

Trends in Spam

Unsolicited bulk email, phishing campaigns, and other forms of spam attacks constitute big business for e-marketers, hackers, and criminal organizations. Any email administrator knows that once a new email server has its IP address published in an Internet DNS MX record, it becomes an immediate target for spammers. Moreover, the recent legal assault on Spamhaus (<http://www.spamhaus.org>), a global non-profit organization hosting an external block list of known bad email servers, is just the latest indication that spam is serious business.

In protecting email infrastructure against spam, there are three trends in particular that email administrators should be aware of:

- Increased volume of spam on the Internet
- The changing nature of reputation checks
- Image spam and its effect on email infrastructure

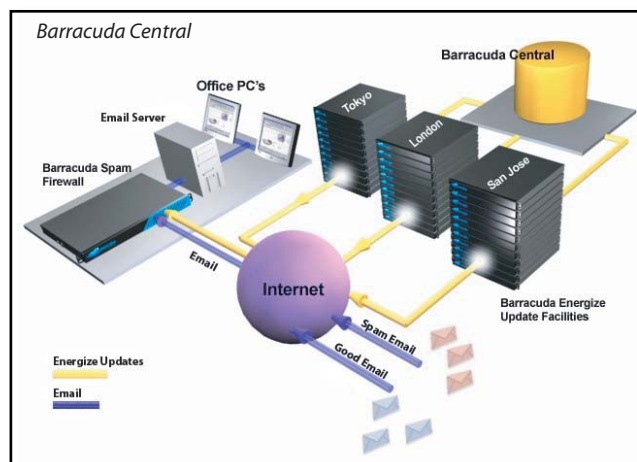
By carefully planning with these trends in mind, email administrators can keep their email infrastructure protected as the changing landscape on the Internet continues to evolve.

Increased Volume of Spam on the Internet

Barracuda Central is an advanced technology operations center where engineers continuously monitor spam trends as they emerge and collect aggregate statistics of blocked email from over 35,000 Barracuda Spam Firewalls deployed worldwide. Based on recent analysis at Barracuda Central, it is estimated that spam constitutes over 85 percent of all email traffic on the Internet.

While for years it has been common for the largest Internet Service Providers (ISPs) to experience levels of spam at 99 percent of total email volume, what is more surprising is the extent to which spam has hit smaller ISPs, enterprises, educational institutions, government organizations, and even small businesses. While in aggregate, the overall level of spam on the Internet has increased about 15 percent in the last six months (May 2006-November 2006), the increase in spam has had a far greater impact on smaller organizations. Barracuda Central has observed many systems reporting back 100 to 200 percent increases in overall email traffic levels, even when adjusted for increases in the number of valid email recipients.

In addition to increases in the overall levels of spam, Barracuda Central has also observed traffic patterns “bursts” that are indicative of spammers concentrating distributed attacks at discreet times. While it is typical that a large organization might receive between 25 to 50 percent more traffic during a burst, Barracuda Central has frequently observed volumes of spam at an order of magnitude or more above normal during peak times. It is believed that spammers concentrate their traffic during peak times in an effort to add latency to delivery of good email while spam filters scan large volumes of spam. By adding this latency, spammers hope to coerce administrators to temporarily disable or bypass their spam filtering solutions in order to deliver email to their users on a more timely basis.



Evolving Reputation Analysis Strategies

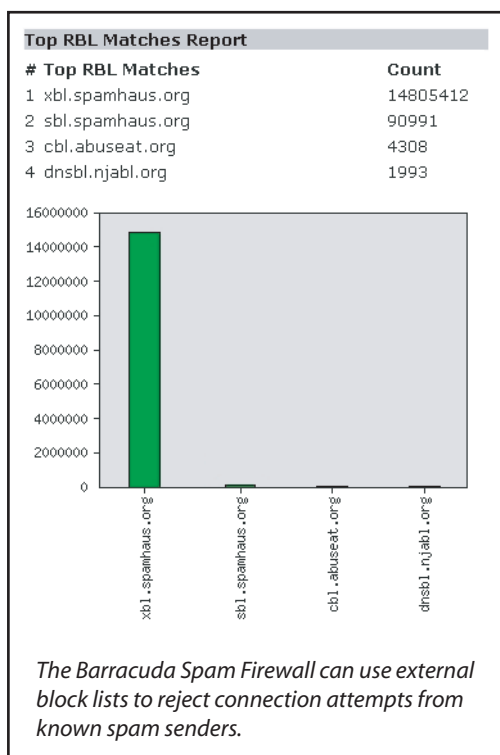
Reputation analysis involves identifying bad email senders based on their history, their current behavior, and their intent. Of many defense strategies used to identify spam, reputation is one particular category that has been evolving over time. Because the Barracuda Spam Firewall utilizes multiple layers of defense, it has maintained its efficacy in fighting spam. However, because of the changing nature of reputation analysis, there is a growing reliance on layers that require additional sophistication and processing power.

RELEASE 1
NOVEMBER 2006

Historical Behavior and Block Lists

Reputation analysis based on historical behavior has been a mainstay of spam protection. Barracuda Central and other external organizations, such as Spamhaus, maintain databases, or “block lists,” of known IP addresses of both email servers and relays used by spammers. Using these block lists, spam protection solutions, such as the Barracuda Spam Firewall, can reject connection attempts from known senders without actually receiving the email, thus saving processing power. Historically, typical Barracuda Spam Firewalls in the field blocked between 50 and 60 percent of all email connection attempts based on simple block list checks.

Because of the popularity and efficacy of this approach to block spam, spammers use multiple techniques to circumvent block lists. The simplest technique to circumvent block lists is simply to change ISPs or to rotate the IP addresses of their email servers. While it’s relatively straightforward to catch the spammer on their new IP address, this very simple technique does provide short-term windows to send out spam, and given the money involved in the business, it is often worth doing.



A more sophisticated technique to circumvent block lists is to use a botnet – or a network of computers (“zombies”) infected with malware. Spammers can use zombies, most of which do not yet have a historical reputation for spamming, to send email on their behalf. Moreover, by distributing the bad email across thousands of zombies’ IP addresses, they can stay under the radar for longer periods of time. Recently Barracuda Central has observed a new trend of “pulsing zombies” in which sophisticated spammers send out a burst of email through a particular zombie, allow it to stay dormant for a time period in an attempt to “expire” its IP address from block lists, and then resume its use.

Barracuda Central estimates that the combination of these and other approaches have eroded some of the efficacy of block lists by as much as 10 to 20 percent in some customer deployments. While these messages are generally blocked by subsequent protection layers, these layers generally require actually accepting and scanning the message, thus requiring more processing power.

Current Behavior and Rate Controls

Establishing reputation of new spammers has always been an issue, so it is important to also establish reputation based on current behavior. Blocking connections from email servers and relays attempting to send too many emails with invalid or non-existent recipients is an excellent means of profiling directory harvest attacks (DHAs). Blocking connections from untrusted email servers or relays that exceed rate control limits by sending too many total messages in a given time period is an excellent means providing resilience against denial-of-service (DoS) attacks. Recipient verification and rate controls are just two examples of creating reputation profiles based on behavior.

That said, zombies have created an additional challenge with profiling current behavior. Because of the scale of the botnets, sophisticated spammers ensure that they distribute email to the same email domains across different zombies to stay below rate control thresholds for any particular recipient’s email infrastructure. And, counter to the “pulsing zombie” strategy, some even trickle the email from each zombie in an attempt to stay off the radar screen of rate controls and block lists altogether.

While Barracuda Central continues to observe high levels of spam blocking due to rate controls and recipient verification, the efficacy of these approaches is likely a function of increasing attack activity on the Internet. Through botnets, more sophisticated spammers appear to be intentionally slipping email past rate controls in particular.

Intent Analysis

The motivation behind most spam is to get the recipient to do something, such as to visit a Web site, to call a phone number, or to buy a stock. This motivation is called the “intent” of the email, and inspecting email for these traits is called “intent analysis.” By far, the most common intent of spam email is to get a user to click on a Web link or URL.

Even if senders try to mask their reputation through botnets or new IP addresses, they ultimately still need to drive a user to a particular Web site. Barracuda Central maintains a database of known Web site addresses used by spammers, and the Barracuda Spam Firewall can block messages based on the Web site addresses embedded in these emails.

While most of the sophistication behind maintaining up-to-date intent analysis databases resides at Barracuda Central, certain spam trends have required some additional processing on the Barracuda Spam Firewall itself. For example, one trend has been the use of free hosting providers, such as Yahoo! Geocities, to redirect users to spam Web sites. The suspect email may contain a link to a free Web site that simply redirects users to a spammer’s Web site. The advantage to this approach is that spammers can rotate the Web site addresses used in their emails very quickly at no cost. Through multi-level intent analysis, the Barracuda Spam Firewall can follow the links and attempt to determine whether links to free sites ultimately resolve to spammer Web sites.

Intent analysis is an extremely effective means of blocking spam and it has increased its efficacy with the commensurate decrease in efficacy of block lists. Barracuda Central has observed that intent analysis typically blocks between 10 and 20 percent of email on Barracuda Spam Firewalls in the field.

While these effects of these trends do not generally decrease spam accuracy, they can increase the processing load on Barracuda Spam Firewalls in protecting email infrastructure.

To bypass traditional and effective text scanning methods, including intent analysis, Barracuda Central has observed an increasing trend in the use of inline images in HTML emails and image attachments. For typical email recipients, image spam can represent 15 to 25 percent of all spam received.

Barracuda Networks has developed several methods to effectively block image spam including fingerprint analysis, optical character recognition (OCR), and animated GIF analysis.

The screenshot displays the Barracuda Spam Firewall's log display interface. At the top, there's a navigation bar with tabs for 'Status', 'Message Log', 'Spam Scoring', and 'Virus Checking'. Below this, a search filter is applied: 'Reason Contains: Intent'. The main area shows a table of blocked messages with columns for 'Time Received', 'Subject', 'Action', 'Reason', and 'Source IP'. The table lists several blocked messages, all with the action 'Blocked' and reasons related to 'Intent' from various domains like 'http://aplicus.com' and 'http://www.geocities.com'. The interface also includes a 'Log Display' section with a date range from 12/14/2006 15:05 to 12/14/2006 15:05, and a 'Current Message Log Count: 498459'.

The Barracuda Spam Firewall can block messages based on Web site addresses embedded in emails.

For all message body parts, including images, the Barracuda Spam Firewall creates a fingerprint. Barracuda Central maintains a database of known image spam fingerprints that are used to block messages. Matching images against known fingerprints is effective in capturing frequently used images with minimal processing.

In addition, Barracuda Central has developed an OCR defense layer, which is effective against a variety of fingerprint obfuscation techniques, including breaking images into component parts, varying pixels, or changing encryption settings. To combat these techniques, the Barracuda Spam Firewall restitches images, normalizes them, and runs OCR to extract the text for further analysis. This technique has been extremely effective at capturing images with text.

Some spammers also use animated GIFs in an attempt to bypass OCR by building resultant images incrementally and then swapping out those images. The Barracuda Spam Firewall reconstructs frames from animated GIFs and analyzes resultant images and behavior. Using special rules tailored to animated GIFs, the Barracuda Spam Firewall also effectively blocks these new forms of spam.

While the Barracuda Spam Firewall has effectively blocked image spam, techniques such as OCR do require more processing power. Some of this processing power burden is offset by the savings in bypassing certain text scanning processes. Still, customers of Barracuda Spam Firewalls should expect increased processing demands on their systems as image spam and other new forms of spam emerge.

Sizing the Barracuda Spam Firewall

Even with increased processing demands with the changing email landscape and proliferation of image spam, most organizations that size their systems based on the recommended Active Email Users specification should continue to observe excellent performance out of their systems.

BARRACUDA SPAM FIREWALLS

**Capacity depends on environment and selected options*

Model	Active Email Users*	Peak Email Capacity
200	1-500	1 million messages / day
300	300-1,000	2 million messages / day
400	1,000-5,000	5 million messages / day
600	3,000-10,000	10 million messages / day
800	8,000-22,000	20 million messages / day
900	15,000-30,000	30 million messages / day

The Active Email Users specification is appropriate for most end user organizations with typical email scanning configuration options set. Organizations with high-volume ecommerce applications, unique Internet email traffic patterns, or a highly visible Web presence may require larger systems.

After applying the Active Email Users specification, organizations should then check the Peak Email Capacity specification that describes the resiliency of the connection management layers of the Barracuda Spam Firewall. Robust connection management layers are important not only for the highest volume ISPs but also for smaller organizations while under attack. In these situations, it is typical that over 99% of email connection attempts can be blocked through techniques such as rate controls, IP block lists, and recipient verification. While most organizations that conform to the Active Email Users specification rarely reach Peak Email Capacity numbers, organizations should remain aware of upper limits for resiliency against attacks.

While some Barracuda Spam Firewall customers have successfully exceeded the Active Email Users specification in the past, many may need to re-evaluate their system sizing based on changes in email mix. With greater overall email volume, fewer messages processed at earlier connection layers, and increased processing associated with image spam, those customers previously operating beyond published specifications may need to increase their system capacity.

For more information about sizing systems, contact your Barracuda Networks systems engineer.

About Barracuda Networks, Inc.

Barracuda Networks is the leading provider of enterprise-class application security appliances for comprehensive email, Internet and IM protection. Its products protect over 35,000 customers around the world, including Adaptec, Caltrans, CBS, Georgia Institute of Technology, IBM, NASA, Pizza Hut, Union Pacific Railroad Company, and the U.S. Treasury Department. The Barracuda Spam Firewall and Barracuda Spam Firewall - Outbound protect organizations against spam, viruses, and violations to e-mail security policy. The Barracuda Web Filter, formerly known as the Barracuda Spyware Firewall, offers comprehensive content filtering and complete network protection against spyware, malware and viruses. The Barracuda IM Firewall, is the only all in one gateway solution for IM traffic management and security. Barracuda Networks is a privately held company with headquarters in Mountain View, California. Barracuda Networks has offices in eight international locations and distributors in over 43 countries. More information is available at www.barracuda.com.



Barracuda Networks
 385 Ravendale Drive
 Mountain View, CA 94043
 United States
 +1 408.342.5400
www.barracuda.com
info@barracuda.com