CIDR House-Rules: Use of BGP router data to identify and address sources of Internet abuse

Karsten M. Self
kmself@ix.netcom.com

1st March 2006

Presented to the Messaging Anti-Abuse Working Group
6th General Meeting
San Francisco, CA
http://linuxmafia.com/~karsten/cidr-house-rules.pdf

Abstract

BGP router data may be used to identify contiguous regions of network space from which significant abuse is observed. Experience suggests a strong power-law relationship in ranking such sources. Applying this knowledge in abuse countermeasures may markedly reduce filtering overhead while minimizing inadvertant blocking and increasing total costs to abuse-tolerant networks.

1 I know where your spam comes from

For typical Internet sites, from a quarter to half or more of all spam and other forms of network abuse may originate from a very small number of sources.

The methods discussed here result from reporting and data analysis on nearly 200,000 spams received at a single ISP POP account since January, 2004. The interest isn't in specific sources, but in the tools used to aggregate information on spam-transmitting peers and the applicability of these methods to large-scale spam mitigation. Several application scenarios are suggested.

The principle is to note sources of spam by IP peer on an aggregated basis. Studying such data over time it has become clear that:

- The bulk of spam originates from a very small subset of network sources.
- These networks are readily identifiable by commonly available tools and methods.

Though based on observations from a single end-user mailbox, trends noted should be similar in character to those seen at the server level. Comparing data from several sources show similar trends. This is not an "ultimate solution", however it may be a useful tool particularly on large sites, sites with large spam loads, or sites in which mitigation methods should incur minimal time, bandwidth, and processor overhead. It would also he helpful to have capabilities directly integrated with standard mail transfer agents.

The intended audience for this discussion includes postmasters, email abuse reporting and mitigation managers, webmail providers, email server developers, email plugin (server or client) developers, blog operators, VOIP vendors, and others dealing with network abuse.

This paper merely introduces the concepts. It is neither a complete solution nor an exhaustive technical analysis.

2 Technical concepts

For presentation calibration: some of the technical concepts covered in this presentation will include: email, SMTP, DNS, CIDR, ASN, BGP, DNSBL, network hygiene, greylisting, proportionate response, and denial of interest. Much of the following discussion assumes a moderate understanding of these terms.

Though initial applications have been for email, and principally based around spam, other abuse for which clear and not readily spoofable peer relationships exist may be appropriate.

3 Existing spam filtering methods

Methods such as whitelist/blacklist, DNSBL, content (rule-based) filters, Bayesian filters, greymilter, and tarpitting are more-or-less widely deployed. They do work and are often effective. Several are strongly endorsed by the author.

However, they share a number of disadvantages:

- Data-lossy, particularly filters, regards spam source. Information gained regarding one IP isn't gainfully applicable to its neighbors, or even (often) itself in subsequent abuse attempts.
- Whack-a-mole, particularly DNSBLs, regards point vs. aggregate source. Rinse, wash, repeat with IPv6: DNSBLs scale very, very poorly in this case.
- Reliance on third parties reliably, accurately, equitably, and expeditiously collect and distribute assessment data.
- CPU and/or wall-clock intensive, particularly for large sites. Often extending to other resources including threads, filehandles, memory, etc.
- Generally fail to impose overhead on spam source.
- Are uniformly applied to mail from both trusted and untrusted sources inducing unnecessary cost.

While not arguing that these methods be disposed of, a method is presented here of taking a large first cut at the spam problem before incurring the cost and uncertainty of other filtering methods.

What If You Could...

- Tie an IP address to the organization responsible for it.
- And a network address space (CIDR block)
- In a manner leveraging existing spam detection / filtering tools for single-point IPs
- Quickly, cheaply, accurately
- And could develop policies for email and network traffic management
- Oh, and could also identify your good / trusted network peers

The answer, of course, is, "You can".

4 BGP, Routeviews.org, and you

Border gateway protocol (BGP), to quote Cisco:

> is an interautonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm]

The key points to recognize are:

- BGP is fundamental to the nature of the Internet. It defines the relationships between autonomous systems – the networks the Internet internetworks between.
- It ties directly to an organization: the AS owner, identified by ASN.
- It ties directly to network data, the CIDRs which BGP peering rules are applied to.
- Though IP space is large, and will likely get vastly larger as IPv6 is widely adopted, pragmatic constraints suggest that ASN proliferation will not change as markedly. Currently there are some 39,500 assigned ASNs with a total namespace of 65,535.

In other words: you've found the folks in charge, where they are, and how they relate to you. Since SMTP deliveries are stateful TCP transactions with defined IP peer relationships (and spoofing is not practically significant), we have a known IP.

Now all you need is something which can return ASN and CIDR data for a given IP address.

The Routeviews project (http://www.routeviews.org/) provides just such a capability, though others exist. It is "a tool for Internet operators to obtain real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet", and was first noted by Joe St. Sauver of University of Oregon. Routeviews provides zonefiles, updated twice daily, and queryable at:

```
host -t txt
   <reversed-IP>.asn.routeviews.org
```

To determine, for example, the ASN and CIDR for the AOL mailserver mailin-01.mx.aol.com at 64.12.137.249:

```
$ host -t mx
   249.137.12.64.asn.routeviews.org
249.137.12.64.asn.routeviews.org
   descriptive text "8176"
   "64.12.0.0" "16"
```

This tells us that the server is in ASN 8176, CIDR 64.12.0.0/16.

For use in mitigating spam, you want to find which ASNs are principally associated with spam traffic, noting volumes of both spam and ham (non-spam) mail received from various sources. Ideally both total mail volume and spam proportion would be noted.

Routeviews.org makes the zonefiles available via rsync to allow large sites to run queries against a local nameserver for increased performance.

5 Pareto's law and spam sources

A power distribution is very evident in monthly data seen to date.

- **Over a two year period, 3-5 ASNs contribute 25% of all monthly spam.**
- **50% of all spam comes from 9 to 35 sources.**

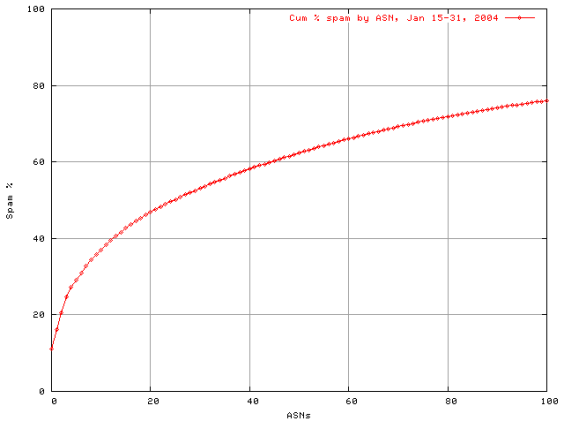

The plot shows the total percent of spam contributed (vertical axis) by ASNs (incremented along horizontal axis).

CIDR data show a similar, though less concentrated, power distribution. Specific ASNs involved vary, though gross abusers have been fairly stable over time. Typical among them are ASNs from China, Korea, large webmail providers (usually 419/Advanced-Fee fraud spam), large European or Middle-eastern ISPs (often quasi-governmental monopolies), blowback/backscatter sources (which would be specific to a given email address at any one time), and occasionally larger US commercial ISPs. Specific trends are highly idiosyncratic. *You are very strongly encouraged to note trends from your own experience, not other sites'*. Sharing data is possible and may be useful but should not be principally relied on.

In conjunction with numerous spam reports sent to the organizations associated with domains, IPs, and/or ASNs, it's further noted that network organizations can be seperated into two classes:

- Those which deal preemptively or reactively in a way which minimizes abuse problems
- Those which don't, can't, or won't.

This observation gives rise to the concept of network hygiene, namely that there are neighborhoods which are well policed and those which aren't. Methods for increasing the accountability of a network's own hygenic practices would be a net benefit.

Additional statistics, tables, and plots follow at the end of this paper.

6 Application

ASNs by themselves don't tell you whether or not traffic is abusive, or if a given IP range is spammy. What's necessary is to identify sources of undesireable traffic, map these to an ASN and / CIDR, and determine your house rules for handling traffic from that CIDR. Two steps are necessary: data aquisition, and policy enforcement.

6.1 Data aquisition

Aquire a list of IPs doing things you don't (or do) like: spam, viruses, open proxies, portscans, blog / comment spam, referrer spam, business partners, friends, vendors, bad breath, drinking white zin. Look up the associated ASN / CIDR. Note which are naughty or nice. This could be accomplished in the case of spam by dropbox accounts, honeypots, server logs, end-user submissions, or other means. Because of the power of aggregation allowed by ASN/CIDR lookups, a reasonably constructed spam provider sample may be very small. On the order of 1:1,000,000 or fewer mails for a very large provider.

6.2 Policy enforcement: CIDR house-rules

Implement a policy at the service (eg: email, web, messaging) or firewall (eg: iptables) level. These are your house rules for interacting with a given CIDR, ASN, IP, or other defined network block.

While blocklisting is one possible option, I'd very much like to see the discussion move beyond that point. A preferred approach is what I term "proportionate response". First: you'll likely want rules to expedite known-trusted mail, or high priority mail from remote organizational sites, peers, clients, vendors, or other established relationships. Secondly, many peers will either have small overall volumes, or not have a clearly identifiable nature. This leaves the set of networks which are both high-volume and overwhelmingly spammy in nature. Of course, any such implementation would have to be evaluated in a business and organizational context.

In proportionate response, a certain level of abuse would be met by a proportionate level of response. For example, a network from which 90% of email was found to be spam, 90% of traffic originating from that network would be denied or dropped, either at the service (protocol) or IP level, at random. If done at the SMTP transaction level, either as a timeout (without 250 OK) or non-permenant rejection, this would mean legitimate mail still has a fighting chance to get through. A 90% reject rate would allow half of mail through on 5 retries, for a typical 2 hour delay. A spam server without retry rules would fail delivery of 90% of its mail, with retries it would suffer large mail spools and possible other resource starvation.

The site implementing such a policy will receive immediate benefit to itself. Widespread adoption is not necessary to be locally beneficial. As multiple and large sites adopt such measures, impacts on abuse-tolerant networks would be significant. The approach is to be both non-invasive and non-retaliatory. You are not taking any action which in any way directly changes or affects a remote system: but are subjecting it to a denial of interest.

As a proportionate response, reject rates could vary with total traffic volume, abusive traffic percentage, and severity of abuse, as suited specific needs. Fine levels of control are therefor possible, operators are not reduced to all-or-nothing responses to abuse.

7 Data and additional references

Some additional information and references on use of BGP and ASN data in spam mitigation.

7.1 Related third-party discussions of spam and ASN data:

- The Routeviews project:
  `http://www.routeviews.org/`
- Chris Siebenmann's blog describing spam combat at the University of Toronto, Canada, including use of BGP and ASN data at the server level:
  `http://utcc.utoronto.ca/~cks/space/blog/spam/SpamByASN`
- Michael Greb's blog on spam, including data on spam by ASN, collected from several spamtrap addresses:
  `http://spam.thegrebs.com/`

7.2 Summaries of spam by ASN & CIDR

Full online reports of my own data are frequently updated at:

```
http://linuxmafia.com/~karsten/monthly-asn-report
http://linuxmafia.com/~karsten/monthly-cidr-report
```

Historical data from January 2004 through present, with some gaps, are saved by year and month in YYYYMM form available at:

```
http://linuxmafia.com/~karsten/monthly-asn-report-YYYYMM.txt
http://linuxmafia.com/~karsten/monthly-cidr-report-YYYYMM.txt
```

From current data, ASNs and CIDRs with most reported spams. Note that report classification isn't entirely accurate though trends are generally well presented.

Report date: Mon Feb 27 23:37:48 PST 2006
Total spams: 11249
Total ASNs: 955

Rank	Cumulative %	%	Spams	ASN	Description
1	9.9%	9.9%	1113	8176	NETSCAPE-ASN
2	18.5%	8.6%	968	4135	CHINANET-BACKBONE
3	24.5%	6.0%	673	4814	CHINA169-BBN CNCGROUP
4	28.3%	3.8%	432	8176	NETSCAPE-ASN
5	31.6%	3.3%	373	4837	CHINA169-BACKBONE
6	34.5%	2.9%	322	4755	KIXS-AS-KR Korea Telecom
7	36.7%	2.0%	248	3269	ASN-IBSNAZ TELECOM ITALIA
8	38.8%	2.0%	230	17858	KRNIC-ASBLOCK-AP KRNIC
9	40.7%	1.9%	217	1668	AOL-ATDN
10	42.1%	1.4%	161	17849	GINAMHANVIT-AS-KR

Report date: Mon Feb 27 23:39:24 PST 2006
Total spams: 11248 Total CIDRs: 2251

Rank	Cumulative %	%	Spams	CIDR	AS & Description
1	9.9%	9.9%	1113	64.12.0.0/16	8176 NETSCAPE-ASN
2	13.7%	3.8%	432	16/	8176 NETSCAPE-ASN
3	15.7%	1.9%	217	205.188.0.0/16	1668 AOL-ATDN
4	17.5%	1.9%	211	212.216.128.0/17	3269 ASN-IBSNAZ TELECOM ITALIA
5	19.2%	1.6%	183	220.163.0.0/17	4134 CHINANET-BACKBONE
6	20.5	1.4%	155	4755/61.17.128.0	VSNL-AS
7	21.9%	1.5%	152	221.220.128.0/18	4814 CHINA169-BBN
8	23.1%	1.2%	139	4755/61.17.176.0	4755 VSNL-AS
9	23.4%	1.2%	136	218.63.0.0/17	4134 CHINANET-BACKBONE
10	25.4%	1.0%	115	61.148.128.0/18	4814 CHINA169-BBN
10	26.4%	1.0%	115	222.129.64.0/18	4815 CHINA169-BBN