IPv6 Pollution Traffic Analysis

Manish Karir
(DHS S&T Cyber Security Division)

Jake Czyz, Kyle Lady, Sam Miller, Michael Kallitsis, Michael Bailey
(University of Michigan)

Internet Pollution

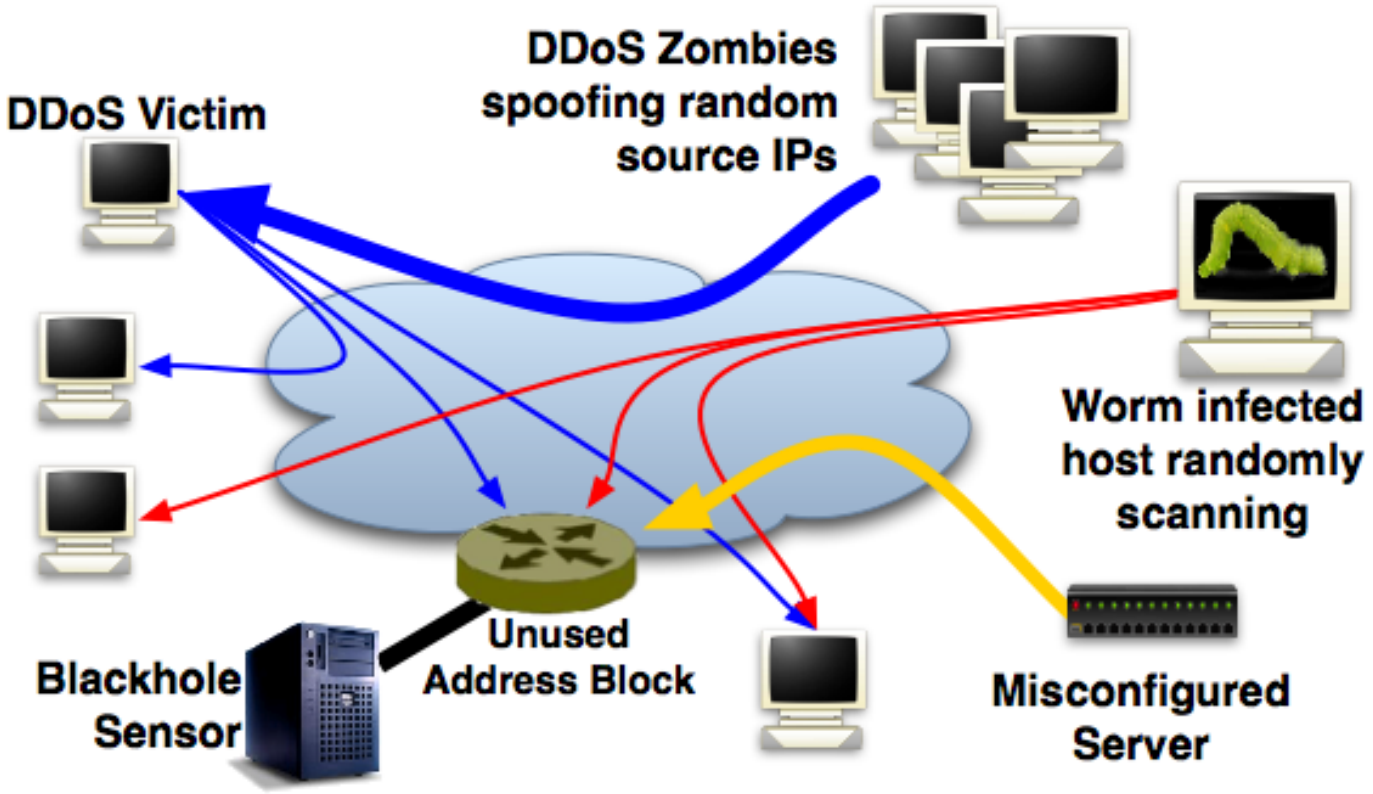

- Darknet sensors monitor unused address block
  - Receives traffic from DDoS backscatter, worm propagation, mis-configuration, and other scanning activity

Internet Pollution

- Traditional Internet Pollution
  - Worm scanning
  - DDoS backscatter
- Modern view of Internet Pollution

  (See Previous talk at NANOG 51)
  - Misconfigurations
  - Topology mapping scans
  - Software coding bugs
  - Bad default settings
  - Routing instability
  - Internet Censorship

IPv4 Previous Work

- We had previously conducted large scale Internet pollution studies for the following /8 network blocks:
  - 107/8,14/8,176/8,1/8,31/8,36/8/42/8,50/8
  - 100/8,101/8,105/8,177/8,181/8,23/8,37/8,45/8 49/8
  - 104/8,185/8
- Not all at the same time but in some cases as many as 5-6 /8 blocks at a time
- Well established processes/systems/techniques
- Long standing network telescope studies (Merit and CAIDA)

Internet Pollution in IPv6

- Previous Work:
  - Sandia Labs/APNIC: 2600::/12
  - Beginning 24 April 2012
  - “Turning Down the Lights” – DUST 2012

- How could we scale this up?
- Are there regional effects?
- Are there differences between unallocated and used address space?

Methodology: Understanding IPv6 Pollution Traffic

- Announcing 5 /12 prefixes(*)
- These are covering prefixes
  - Different from the previous work in IPv4
- Determine announcement visibility
- Determine data plane effects (port blocking?)
- Data analysis -> Report results to community

- Check to see if we broke the Internet (do this first!)

Coordination with RIRs

- Letters of Authority (LoAs) acquired from each RIR
  - 2400::/12 - APNIC
  - 2600::/12 - ARIN
  - 2800::/12 - LACNIC
  - 2A00::/12 – RIPE
  - 2C00::/12 – AFRINIC
- Permission to announce the covering /12 address blocks
  - Initially through 31 Dec 2012
  - Started announcing all five routes on 7 Nov 2012
  - Extension for observing long term trends

The Datasets

- Weekly data starting Nov 12 -Present
- Here: different subsets of this data
- 5 IPv6 /12 blocks – one for each RIR
  - 2400::/12 - APNIC
  - 2600::/12 - ARIN
  - 2800::/12 - LACNIC
  - 2A00::/12 (*) – RIPE
  - 2C00::/12 – AFRINIC
- Announced from AS 237 – Merit Network
- Coordinated with AS 7018 (ATT) and AS 6939 (Hurricane Electric)

*After an initial announcement, RIPE announcement was reduced to 2a04::/14 and 2a08::/13 (reduction of 25% of address space)

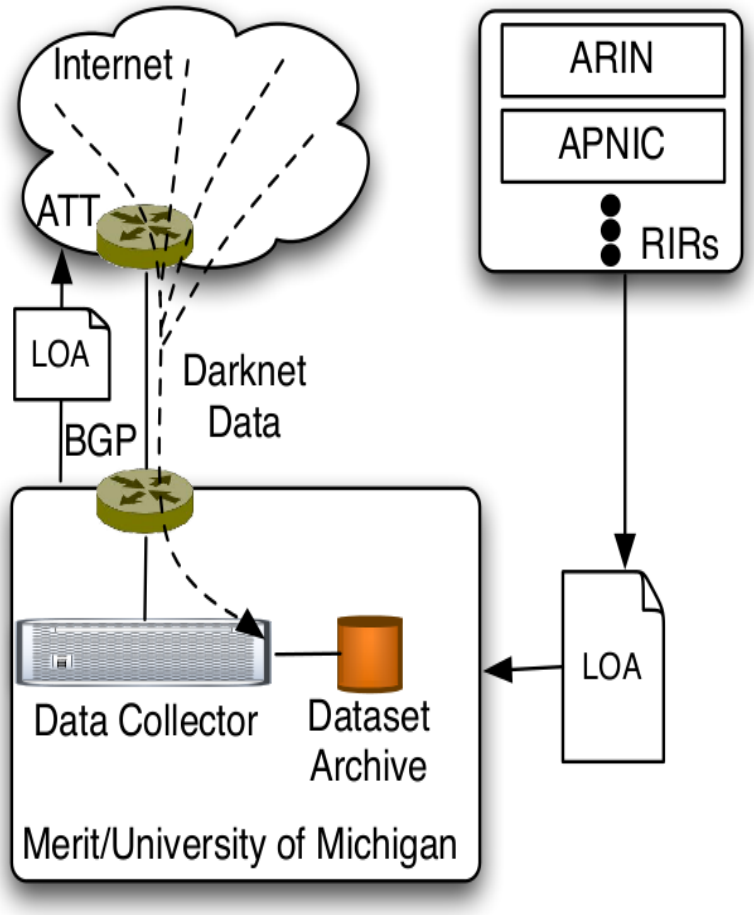

Validating Routing Visibility

- The announcements were visible from 8 of the 9 IPv6-capable monitors from the routeviews project
  - On average 74 out of 93
  - Not visible: KIXP in Kenya
- Also visible from 9 of the 12 v6-capable monitors maintained by RIPE
  - Not visible: MSK-IX in Russia, PTTMetro-SIP in Brazil
  - Partial visibility: DE-CIX in Germany saw 2 of the 6 routes
- Diminished visibility of RIPE / 12 starting in mid-January
  - Unclear why

Route Server	LACNIC 2800/12	ARIN 2600/12	APNIC 2400/12	RIPE 2a04/14+ 2a08/13	AFRINIC 2c00/12
r-v					
r-v.eqix	✓	✓	✓	✓✓	✓
r-v.isc	✓	✓	✓	✓✓	✓
r-v.jinx	✓	✓	✓	✓✓	✓
r-v.linx	✓	✓	✓	✓✓	✓
r-v.kixp					
r-v.saopaulo	✓	✓	✓	✓✓	✓
r-v.sydney	✓	✓	✓	✓✓	✓
r-v.telxatl	✓	✓	✓	✓✓	✓
r-v.wide	✓	✓	✓	✓✓	✓
rrc00	✓	✓	✓	✓✓	✓
rrc01	✓	✓	✓	✓✓	✓
rrc03	✓	✓	✓	✓✓	✓
rrc04	✓	✓	✓	✓✓	✓
rrc05	✓	✓	✓	✓✓	✓
rrc06					
rrc07	✓	✓	✓	✓✓	✓
rrc10	✓	✓	✓	✓✓	✓
rrc11					
rrc12		✓	✓		
rrc13					
rrc14	✓	✓	✓	✓✓	✓
rrc15					

Validating Data Path Continuity

- Goal: Ensure live hosts weren't affected by route announcements
- Ping 12k v6-capable hosts in diverse ASes during initial announcements (derived from Alex Top N lists)
- Confirmed no change in reach-ability of hosts

Probed IPs by Region

	IPs	ASNs
AfriNIC	9	8
APNIC	1622	603
ARIN	1219	530
LACNIC	159	62
RIPE	9409	3654

Validating no - Port Filtering

- nmapped dark addresses from ~5 hosts distributed around the world
- Occasional packet loss, as expected
- No ports consistently filtered
- Very different from v4
  - Windows-specific ports (e.g., dcom-scm on 135) are frequently filtered

Does the covering prefix matter?

Volume Differences w/o 2a00::/14

2a00::/12 12 Nov-13 Nov

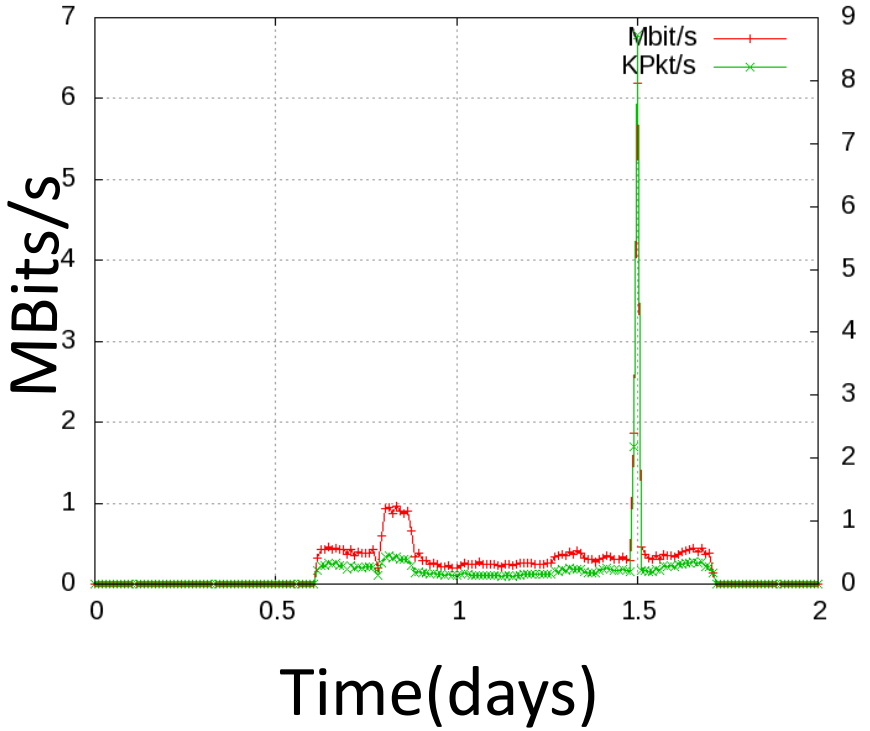

2a08::/13, 2a04::/14 3 Dec-9 Dec

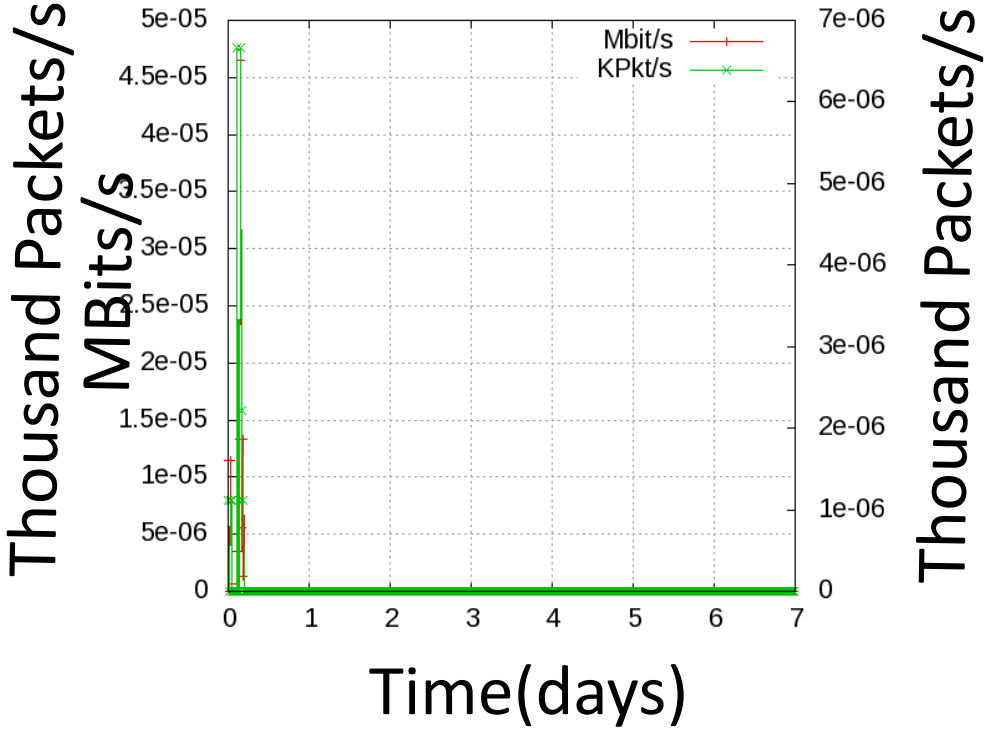

Withdrawing 25% of routed space resulted in orders of magnitude decrease in volume

Spatial Analysis
(week of 2012-11-19)

Traffic Volume:

APNIC and ARIN dominant (higher IPv6 adoption)

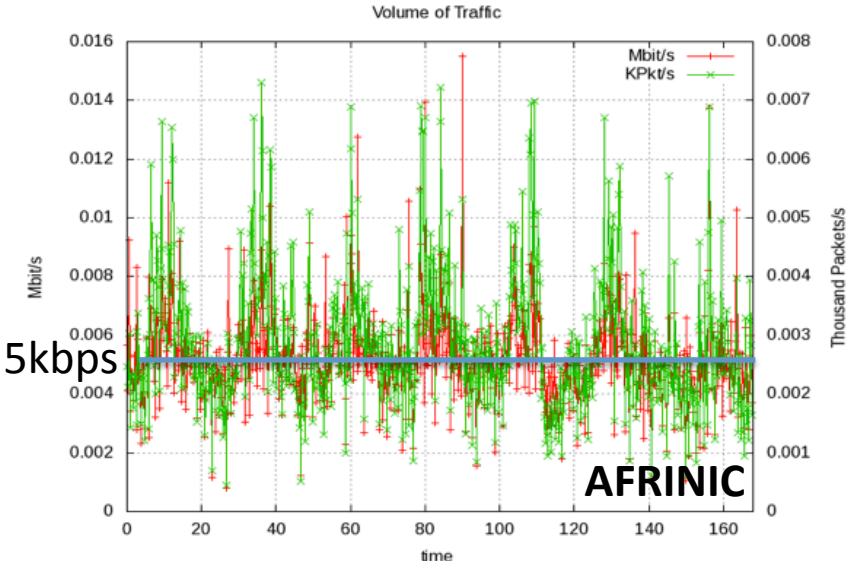

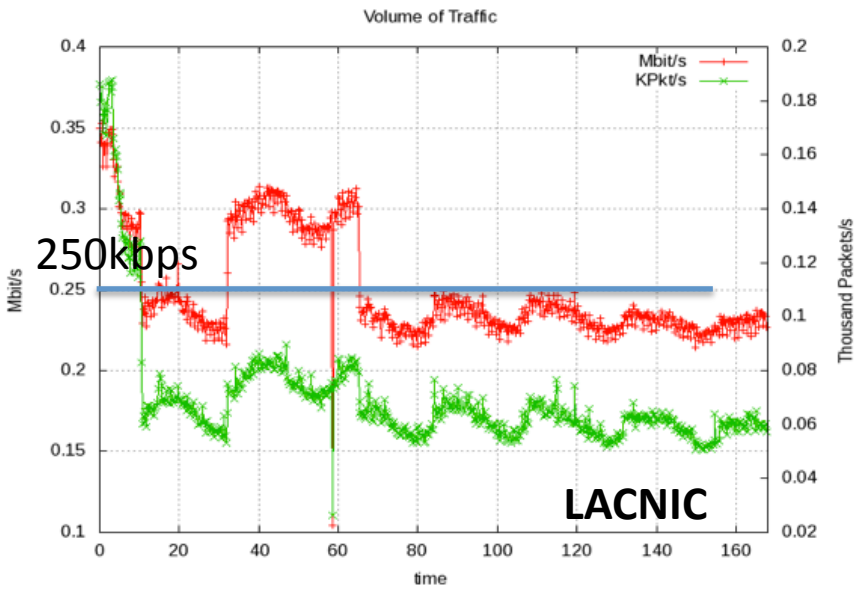

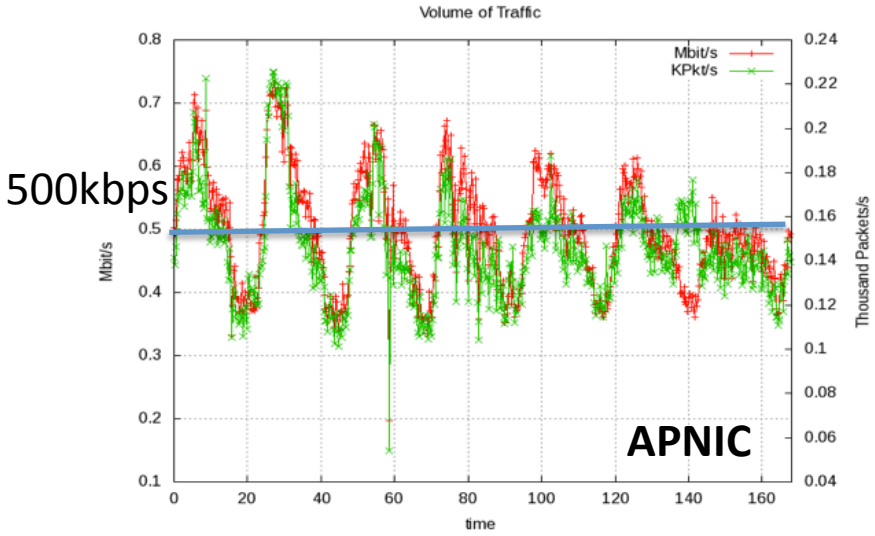

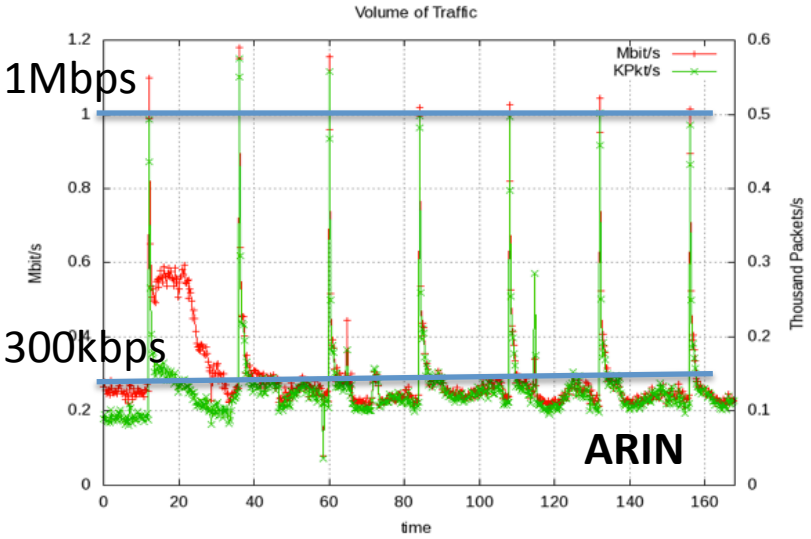

Traffic Breakdown by Protocol

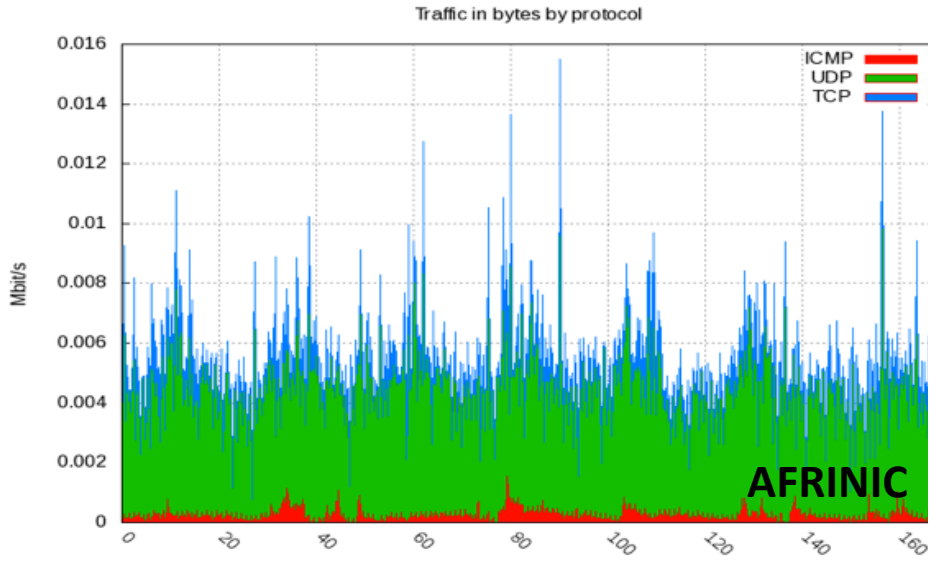

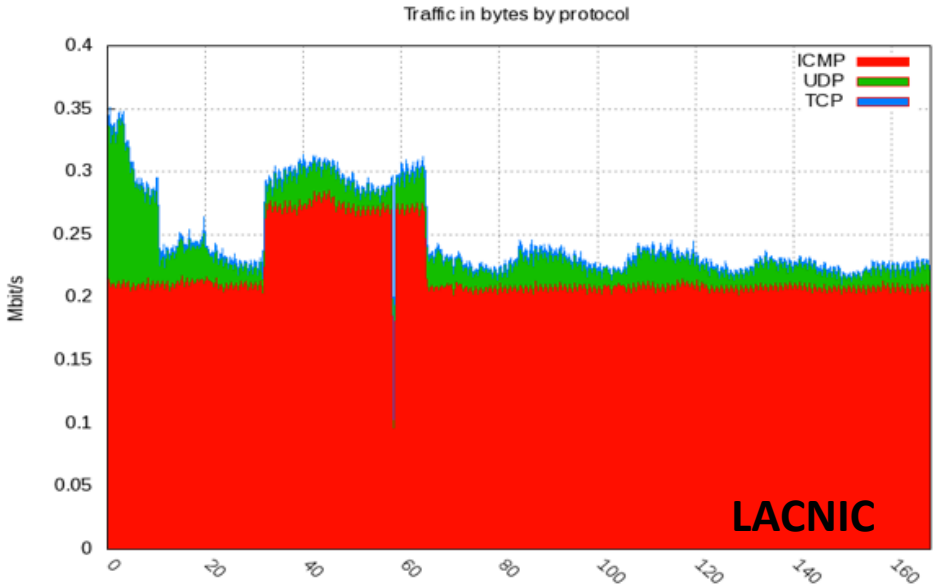

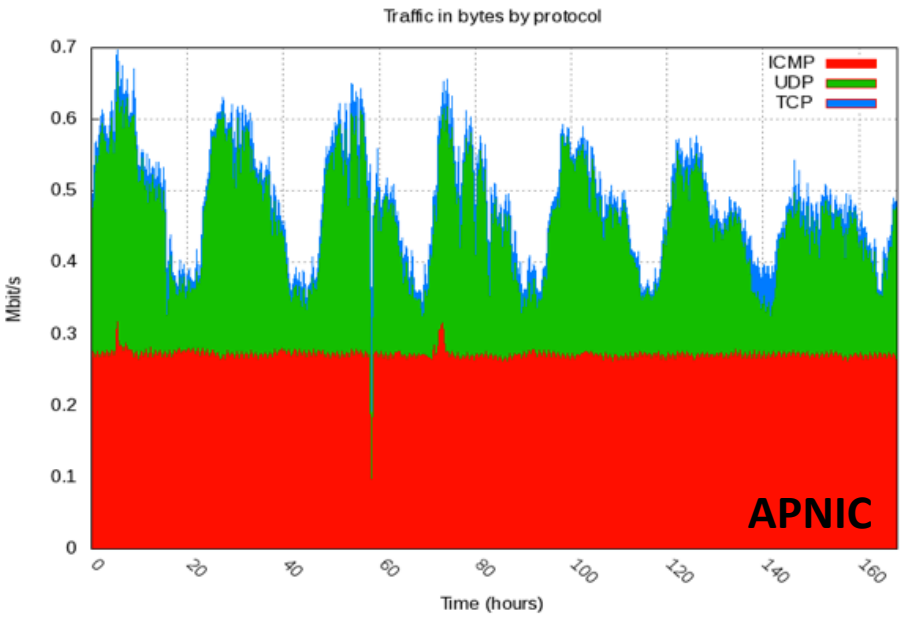

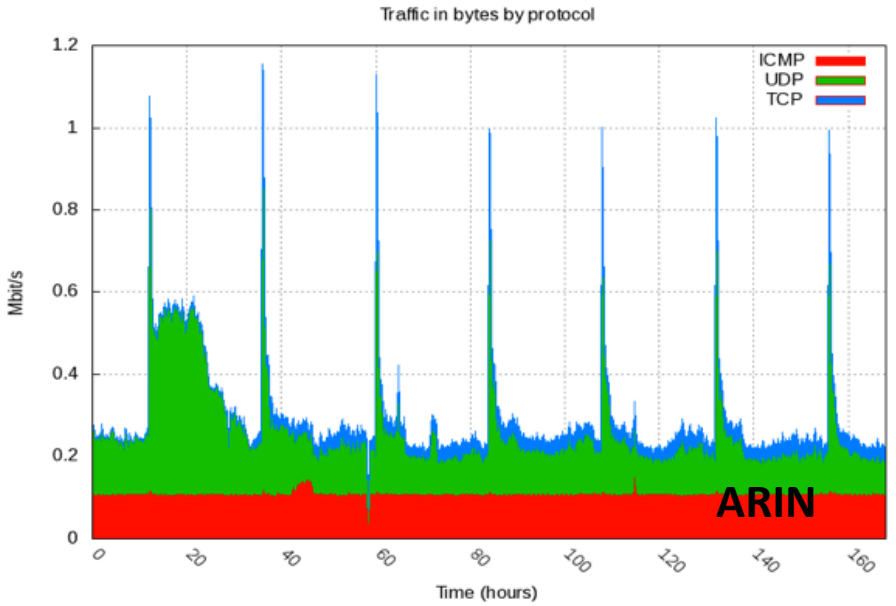

Long-term Trends

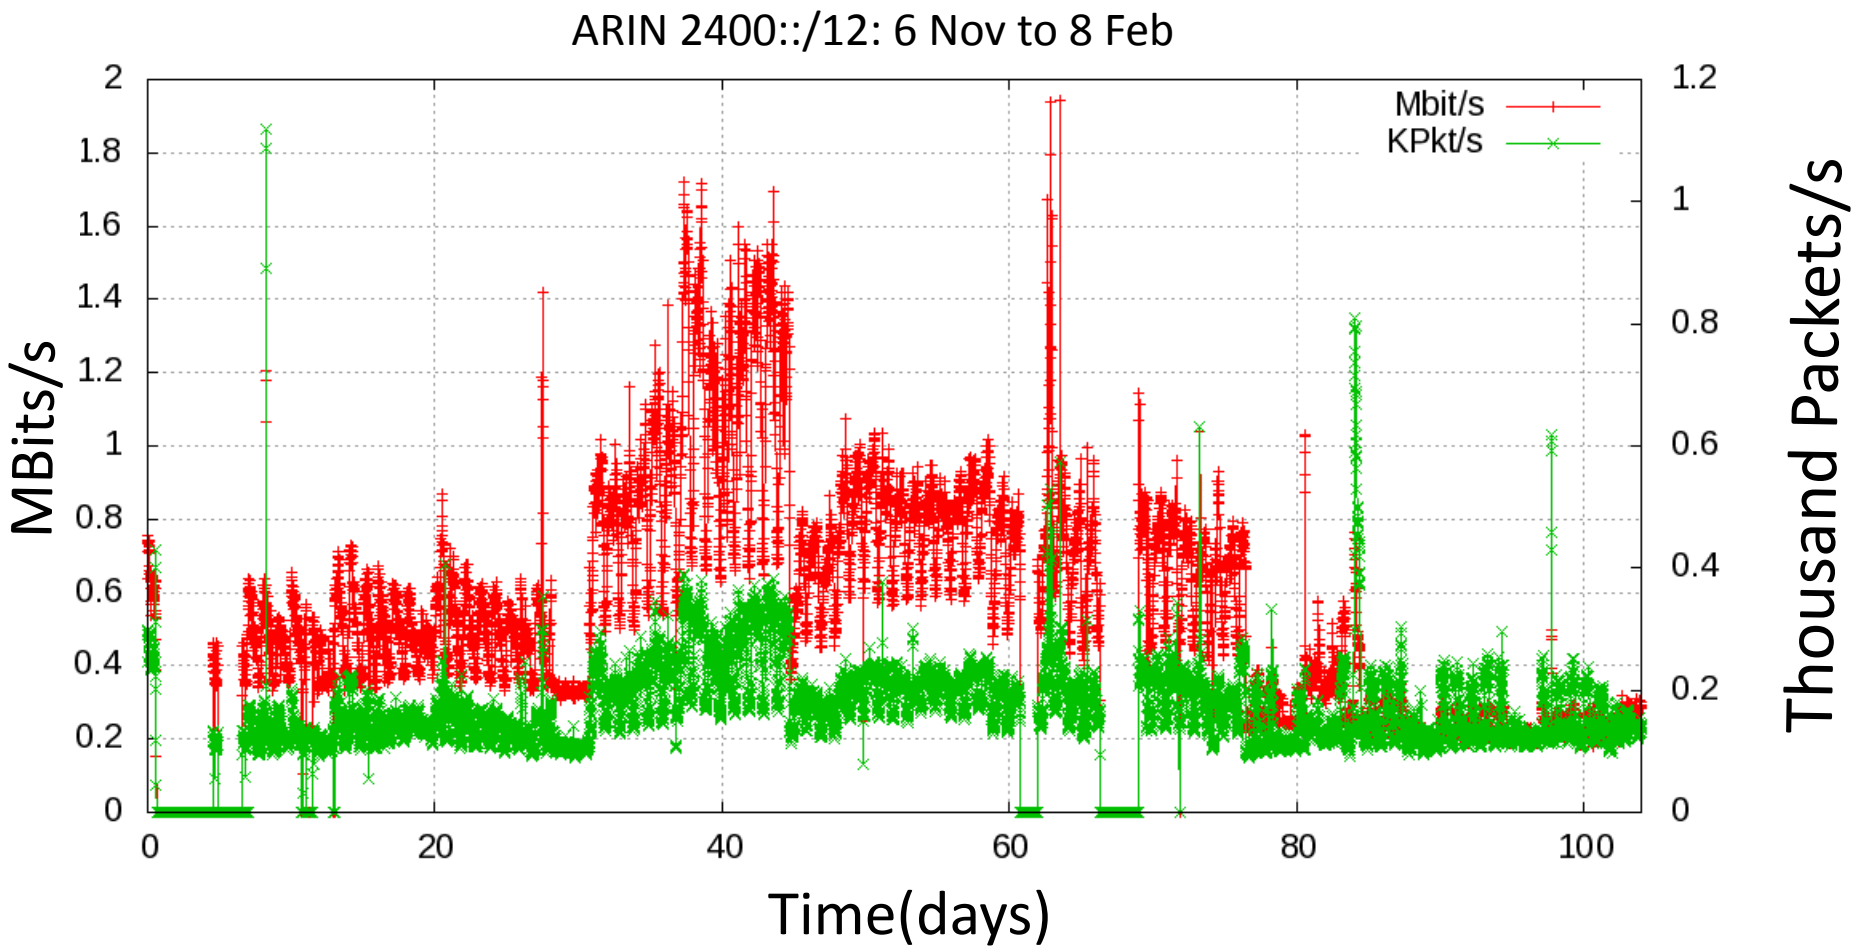

Top Destinations in the Traffic

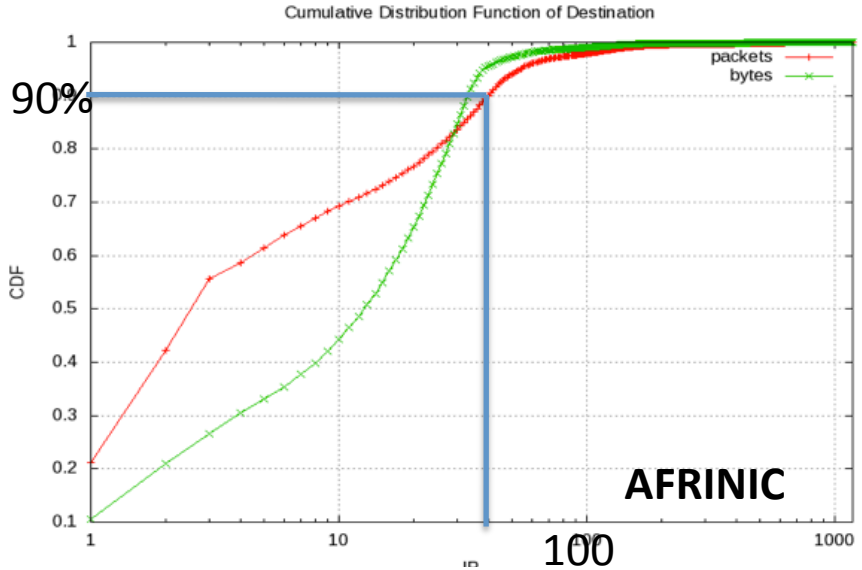

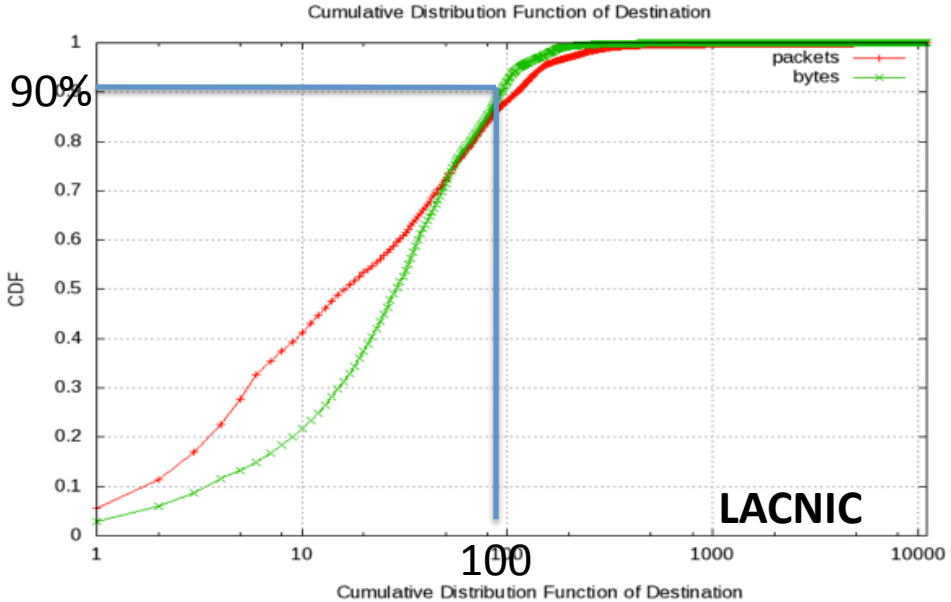

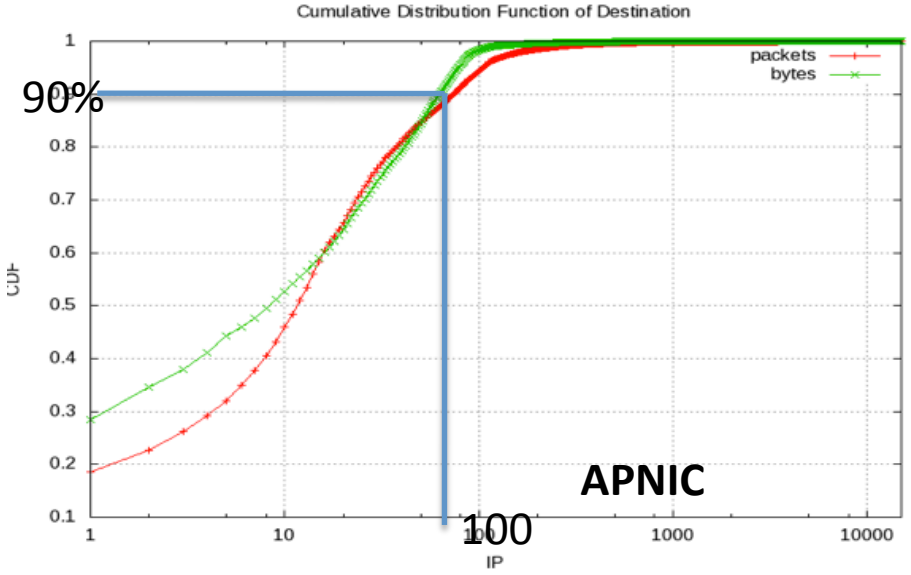

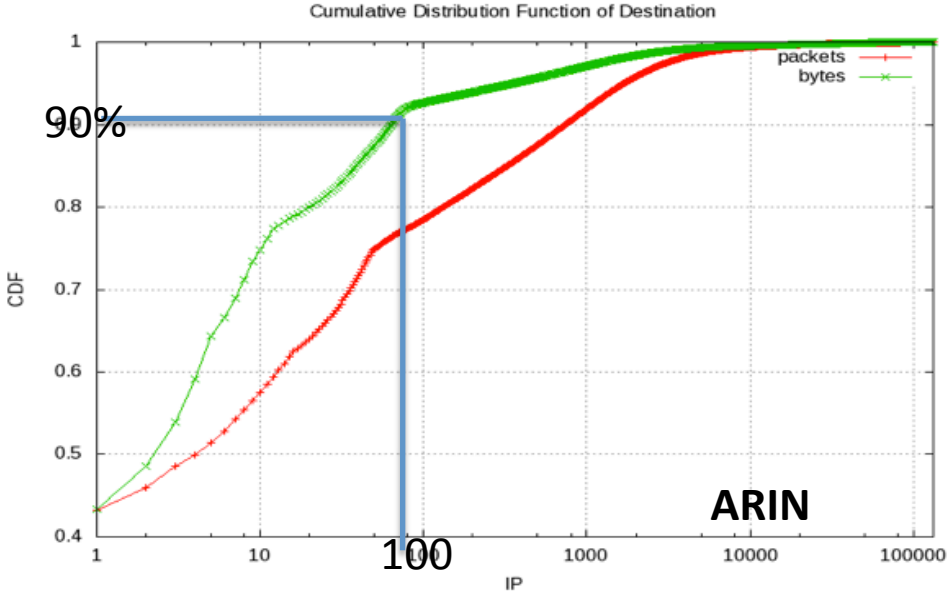

Top Sources

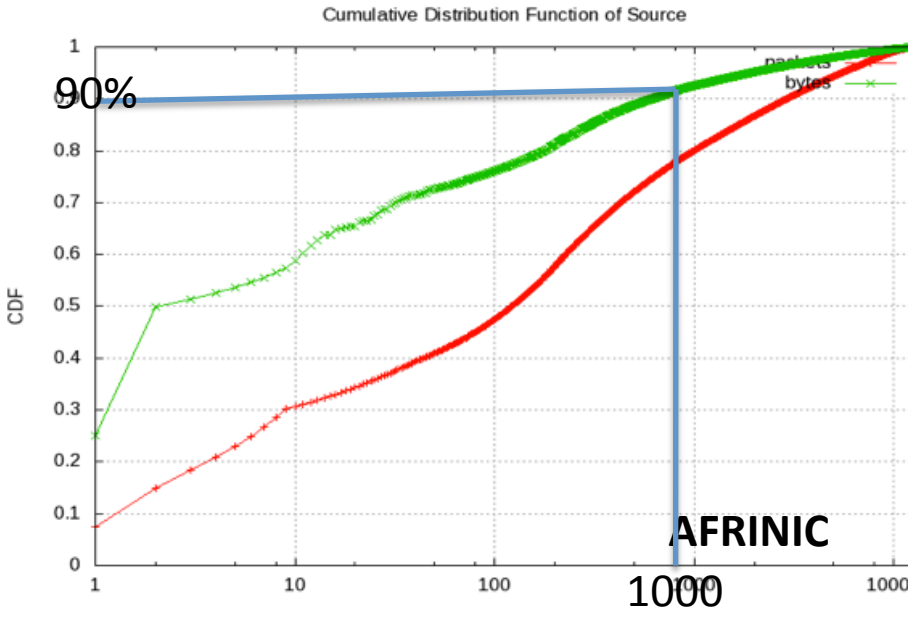
Cumulative Distribution Function of Source
CDF
packets
bytes
90%
AFRINIC
1000

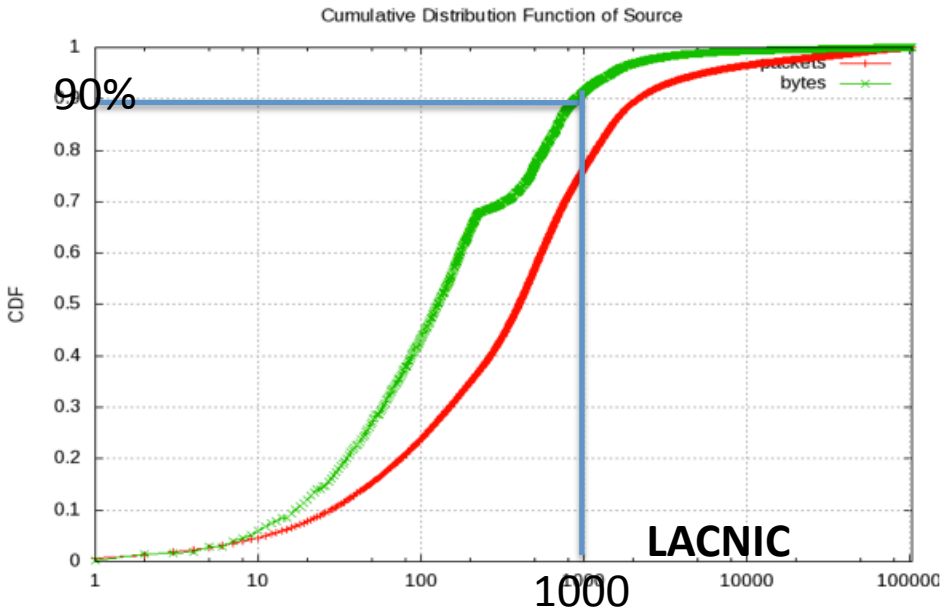
Cumulative Distribution Function of Source
CDF
packets
bytes
90%
LACNIC
1000

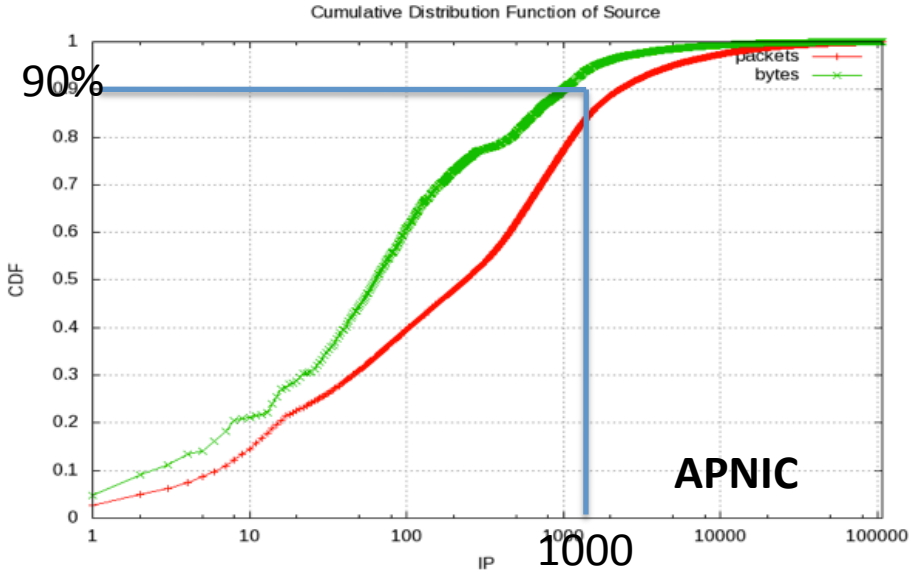
Cumulative Distribution Function of Source
CDF
IP
packets
bytes
90%
APNIC
1000

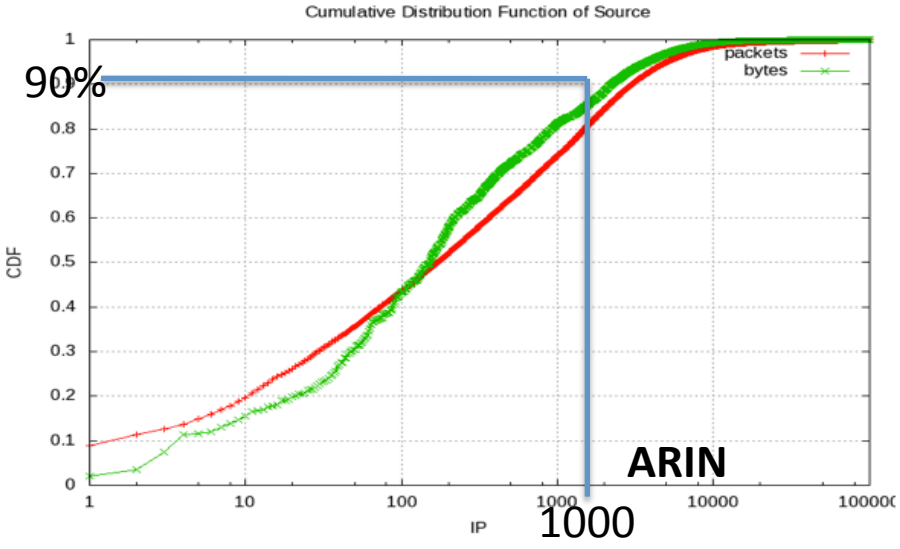
Cumulative Distribution Function of Source
CDF
IP
packets
bytes
90%
ARIN
1000

Time-to-live values for UDP

Most traffic from Linux sources

(default TTL values for Windows / Linux /Solaris = 128 / 64 / 255)

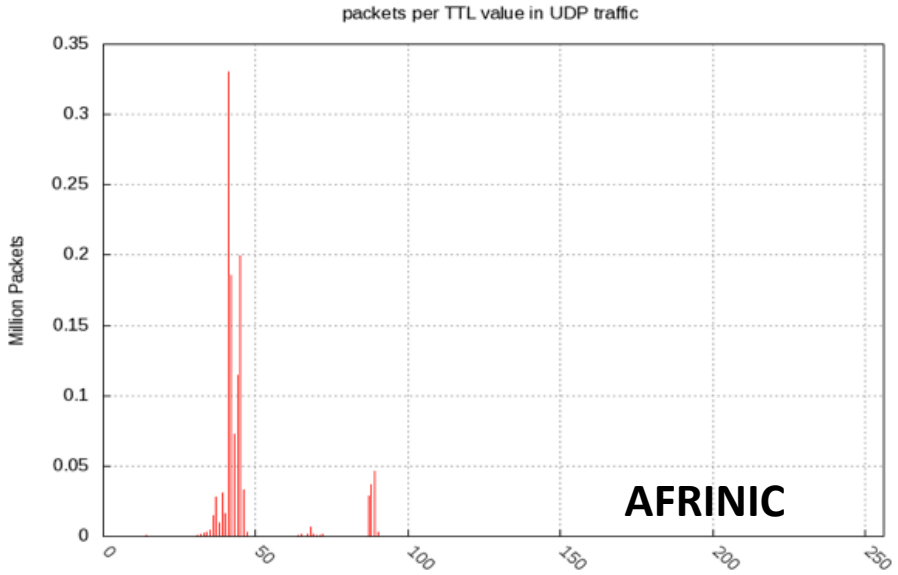

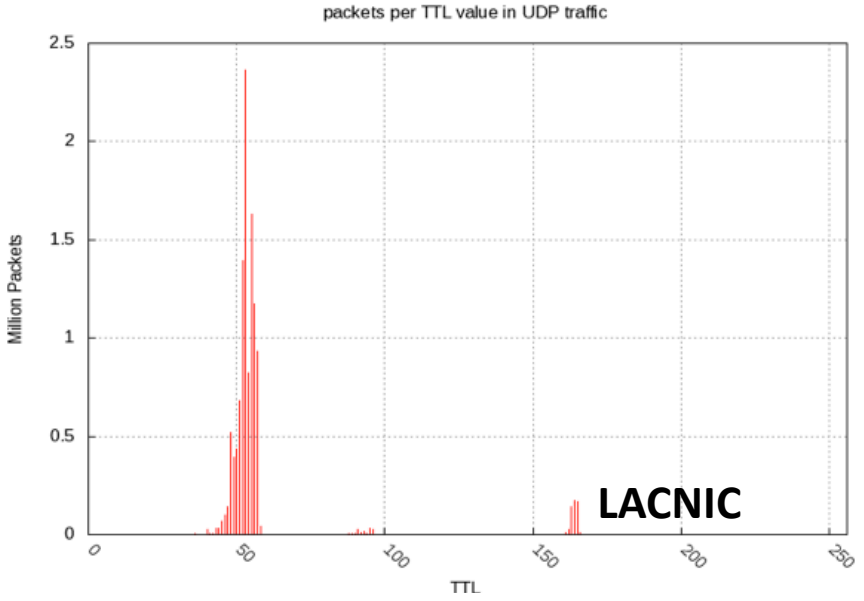

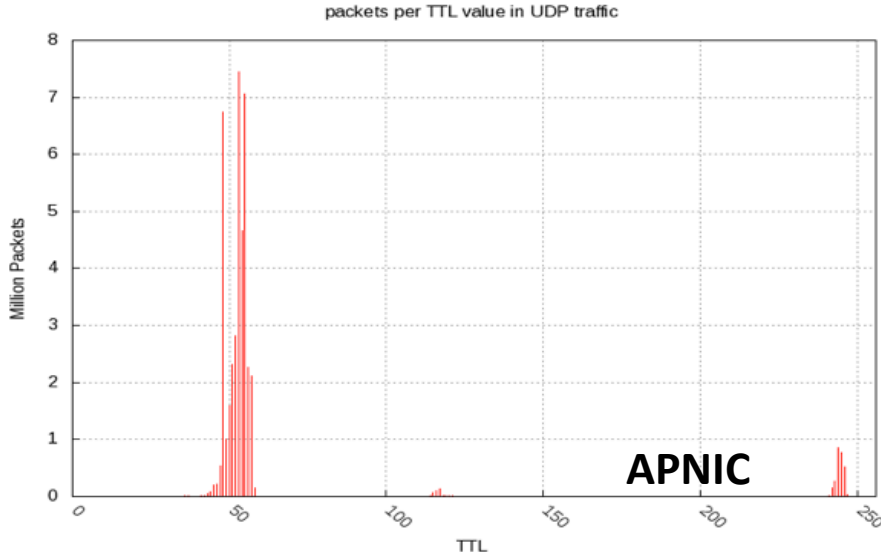

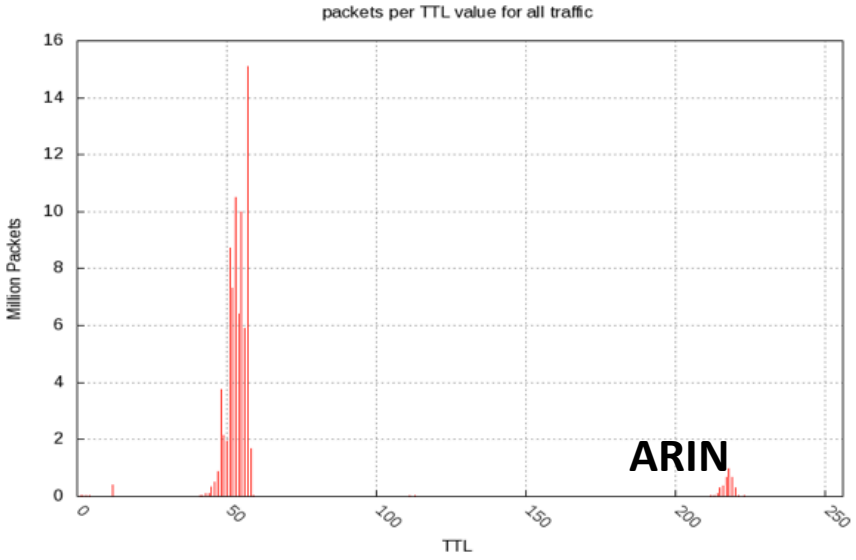

APNIC

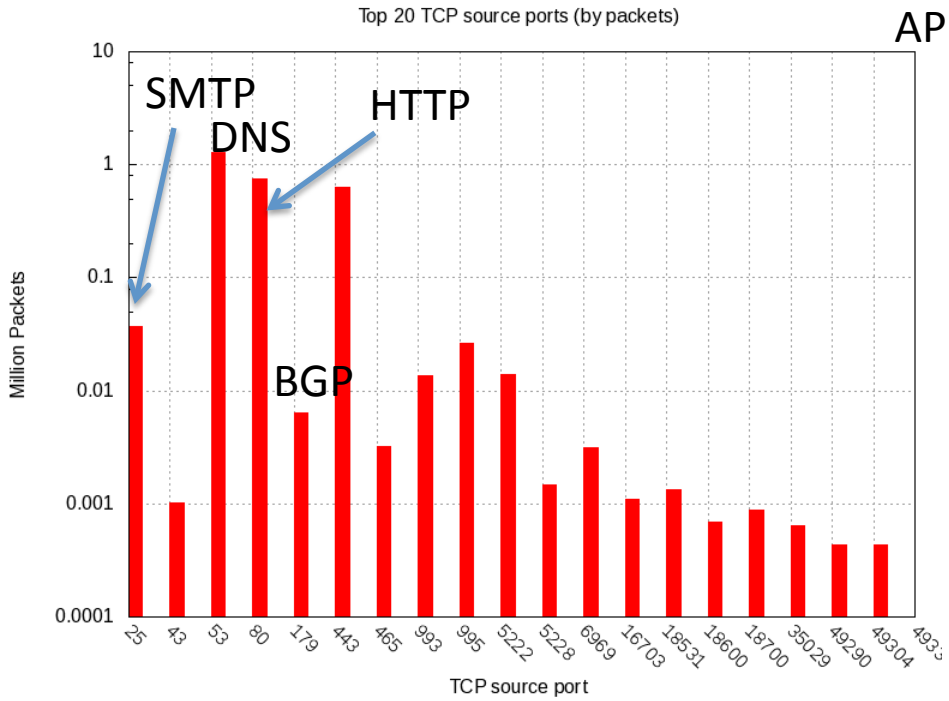
Top 20 TCP source ports (by packets)
SMTP
DNS
HTTP
BGP
Million Packets
10
1
0.1
0.01
0.001
0.0001
25 43 53 80 179 443 465 993 995 5222 5228 6969 16703 18531 18600 18700 35029 49290 49304
TCP source port

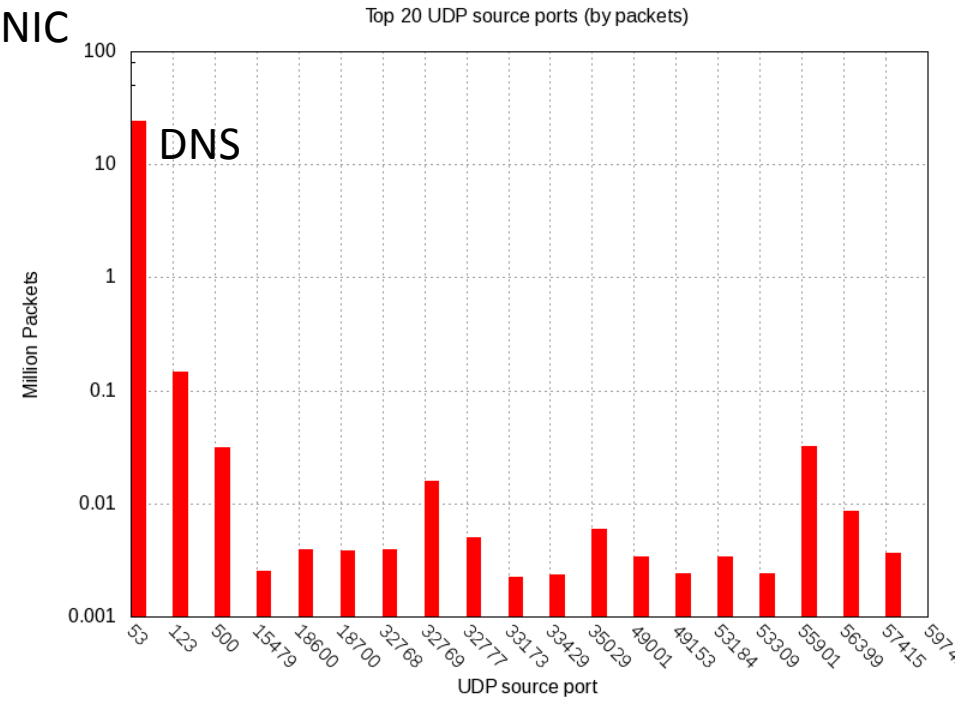
Top 20 UDP source ports (by packets)
DNS
Million Packets
100
10
1
0.1
0.01
0.001
53 123 500 15479 18600 18700 32768 32769 32777 33173 33429 35029 49001 49153 53184 53309 55901 56399 57415
UDP source port

ARIN

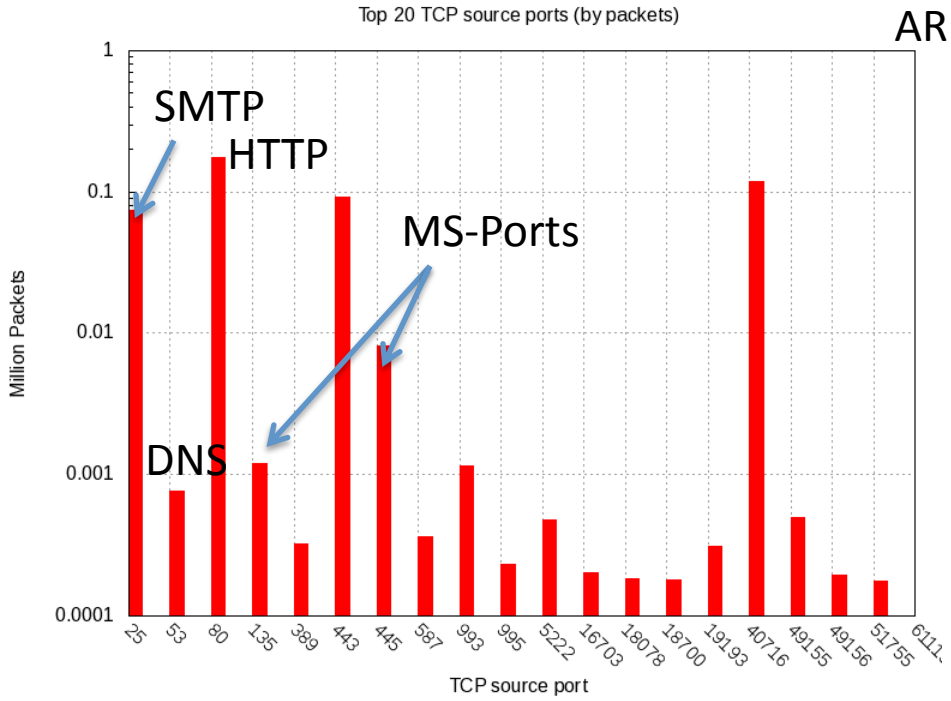
Top 20 TCP source ports (by packets)
SMTP
HTTP
MS-Ports
DNS
Million Packets
1
0.1
0.01
0.001
0.0001
25 53 80 135 389 443 445 587 993 995 5222 16703 18078 18700 19193 40716 49155 49156 51755
TCP source port

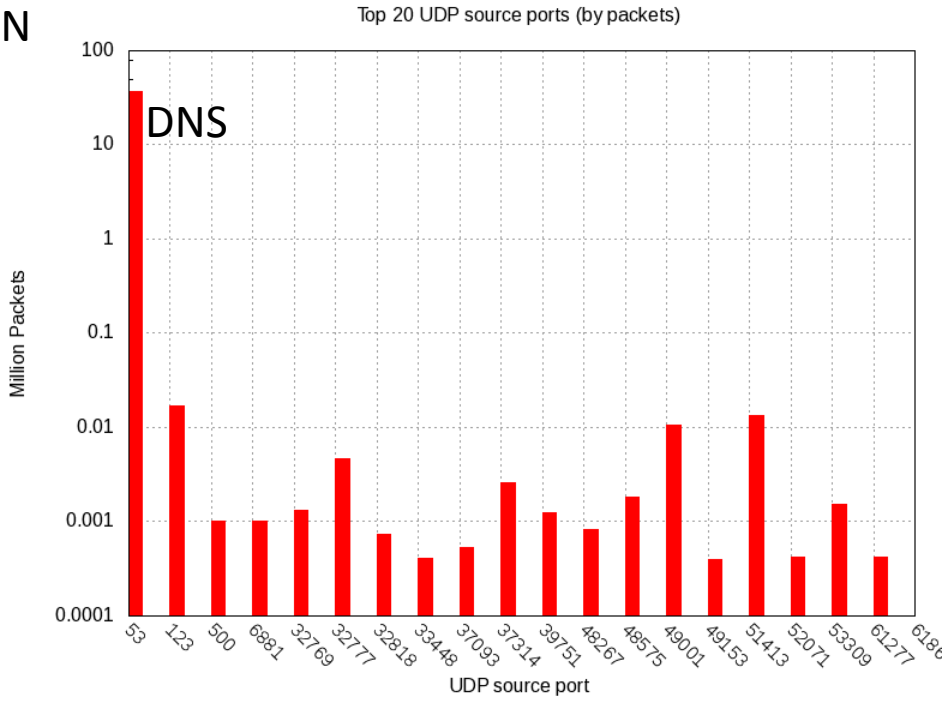
Top 20 UDP source ports (by packets)
DNS
Million Packets
100
10
1
0.1
0.01
0.001
0.0001
53 123 500 6881 32769 32777 32818 33448 37093 37314 39751 48267 48575 49001 49153 51413 52071 53309 61277
UDP source port

LACNIC

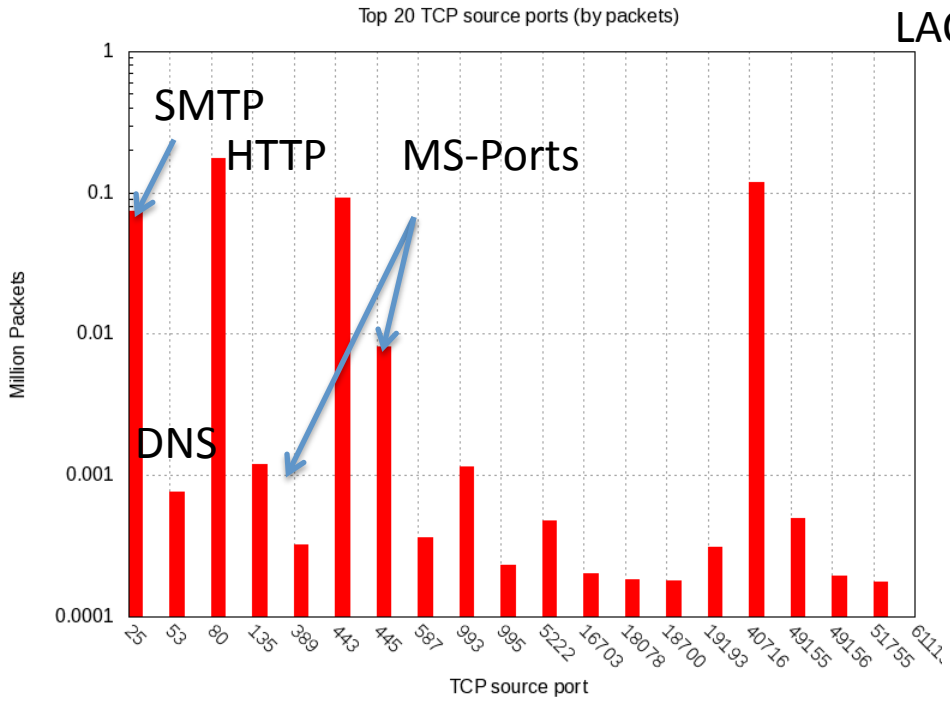

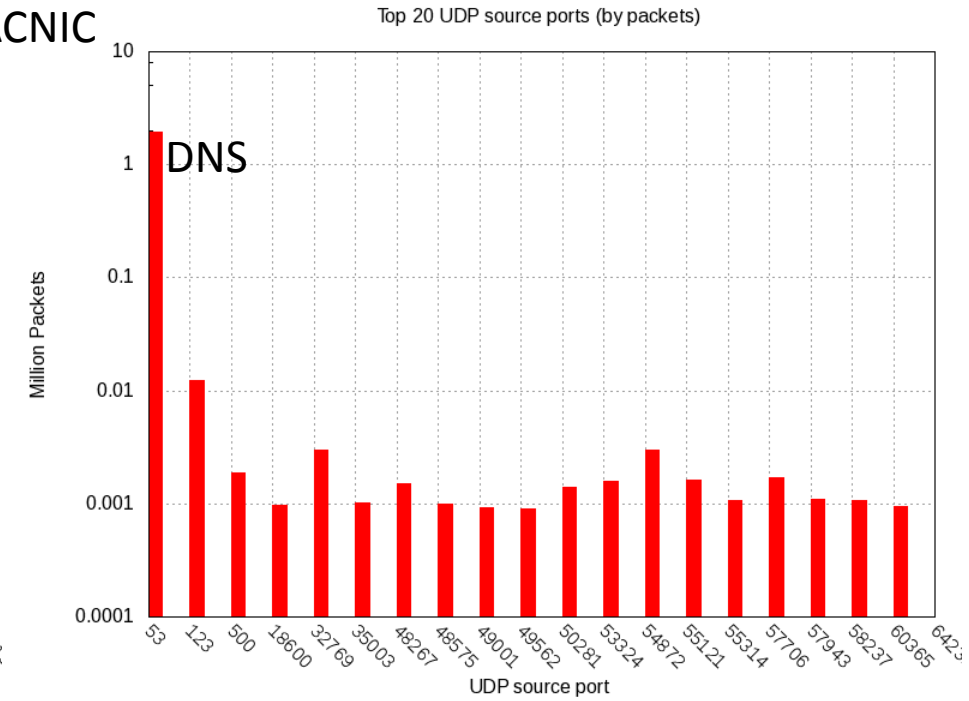

AFRINIC

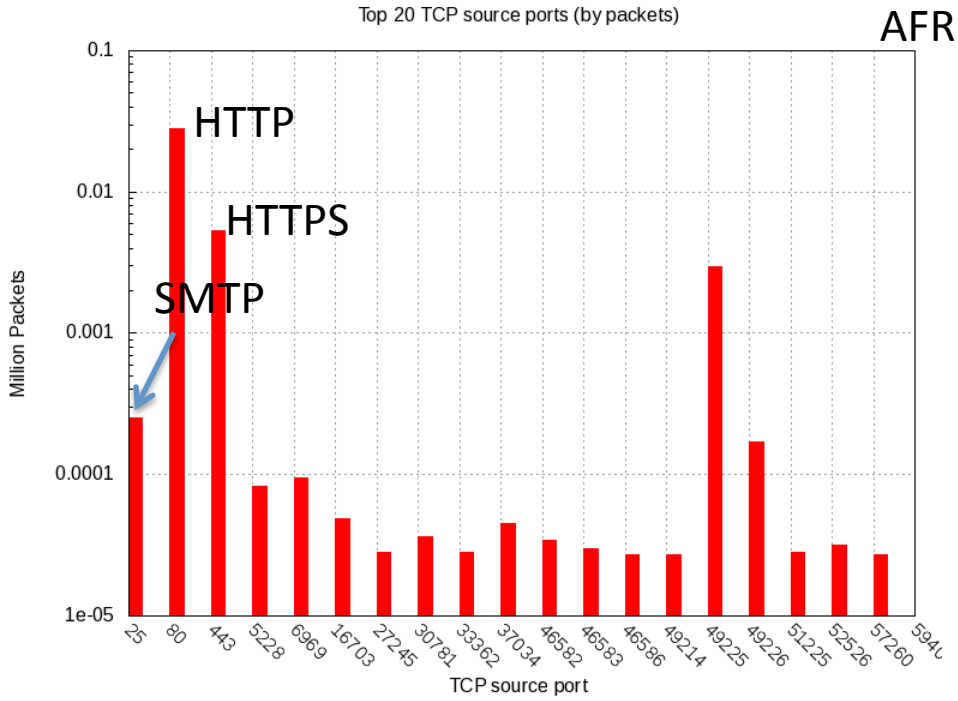

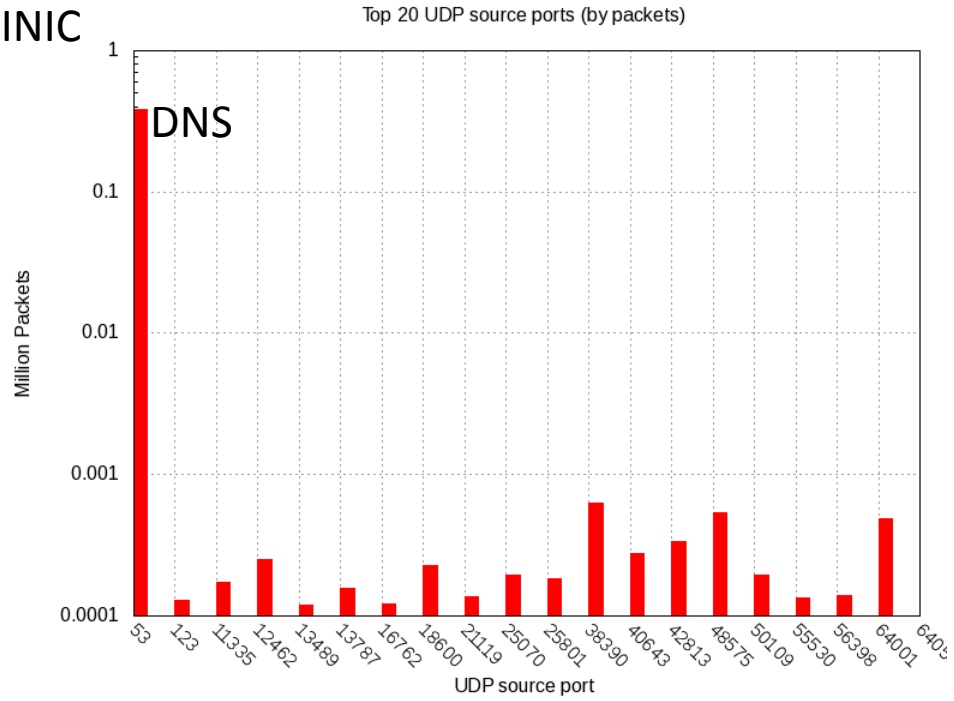

Case Studies

Worm Activity/Scanning?

- Some minor amounts of traffic on slammer/conficker ports (3 month dataset)
- Slammer signature does not match the traffic
- No signs of varying destinations for port 445 traffic single src and destination
- ICMP Probing/Scanning
  - Over 6K unique sources sending >1K ICMPv6 (APNIC), 3.2K (ARIN), 3.9K (LACNIC), 0.8K (AFRINIC), 0 (RIPE)
  - Clear evidence of sequentially scanning but generally limited to smaller subnets rather than /0 or /12
  - Akamai sourced ICMPv6 activity also visible e.g. a single IP send 2.5M packets to 141 unique destinations

Link-local addresses?

- We see over 800 unique link-local addresses as the source address in our dataset (3 month dataset)
- In one case we see a single IP address send over 71M ICMPv6 packets to roughly 27 unique destinations (cycle)
- If we see link-local addresses it is likely IPv6 address spoofing will work from those networks as well
- Check your filters (BCP 38 for IPv6?)

NTP/BGP Services

- We are able to identify data for both NTP and BGP in our datasets (3 month dataset)
- NTP traffic from over 4.7 unique sources – but in clusters
  - 800 from AT&T, 750 from Verizon Wireless, 870 from Edgecast
  - In all three of these cases clients are attempting to reach lara.nono.com (ARP networks Inc operated time-server in IPv6 pool.ntp.org)
- BGP traffic from over 330 unique sources
  - Appear to be legitimate BGP traffic as the addresses usually belonged to loopback interface Ips

SMTP Traffic

- SMTP traffic from 4.3K unique email servers (3 month datset)
- 2.4K in APNIC, 0.9K ARIN, 1.2K in LACNIC, 0.13K in AFRINIC, 5 in RIPE data
- Email severs attempting to reach other email servers (Google/comcast email servers)

DNS Traffic

- One of the largest contributors to pollution traffic (3 month dataset)
- Roughly 50% of ALL IPv6 announcing ASNs appear to be sending some DNS traffic to our darknet monitor
- AS6939 (HE) tops the list with 55K unique sources, ATT (AS7287) – 23K, Edgecast -13K, PROXAD – 9K, and OVH – 8K are in the top 5 with over 5K unique IPs each
- We observe both DNS queries as well as responses

DNS Queries

- Number of queries:
  - 176M – APNIC
  - 75M – ARIN
  - 71M – LACNIC
  - 6.9M - AFRINIC
- Sources of queries:
  - 85K – APNIC
  - 59K – ARIN
  - 30K – LACNIC
  - 7.6K – AFRINIC
- Only 134 queries in the RIPE region dataset

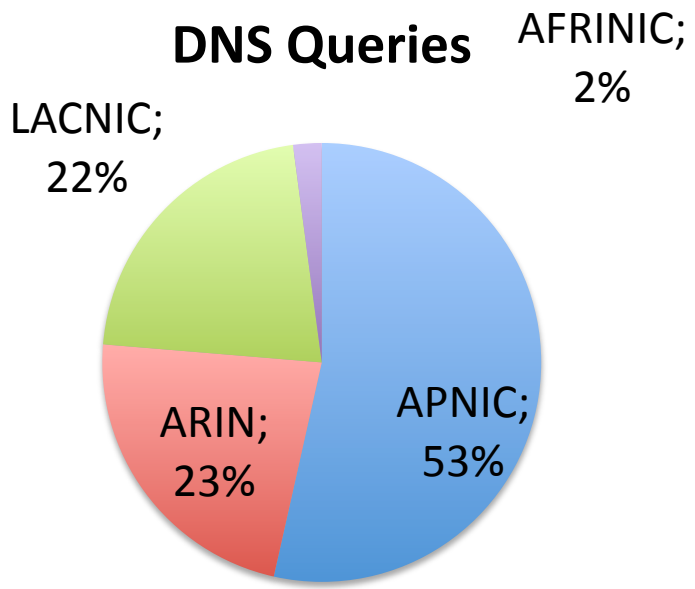

DNS Responses

- Number of response packets:
  - 450M – APNIC
  - 365M – ARIN
  - 73M – LACNIC
  - 3.9M – AFRINIC
- Sources
  - 16K – APNIC
  - 16K - ARIN
  - 9.8K - LACNIC
  - 3.3K - AFRINIC

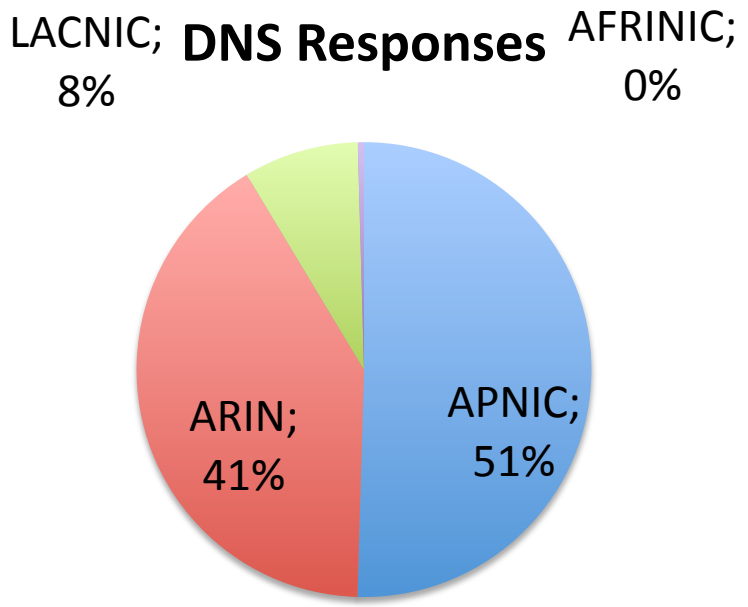

- We observe no responses in the RIPE region dataset
- 54% of APNIC region responses are from DNS root servers
- 5% of all ARIN region responses are from a single resolver operated by RIPE, 4% from 2 resolvers operated by Comcast
- 18% of LACNIC region DNS responses are from servers operated by ARIN
- Some are DNS-based block list traffic from bit.nl (22M – APNIC, 2.5M ARIN, 6.4M LACNIC)

Periodic spikes in UDP DNS traffic

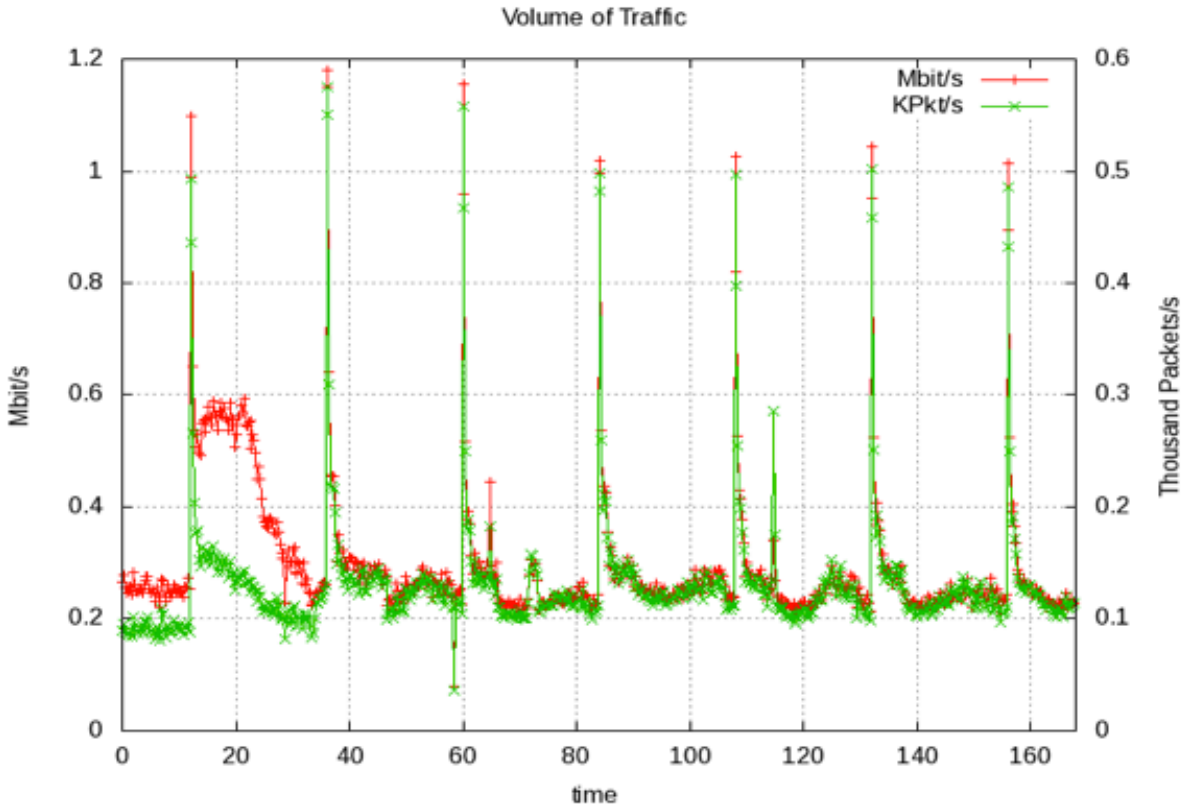

•Spikes are all UDP, port 53 DNS responses from either ns.ripe.net or a handful of comcast.net resolvers.

•All of the packets have destination set to the same value: 2607:fad0::1 which is an IP address based out of Liquidweb IP address space. AS 32244.

Routing Related Issues and IPv6 Pollution

- Near Misses
  - Darknet traffic destinations “near” routed prefixes
  - Used edit-distance analysis
  - 40-80% of all packets within 1 hex character from a routed prefix
  - Explains partially why we see negligible RIPE region traffic
- Route Instability
  - A key factor in our study is the covering prefix announcement
  - Routing instability can result in additional pollution traffic
- Partial visibility
  - Pollution traffic can also be caused by prefixes that are partially visible
  - We also noted that:
    - Partially visible prefixes are also 10 times more unstable than an average prefix
    - These partially visible prefixes are generally at the edges of the Internet
    - They are much more common in IPv6 than IPv4

Conclusion

- First large-scale study of IPv6 Internet Pollution
  - Some amount of route filtering
  - Minimal or no port filtering
  - Significantly lower volume of background traffic in v6
  - Significant change in protocols and ports over v4
- Highlight key contributors to this traffic
- Case studies highlight the highly unpredictable nature of Internet pollution traffic – you never know what you are going to get

Conclusion

- Future: long-term collection
  - Observe and explain trends
  - Understand how the IPv6 ecosystem operates
  - Aid operators
    - Sharing information with the operational community
    - Diagnosis of network problems
    - Early warning of misconfigurations
    - Notification of malicious clients
  - Re-introduce the RIPE prefixes into our study