/articles/2006/04/dirtydozapr06.html.

[13] SpamCop stats: http://www.spamcop.net/spamstats.shtml.

[14] Dlink SECURESPOT: http://www.dlink.com/products/?sec=0&pid=486.

[15] Bsecure Technologies: http://www.bsecure.com/.

[16] Symantec Norton Internet Security: http://www.symantec.com/home_homeoffice/products/overview.jsp?pcid=is&pvid=nis2006.