IPv6 Source Addresses

What Could Possibly Go Wrong?

Jen Linkova, furry@google.com

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

Data Set:

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

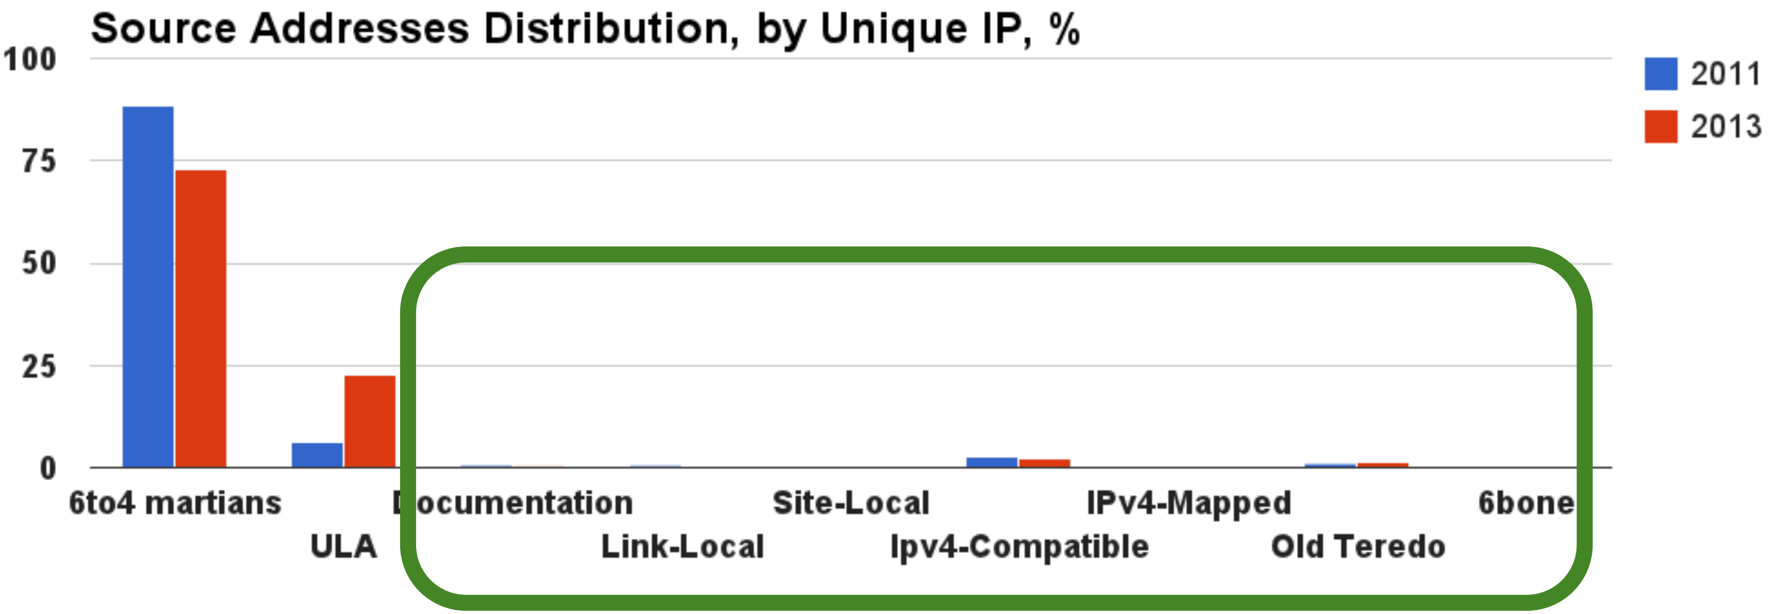
Source Addresses Distribution, by Unique IP, %

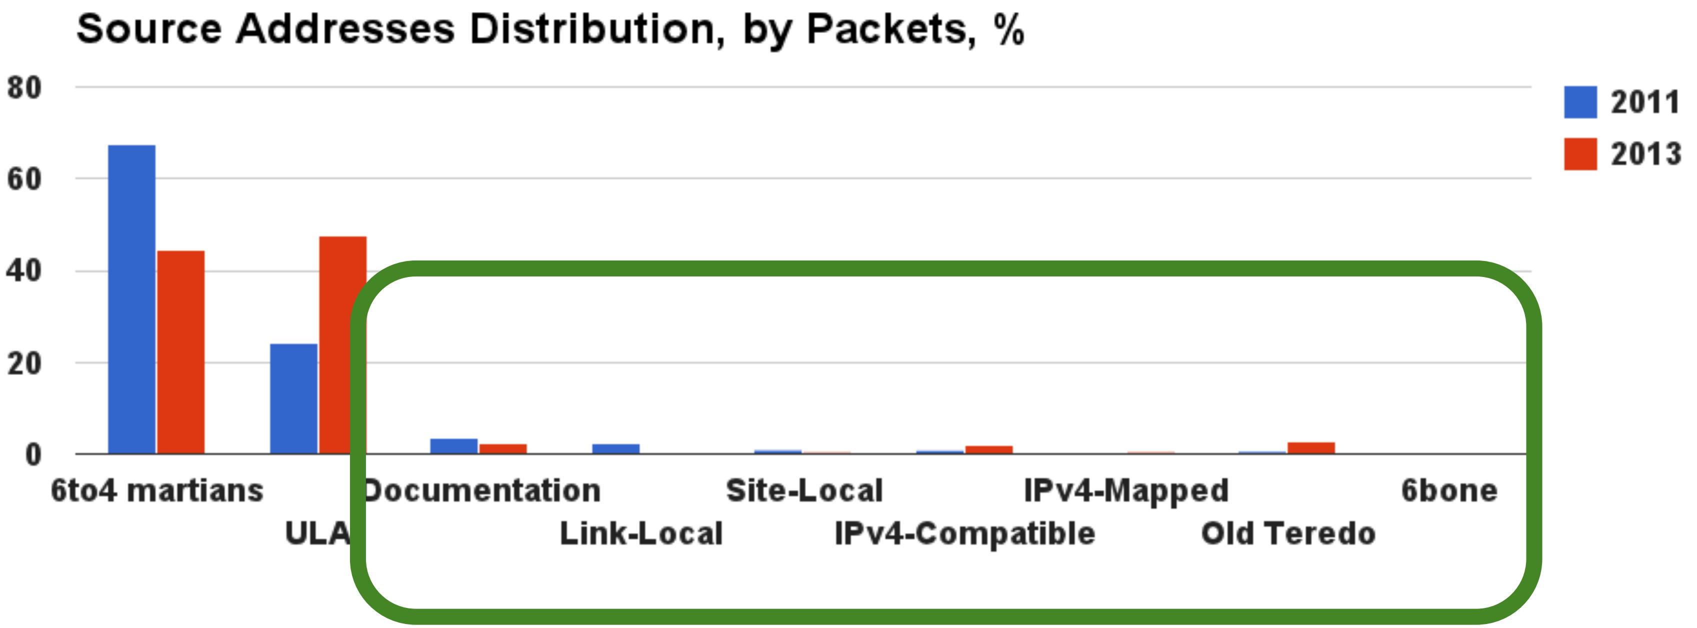
Source Addresses Distribution, by Packets, %

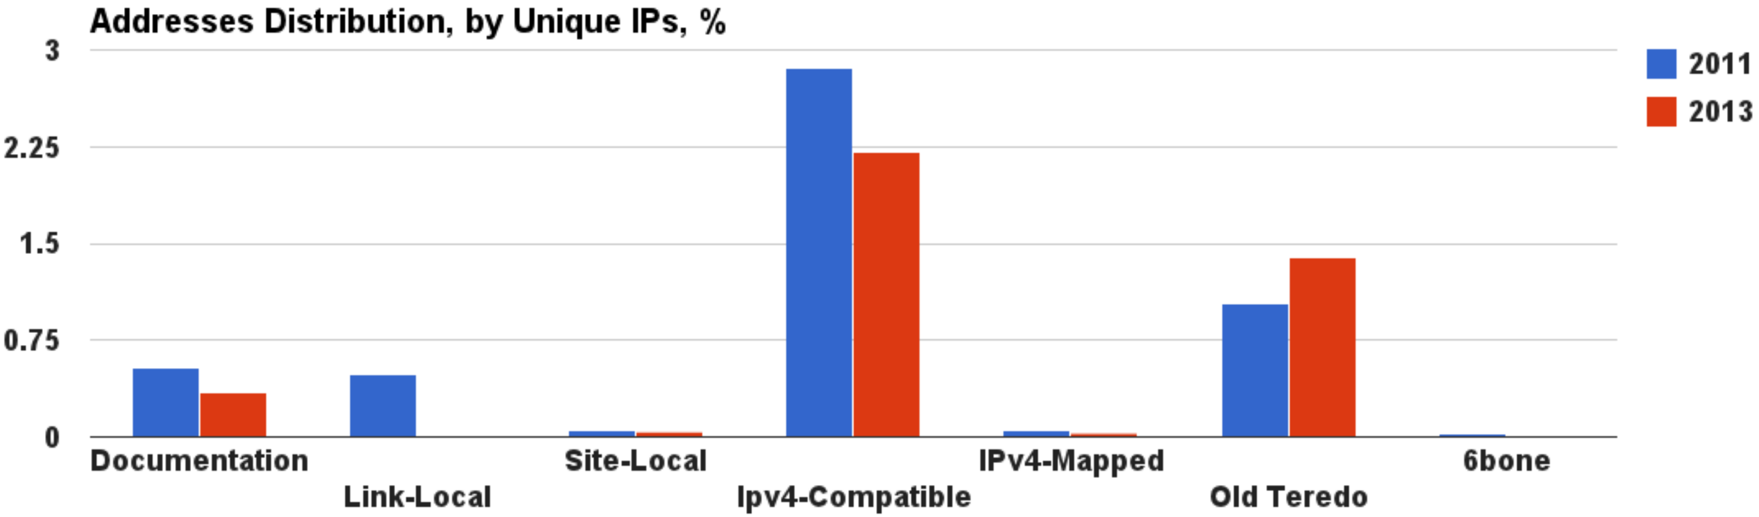
Addresses Distribution, by Unique IPs, %
3
2.25
1.5
0.75
0
2011
2013
Documentation
Link-Local
Site-Local
Ipv4-Compatible
IPv4-Mapped
Old Teredo
6bone

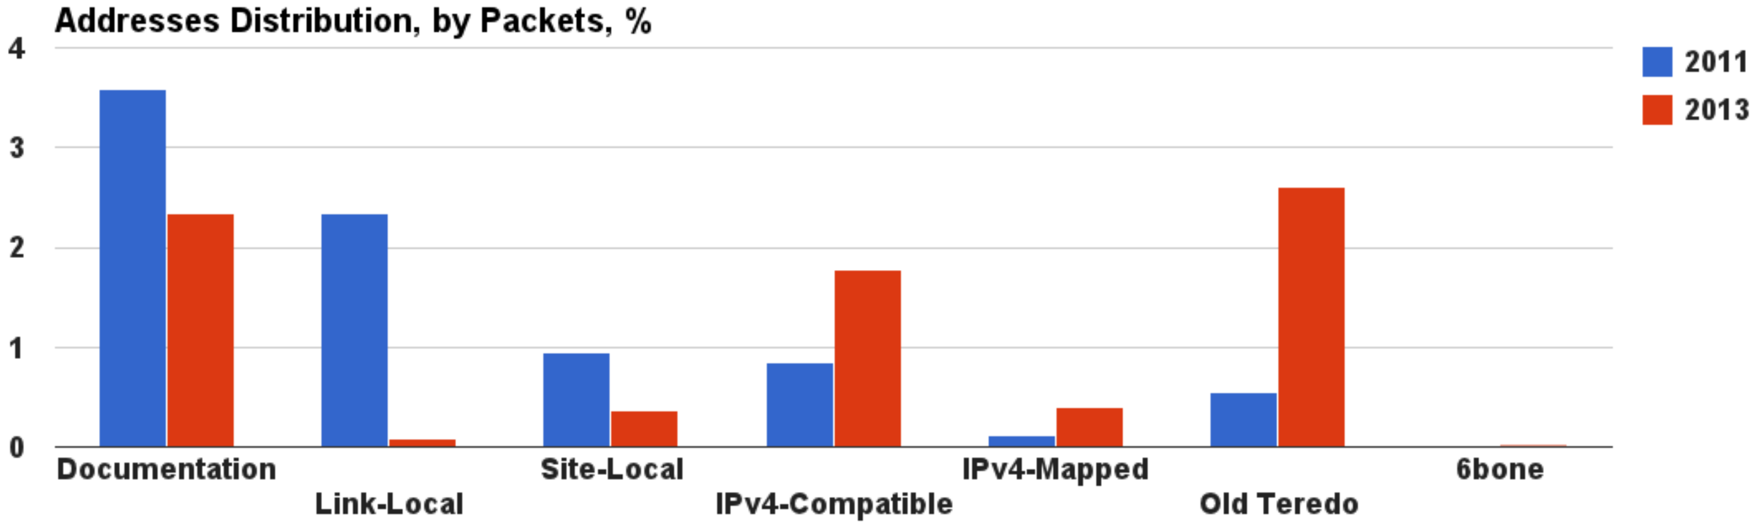
Addresses Distribution, by Packets, %
4
3
2
1
0
2011
2013
Documentation
Link-Local
Site-Local
IPv4-Compatible
IPv4-Mapped
Old Teredo
6bone

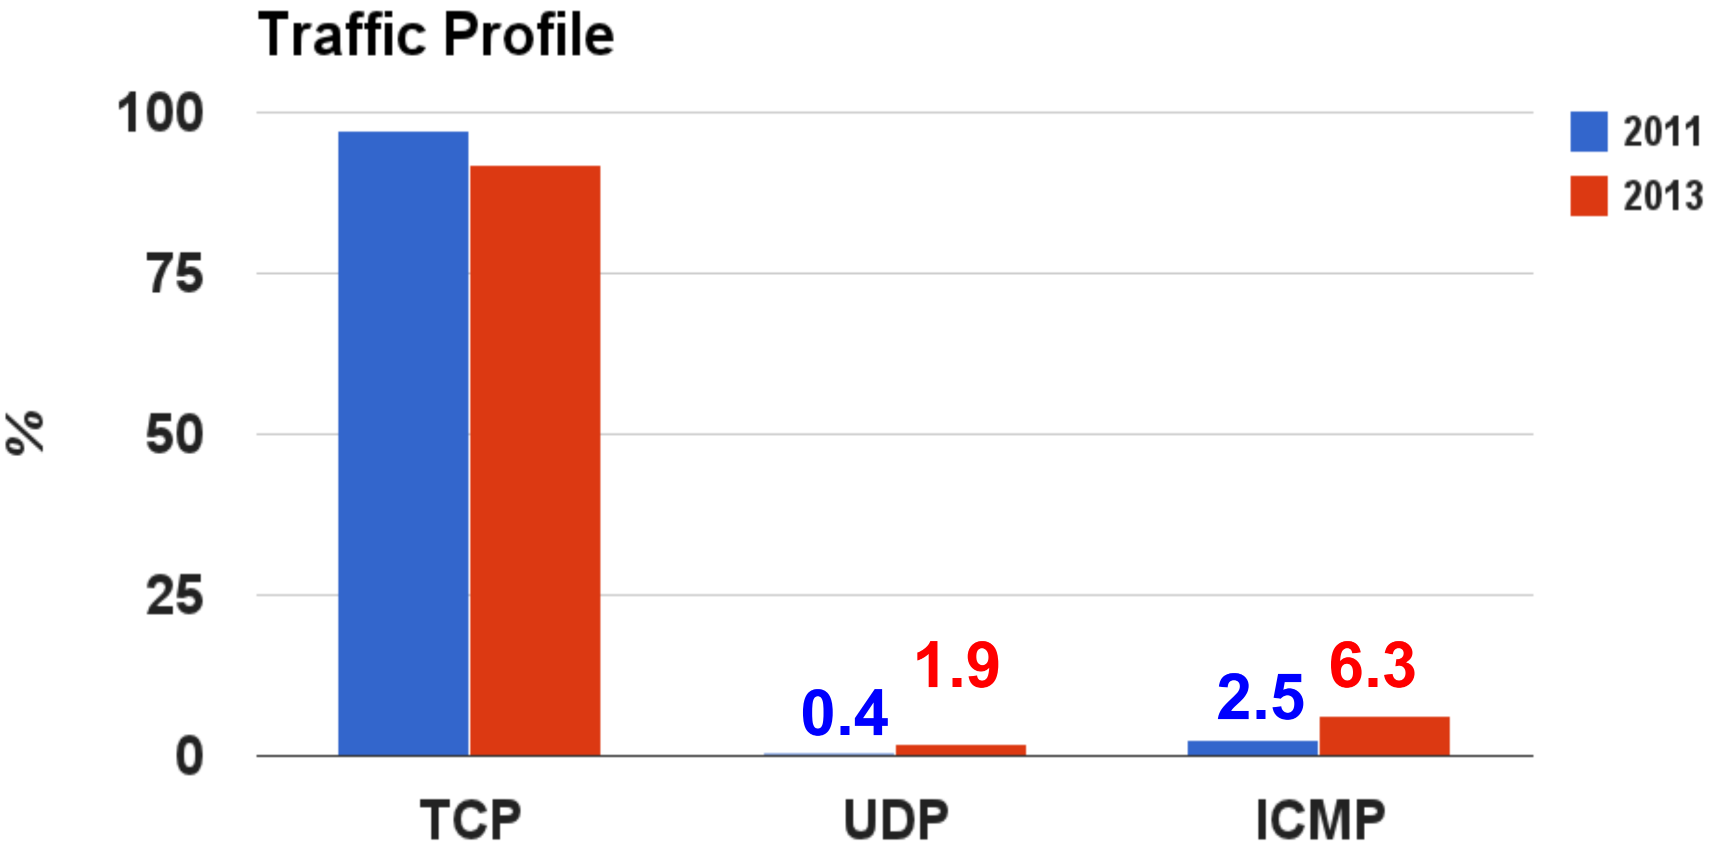
Traffic Profile
%
100
75
50
25
0
2011
2013
0.4
1.9
2.5
6.3
TCP
UDP
ICMP

ICMP Traffic Profile

- Users' Traffic
  - Echo Requests
- Infrastructure
  - Time Exceeded
  - Packet Too Big
  - Destination Unreachable
    - > 99% - 'Address Unreachable'

    *** Neighbor Discovery Redirects**

Link-Local Unicast

fe80::/10

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

* “Based on MAC-48”: “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 TCP (~90%)
- Non-TCP traffic:
  - 2011: mix of ICMP
    - destination unreachable
    - packet too big
    - time exceeded
    - **ND redirects**
  - 2013: traffic from TWO routers only
    - **ND redirects** to Google frontends IPs.

Neighbor Discovery Redirects

RFC 4861 - Neighbor Discovery for IP version 6 (IPv6)

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 - MUST identify a neighbor

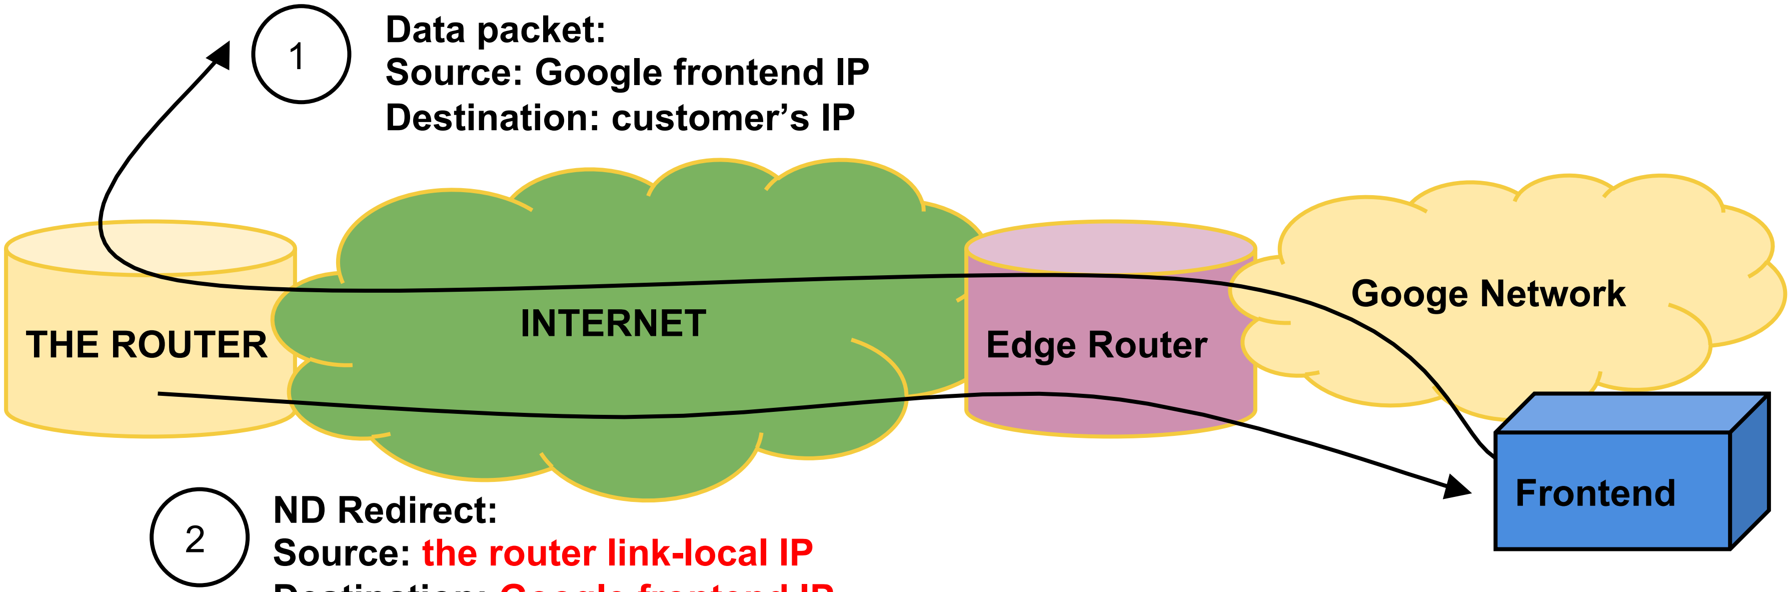

1 **Data packet:**
Source: Google frontend IP
Destination: customer's IP

2 **ND Redirect:**
Source: the router link-local IP
Destination: Google frontend IP

How Did They Get There?

- None of those packets are from devices directly connected to Google routers
- Packets with link-local source came from Internet - successfully routed
- 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
ULA
fc00::/7

	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

IPv6 is hard: There is some confusion between fc00::/7, fc::/7 and fc0::/7!

‘U’ Stands For ‘Unique’...Really?

- What is the proper way to detect non-random GID?
  - highest octet is ‘0’ or ‘1’ OR
  - hex representation contains [a-f] or [0-9] only OR
  - hex representation contains 3 or less different symbols (excl. ‘:’)
  - two octets are ‘0’
- Non-Random Prefixes Top List:
  - fc00::/48
  - fd00::/48
  - fdfd:cafe:cafe::/48
- Non-random ULA prefixes:
  - 2011: 2.8%
  - 2013: **0.7%**

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

	Addresses (% of all unique IPs)	Prefixes	Packets (% of total packets)	Traffic Profile			
				TCP	ICMP Dest. 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

Anomalies

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

- ~1% of all logged packets: **3ffe:831f::/32**
    - Was used by Teredo on Windows machines
    - 100% of traffic is ICMP Echo Requests

- 0.01% of all logged packets are from actual 6bone block
    - 7 IP addresses detected
    - 100% of traffic is TCP

IPv4-Mapped ::FFFF:0:0/96

- Used in the IPv6 basic API to denote IPv4 addresses
- Should NOT appear on the wire
- 2011/2013 - ~0.1% of analyzed traffic

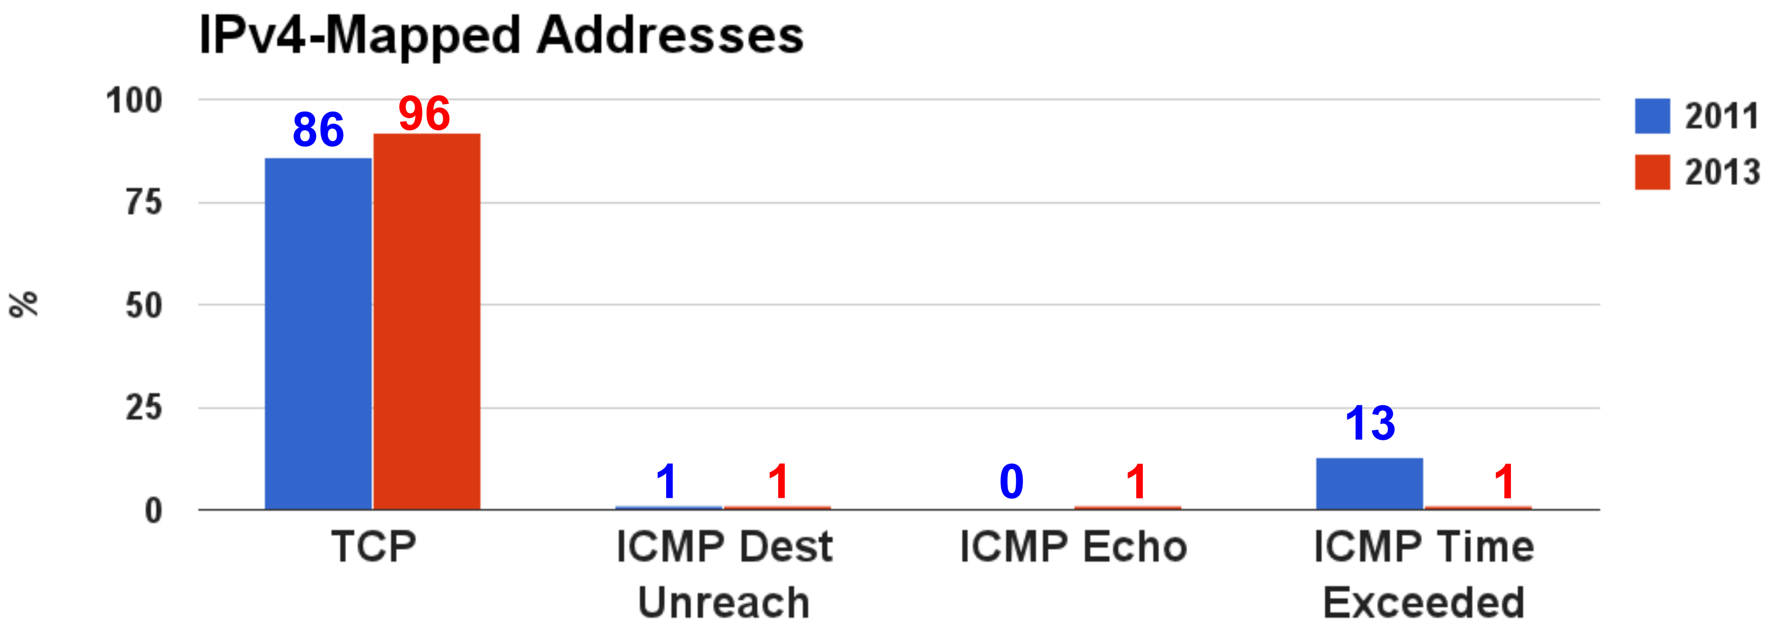

IPv4-Compatible ::/96

- Deprecated since 2006
- Should NOT appear on the wire
- 2011/2013 - ~2% of analyzed traffic
- Most of encoded IPs are private
- Mostly (97%): ICMP Destination Unreachable

::/64 Subnet

- Very few packets from
  - ::/1
  - :: (unspecified)
- Mystery Traffic:
  - Interface ID: 64 non-zero bits, NOT based on MAC48

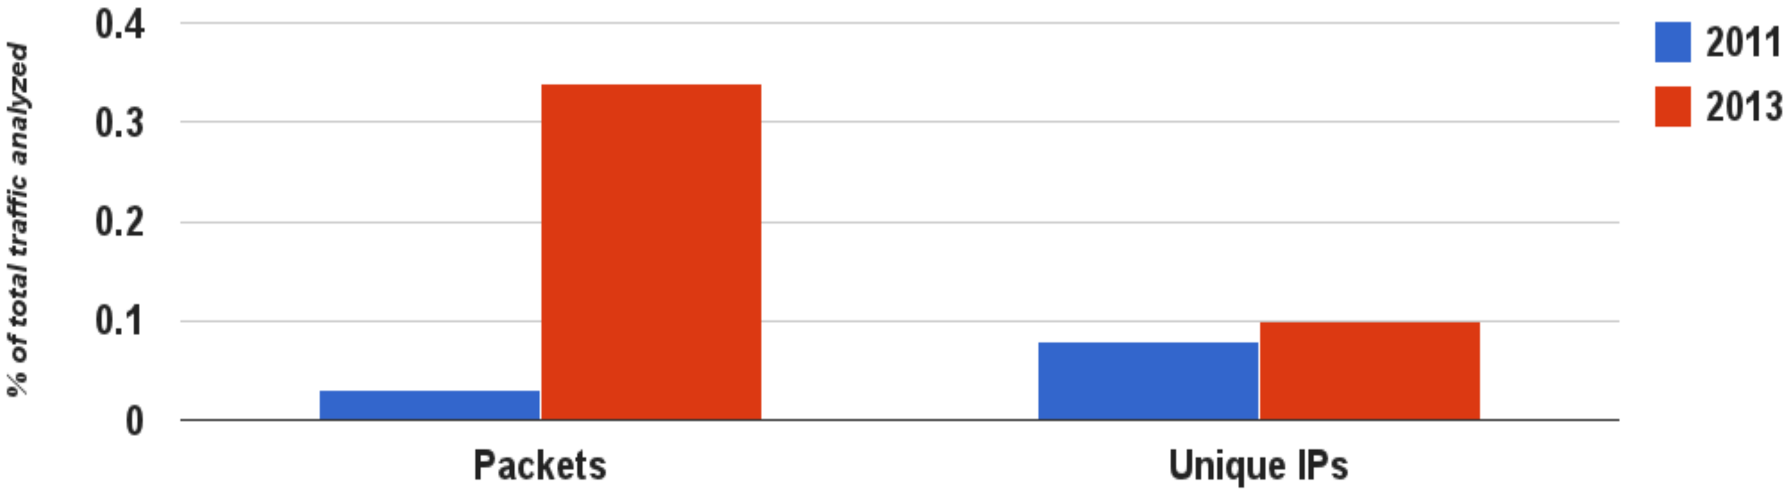

What We DID NOT See

- Multicast Sources
- Very little traffic from random blocks
  - addresses like ‘a:a:a:a:a:a:a:a’ are popular

Summary

- Address selection is still broken
- Things are getting better
- No explanation for some mystery packets
- Scoped Address Architecture is ignored ;(
- ..let alone BCP38….:-((

QUESTIONS?