

What's going on in 1.0.0.0/8

George Michaelson ggm@apnic.net

Geoff Huston gih@apnic.net



“Standard” Address Testing

- IANA assigns /8 to APNIC
 - RIPE NCC, on APNIC’s behalf, announces selected subnets to test “reachability”
 - RIS, other tests applied.
 - Encourage operational community to test reachability
- APNIC releases /8 to registration services
 - Assignments and allocations proceed


APNIC gratefully acknowledges the assistance of the RIPE NCC and community supporting tests on AP ranges!


Except..

- IANA allocates Net 1.0.0.0/8 to APNIC in January 2010
 - We had some sense this was going to be different
 - Just how different wasn't clear
 -

(Not unexpected) First Warning

- RIPE announces 4 /24s for normal testing at AMS-IX
 - Link floods
 - announcements withdrawn.
 - Report written up on RIPE labs.





Lets Get Serious about Bad Traffic

- There is an issue here about 1/8 acting as a traffic magnet for unsolicited traffic
 - Just how “bad” is 1/8?
 - Are some bits REALLY bad?
 - What sort of badness are we seeing in the traffic?
- So we commenced a program to analyze the “badness” in 1/8


Bigger Badder Faster

- Need multi-gig collectors and large disk space
 - Exceeds APNIC's transit capacity
- Sought collaborators in R&D & Ops community
 - Many responses, for which we thank everyone
 - We worked with Merit, AARNet, Google and Youtube for this exercise
 - Rapid deployment possible, Tb file systems available with good IX connectivity within a couple of days – thanks!
- Long announcements. 1 week+
 - Explore >24h traffic patterns, diurnal behaviours


Tickling Badness

- Simple code to ACK all incoming TCP SYNs
 - If any follow up packet sent, that's interesting!
 - See if we can 'draw traffic out of the woodwork'
 - Distinguish one-way probes and DDoS engines, scanners from 'real' uses of the network
 - Based on Geoff's lightweight TCP experimental 'very bad idea' code
 - Not such a bad idea after all then!
- PCAP filter to collect all traffic, dump to disk
 - Applicable to passive and active/tickled capture


Traffic to 1.0.0.0/8



Traffic to 1.0.0.0/8




Packet Rate to 1.0.0.0/8



Packet Rate to 1.0.0.0/8


Peak Burst
at 220Kpps




Per subnet

- Sum packet counts seen per second, running average, promote to /16 and /24 counts
 - Rapidly identifies sub-spaces of the /8 range which have high traffic
- Establishes baseline load across entire net
 - But is it uniform?


Traffic Spread by /16



Traffic Spread by /16



Traffic Spread by /16



Yes, that's
a Log Scale!


The "hot spots" appear to lie in the low /16s

What is in all these packets?

Packet Size Distribution

- Most packets are very small (< 120 bytes)
- BUT 31% of the packets are exactly 200 bytes in length

Packet Size Distribution



IP Protocol Distribution

76%	UDP
20%	TCP
2.5%	ICMP
0.6%	6in4 (proto 41)
0.1%	GRE

This high concentration of UDP is unusual. Other networks see 55% TCP and 40% UDP in their levels of unsolicited incoming traffic.

UDP Port Distribution

<i>Port</i>	<i>Count</i>	<i>Description</i>
15206	45%	SIP response with RTP payload
33368	12%	some form of DNS?
0	6%	huh?
514	4%	syslog
80	3%	looks like firewall probing
33528	3%	pseudo-DNS again
3072	1.5%	and more of the same
53	1.4%	real DNS

port 15206?

```
08:48:36.000111 IP (tos 0x8, ttl 55, id 0, offset 0, flags [DF], proto UDP (17), length
200) 208.48.241.3.36670 > 1.1.1.1.15206: [udp sum ok] UDP, length 172
 0x0000:  4508 00c8 0000 4000 3711 7fe7 d030 f103  E.....@.7....0..
 0x0010:  0101 0101 8f3e 3b66 00b4 878c 8008 dbba  .....>;f.....
 0x0020:  cc7b 0288 55dd 8ce2 7a63 677b 7e66 6b14  .{..U...zcg{~fk.
 0x0030:  6962 1517 1613 1d05 0605 12d4 9d8c 8dea  ib.....
 0x0040:  6617 ef83 8d9e eee5 f85d 6050 919a 9758  f.....]`P...X
 0x0050:  6c66 49d6 5b4d dac5 c3d9 4453 c2d5 4d7a  lfI.[M....DS..Mz
 0x0060:  647f 7966 6f67 7360 1510 1d14 111f 0404  d.yfogs`.....
 0x0070:  6490 8e8f 9566 16ce 9b84 859b 93ef 6510  d....f.....e.
 0x0080:  4491 859c 5b6e 626d 7b4b 4ece d64d 4f7f  D...[nbm{KN..MO.
 0x0090:  5ac4 555d 4976 7b67 7b7c 7073 6e15 6c15  Z.U]Iv{g{|psn.l.
 0x00a0:  141b 0619 1b15 dd86 8e9b c514 d887 8399  .....
 0x00b0:  9e9c 9de0 637a e693 91d9 617b 7f4c 7764  ....cz....a{.Lwd
 0x00c0:  47d5 5e45 7c46 f7c4  G.^E|F..
```

UDP packets of 172 bytes in size, appears to point to some kind of audio streaming going on here

And SIP as well..

```
08:48:36.003010 IP 77.165.37.131.5060 > 1.1.1.1.5060: SIP, length: 486
 0x0000: 4508 0202 bf3e 0000 3411 507b 4da5 2583 E....>..4.P{M.%.
 0x0010: 0101 0101 13c4 13c4 01ee b101 5245 4749 .....REGI
 0x0020: 5354 4552 2073 6970 3a31 2e31 2e31 2e31 STER.sip:1.1.1.1
 0x0030: 2053 4950 2f32 2e30 0d0a 4672 6f6d 3a20 .SIP/2.0..From:.
 0x0040: 3c73 6970 3a47 6c6f 6261 6c55 4131 4031 <sip:GlobalUA1@1
 0x0050: 2e31 2e31 2e31 3a35 3036 303e 3b74 6167 .1.1.1:5060>;tag
 0x0060: 3d38 3061 6537 6530 302d 6330 6138 3031 =80ae7e00-c0a801
 0x0070: 6665 2d31 3363 342d 3435 3032 382d 3533 fe-13c4-45028-53
 0x0080: 6264 6234 2d34 3036 6638 3837 332d 3533 bdb4-406f8873-53
 0x0090: 6264 6234 0d0a 546f 3a20 3c73 6970 3a47 bdb4..To:.<sip:G
 0x00a0: 6c6f 6261 6c55 4131 4031 2e31 2e31 2e31 lobalUA1@1.1.1.1
 0x00b0: 3a35 3036 303e 0d0a 4361 6c6c 2d49 443a :5060>.
```

TCP Port Distribution

<i>Port</i>	<i>Count</i>	<i>Description</i>
21	40%	ftp
80	9%	http
1433	4%	ms-sql – (slammer lives!)
455	3%	ms-ds – (slammer again!)
6112	2%	?
25	2%	smtp

Who's Bad?

<u>/16 Address Prefix</u>	<u>Average Traffic(AS35361)</u>	<u>Average Traffic(AS237)</u>
1.1.0.0/16	86,757 kbps	79,981 kbps
1.4.0.0/16	19,714 kbps	12,564 kbps
1.0.0.0/16	10,241 kbps	8,816 kbps
1.10.0.0/16	3,656 kbps	3,320 kbps
1.2.0.0/16	3,611 kbps	12,010 kbp

?



Bad, or ...

Hanlon's Razor:

“Never attribute to malice that which can be adequately explained by stupidity.”

(or “cock-up before conspiracy!”)

A lot of this traffic appears to be leakage from private network domains

Some traffic is scanning, some is virus and worms, but the majority of traffic is leakage

Outcomes


- Holdback on the worst 5 /16s of net-1 recommended for the moment
 - Subject to ongoing testing
 - Parts of these blocks may become viable to release to community
 - Some parts clearly unusable for foreseeable future
- Ongoing tests of all new nets now part of APNIC's process

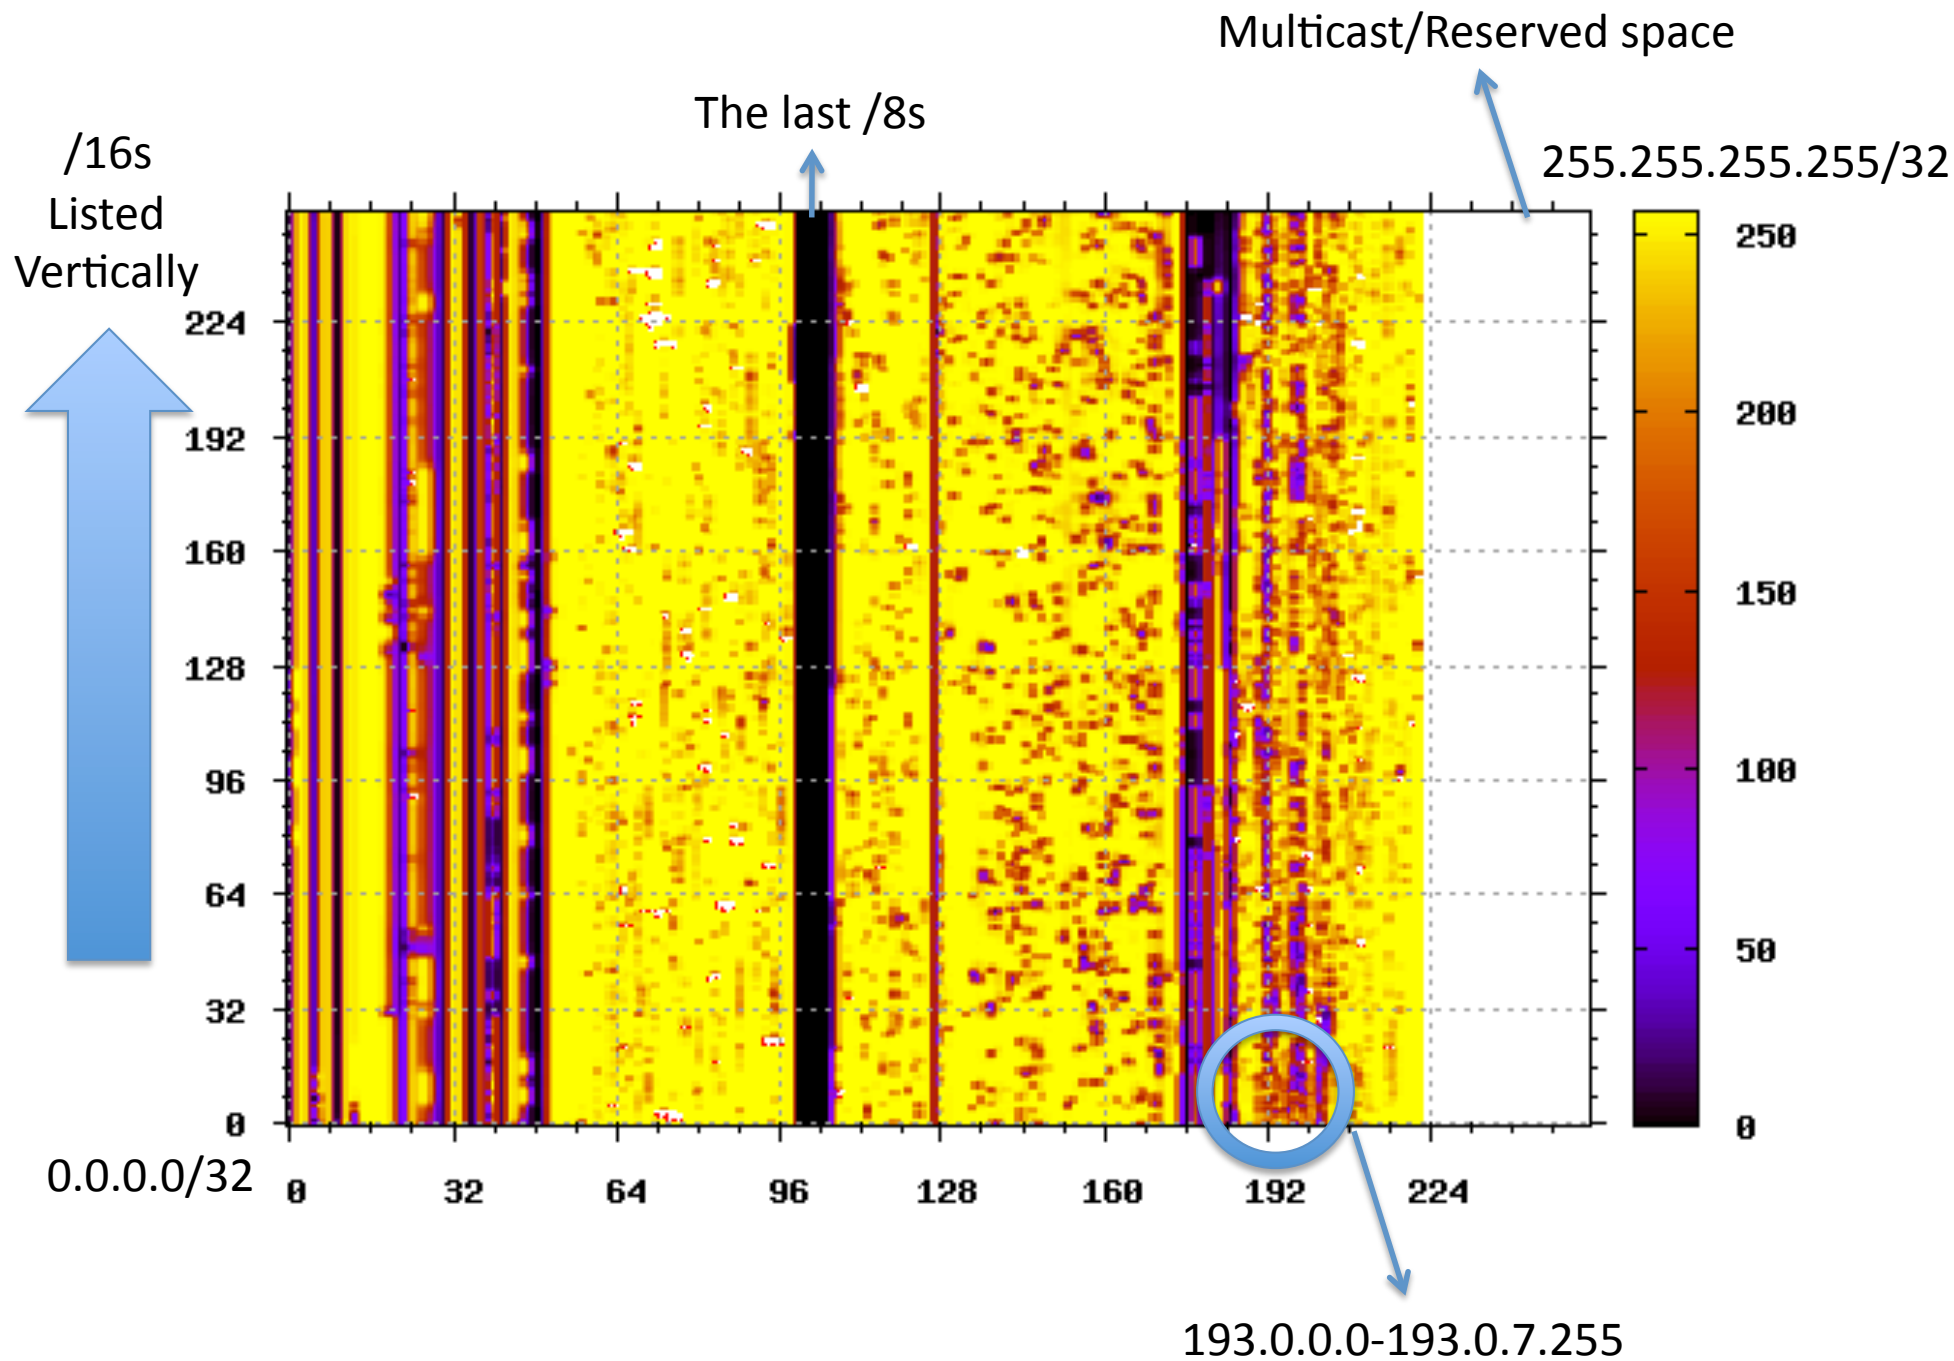
Visualization

- Look at the data in time-series, convert to movie
 - “see” the patterns of usage, identify subnets for further work
 - Applicable to net address and port
 - Allows side-by-side comparison of src, dst behaviours
- Easy to do, easy to understand

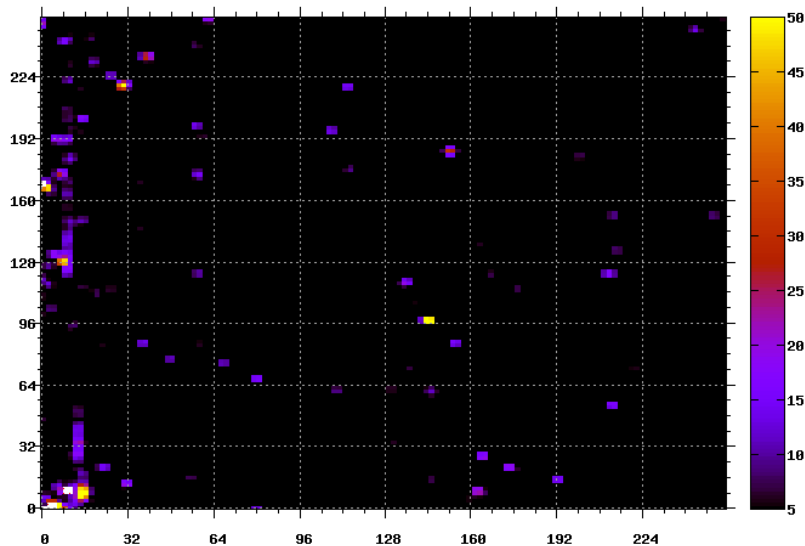
What does the net look like?

- Map of /16s, as {x,y} points
 - Colour denotes density of the /24 in use for that /16
- Thick Black stripes:
 - IANA reserve
- White stripe 224-255
 - Multicast/Reserved






What does net 1.0.0.0/8 look like?



- Map of /24s
 - as {x,y} points.
 - /16 run vertically
 - Colour denotes intensity
 - of traffic directed to the /24

Ranges of /24 inside /16 under higher traffic

/24s
Listed
Vertically



Hotspot: 1.0.0.0/24, 1.2.3.0/24 etc

Distinct /24 under high traffic

Lets go to the movies...

