The ISP Column
A monthly column on things Internet

March 2010

Geoff Huston
George Michaelson

Traffic in Network 1.0.0.0/8

Background

The address plan for IPv4 has a reservation for "Private Use" address space. This reservation, comprising three distinct address blocks, namely 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16, is intended for use in private contexts where networked devices do not have any requirement to be visible to the public Internet. However, it has long been recognised that other addresses have also been used in private contexts. Some of these uses are entirely informal and remain entirely within a particular private network, while other uses have been a little more systematic. One recent study of this form of address use noted that the address block of network 1, or 1.0.0.0/8, was "widely used as private address space in large organizations whose needs exceed those provided for by RFC 1918" [1].

Network 1.0.0.0/8 was assigned by IANA to APNIC on 19 January 2010, for use for as public unicast space for further address allocations and assignments. Before APNIC commences distribution from this address block, we have undertaken a study into 1.0.0.0/8. The particular question under investigation here is the extent to which addresses in network 1.0.0.0/8 are an "attractor" for unusually large quantities of unwanted traffic. In particular, is there a significant level of "leakage" of supposedly private-use traffic directed to addresses in 1.0.0.0/8 that "leak" into the public Internet?

While network 1.0.0.0/8 was an unallocated and (officially at least) unadvertised network any traffic directed to an address in this network that "leaked" into the public Internet would follow a "default" routing path to the point where there was a "default-free" routing element, where the traffic would be discarded. As soon as any address in 1.0.0.0/8 was advertised as reachable into the public Internet then instead of being discarded at the boundary of the default-free zone (DFZ), the packets would be passed towards the advertised destination.

Initial Experience with 1.0.0.0/8

On 27 January 2010 the RIPE NCC, in collaboration with APNIC, announced 4 prefixes in 1.0.0.0/8, namely 1.1.1.0/24, 1.2.3.0/24, 1.50.0.0/22 and 1.255.0.0/16. This exercise recorded an immediate effect in terms of a jump in traffic to the announcement point that saturated the 10Mbps port (Figure 1) [2].

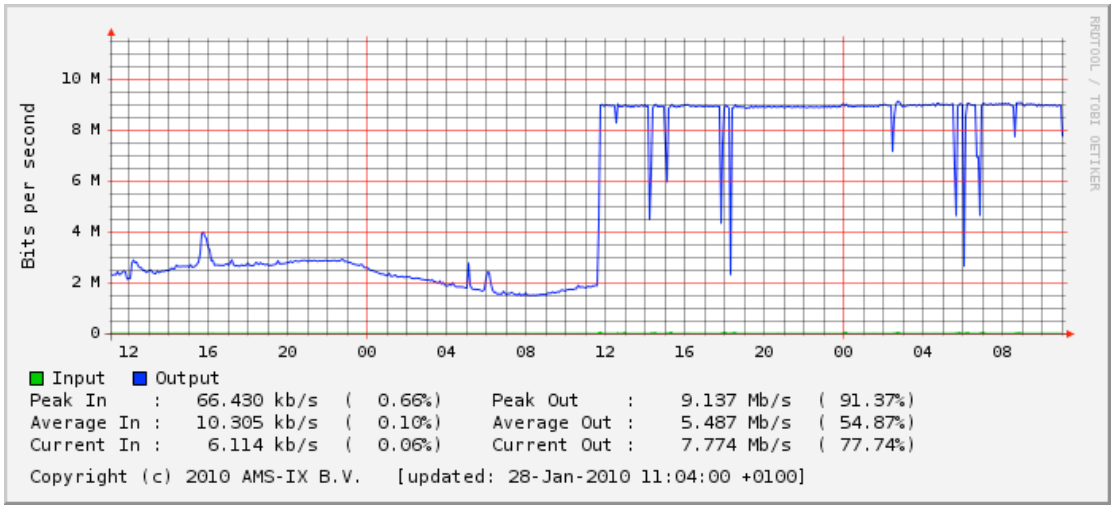

Figure 1 – Traffic load at the RIPE NCC announcement point of 1.0.0.0/8 more specifics [2].

The announcements of 1.1.1.0/24 and 1.2.3.0/24 were dropped on 2 February, and the announcements of 1.50.0.0/22 and 1.255.0.0/16 were dropped on 9 March 2010. Some packet capture of the traffic that was being sent to these prefixes was performed at this time, and the results of an analysis of this traffic are described in [2].

Capturing a Sample of Traffic in 1.0.0.0/8

It is clear that there is a significant amount of traffic that is being directed to addresses in 1.0.0.0/8, and it is well in excess of 10Mbps of sustained load. This traffic is likely to be a combination of leakage of traffic from private use domains, potential leakage from mis-configured equipment, and a certain amount of scanning activity that passes across 1.0.0.0/8 as part of a walk across the entire network address range. Following the earlier work undertaken by the RIPE NCC, we commenced further investigation of the traffic in network 1 in February 2010.

The primary objective of this work was to quantify the extent to which all networks in 1.0.0.0/8 attract "pollution" or "unwanted" traffic.

A related objective was to investigate the distribution of traffic directed to addresses in 1.0.0.0/8 to determine if there are particular ranges of addresses within this block that are "hot spots" for attracting unwanted traffic, and to quantify the additional traffic that such addresses attract.

In February 2010 we solicited assistance from potential collaborators to perform a comprehensive packet capture for the entirety of 1.0.0.0/8. Collaborative experiments have been undertaken with Merit and with YouTube, and this report summarizes the initial findings of these studies. We would like to acknowledge their assistance here, as their generous support and flexible responses to our requirements have been vital in assembling this report.

AS237 Announcement of 1.0.0.0/8

AS237 (Merit) announced network 1.0.0.0/8 from 22 February 2010 until 1 March 2010. A single system was used as the packet capture device and the configuration was entirely passive (i.e. the system did not respond to any packets that it received in 1.0.0.0/8). The system performed a full packet capture of all packets received, including all payload.

The analysis reported here is based on the 6 day period from 0000 23 February 2010 UTC -6 through to 2400 28 February 2010 UTC -6.

Traffic Profile

Figure 2 shows the traffic received by the collection point for this 96 hour period on a second-by-second basis.

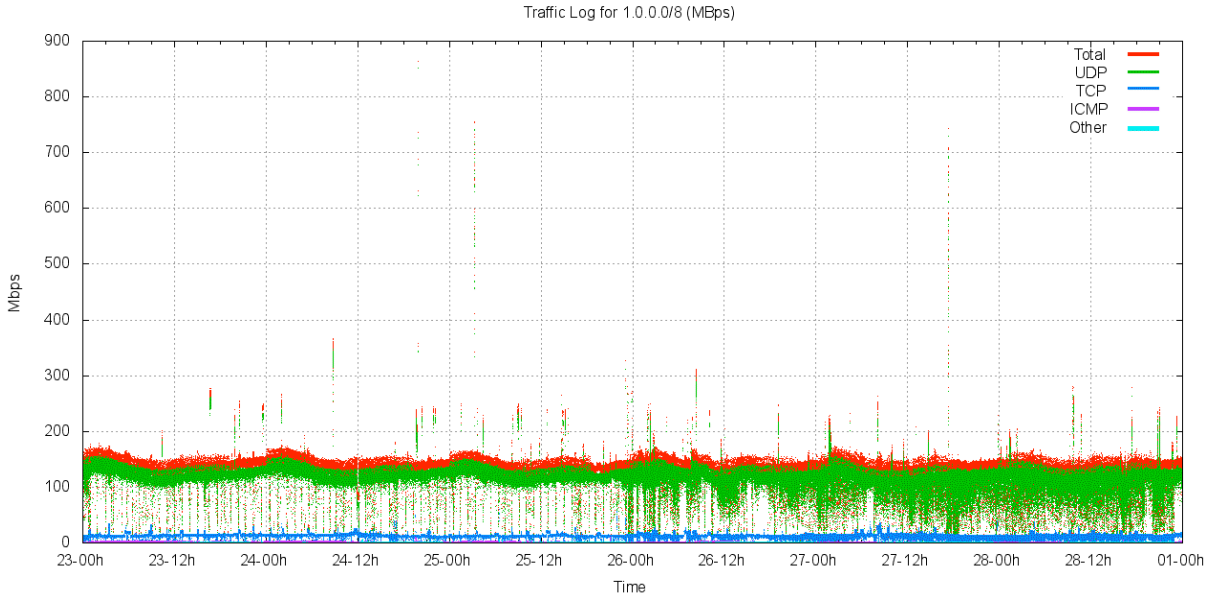

Figure 2 – Traffic received by AS237

The traffic logs show that the traffic sent to 1.0.0.0/8 is a relatively steady 160Mbps for the period. There is a slight element of a 24 hour diurnal cycle in the data, but it is not a pronounced cycle in terms of total received traffic levels.

There is a regular interval of reduced traffic in these logs, that we believe corresponds to the file cycle interval for the packet capture system. Other instances of short term reduced traffic incidents appear to reflect packet loss by the packet capture system, but there is no direct way to substantiate this assumption from the gathered data.

There are short bursts of between 1 and 30 seconds of elevated traffic levels. There are 20 or so incidents of burst traffic levels of between 200Mbps and 300Mbps. There is a 3 second isolated burst at 860Mbps and a 10 second burst at 750Mbps in this period.

Protocol Profile

The profile of traffic in each protocol is shown in the following collection of figures.

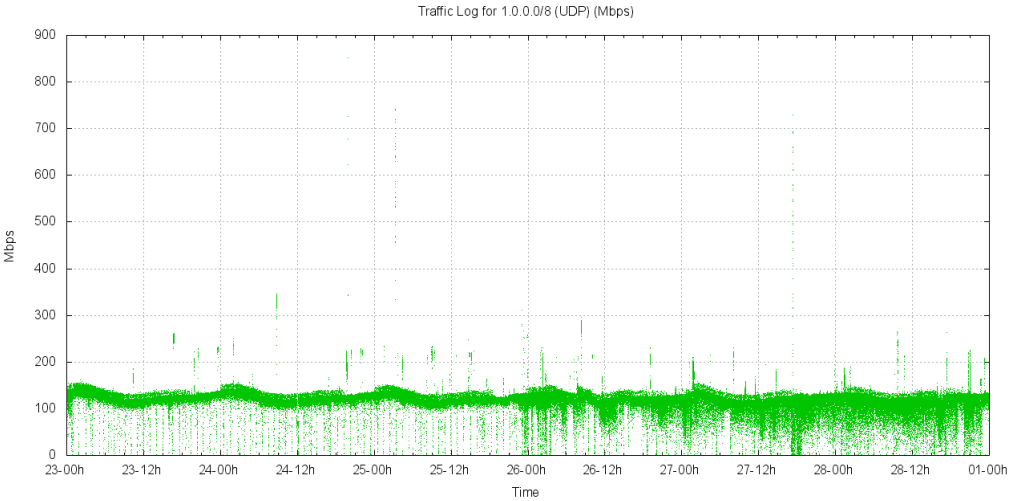

Figure 3 – UDP Traffic received by AS237

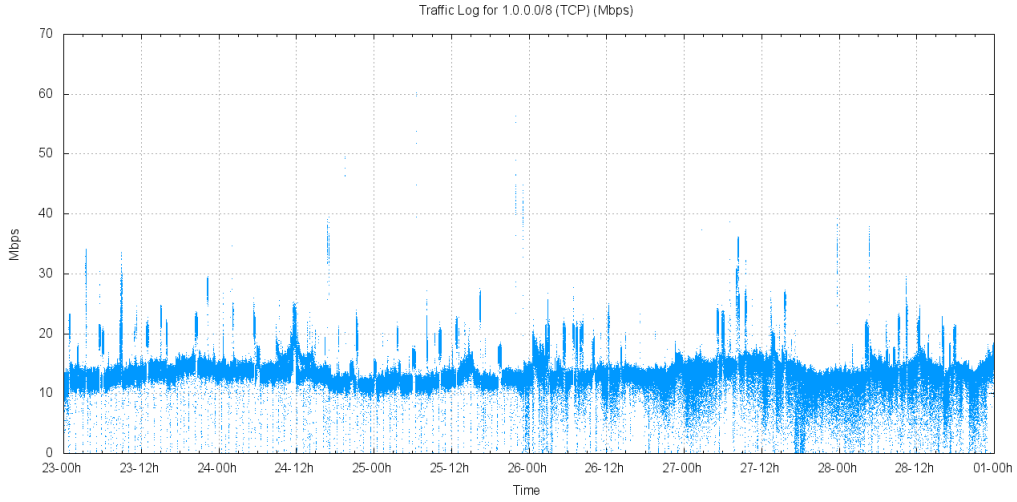

Figure 4 – TCP Traffic received by AS237

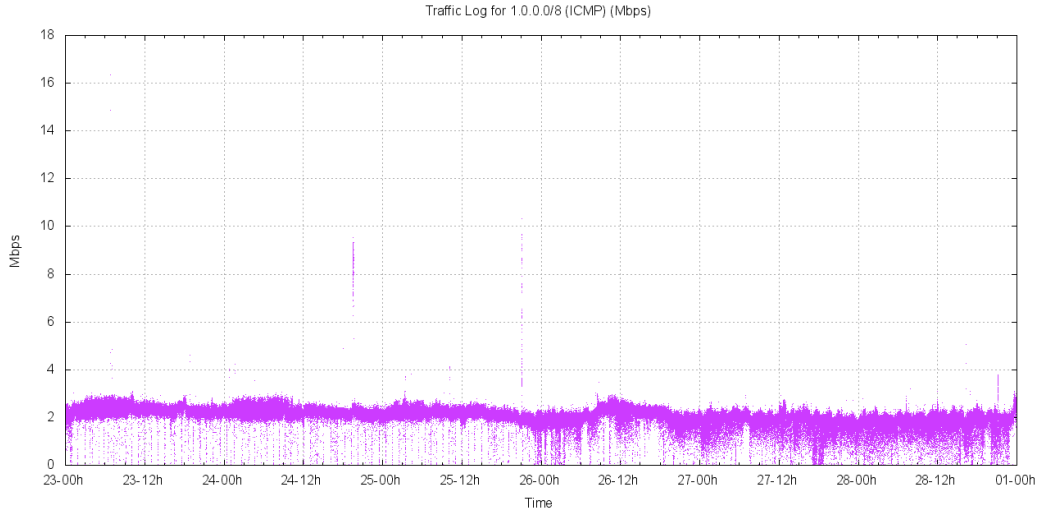

Figure 5 – ICMP Traffic received by AS237

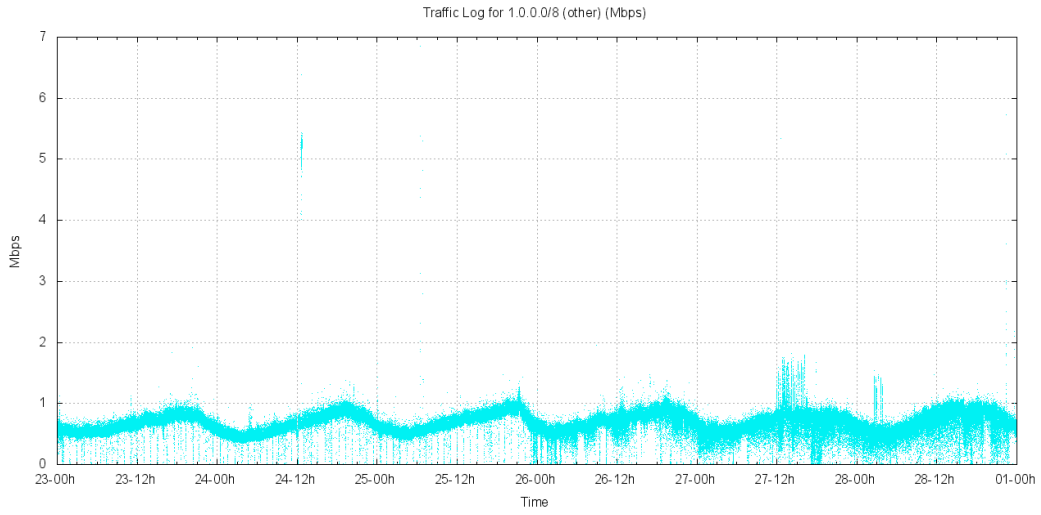

Figure 6 – Other Traffic received by AS237

Figure 7 shows the received packet rate for the period. Of note is the pronounced diurnal pattern in the data, particularly in the UDP data. The peaks are not as pronounced, indicating that the peaks in traffic rate are due to a burst of larger UDP packets.

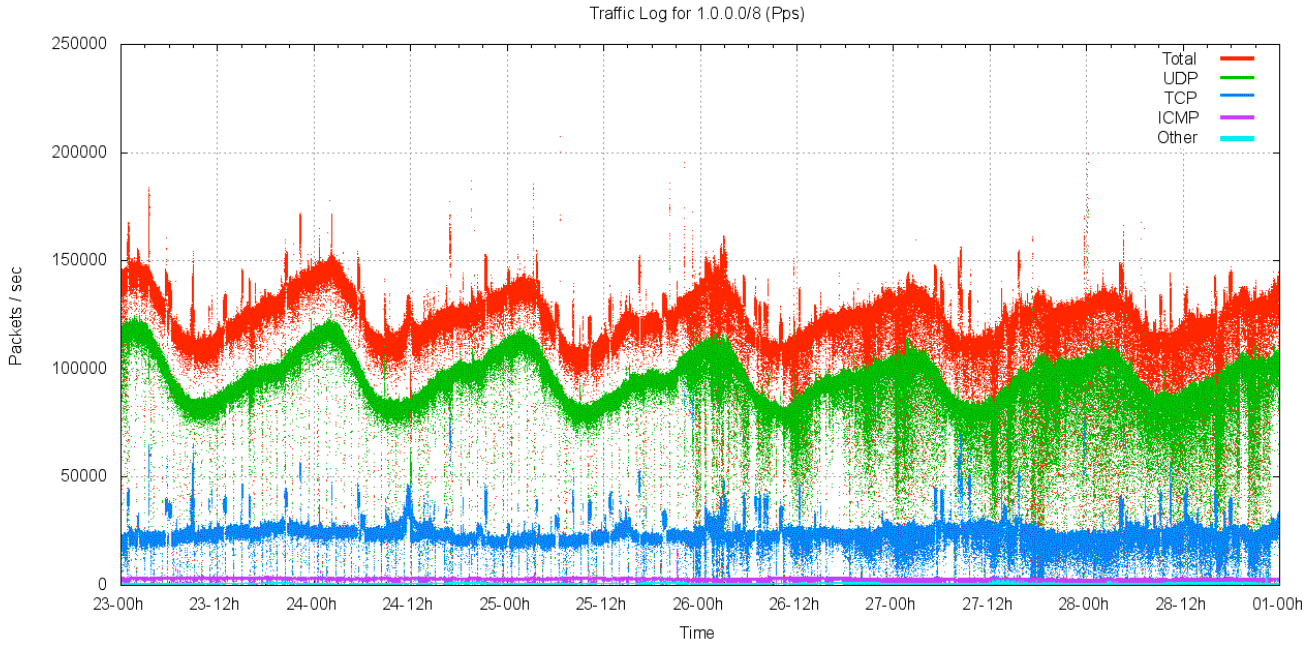

Figure 7 – Received Packet rate

In terms of protocol distribution as measured by bytes of traffic the distribution according to protocol is shown in Table 1.

Protocol	Proportion of Traffic	Proportion of Packets
UDP	88.1%	76.9%
TCP	9.8%	19.8%
ICMP	1.6%	2.5%
Other	0.5%	0.7%

Table I – Distribution of traffic by Protocol

The distribution of packet sizes by protocol is shown in Figure 8. TCP packet sizes were consistently in the range 69-70 bytes, indicative of a TCP SYN packet header of 16 bytes of framing, 20 of IP, 20 of TCP, and 13-14 bytes of TCP options.

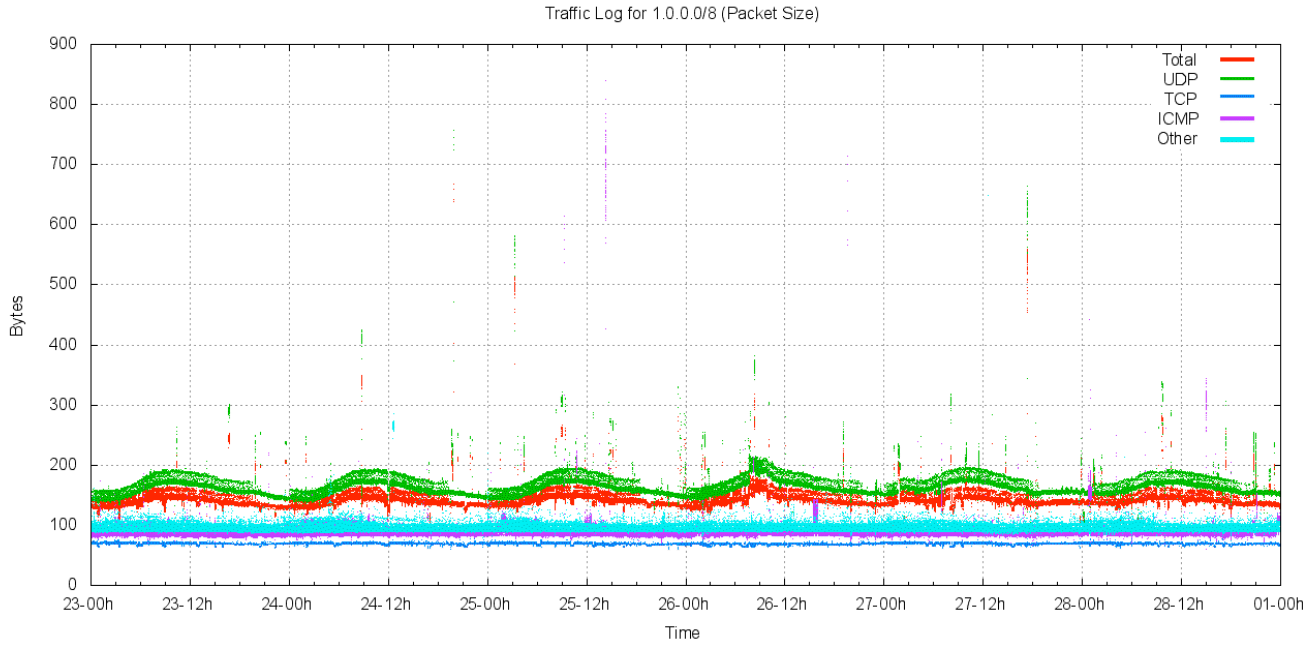

Figure 8 – Distribution of Packet Sizes by Protocol

Distribution of Traffic by /16s

The traffic in network 1 is not evenly distributed across all destinations. Figure 9 shows the distribution of traffic within 1.0.0.0/8 divided up into address blocks of a /16 in size. Note that this is a logarithmic scale of traffic levels shown in this figure.

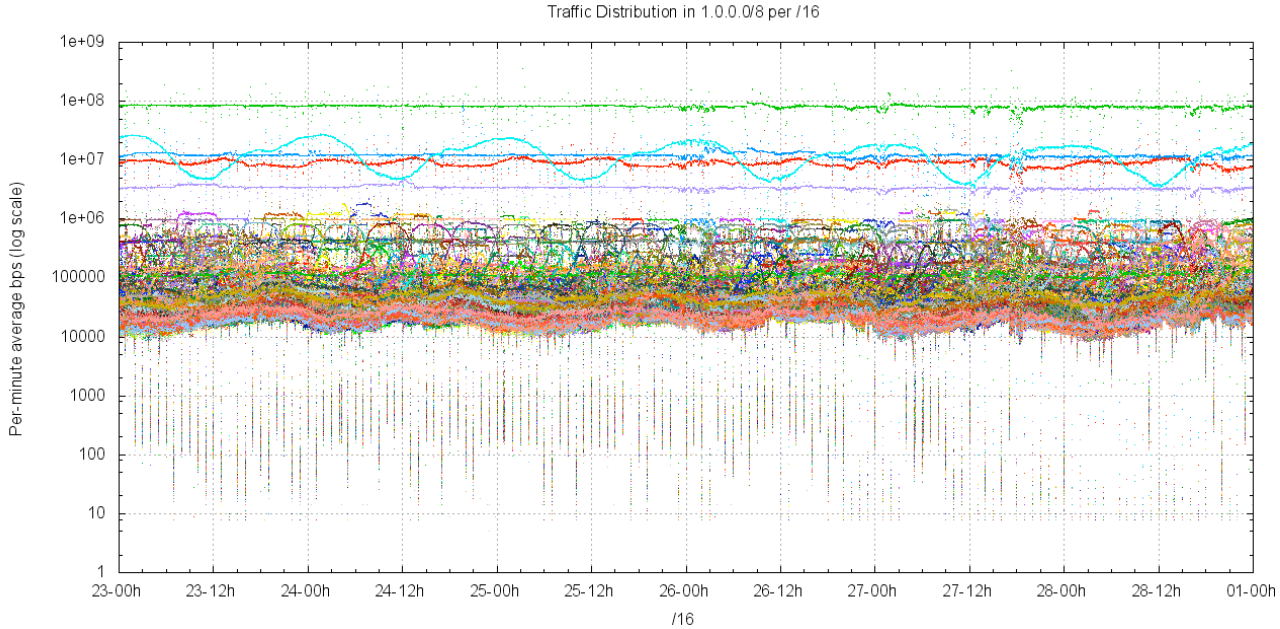

Figure 9 – Traffic levels per /16

Is this level of traffic being directed to addresses in network 1.0.0.0/8 "normal"? What is a "normal" expectation in terms of traffic?

To provide an answer to this question we advertised a smaller prefix, namely 27.128.0.0/12, under similar conditions to the advertisement of 1.0.0.0/8, using a similar packet collection process. This prefix was drawn from other parts of the unallocated address pool managed by

APNIC. This advertisement of this address block was undertaken in collaboration with YouTube, and 27.128.0.0/12 was advertised by AS36351 on the 19th of February 2010.

The traffic attracted by this advertisement was collected by four passive collectors in a load-sharing configuration. The aggregate traffic levels, per /16 is shown for the 24 hour period in Figure 10. It is evident that there is a distinct correlation across the /16s of this address block with an average traffic level of slightly lower than 10Kbps per /16.

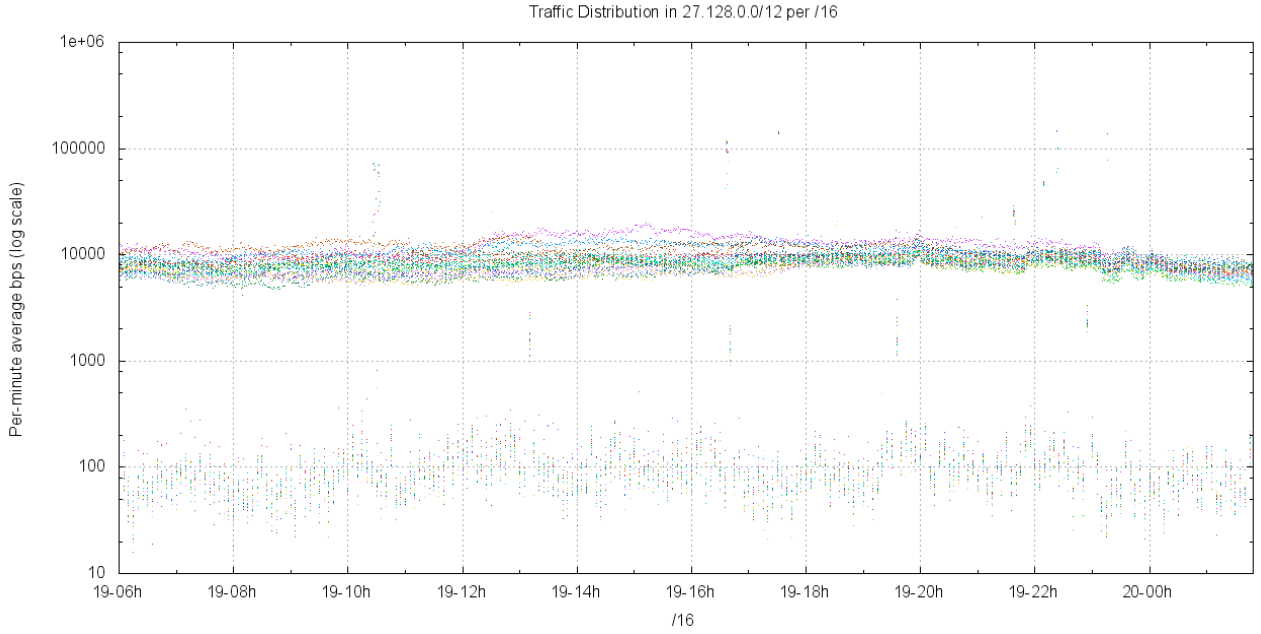

Figure 10 – 27.128.0.0/12 traffic levels per /16

The comparable levels of traffic per /16 in 1.0.0.0/8 are slightly higher than those observed in 27.128.0.0/12. Informally, Figure 9 indicates traffic levels per /16 in 1.0.0.0/8 between 10kbps and 100kbps. However, it is still that case that a /16 is a large address block, and heavy traffic levels heading to individual addresses within a /16 address block can skew the average traffic for the entire /16.

Distribution of Traffic by /24s

A similar analysis can be performed at the granularity of /24s. Figure 11 shows the profile of traffic to each /24 in the address block 1.1.0.0/16

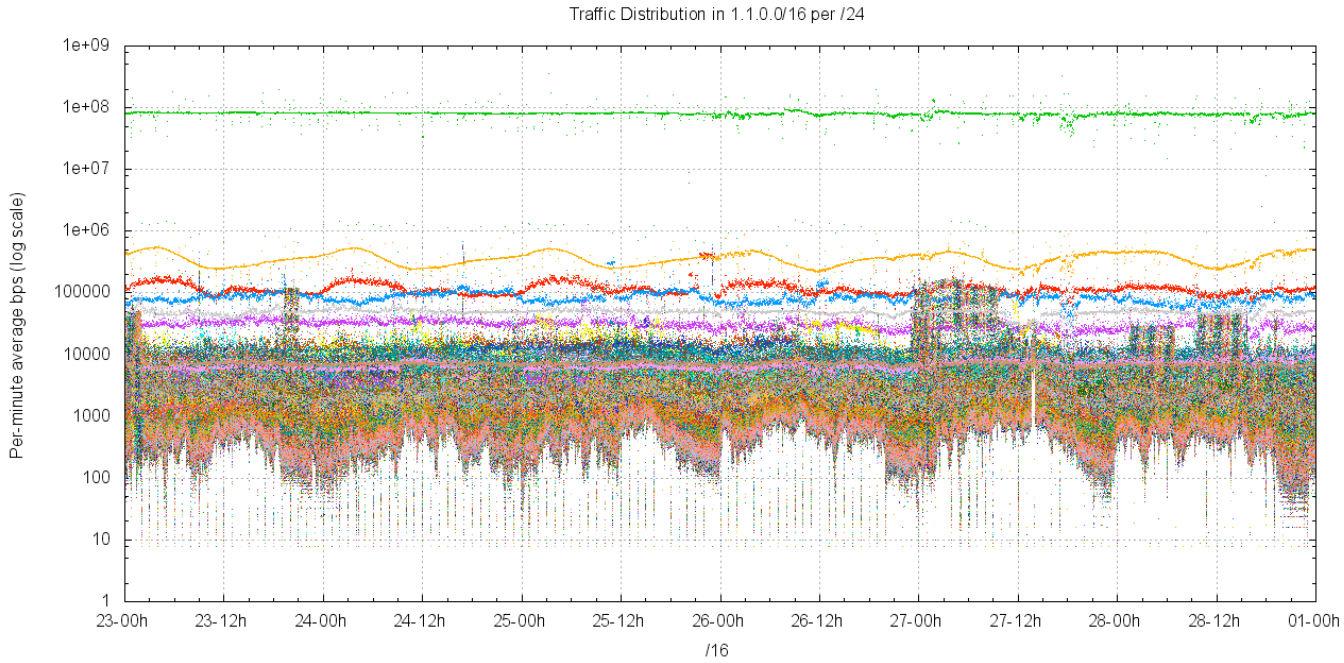

Figure 11 – Traffic levels per /24 in 1.1.0.0/16

One /24 within this address block dominates all others, namely 1.1.1.0/24. The traffic profile for this particular /24 is shown in Figure 12. There is a steady load of between 80 to 100Mbps of traffic directed to this prefix, peaking at more than 800Mbps for one second, with the predominate volume being directed to the single address 1.1.1.1.

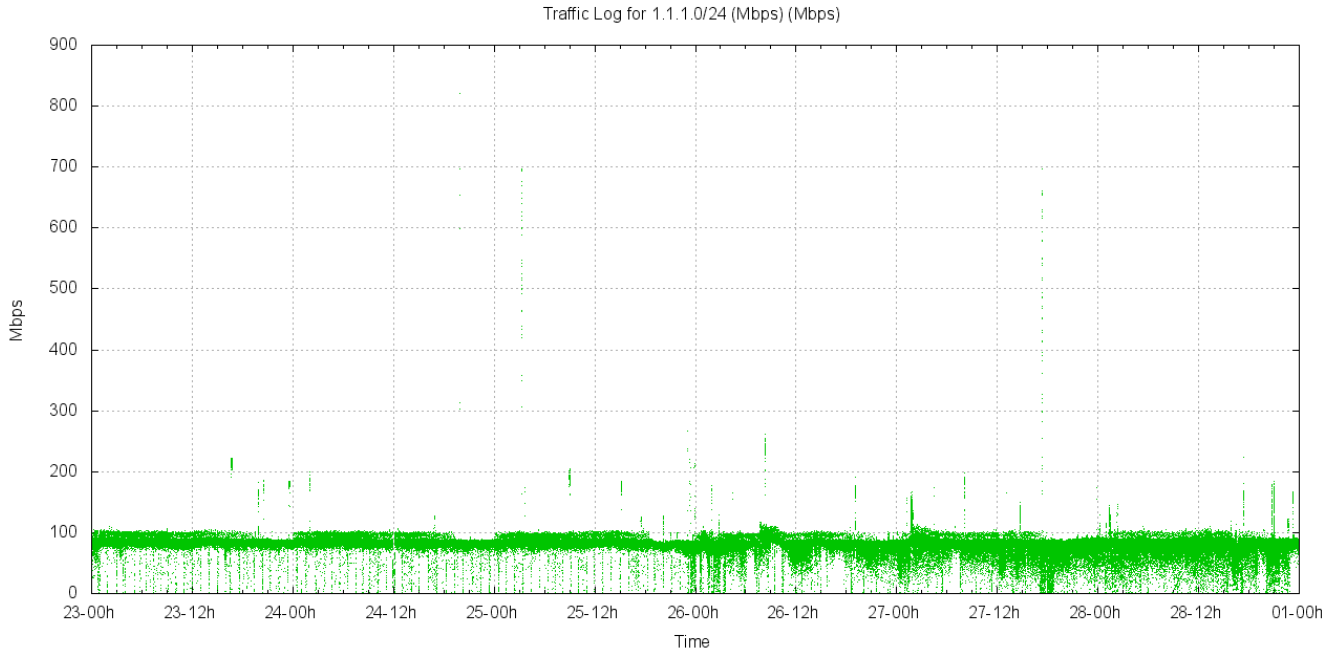

Figure 12 – Traffic levels for 1.1.1.0/24

In order to provide an overview of the traffic sent into each of the /24s in 1.0.0.0/8, the average traffic in each /24 has been divided into "bins" of 50bps increments and the number of /24 prefixes that fit into each of these "bins" is shown in Figure 13.

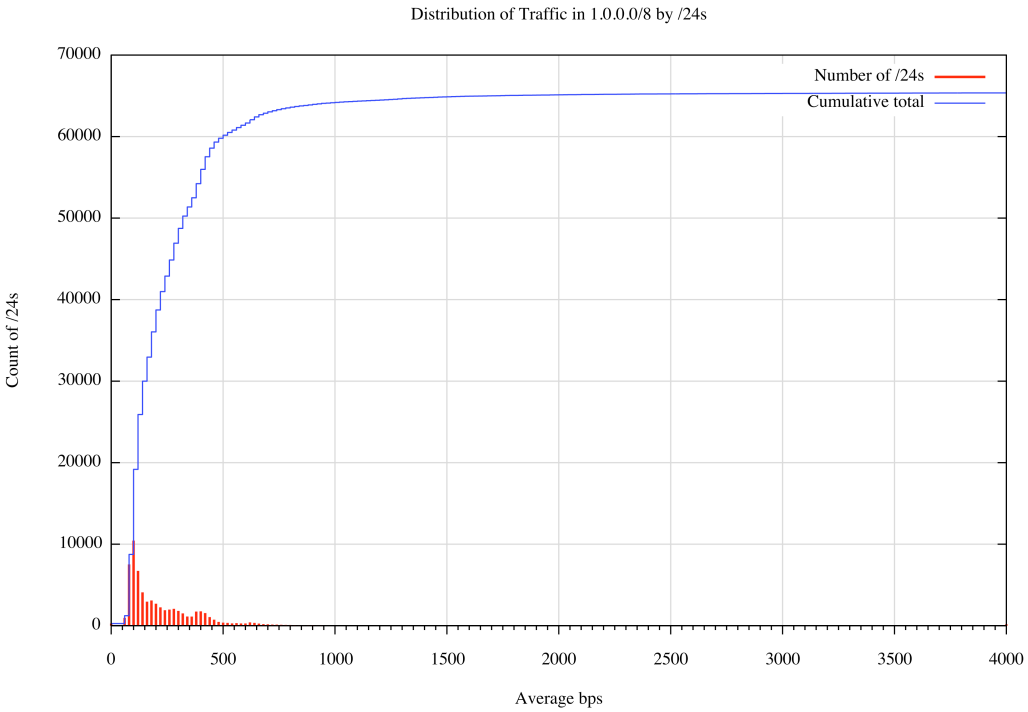

Figure 13 – Distribution of traffic in 1.0.0.0/8 per /24

The distribution has some similarity to an exponential decay function (which would be the case if the log of the distribution was linear). Figure 14 shows the same data presented using a log scale for the count of /24s in each average traffic level bin.

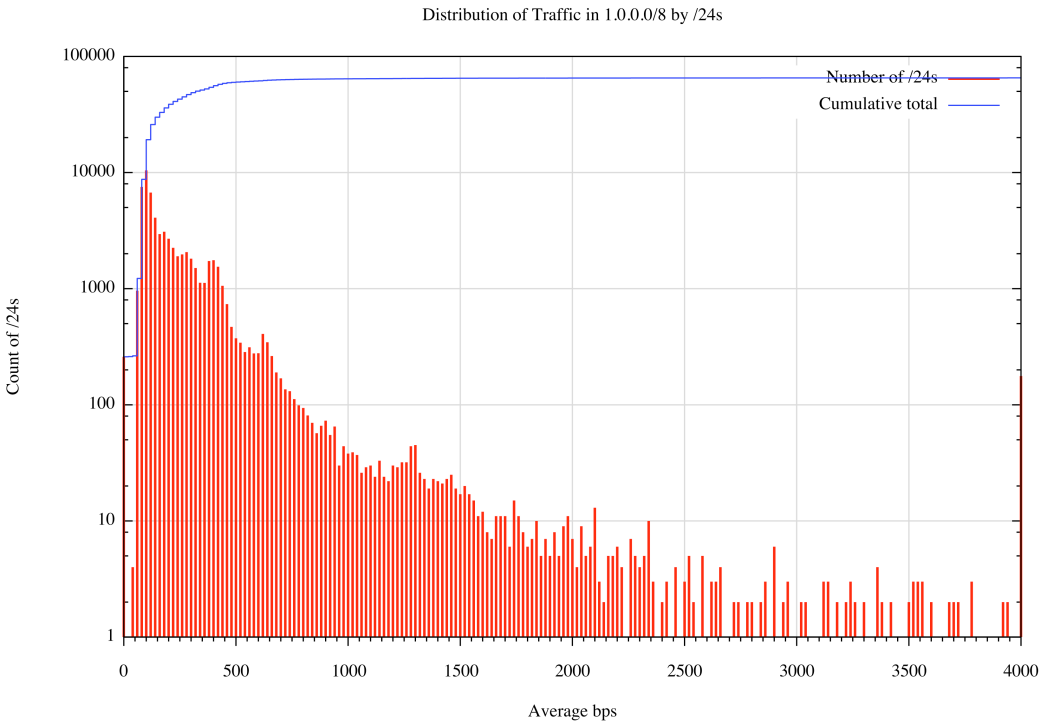

Figure 14 – Distribution of traffic in 1.0.0.0/8 per /24 (log scale)

In this distribution 98% of the /24s in 1.0.0.0/8, or 64,252 /24s received an average traffic load less than 132 bytes per second (or an average of <= 1 average-sized packet per second). The

following table shows the distribution of /24s by average-sized packet rates, using an average packet size of 132 bytes (as measured in the packet capture for 1.0.0.0/8).

Average Pkt ps	*Average bps*	*Number of /24s*	*Cumulative Total*	*% of total*
<= 1	<= 1056bps	64,252	64,252	98.0%
1 .. 2	<= 2112bps	916	65,168	99.4%
2 .. 3	<= 3168bps	195	65,303	99.6%
>= 3	> 3168bps	233	65,536	100%

Table III – Distribution of traffic by /24s

Control Point Comparison

A similar exercise has been undertaken for traffic in 27.128.0.0/12. The distribution of traffic per /24 is shown in Figures 15 and 16, and in Table IV.

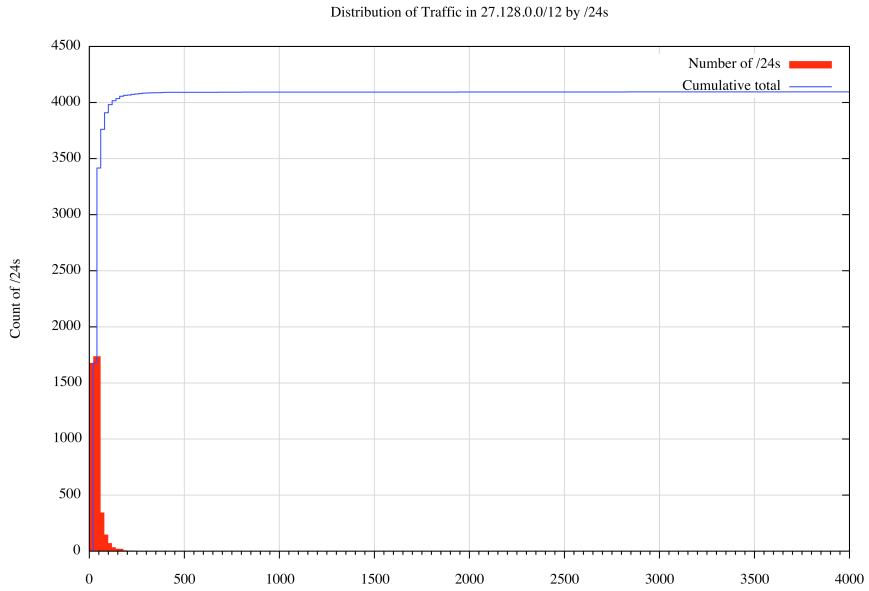

Figure 15 – Distribution of traffic in 27.128.0.0/12 per /24

Average Pkt ps	*Average bps*	*Number of /24s*	*Cumulative Total*	*% of total*
<= 1	<= 776bps	4,092	4,092	99.9%
1 .. 2	<= 1,552bps	1	4,093	99.9%
2 .. 3	<= 2,328bps	1	4,094	99.9%
3 .. 4	<= 3,104bps	1	4,095	99.9%
>= 3	> 3,104bps	1	4,096	100%

Table IV – Distribution of traffic by /24s in 27.128.0.0/12

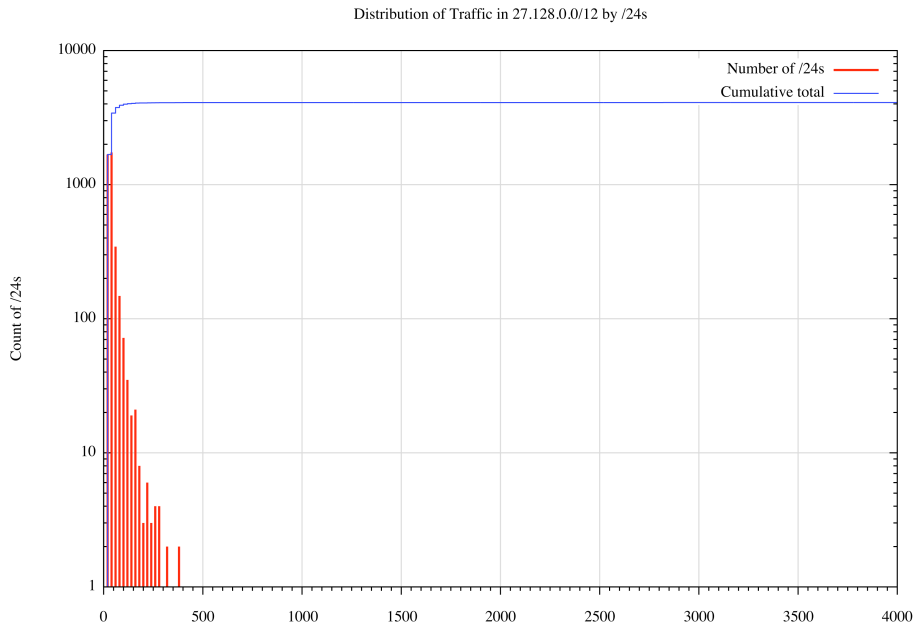

Figure 16 – Distribution of traffic in 27.128.0.0/24 per /24 (log scale)

In this case 4,087 /24s (or 99.8% of the /24s) received less than 258 bits per second, or the equivalent of 1 packet every 3 seconds (the average packet size for the 27.128.0.0/12 packet capture was 97 bytes).

AS36351 Announcement of 1.0.0.0/8

The experiment of passive listening of incoming packets addressed to network 1.0.0.0/8 was repeated for a further 6 hours on the 21st March 2010. This was undertaken by YouTube, and the prefix was originated by AS36351, collected using 4 systems and a load balancing front end. The profile of traffic gathered in this interval is shown in Figure 17.

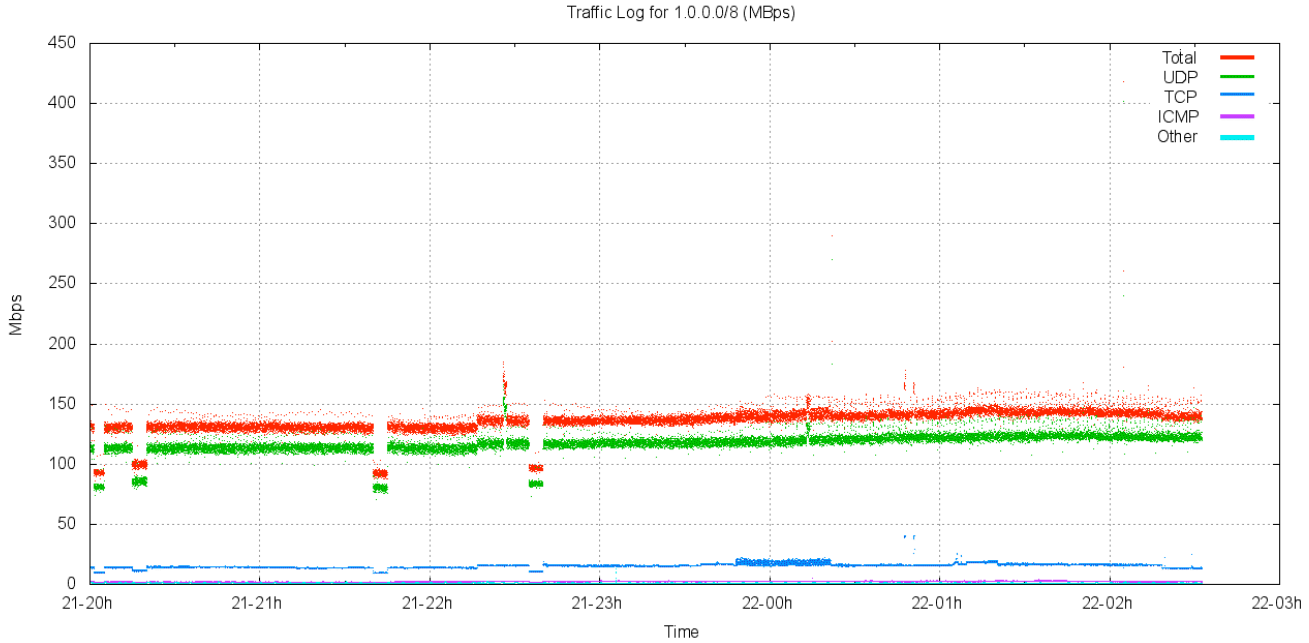

Figure 17 1.0.0.0/8 traffic to AS36351

The profile of traffic for this 6 hour period is comparable to the profile of the original AS237 announcement, with the traffic level positioned at some 150Mbps, with the majority of the traffic using the UDP protocol. This independent announcement and data collection by AS237 and AS36351 has confirmed the consistency of the traffic directed to this network, which allows us to have a higher level of confidence in the results of this analysis.

Summary of Results

Network 1.0.0.0/8 currently attracts an average of some 140Mbps - 160Mbps of incoming traffic, as a continuous sustained traffic level, with significant overload peaks. The traffic is largely UDP protocol traffic, with a smaller component of TCP, ICMP and other protocols. The TCP traffic is largely composed of TCP SYN packets, and not "leaks" from the interior of TCP sessions.

Traffic appears to be a combination of incremental scans that pass across part or the entirety of the addresses in this block and streams of UDP traffic addresses to particular individual addresses in the block.

The traffic in 1.0.0.0/8 is not evenly distributed. The majority of the traffic is directed at the single address 1.1.1.1, and the covering /24, 1.1.1.0/24, receives some 90 – 100Mbps of traffic as a continuous load, with isolated peaks of 1 second intervals in excess of 800Mbps.

Using a control point of traffic sent to a passive announcement of 27.128.0.0/12 to establish a "normal" background traffic level, it appears that 90% of the /24s in 1.0.0.0/8 receive the equivalent of a single average-sized packet every 3 seconds or longer.

Using a benchmark threshold level of 1 average-sized packet per second, or 1.056 Kbps, 98.8% of all /24s in 1.0.0.0/8 receive less than this threshold.

It appears that while individual addresses within 1.0.0.0/8 are heavily "polluted" with unwanted incoming traffic, this traffic is not evenly spread across the entire /8, and much of this address block sees a traffic level that is comparable with that directed to any other part of the Internet's address space.

Acknowledgements

This work would not have been possible without the generous assistance from our collaborators in this exercise. We wish to thank staff from the RIPE NCC, Merit, AARNet, Google and YouTube for both their patience and enthusiasm to provide us with the necessary facilities to undertake this work. Particular thanks are due to Nathan Hickson of YouTube and Steve Padgett of Google for their generous assistance.

References

[1] Leo Vegoda, "Awkward /8 Assignments", Internet Protocol Journal, September 2007. (http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_10-3/103_awkward.html)

[2] "Pollution in 1/8", RIPE Labs Report, RIPE NCC, 3 February 2010. (http://labs.ripe.net/content/pollution-18)

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

Authors

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

George Michaelson has a B.Sc from the University of York and is a research scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. George explores issues in Internet Number Resource management, Internet standards, and network measurement by collaborative research. George has over 28 years experience in Computer Science, Networking, ICT Administration and Research conducted in Australia and the UK. He participates in standards development in the IETF and has been a working group chair as well as an RFC author. He is a member of the BCS.