How to Stop the Latest Internet Email Traffic Emergency

Spam "Bounce" Messages are Compromising Networks

A REPORT FROM

THE IRONPORT THREAT OPERATIONS CENTER

WITH A FOREWORD BY PETER CHRISTY

A MESSAGING NEWS PRESS PUBLICATION

Table of Contents

“The misdirected bounce problem . . . is a good example of the complexities of email control problems and the value in comprehensive solutions.”

PETER CHRISTY
Principal, Internet Research Group

Foreword

by Peter Christy

The control of spam and other forms of email abuse and fraud is a complicated and evolving topic. For most companies the goal of email operations is best practice competence, not competitive differentiation. In these cases, it makes more sense to count on a trusted vendor or service provider to deliver comprehensive email control and protection systems, rather than trying to create them locally. Unless email is a key competitive differentiator, creating these solutions internally is mostly an opportunity to fail rather than a means to succeed—the definition of something that is best delegated to outside experts. The issue of high-volume spam traffic caused by “misdirected bounces” discussed here is a great example of a problem that has to be solved, but one that you probably don’t want to have to solve yourself.

Modern spam methods can cause secondary traffic problems if the mail systems receiving the spam reject (or “bounce”) misaddressed messages—creating a onslaught that can be much more of a problem to the party receiving the bounces than the spam was in the first place. Mail protection solutions must anticipate this kind of second-order effect, and not just solve the simple parts of the problem.

Let’s briefly consider the history of email to understand how we got here. The Internet originated with a network research project to improve collaboration between computer science research groups. Network email was an obviously valuable tool within that small community. Email provides instantaneous, distance-independent and essentially free communication. As the Internet community grew, new users discovered these benefits. Additionally, as more and more people, businesses and organizations connected to the Internet the importance and value of email grew rapidly.

In the end, it isn’t surprising that peddlers, con artists and criminals also found value in free, instantaneous worldwide communications. But that

potential certainly wasn't in the minds of the creators of the early network and mail protocols. Sending huge volumes of mail for limited responses is a sensible business because it costs next to nothing to send those messages. Now that the issues of spam and fraud are evident, it is regrettable that the underlying protocols provide little help solving the problem (all of the information in an email message can be forged by the sender). This just reflects the amazing and unanticipated growth and value in today's network. Now we have to live with the consequences.

The misdirected bounce problem (and solution) described in the following report is a good example of the complexities of email control problems and the value in comprehensive solutions. When a spammer generates large volumes of low-value email, the first goal of an email administrator is to keep that undesired mail off of expensive servers and out of users' mailboxes. But the problems don't stop there. A lot of the spam generated doesn't have valid user addresses (e.g., is addressed to BadName@company.com). What should the receiving mail system do it that case? Standard mail etiquette says you should send a "bounce" message to the sender — so they know the address is invalid and remove it from their lists. Sending a lot of these misdirected bounce messages back to a spammer sounds like a reasonable form of retaliation, until you take into consideration the likelihood that the return address is forged.

We've all experienced messages in our inbox saying that someone you don't even know didn't get mail you sent. When, of course, you know that you didn't send the message in the first place. That is an example of this kind of bounce message to a forged return address. Let's call it spam "secondhand smoke." A spammer sent out messages forging your return address because they believed that would increase the probability of the message getting through (using the return address of one individual in a company to send spma to another for example). But, if the address is forged, the bounce is sent to an innocent party — the owner of the forged sender address. So more parties are impacted by the spam barrage then we might at first have suspected and we have a secondhand smoke problem to worry about as well.

How can this secondhand smoke be mitigated? The broad answer is comprehensive spam control solutions, but the details are of interest as well. Discontinuing the practice of sending any bounce messages would solve this problem, but at the cost of losing a key tool for email address management. It would be like "throwing out the baby with the bath water." Fortunately, we have other possible tools to apply — the same tools we use to keep spam out of the server and the inbox. To grossly oversimplify the solution, rather than bouncing messages due to invalid recipient address before doing spam filtering, perform spam filtering first — and use that analysis to condition whether or not a bounce message should be sent.

What lesson should an email manager learn from this interesting little story? Misdirected bounces can be a very important problem, but are unlikely to have been anticipated by most knowledgeable email administrators. This is not the first, nor will it be the last, such problem that will arise in the ongoing war against inappropriate use of email. Rather than hire an email team that has to know every nuance about email abuse, for most a smart alternative is to use solutions from a trusted expert vendor—who can provide protection not only from today's threats, but from those certain to evolve in the future.

Peter Christy
Principal, Internet Research Group

The sheer volume of bounce messages generated by spam has grown to the point where it consumes an estimated (US) $5 billion per year in IT resources.

Introduction

Bounce messages are notifications of undeliverable email messages that are returned to their sender. When a receiving mail server gets a message with an undeliverable address, it will generate a new message back to the purported sender—notifying them that, "the email you tried to send was undeliverable". This email notification is often referred to as a "bounce" message.

When a spammer is sending out ten million spam messages per day, 20 percent or more will bounce because of invalid addresses. Since the spammers don't want to deal with two million incoming bounce messages, they typically forge the return address and the bounces become "misdirected" or returned to an innocent third party that had nothing to do with the spam in the first place.

The IronPort Threat Operations Center has the ability to measure global email traffic patterns using IronPort's SenderBase® traffic monitoring network. This network samples an astounding 25 percent of the world's email, providing unique insight into trends. SenderBase also has a unique capability to measure the composition of this email, as illustrated in Figure 1. This study shows that, in aggregate, global email is made up of only 20 percent legitimate messages, spam makes up 67 percent, misdirected bounces make up 9 percent, viruses make up 3 percent, and phishing emails make up less than 1 percent.

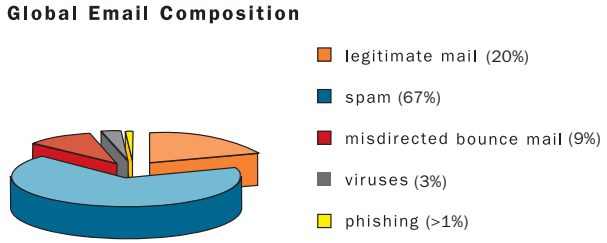

FIGURE 1: *Spam bounce volume is approaching the same level as legitimate mail volume*

The sheer volume of bounce messages generated by spam has grown to the point where it consumes an estimated (US) $5 billion per year in IT resources. This huge expenditure is also growing at the same rate as spam volume, which has historically grown at 100 percent annually— an alarming prospect for the future.

An even more troubling trend has misdirected bounces becoming more than an annoyance. Bounces can actually cause a massive distributed denial of service (DDoS) attack, which can knock even the largest email systems offline for days. IronPort® has found that over half of Fortune 500 enterprises have experienced a disruption of service or a total denial of service due to misdirected bounces, creating costs that eclipse those quantified in this study. ■

The Bounce Problem

Bounce messages are an inherent part of SMTP. Similar to a postal envelope address, an SMTP email has an "envelope" to and from address that is not exposed to the end-user, but that email gateways use to properly route mail. It is not uncommon to have an envelope return address be different then the actual from address that is exposed to the end-user in their mail client. For example, if a company were to use an email service provider for their monthly newsletter, the emails might appear to the end-user to be from "user@from.com" but would have an envelope return address of "bounce@isp.net". The anatomy of an SMTP email is illustrated in Figure 2.

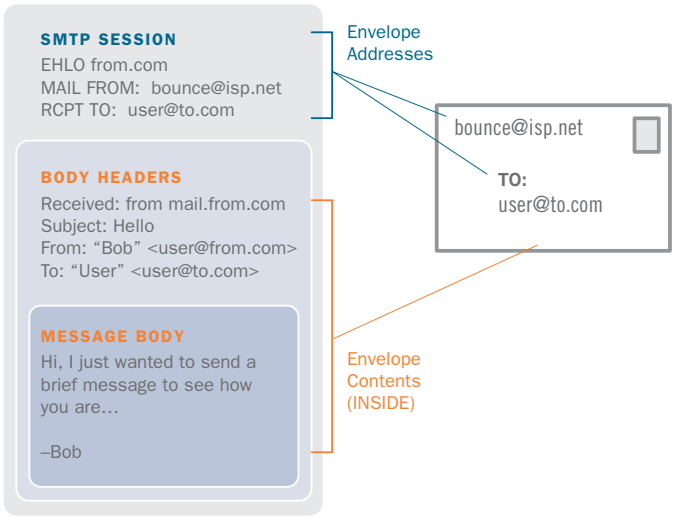

FIGURE 2: *Anatomy of an SMTP email message*

Email was designed at a time when the Internet was a collection of trusted scientists and academics sharing information. At the time the protocols were designed, it was inconceivable that anyone would forge the addressing

scheme. As a result, there is virtually no way to verify that a return address is valid or spoofed. Modern spammers have taken advantage of this loophole in the email infrastructure and created a situation where bounce messages are polluting the Internet—in some cases disabling entire email networks.

Spam works on volume. Since spam response rates are very low, spammers increase their revenues by blasting out ever-increasing volumes. It is not uncommon for a spam attack to involve ten million messages or more. Since many of the addresses the spammers are mailing to are invalid, a bounce rate of 20 percent or more is also typical. So a ten million message spam attack will create two million bounce messages in return. Spammers don't want that type of volume returning to their networks, so they typically forge the return address with a random address— causing billions of bounce messages to scatter across the Internet. These bounces, aimed at forged return addresses, are know as "misdirected bounces" and have grown to massive proportions. A misdirected bounce is illustrated in Figure 3.

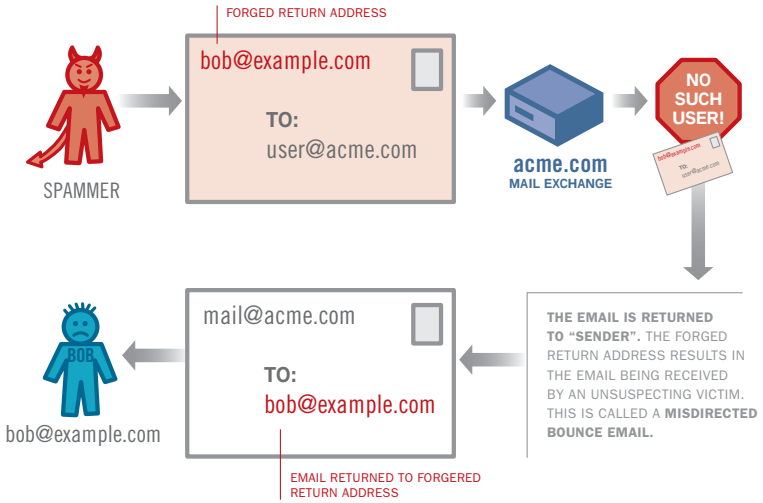

FIGURE 3: *Legitimate mail is sent and accepted. If the address is incorrect, the message is returned to its sender. Spam with a forged email address is also returned—to an unsuspecting victim of misdirected bounces.*

Misdirected bounces have a few variants. The most common is the text reply notifying the user that, the message they sent was not delivered, due to invalid address. Another form of misdirected bounce is caused when an end-user sets up an "out of office" autoreply. This out of office notice gets misdirected to the unwitting third party that was unlucky enough to appear in the return address of a spam message. A particularly troublesome variant occurs when a receiving mail server is configured to notify the sender that their message contained a virus. This helpful notice is misdirected to a random third party with a message stating, "the message you sent to user@acme.com could not be delivered because it contains the mytob virus". When, in fact, the end-user never sent an email to user@acme.com, it just happened that the mytob virus forged the return address—causing a misdirected virus notification. ■

"We saw the beginnings of the misdirected bounce problem and it just kept getting worse. It is the Internet age equivalent of all the world's junk mail being marked 'return to sender' and then delivered to an arbitrary, or strategically selected, post office."

ALLISTAIR SCOTT
KITG Gateway Team Leader, Arup

A Constant Drain on IT Budgets—or Worse

The sheer volume of misdirected bounces have created a nuisance for IT teams around the world, depleting valuable resources. Misdirected bounces consume system capacity and bandwidth used to process and store these messages. They also generate expensive IT trouble tickets from end-users, who are confused by incoming misdirected bounces. Misdirected virus notifications cause a significant amount of end-user confusion. End-users may get a message from an address they recognize or one they don't, but in either case the end-user will assume their PC is infected—when in fact they are just the unlucky recipient of a misdirected bounce or virus notification.

The IronPort Threat Operations Center measured global volume of misdirected bounces at an astounding 4.5 billion messages per day. Typically 10 percent of these misdirected bounces have valid addresses, yielding 450 million misdirected bounces that make it through to end-user mailboxes every day. If only 0.2 percent of these messages generate an IT trouble ticket (a very conservative assumption) that corresponds with 900,000 tickets per day. At a global ticket cost of (US) $20 per ticket, this equals (US) $4.5 billion annually consumed by misdirected bounces. Adding the cost of system resources, bandwidth and service outages could easily increase this number by 3-5x. But these costs are harder to quantify and were, therefore, not addressed by this study.

There is another, far more significant, cost associated with misdirected bounces. If a spammer sends out a 100 million message spam campaign, and (instead of rotating through random return addresses) uses a single forged return address of, for example, "postmaster@victim.com", then postmaster@victim.com is going to receive 20 million bounce messages from 20 million different legitimate mail gateways across the Internet.

These 20 million legitimate mail gateways are performing what should be an electronic courtesy, but they are actually unwittingly participating in a massive DDoS attack.

These unintentional denial of service attacks are not uncommon. In a survey associated with this study, IronPort found that 55 percent of the Fortune 500 have had a disruption in service or a full scale outage due to misdirected bounce attacks. Consumer brands and financial entities are often targeted as spammers will use their return addresses in an effort to appear more legitimate. Figure 4 illustrates the full impact of one such denial of service attack (resulting from an exceptionally large spam mailing). This is an actual screen shot taken from the IronPort appliance protecting a Fortune 500 insurance company. During a 24 hour period, this firm saw their typical mail volume of 10,000 messages leap to 3,653,201 messages—a 360x increase

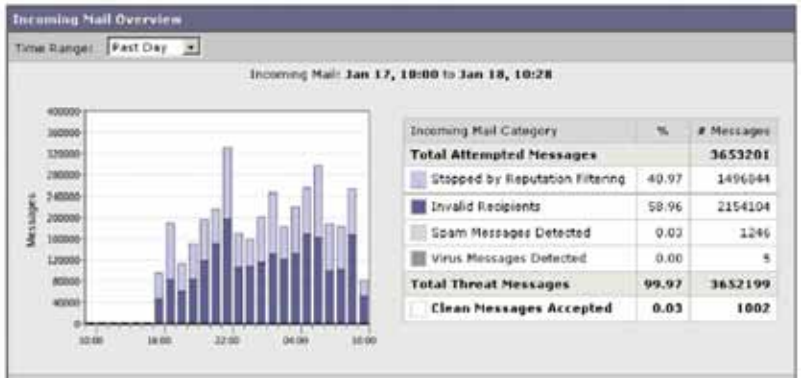

FIGURE 4: *A spam bounce attack resulting in a denial of service attack*

in volume—driven by a misdirected bounce attack. The good news shown in Figure 4 is that the IronPort appliance was able to withstand this massive volume surge. It dropped 59 percent of the mail prior to accepting it, by utilizing high performance address validation, and dropped the remaining 41 percent based on the reputation of the sender.

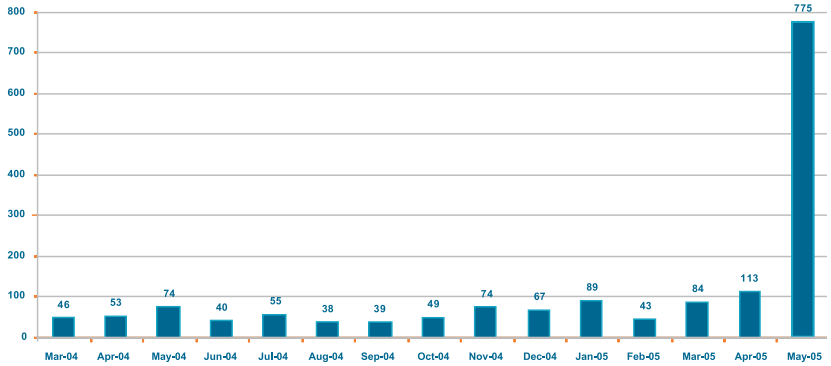

FIGURE 5: *Spam bounce attacks*

Networks of all size and scale are potential victims. Figure 5 shows another real world graph, this time from a large ISP. The graph shows a surge from the ISP's typical 175 million messages, to 775 million messages in a single day.

Figures 4 and 5 illustrate a profound point. Misdirected bounces can target any network with a very large increase in volume—on the order of a 50-100x increase. It is very difficult for any messaging system to have sufficient capacity to withstand this type of volume surge. Misdirected bounce attacks can also reach enormous scale. Figure 5 shows one attack generating nearly 800 million messages. The goal of most spammers is not to disrupt, but rather to avoid detection. However, there are groups and organizations whose sole function is to disrupt western economies. Imagine if these organizations launched a misdirected bounce attack on every government organization? With simple spammer tools and a few common servers, anyone can launch a massive DDoS attack that will knock a traditional mail server offline for days with a 10-100x increase in message volume. ■

55 percent of the Fortune 500 have had a disruption in service or a full scale outage due to misdirected bounce attacks.

Secure Bounces: The Solution

Misdirected bounces are expensive and extremely dangerous. While there is little that can be done to prevent these attacks using traditional mail infrastructure (such as the popular open-source "sendmail" program), IronPort has developed unique technology that can withstand a DDoS attack. And, better yet, prevent it from happening in the first place.

Performance: Connections and Message Processing

Withstanding a highly distributed denial of service attack with misdirected bounces is extremely challenging, because the sources of the bounces are legitimate mail servers that cannot simply be blocked. IronPort email security appliances have extremely high concurrency, allowing them to support up to 10,000 simultaneous connections. A single IronPort appliance can process incoming mail at rates of up to 1 million messages per hour. This breakthrough performance represents a 10x increase in processing versus traditional UNIX-based systems, with raw capacity that can shrug off a DDoS attack.

Email Address Validation

The architecture of the IronPort appliance uses LDAP-based address validation early in the email processing pipeline. Consequently, it can validate the address of incoming mail — prior to actually accepting the message. This was a key attribute in allowing the appliance to scale, while under the real world attack shown in Figure 4. 59 percent of the incoming misdirected bounces were discarded because of invalid addresses, a ratio not at all uncommon in a misdirected bounce. The remaining 41 percent of the incoming volume was originating from illegitimate sources and was thus blocked by IronPort Reputation Filters™. By intelligently processing mail prior to actually accepting the SMTP message body, the system can increase efficiency and shrug off even a large volume spike like the 360x increase shown in this example. However, there is a concern with this approach, because the corporate directory is now exposed to a spammer's directory harvest attack.

Throttling Traffic with Email Sender Reputation

To protect the corporate directory, IronPort uses sender reputation and secure bounce logic. The IronPort appliance keeps track of the number of invalid addresses received from a given sender over time. At a certain threshold, the appliance assumes the sender is just guessing at addresses and drops mail from that sender. This threshold varies by the reputation of the sender. A sender with a poor reputation might get one or even zero attempts to deliver a message. A slightly stronger reputation will be allowed five invalid delivery attempts before messages are dropped. A sender with a long history or reliable mail patterns will be allowed 20 or more attempts. This type of graduated response allows the IronPort appliance to intelligently determine whether or not to issue a bounce message, based on the reputation and behavior of a given sender.

IronPort Technology Solves the Problem at its Core

IronPort has developed a secure bounce technology that prevents misdirected bounces from starting in the first place. IronPort appliances can be configured to issue a bounce message during the SMTP conversation. This means that the appliance holds the connection open with the sending mail server and validates the recipient address prior to accepting. Validating the address during the conversation has the advantage of never issuing a misdirected bounce, since the sender gets the bounce notice directly — instead of relying on a separate email notice that is sent to the (frequently forged) email envelope return address.

Email administrators need to be aware of the problems of misdirected bounces. Those with traditional mail gateways need to begin using conversational bounces, instead of the delayed bounces that are frequently misdirected (and thus are polluting the Internet). Making this change can expose the corporate directory to harvesting, but it is a fair trade-off to prevent the DDoS attacks associated with misdirected bounces. IronPort customers can rest assured that they will not be misdirecting bounces and contributing to the Internet email traffic emergency. ■

"IronPort created a robust bounce solution that actually prevents misdirected bounces — solving the problem at its core, instead of just reacting to it."

ALLISTAIR SCOTT
KITG Gateway Team Leader, Arup

IRONPORT EMAIL SECURITY APPLIANCES

- IronPort C10/C100
- IronPort C300
- IronPort C600
- IronPort X1000

NOW WITH
BOUNCE VERIFICATION
Don't let your network be compromised by email bounce attacks.

www.ironport.com / email

IRONPORT. A SPAMMER'S WORST NIGHTMARE.

The *IronPort C-Series*™ and *IronPort X-Series*™ email security appliances are in production at eight of the ten largest ISPs and more than 20 percent of the world's largest enterprises. These industry-leading systems have a demonstrated record of unparalleled performance and reliability.

The same technology that powers and protects IronPort's most sophisticated customers is available for companies of all sizes, starting with the entry-level *IronPort C10™/C100™.*

By reducing the downtime associated with spam, viruses and blended threats, IronPort® email security appliances vastly improve the administration of corporate email systems, reduce the burden on technical staff and provide state-of-the-art network protection.

IRONPORT®

MAKING THE INTERNET SAFE.™

PETER CHRISTY is a principal at the Internet Research Group and a recognized industry expert. He has been involved with the computer and communications industries since the late 1960's. A graduate of both Harvard and Berkeley, Christy has held positions at a wide variety of technology companies, including Hewlett Packard, IBM/Rolm, Sun Microsystems and Apple, among others. His experience is brought to bear in the current Internet Research Group activities, strategy and research on diverse topics related to "infrastructure"—the interaction of networks with systems and applications.

THE IRONPORT THREAT OPERATIONS CENTER (the TOC) is IronPort's 24x7 view into global email activity. TOC analysts use sophisticated tools to easily visualize complex real-time and historical traffic patterns, analyze anomalies to uncover new threats and track email traffic trends. Automatically generated alerts are verified and updates issued to IronPort's security appliances on a constant, rapid basis—successfully countering threats, such as new virus outbreaks. TOC tools are powered by a series of proprietary algorithms that process data from SenderBase, the world's first and largest email and Web traffic monitoring network.

MNP MESSAGING NEWS PRESS
www.messagingnews.com

6/06 P/N 451-0305-1