Fighting Spam and Winning

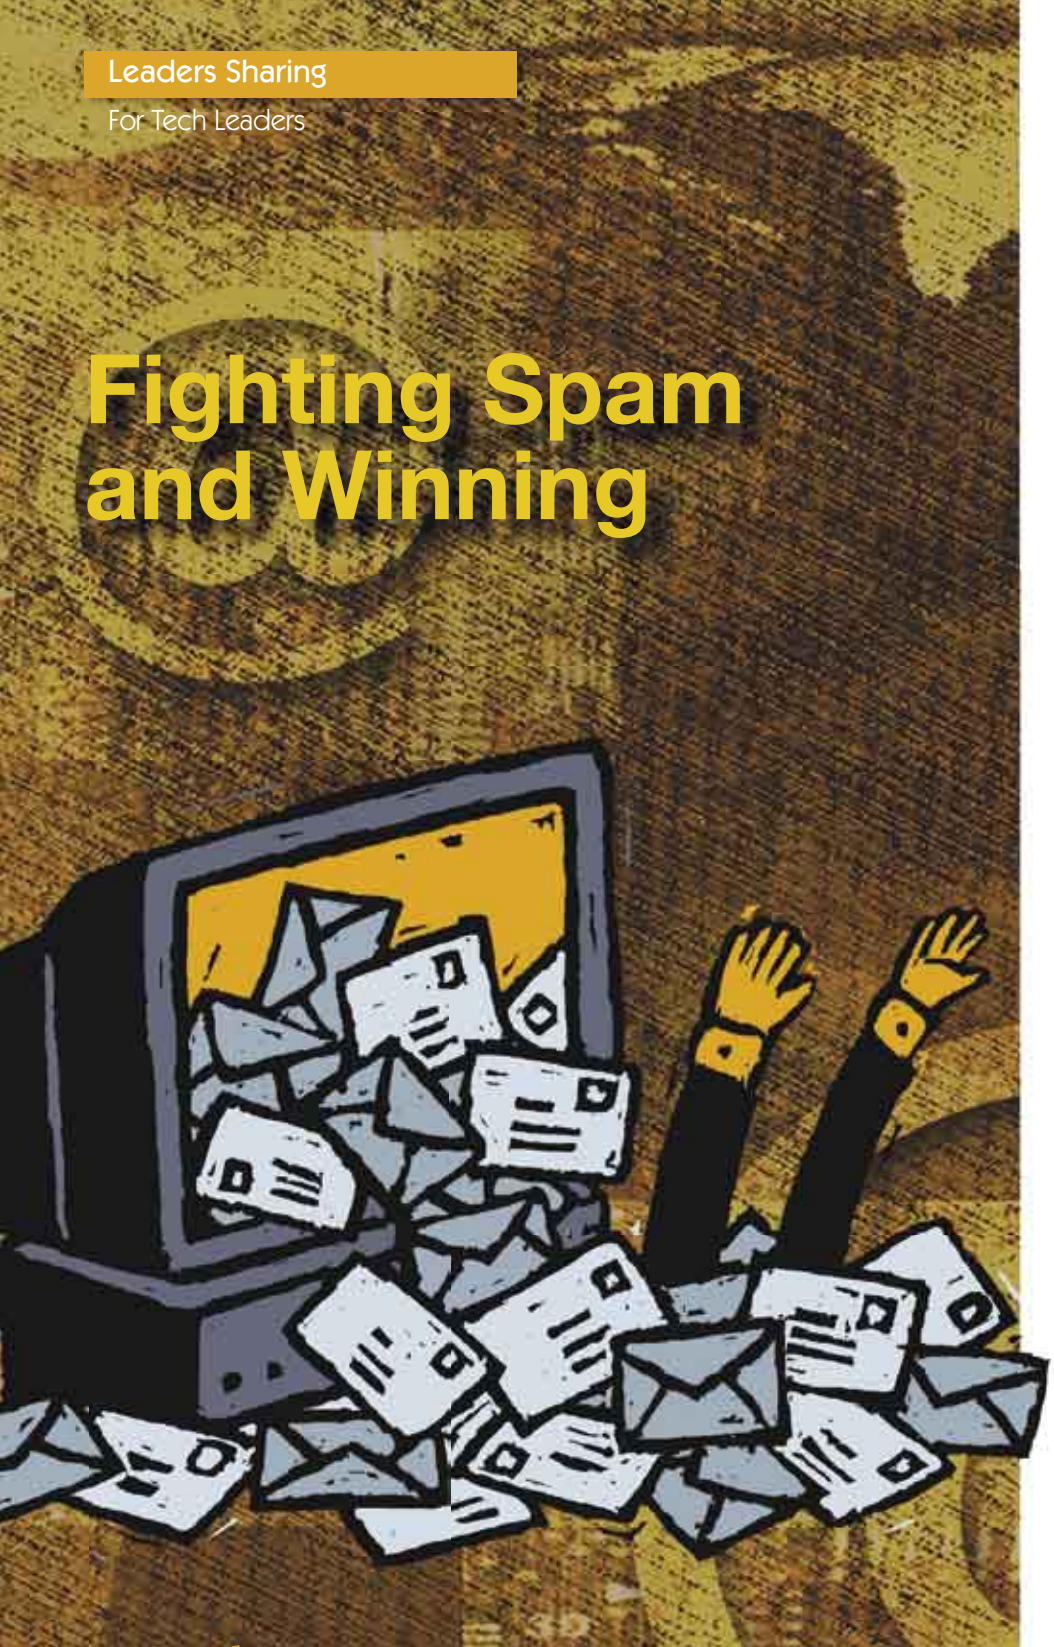

By Cornelius B. DuBose

Subject: Filtering unsolicited e-mail messages (spam)

Technology: E-mail

Standards: *NETS•A* I (http://www.iste.org/standards/)

E-mail is the lifeblood of communications for our district—Lake Forest (Illinois) School District 67. Teachers use it to communicate with parents. Administrators use it to communicate with faculty. So the issue of how to prevent unsolicited commercial e-mail or spam was serious. As is the case in the corporate world and in our homes, we noticed a substantial increase in the amount of spam our staff was being forced to deal with. Daily, staff members' inboxes were clogged with unwanted messages. Many complained about e-mail related to topics inappropriate for a school environment. Time was wasted deleting these messages, and in the deletion process, legitimate e-mails were sometimes accidentally erased. Something had to be done to put a stop to this.

Our goal was to put a system into place that would prevent unwanted and inappropriate messages from clogging user's inboxes and to maximize our network resources by keeping spam off our e-mail server.

Spam Basics

When a user sends e-mail, the message travels from the sender's computer to the e-mail server. The e-mail server then delivers the message to the recipient's e-mail server. Often, the message will be passed along from one e-mail server to another until it reaches its final destination. The protocol that controls this movement is the Simple Mail Transfer Protocol (SMTP). SMTP requires that every message contain certain standardized identifying information known as the *e-mail header*. The header contains a record of a message's journey across the Internet, including a record of each server that handled the message. It also includes:

- message recipient
- sender

- subject of the message
- message ID
- date sent
- reply address

The header is, as the name suggests, at the beginning of each message. E-mail client software interprets the headers to display appropriate information about the message. Most e-mail clients hide the headers by default.

Here is a sample e-mail header:

```
Received: from mail.somewhere.com
by mail.lfelem.lfc.edu with SMTP
(QuickMail Pro Server for Mac
3.0.1); 24-Nov-2003 08:32:48 -
0600
Date: 24 Nov 2003 08:32:48 -0600
Message-ID: <1005023917@lfelem.l
fc.edu>
From: Sample Spammer <sspam@
somewhere.com >
Subject: RE: Information
Requested
To: Cornelius DuBose
<cdubose@lfelem.lfc.edu>
X-Mailer: QuickMail Pro 3.0 (Mac)
X-Priority: 3
MIME-Version: 1.0
Reply-To: Your Friend
<reply@lfelem.lfc.edu>
Content-Transfer-Encoding:
quoted-printable
Content-Type: text/plain;
charset="US-Ascii"
```

The header is important in the battle against spam not only because it tells where e-mails have been but also because spammers manipulate the headers to hide their identities. One technique is to include incorrect reply addresses, called *spoofing*. Another is to leave the reply header out altogether. Still another is to have the "from" line include the e-mail address of the recipient or someone else in the same organization.

The so-called CAN-SPAM Act, signed into law late in 2003, prohibits the use of deceptive e-mail subject lines, headers, and return addresses. However, spammers can still legally send spam as long as they follow those rules. Worse, the law only affects spammers who reside in the United States. The law doesn't apply to spam coming from outside the country.

Open or third-party relays are another tool spammers use because these servers will allow anyone to send mail through them. Any computer can quickly become an e-mail server. All that is needed is a computer connected to the Internet (even temporarily through a dial-up connection) and e-mail server software, which can be freely downloaded. A server setup in this manner is usually unregistered and therefore anonymous. Most legitimate e-mail servers are registered by their mail exchange (MX) record maintained in Domain Name Servers (DNS) across the Internet. Unregistered servers work fine for sending mail, and that's all a spammer cares about.

Spam-Fighting Tools

Four general approaches exist for filtering spam:

- external services
- client products
- server products
- gateway products

Spam filtering services, as the name implies, are outsourcers that detect spam for a fee. These services work by having any e-mail sent to your district rerouted through their servers. The e-mail is then examined for spam, and an action (determined by the district) is initiated. A spam action can include: adding identifying text to the subject of the message (tagging), putting a header or footer in the message, sending the e-mail to a specific address, or denying delivery altogether. After the spam test, e-mails are routed back to the district's e-mail server. Although there is some lag, the overall delay is minor. To accomplish all of this routing, changes need to be made to the e-mail server's DNS entry. Using such a service requires no new investment in hardware and limits the time burden on district staff.

These services claim very high spam capture rates, some as high as 99%. No new hardware investment is required, district staff time burdens are minor, and spam can be prevented from ever reaching your e-mail server.

Client spam products reside on individual computers throughout the district. These programs screen e-mails as they are delivered to the individual user's inbox. Most of these programs come with a set of built-in filtering rules that attempt to determine what is spam and what is legitimate e-mail. In addition, users can create customized rules. As with the filtering services, users can determine if spam should be automatically deleted, flagged, or placed in a separate folder. This type of spam protection enables individual staff members to be in control of what is considered spam and to determine what spam actions to take.

Although this method can be extremely effective, it does require the most effort and expertise from users and does nothing to keep spam off the e-mail server.

Server filter products deal with spam before it gets to the e-mail server. This method of spam detection requires that an e-mail filtering server be installed between the Internet router (gateway) and the mail server. All mail destined for the mail server is routed to the filtering server instead. The filtering server determines if the arriving mail is spam. The filtering server then passes the mail to the district e-mail server. As with the other methods, various actions can be implemented once spam is detected. A DNS reconfiguration must occur for this method to function.

Message lag from the rerouting should be unnoticable. This method can involve significant financial in-

Thank You!

ISTE members are wonderful and generous. The *L&L* staff would like to especially thank the members who volunteered time from their busy NECC schedules to meet with us. Look for members like this month's subject, Daphne Griffin, in the member profile section each issue.

We'd also like to express our gratitude to **ISTE 100 member Intel** and its Innovation in Education program for providing Intel digital microscopes as gifts for the participants.

If you'd like to be the subject of a member profile this year, please contact us at **letters@iste.org.**

vestment. In addition to the server software, a spam detection service must be purchased. Server filter products are effective in reducing spam and keeping it from reaching your e-mail server. An additional bonus is the potential of adding reporting capabilities such as being able to audit e-mail use in the district. Staff workload is very low because the spam detection on the server is automatically updated.

Gateway spam filtering occurs at the district's Internet gateway (where the local network connects to the Internet) and may be performed by a firewall. A firewall is a security device through which all traffic of any kind (e-mail, Web, FTP, and so on) entering or leaving the district from the Internet must pass. As e-mail passes through the firewall, the firewall determines whether it is spam. Once that determination is made, the firewall can be programmed to handle the spam. Not all firewalls offer this capability, and often those that do are limited to specific software vendors.

The advantages are that no DNS reconfiguration is needed, spam is reduced and doesn't reach the e-mail server, and you may not need to buy a new firewall. It does require a very active role from district technical staff to implement.

Our Solution

After considering the various tools outlined above, our district determined that the gateway solution suited our needs the best because:

- There would be no DNS reconfiguration, which the server or outsourcing options would require.
- Control of the district's mail would remain within the district.
- The district would have no recurring costs, as there would be with the outsourcing option.
- Finally, no new hardware would need to be purchased, because a firewall was already in place.

Our firewall functions only with a product called Spamscreen, provided from the manufacturer, Watchguard. Spamscreen provides us the ability to identify spam using Realtime Blackhole Lists (RBL), filtering rules, or both. An RBL is a list of spam-sending domains stored in a remote DNS server. When e-mail arrives, its header is examined to determine its domain of origin. That domain is then checked against the RBL server list. If the domain is found to be in the RBL, the message is designated as spam. Rule filtering scans a message header to see if the header meets some criteria that would define it as spam. The criteria might be that the subject line contains the term *mortgage* (or some more provocative term). It might be that the reply header is not properly configured or is missing. Or it might be that the message comes from a particular user at a domain. As with RBLs, an entire domain can be identified as spam. Rule filtering enables a more customized approach to spam detection. However, it requires more effort than the use of a RBL.

In addition to the RBL and the filtering rule tests, Spamscreen also checks to see if the mail comes from a registered e-mail server. This is done by checking the domain listed in the header against a DNS server to ensure that the mail server attempting to send to our system has a valid MX record. If the message does not, it is deemed spam.

No system is perfect. If a system is set up stringently enough to effectively stop spam, it will more than likely stop some legitimate e-mail as well. A legitimate e-mail that has mistakenly been identified as spam is termed a false positive. Conversely, an e-mail that should be identified as spam but makes its way to the user is called a false negative.

To determine whether a message is spam, Spamscreen assigns every message a point value—its spam score.

If the assigned score is greater than the threshold weight the message is designated spam. Out of the box, Spamscreen sets the default threshold weight to 1,999 points. We did not change this value. Any message failing an RBL or MX record test automatically gets a spam score of 2,000 points.

Filtering points can be relatively flexible. For example, you might assign the word *sex* a score of 2,500 or more. Another rule might assign negative 1,200 points for messages originating from a listserv at Yahoo groups. In that way, a valid message with the subject "sex education" would not be considered spam because its final score would be 1,300 points.

The Battle Begins

Loading the software took us only a few minutes. Spamscreen comes configured with a host of filtering rules and a default RBL that can be selected. It provided three options for dealing with identified spam. We could add a tag to the subject line, delete the message, or do nothing at all. We set the spam action to tag and elected to have the heading "[SPAM]" added to the beginning of the subject line of any message identified as spam. The message, however, would still be delivered to its intended recipient. The recipient could determine whether messages identified as spam were actually false positives.

We put the system into place on a Monday afternoon. I remember the anticipation I had the next morning when I opened my e-mail. I expected to see the normal 50 or so spam messages I had been receiving all tagged. Unfortunately, that was not the case. Only about 20 or so of the messages were tagged. The rest were sitting in my mailbox as usual. We definitely had more work to do.

We began with spam addresses the district technician and I had gathered. Most of these addresses were spam domains, addresses from which we would receive no messages except spam. We began writing filter rules to block these entire domains. Spamscreen filter rules are written in Perl Compatible Regular Expressions (PCRE). Regular expressions search for patterns within text using a specific syntax. For example the rule to block a domain such as spamsenders.com is written in PCRE using the syntax: ^((?i)From):\s(?i).*\@\S*spamsenders\.com. We spent a great deal of time learning the syntax of PCRE and writing new rules to block domains. We also spent time searching for and finding additional RBL lists to use.

Fortunately, Spamscreen allows for multiple rules to be imported into the software. An imported rule contains the rule, its name, and its spam weight. We developed an Excel spreadsheet to facilitate the process of writing these rules. Domains could be entered into the spreadsheet and importable rules created. It then became easy to go through our messages, create a list of domains, copy and paste the list into the spreadsheet, and generate rules to import. We entered literally hundreds of domain rules during the next few days, and more and more spam was tagged.

It was time to enlist the aid of the staff in the fight against spam. We began by creating an e-mail account that could be used to report spam and to report false positives. We sent an e-mail message to all staff asking for their assistance in both eradicating spam and assuring that the system was working properly. We asked them to forward the e-mail address, singularly or as a list, of any spam messages they received. They were also asked to report any message that was tagged as spam that wasn't. We also explained that the ultimate goal of the system was to deny spam. We hoped this would encourage staff to report false

positives. Finally, we suggested that staff create a folder in their e-mail program to hold spam and a filter rule that routed any message that included [SPAM] in the subject line into this folder. In this manner, any identified spam messages would not appear in staff members inboxes, and false positives could be discovered before deletion.

As staff began to send in spam addresses, the percentage of tagged spam rose. To increase this percentage further, we began to analyze the spam we received, looking for patterns that would allow us to write generic rules that would catch multiple spam domains instead of single domain addresses. We noticed several patterns:

- using numbered servers at the same domain to send the same spam message multiple times
- using a top-level domain ending in .biz
- words such as products, marketing, rewards, discounts, and so on appearing in domain names

We developed rules to designate all messages that fit these patterns as spam. Because school districts quite often have numbers in their addresses, we had to write a counter (negative) rule to allow messages from school districts to pass through as legitimate e-mail.

As we caught more spam, the number of false positives reported also increased. Fortunately, Spamscreen has the capability of excluding either entire domains or specific e-mail addresses from a spam check. We began by entering several of the major domains (i.e., Comcast, AOL, Hotmail) into this list. This however, allowed spam from these domains to pass through the system unchecked. To prevent this, we removed these domains from the exceptions lists and used a rules-based approach to deal with these false positives. This also gave us the ability to write rules that could catch spam from specific e-mail addresses at these major domains.

Where We Stand

As of this writing, our system has been in place for about six months. We believe at this point we are tagging about 95% of the spam entering the district. In the first couple weeks, we would enter 50 or 60 spam rules a day, now we enter two or three new addresses a day.

It is essential to have the cooperation of the staff to make a system such as this successful, and our staff has been cooperative. They continue to send addresses to block as well as false positive notifications. Both categories of notification have drastically reduced over time. Many staff members have made positive comments about the system. Some have even thanked us for installing it.

One of the things we discovered in this process, however, is that what is considered spam by one, is not considered spam by another. For example, a staff member wanted *Sports Illustrated* blocked, but the Physical Education department wanted these messages to pass through. Individual client filters can be used in cases such as this. Each Friday we distribute an e-mail "how-to" newsletter to all staff. We instruct staff on using personal spam filters. We also touch on the topic in every e-mail-related workshop we offer.

We have not yet reached our final goal of freeing up network resources, server time, and server space by denying spam altogether. We will begin denying spam during the summer break. Our test, during spring break 2004, resulted in only a single complaint. For the time being, however, we are extremely pleased with our system. Its benefits are apparent every morning when staff members open their inboxes. They aren't clogged with unwanted and often inappropriate messages.

We learned several valuable lessons during the course of implementing the system:

- The selected solution should match the available time and technical level of the support staff. Some solutions, such as ours, can require large commitments in staffing, especially at first.
- Keep all staff informed and updated. You will need their help and, in some instances, their understanding to make this work.
- Provide a clear definition of what will and won't be considered spam.
- Set-up processes for general staff to participate. No one wants spam, and this is one area where people will gladly lend their assistance.
- Don't begin by denying spam. Despite the temptation to simply turn off the switch on day one, you really need to know and understand the effect the spam filter will have on the ability of your staff to receive legitimate messages.

Resources

Brightmail: http://www.brightmail.com/
CyberLink: http://www.cyberlynk.net/services_solutions/dsp_spam.cfm
SpamSoap: http://www.spamsoap.com
SurfControl E-mail Filter: http://www.surfcontrol.com
WatchGuard Technologies: http://www.watchguard.com

Cornelius DuBose has been an educator for more than 31 years. A graduate of North Park College, he began his career as a sixth-grade teacher. He received an MS in educational technology from National-Louis University. He has taught technology-related courses for Northwestern University's Center for Talent Development; DePaul University's School of Computer Science, Telecommunications, and Information Systems; and National-Louis University's educational technology program. Currently, he is the director of technology for Lake Forest School District 67, in Lake Forest, Illinois.

http://www.iste.org/LL