IPv6 Source Addresses

What Could Possibly Go Wrong?

Jen Linkova, furry@google.com

Methodology and Data Set

- Logging all IPv6 packets from reserved/invalid sources entering Google network from Internet
- Collecting the data for a few days

Data Size:

- 2011:
  - 1.1M packets
  - 32.5K Unique IPs
- 2013:
  - 15M packets
  - 476K Unique IPs

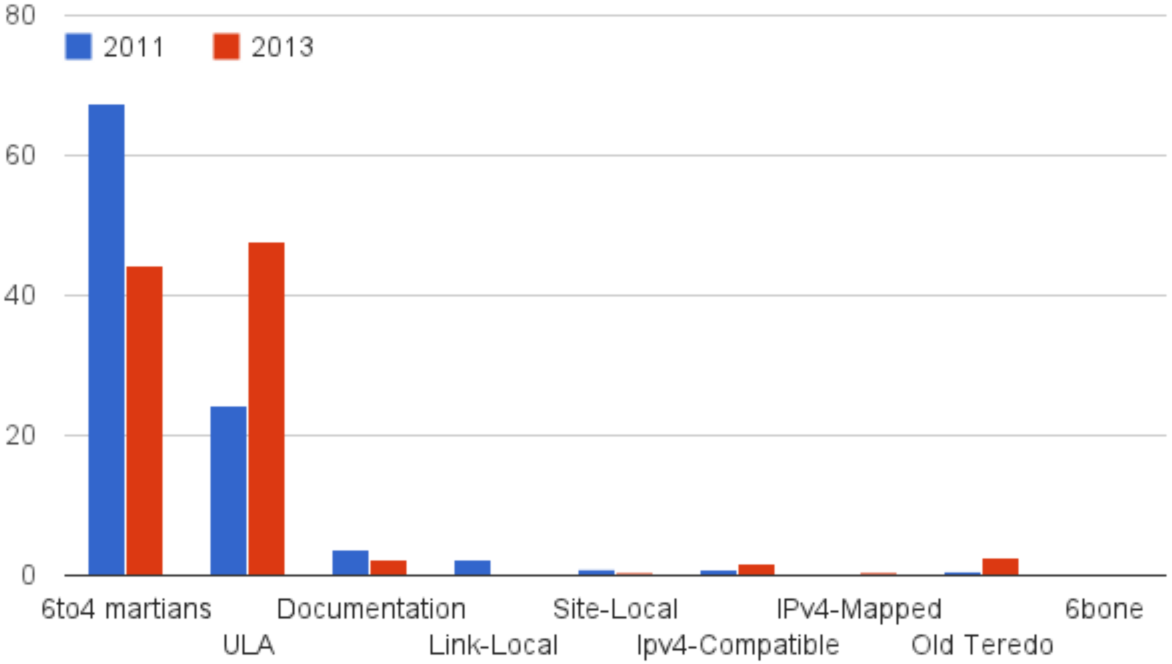

Source Addresses Distribution by Packets Count, %

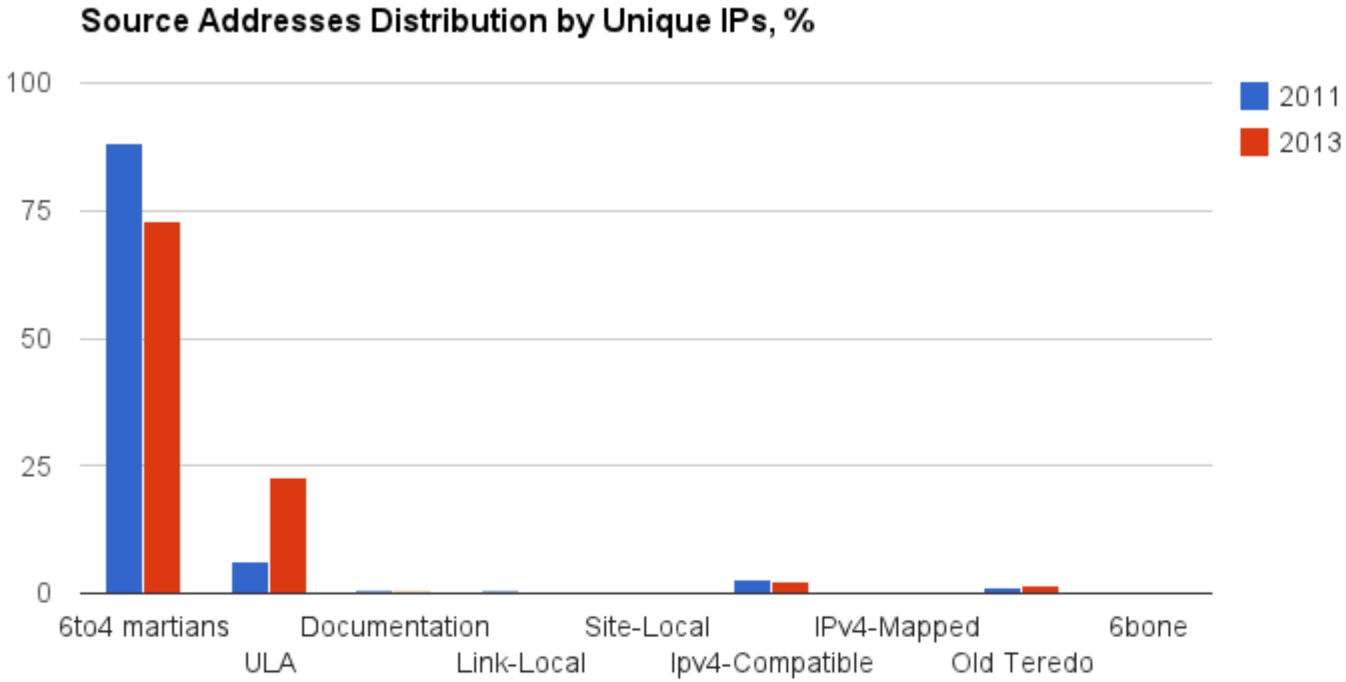

Source Addresses Distribution by Unique IPs, %

Zooming In... (6to4 and ULA Excluded)

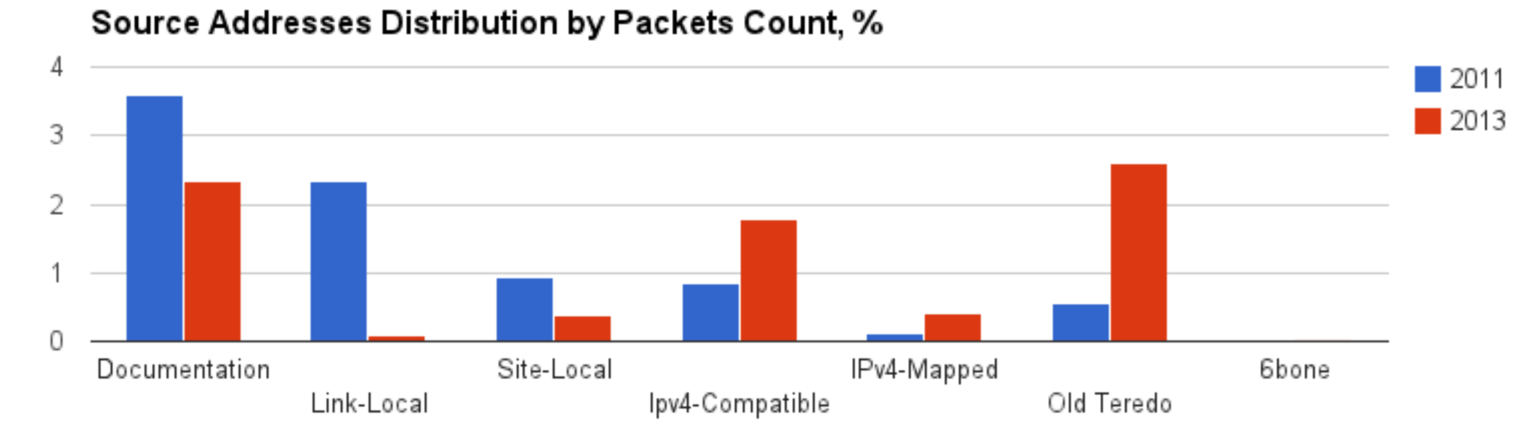

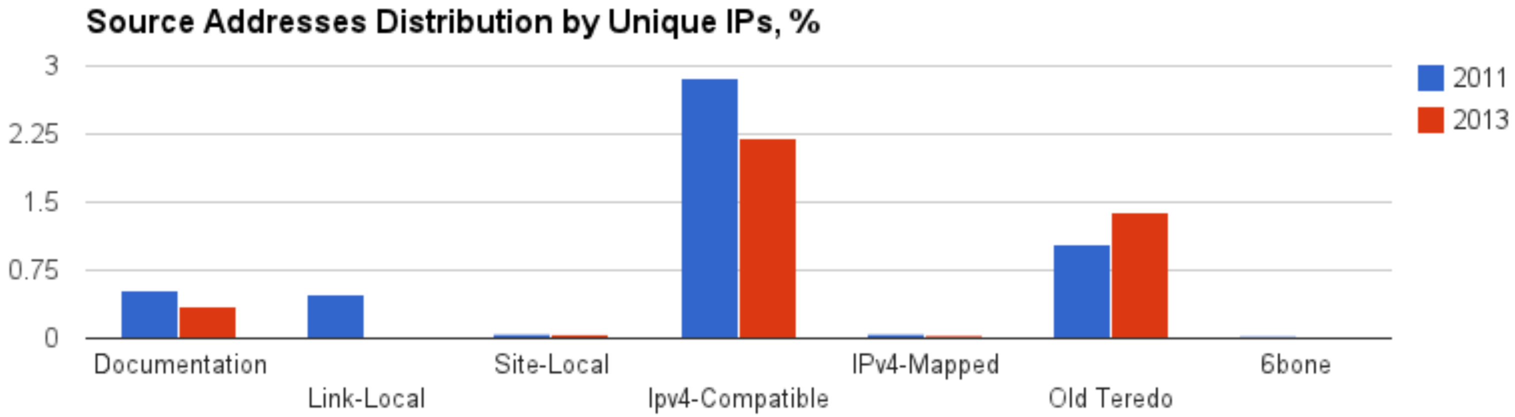

Some Good News

- No multicast Sources
- Very few people are using unallocated/bogon blocks
  - when they do, they choose them randomly
  - although some people like addresses like ‘a:a:a:a:a:a:a:a’

	Packets	Addresses	/64 Prefixes
2011	470 (0.4%)	39 (0.1%)	34
2013	9035 (0.6%)	168 (0.04%)	105

Traffic Profile

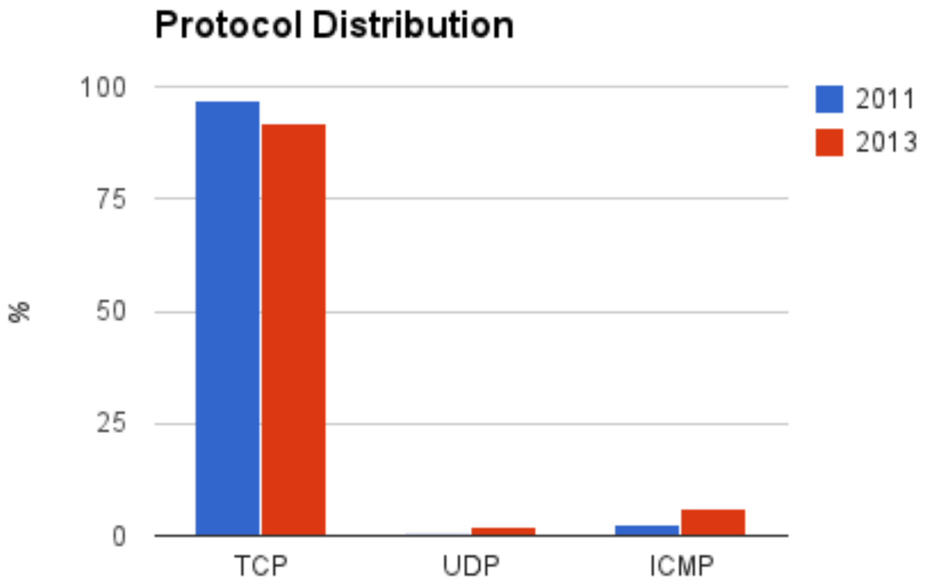

- TCP & UDP dropped from 97% in 2011 to 92% in 2013
- More ICMPv6 (from 2.5% to 6.3%)

ICMP Traffic from Invalid Sources

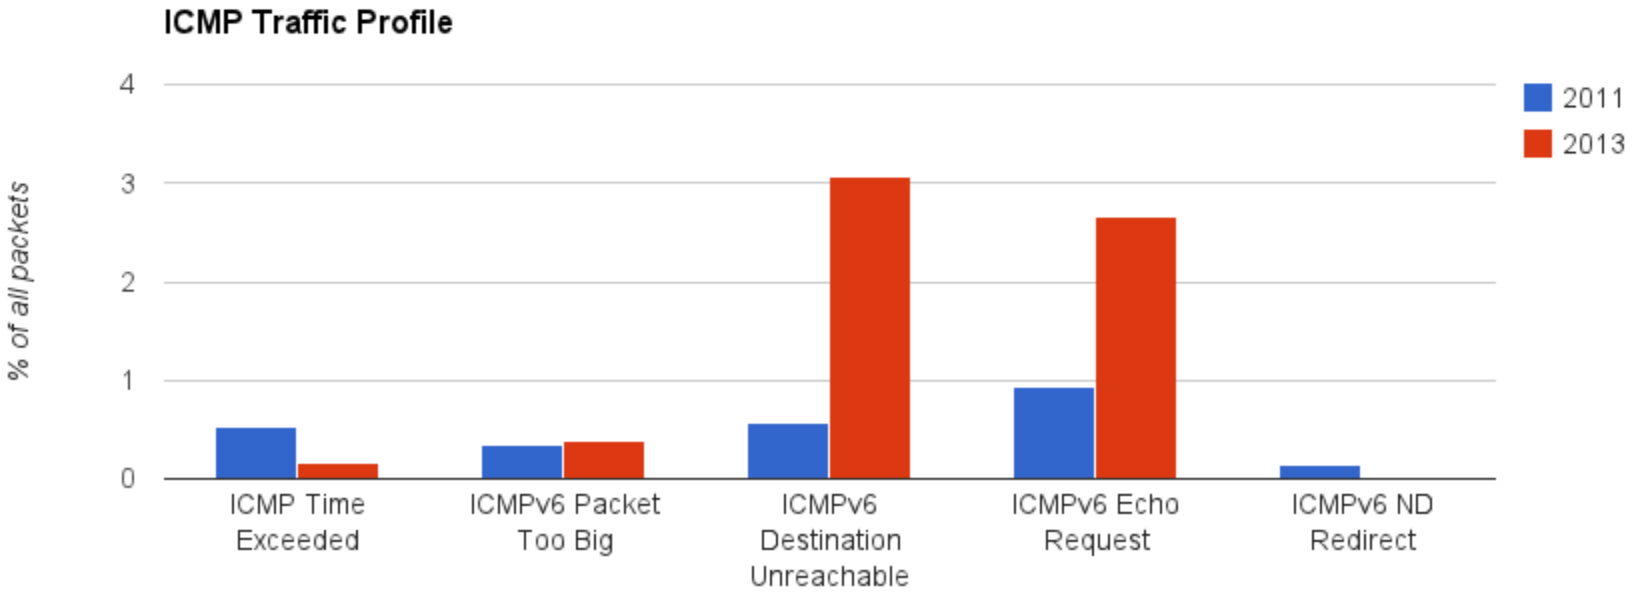

ICMP Traffic from Invalid Sources (contd.)

- Time Exceeded: dropped from 0.52% to 0.17%
  - is routing better now?
- Packet too Big: slight increase (0.35% to 0.38%)
- Destination Unreachable: increased from 0.57% to 3.07%!!
  - > 99% - ‘Address Unreachable’
- Echo Request (0.94% -> 2.66%): users keep pinging us.. from invalid addresses ;)
- and finally…one interesting type of ICMP (see next slide)

Neighbor Discovery Redirects

- Coming from link-local address to Google frontends
- RFC 4861 - Neighbor Discovery for IP version 6 (IPv6) says:

Source Address: MUST be the link-local address *assigned to the interface from which this message is sent.*
Destination Address: The Source Address of the packet that triggered the redirect.

….

A router SHOULD send a redirect message [skip] whenever it forwards a packet that is not explicitly addressed to itself [skip] in which: ***the Source Address*** *field of the packet* ***identifies a neighbor***

- Two routers (from two vendors) somewhere in the Internet keep sending redirect packets...since 2011..

Link-Local Unicast
fe80::/10

Addresses Distribution

	Packets Count (% of all packets)	Unique Address		Vendors	
		Total	MAC48 based (*)	Known	Unknown OUI
2011	26198 (2%)	156	129 (82%)	24	2
2013	11676 (0.08%)	35	32 (91%)	18	1

* "Based on MAC-48" means that "U/L bit is set and "FF:FE octets present".

Other addresses look like privacy extensions or based on locally administered MAC-48.

Traffic Profile

- Majority of traffic is still TCP (~90%)
- In 2013 majority of non-TCP traffic is from those two devices sending ND Redirects.
- None “Packet too big” or “Time Exceeded” anymore, only “Destination Unreachable” (very few)

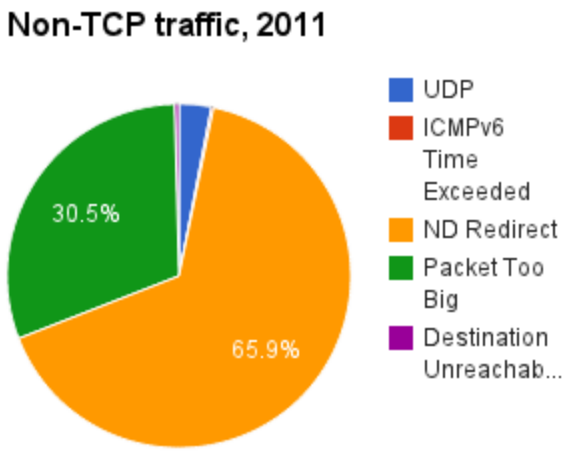

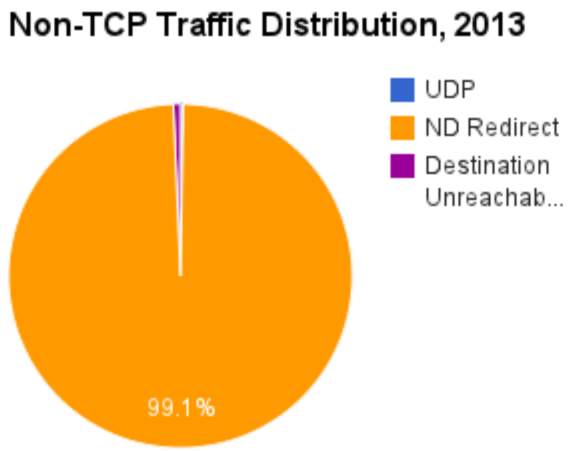

Additional Observations

- None of those packets are from devices directly connected to Google routers
- Packets with link-local source came from Internet - successfully routed
- What about RFC4007 “IPv6 Scoped Address Architecture”?

Section 9, “Forwarding”:

If transmitting the packet on the chosen next-hop interface would cause the packet to leave the zone of the source address, i.e., cross a zone boundary of the scope of the source address, then the packet is discarded.

Unique Local Unicast Addresses
fc00::/7

Addresses Distribution

	Packets (% of total packets analyzed)	Prefixes			Addresses		IPs/prefix (avg)
		Total count	Locally Assigned	Invalid ULAs a.k.a ‘globally assigned’	Total count (% of total packets)	IEEE MAC48 based	
2011	271056 (24%)	652	644 (99%)	8 (1%)	2063 (6.0 %)	88 (4.27%)	~3
2013	7125395 (48.0 %)	15545	15518 (99.8%)	27 (0.2%)	108920 (23%)	1452 (1.3%)	~7

Apparently there is some confusion between fc00::/7, fc::/7 and fc0::/7

Global ID Randomness

- What is the proper way to detect non-random GID?
- Approach chosen:
  - highest octet is ‘0’ or ‘1’
  - hex representation contains only [a-f] or only [0-9]
  - hex representation contains 3 or less different symbols (excl. ‘:’)
  - two octets are ‘0’

	Non-Random prefixes	Packets from non-random addresses			Top 5 prefixes
		Total number	% of all ULA traffic	% of total packets	
2011	18 (2.8%)	65800	24%	5.9%	fc00::/48 fd00:5000::/48 fd00::/48 fc01:a:1::/48 fc00:10:18:/48
2013	112 (0.7%)	801495	11.2%	5.4%	fc00::/48 fd00::/48 fccc:15::/48 fdfd:cafe:cafe::/48 fc00:1000:1010::/48

ULA: Traffic Profile Dynamics

- Less TCP connections:
    - 98% in 2011
    - 94% in 2013
- More ICMP Destination Unreachable
    - < 0.01 % in 2011
    - 2% in 2013

Site Local Addresses fec0::/10 (Deprecated Since 2004)

Addresses/Traffic Distribution

	Addresses (% of all unique IPs)	Prefixes	Packets (% of total packets)	Traffic Profile			
				TCP	ICMP Destination Unreachable	ICMP Time Exceeded	UDP
2011	16 (0.05%)	8	10497 (1%)	64%	1%	35%	< 0.1%
2013	205 (0.04%)	21	55963 (0.4%)	40%	40%	20%	< 0.1%

Traffic profile is different from ULA sources!

Anomalies

6Bone: 3ffe::/16 and 5f00::/8

Almost all traffic is from 3ffe:831f::/32 (old M$ Teredo net)
Shouldn’t be used by Windows since long time ago

	Packets	Addresses	Traffic Profile
			ICMP Echo Request
2011	6135 (1%)	334 (1%)	100%
2013	389920 (3%)	6622 (1%)	100%

6bone traffic:

	Packets	Addresses	Traffic Profile
			TCP
2011	142 (0.01%)	7	100%
2013	3192 (0.02%)	8 (7 from 3ffe:: and 1 from 5f00::/8)	100%

IPv4-Compatible and IPv4-Mapped

- ::FFFF:0:0/96 - IPv4-Mapped
- ::/96 - IPv4-Compatible (deprecated for long time...)
  - most IPv4 addresses encoded in compatible are private

		Packets	Unique IPs	Traffic Profile			
				TCP	ICMP Desti. Unreach	ICMP Echo	ICMP Time Exceeded
2011	v4-mapped	1217 (< 0.1%)	16 (<0.1%)	86%	1%	none	13%
	v4-compatible	9475 (1%)	929 (3%)	41%	58%	< 0.1%	none
2013	v4-mapped	60213 (<0.1%)	145 (<0.1%)	92%	1%	1%	1%
	v4-compatible	**266682 (2%)**	**10526 (2%)**	3%	**97%**	< 0.1%	none

Other Addresses from ::/64

- Very few packets from
  - ::/1
  - :: (unspecified)
- There are other source addresses from ::/64 with interface ID not based on MAC48
  - What are they??

	Packets (% from total packets count)	Addresses	Traffic Profile
			TCP
2011	318 (0.03%)	25 (0.08%)	100%
2013	51047 **(0.34%)**	498 (0.1%)	100%

QUESTIONS?