

# IPv6 Performance

Geoff Huston  
APNIC Labs  
February 2015



# What are we looking at:

- How “reliable” are IPv6 connections?
- How “fast” are IPv6 connections?



# What are we looking at:

- How “reliable” are IPv6 connections?

Do all TCP connection attempts succeed?

- How “fast” are IPv6 connections?

is V6 slower than V4?

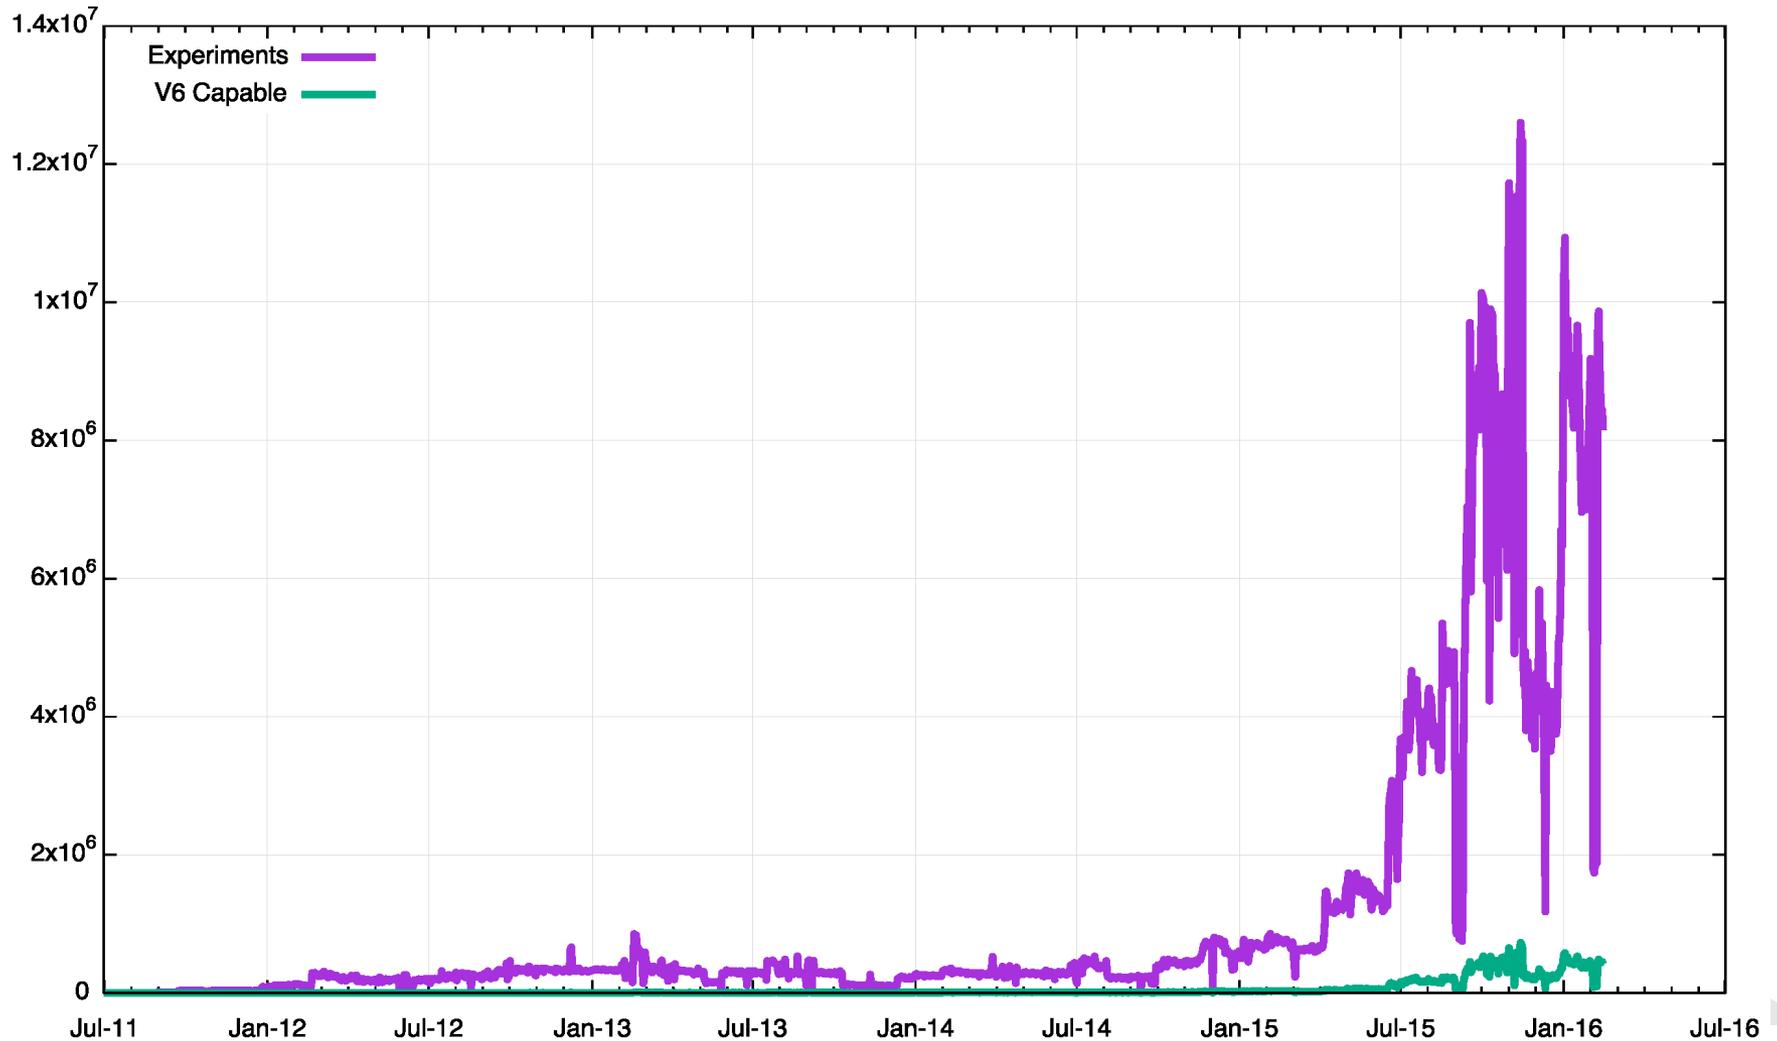


# The Measurement Technique

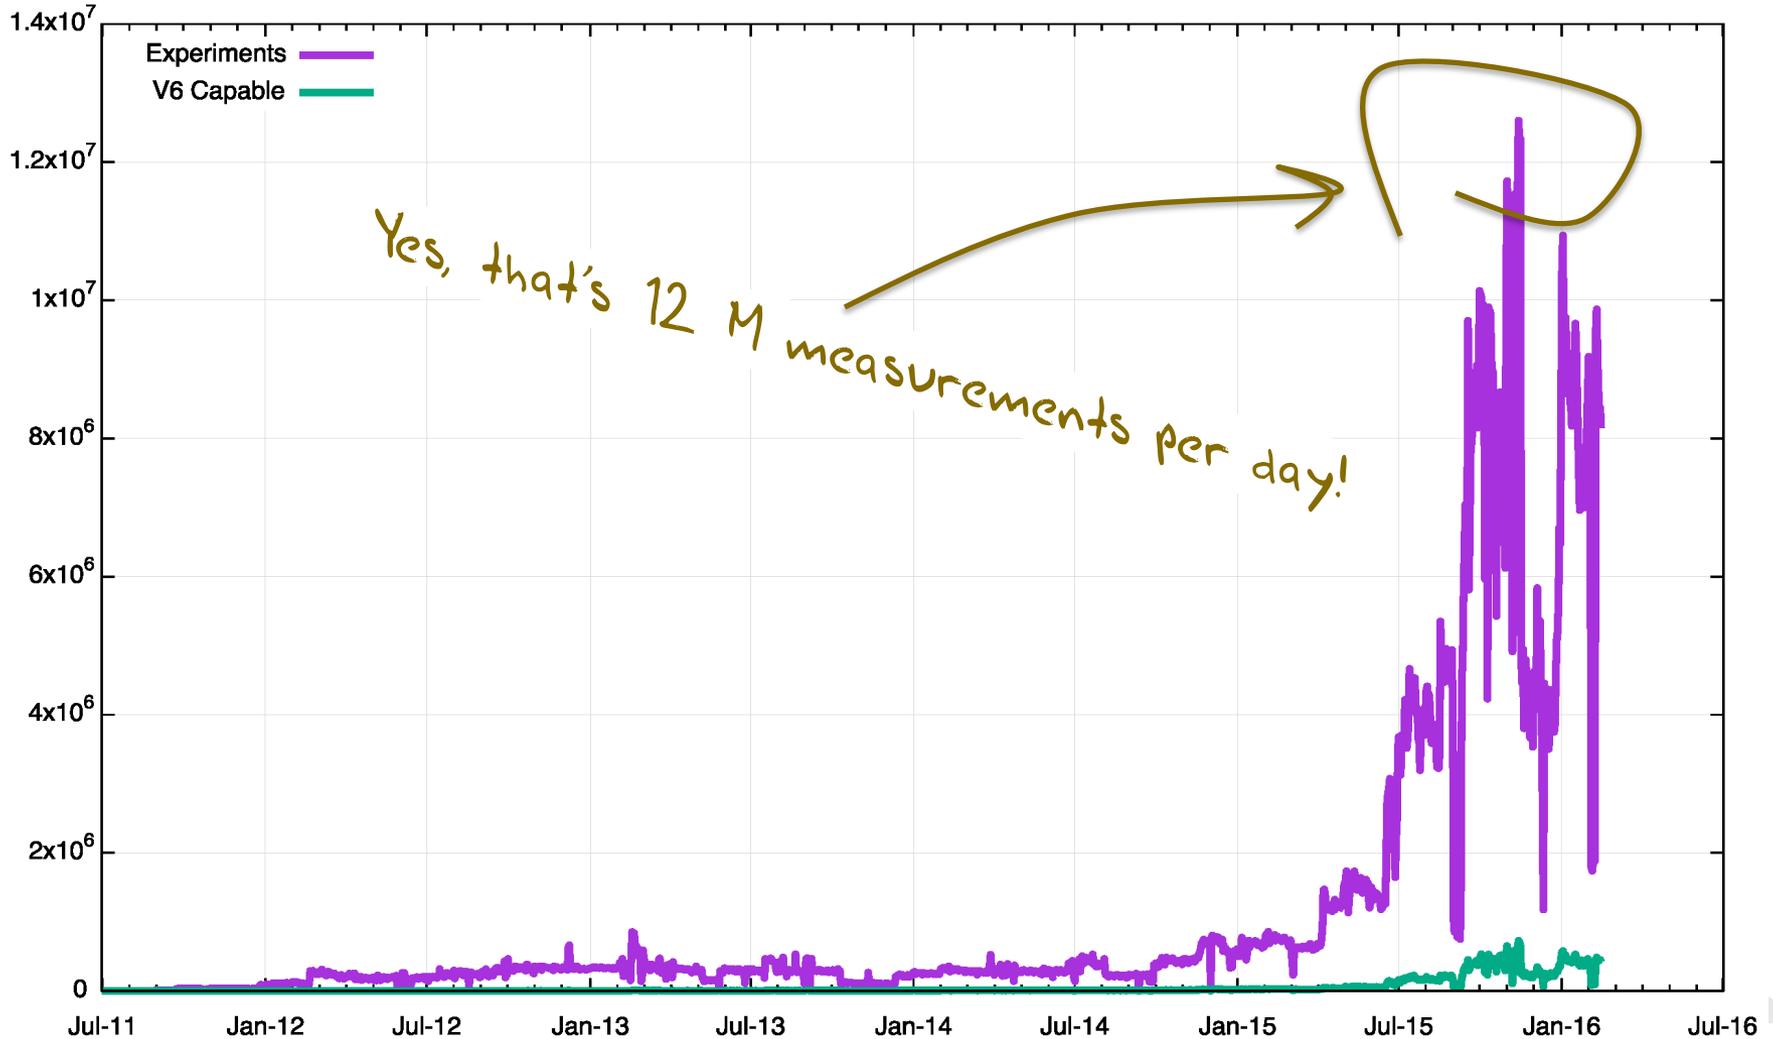
- Embed a script in an online ad
- Have the script generate a set of URLs to fetch
- Examine the packets seen at the server to determine reliability and RTT



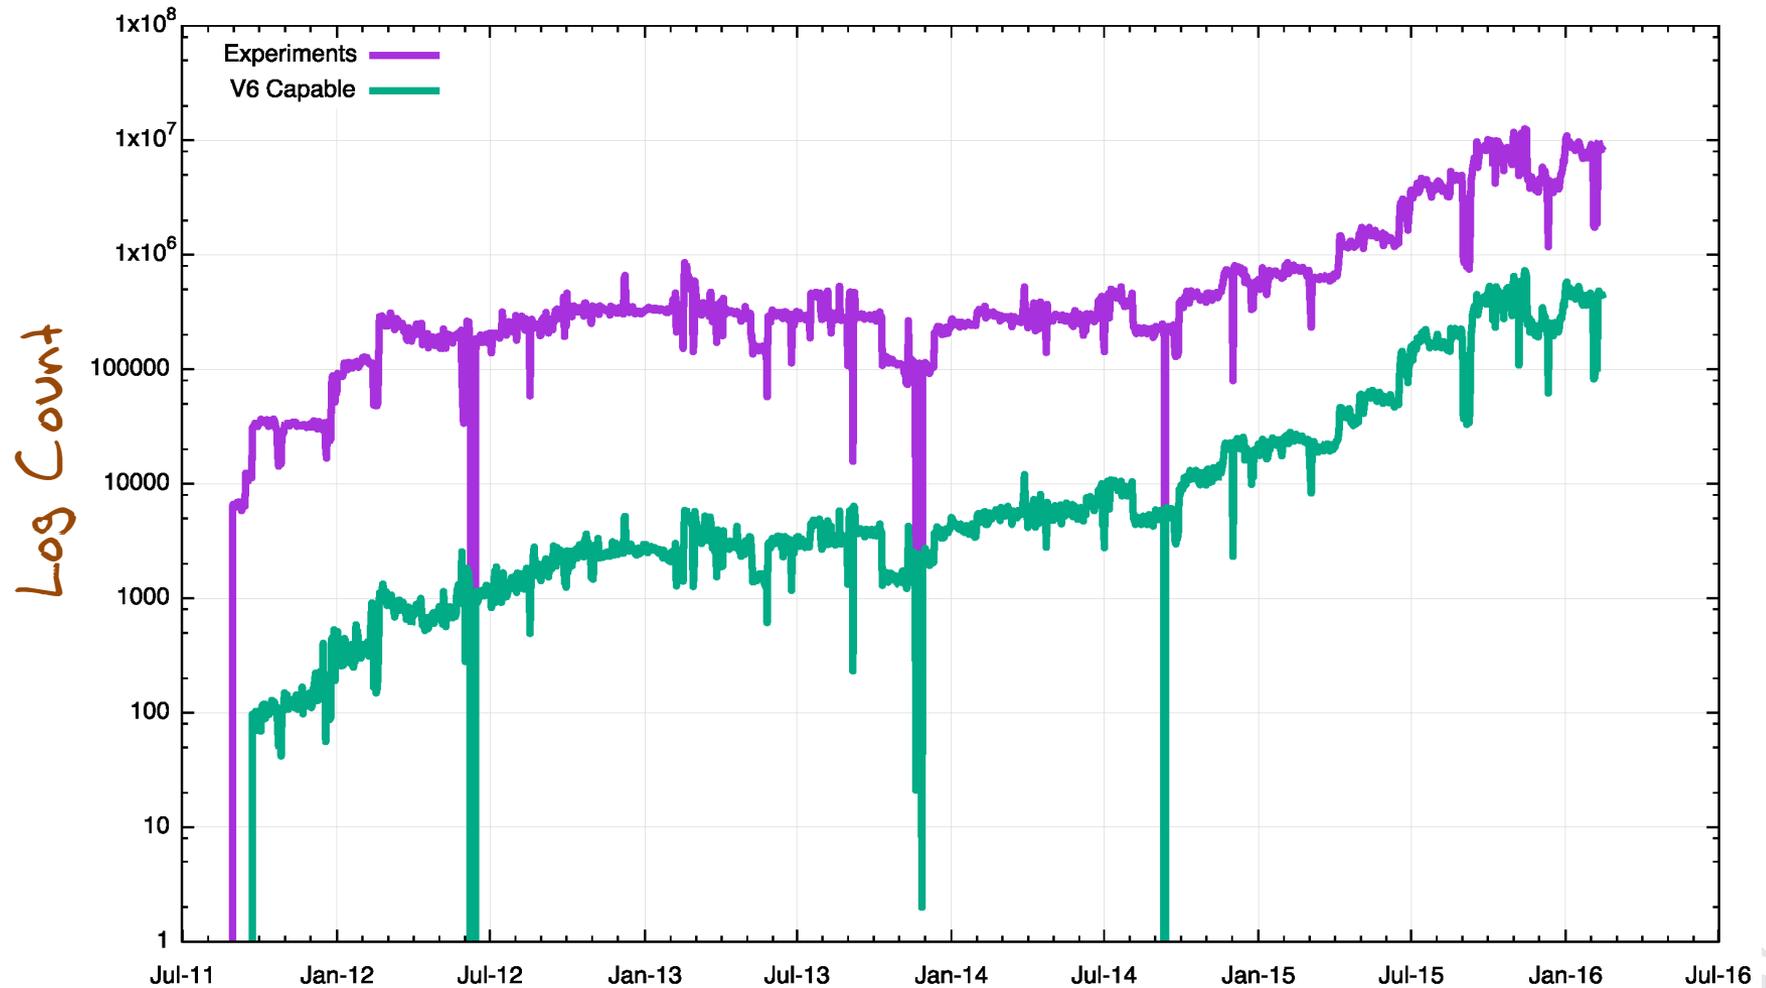
# Measurement Count



# Measurement Count



# Measurement Count



# What are we looking at:

- How “reliable” are IPv6 connections?

Do all TCP connection attempts succeed?

- How “fast” are IPv6 connections?

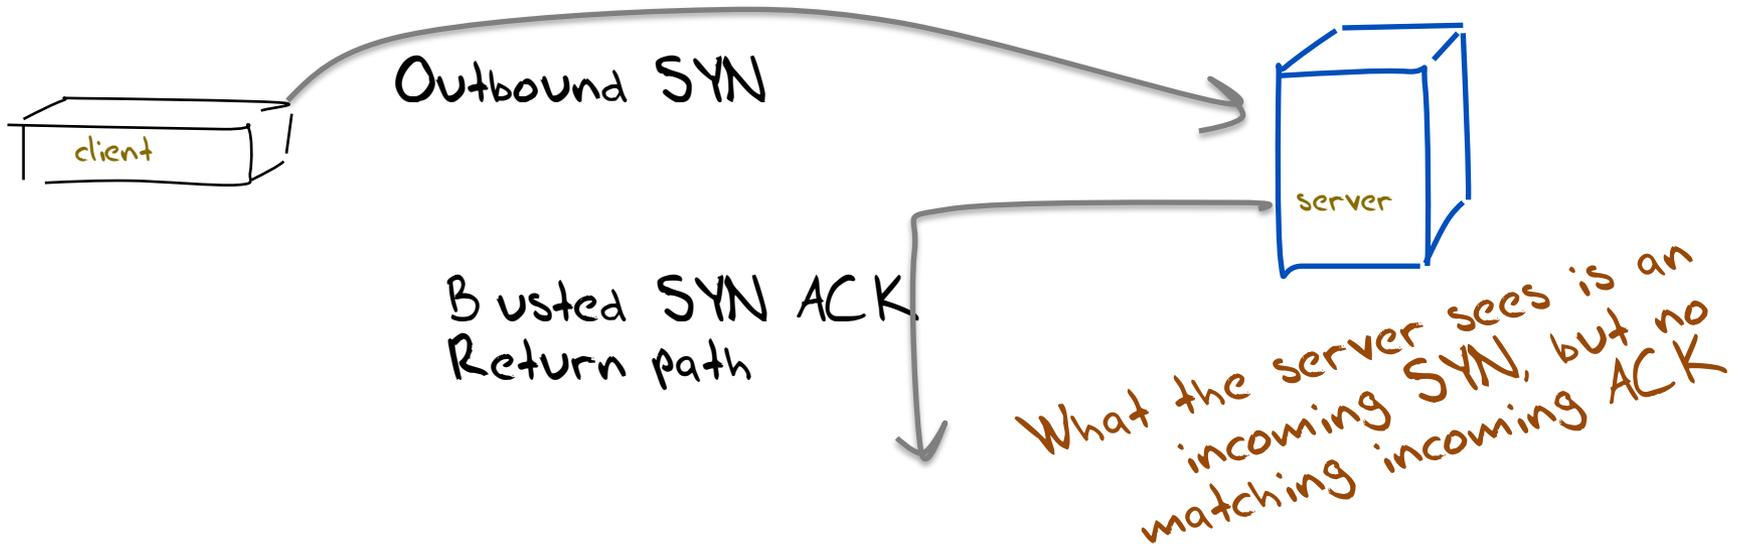
is V6 slower than V4?



Let's dive into Naked ~~SINs~~!



# Connection Failure



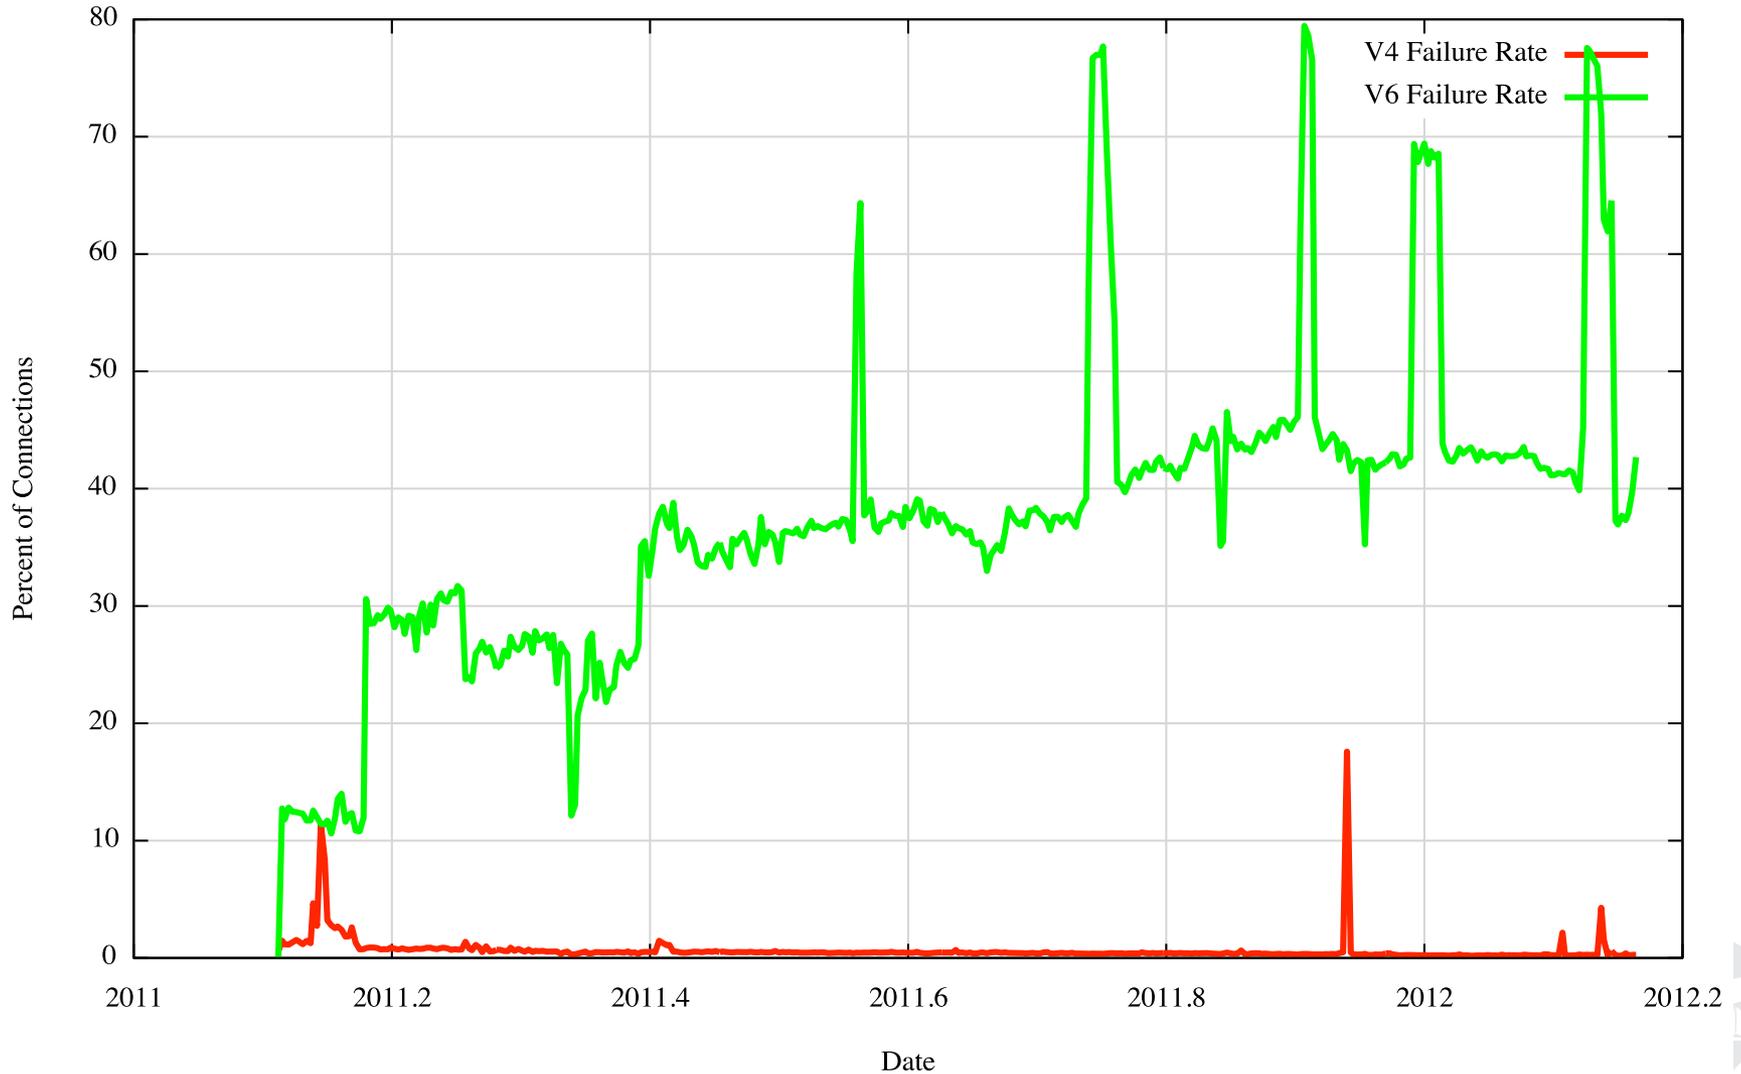
# Compare two data sets

- The first data set has been collected across 2011
  - Teredo and 6to4 were still active as IPv6 mechanisms
  - Little in the way of other IPv6 services
- The second data set has been collected across 2015/2016
  - Missing comparative IPv4 data for the period September – October ☹️



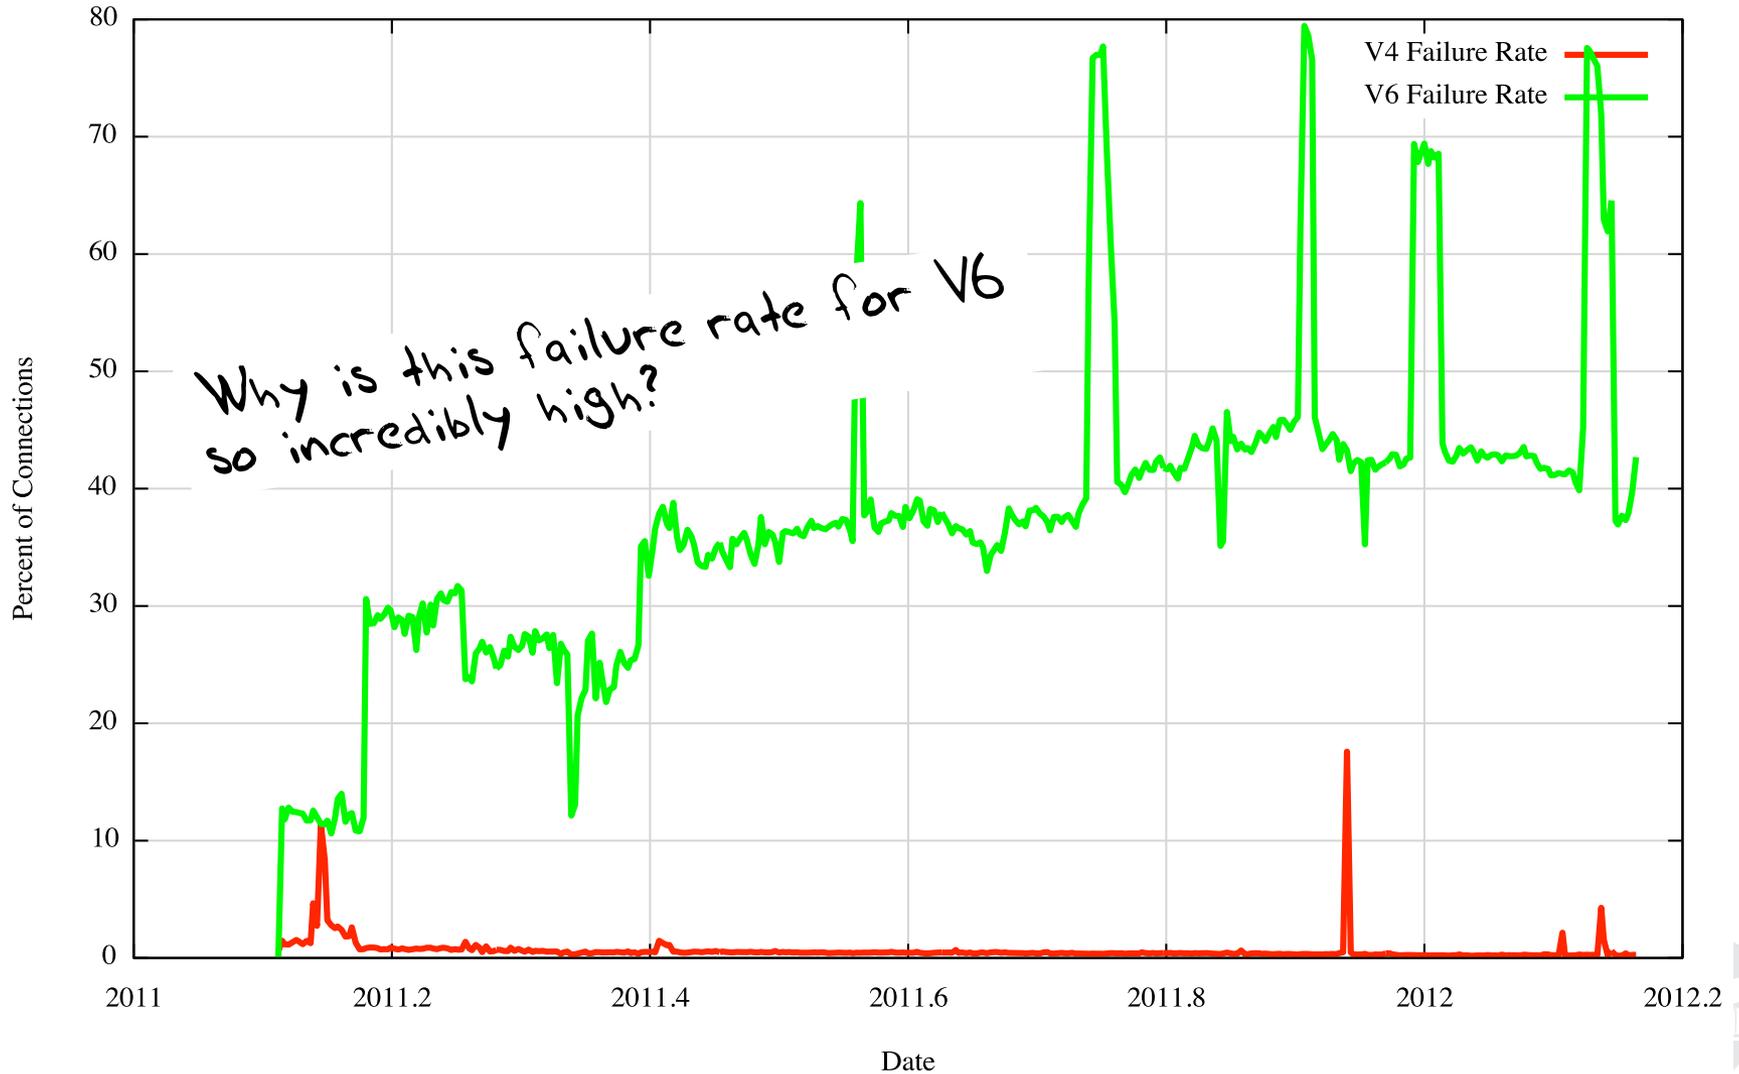
# 2011 - Measuring Failure

Connection Failure Rate



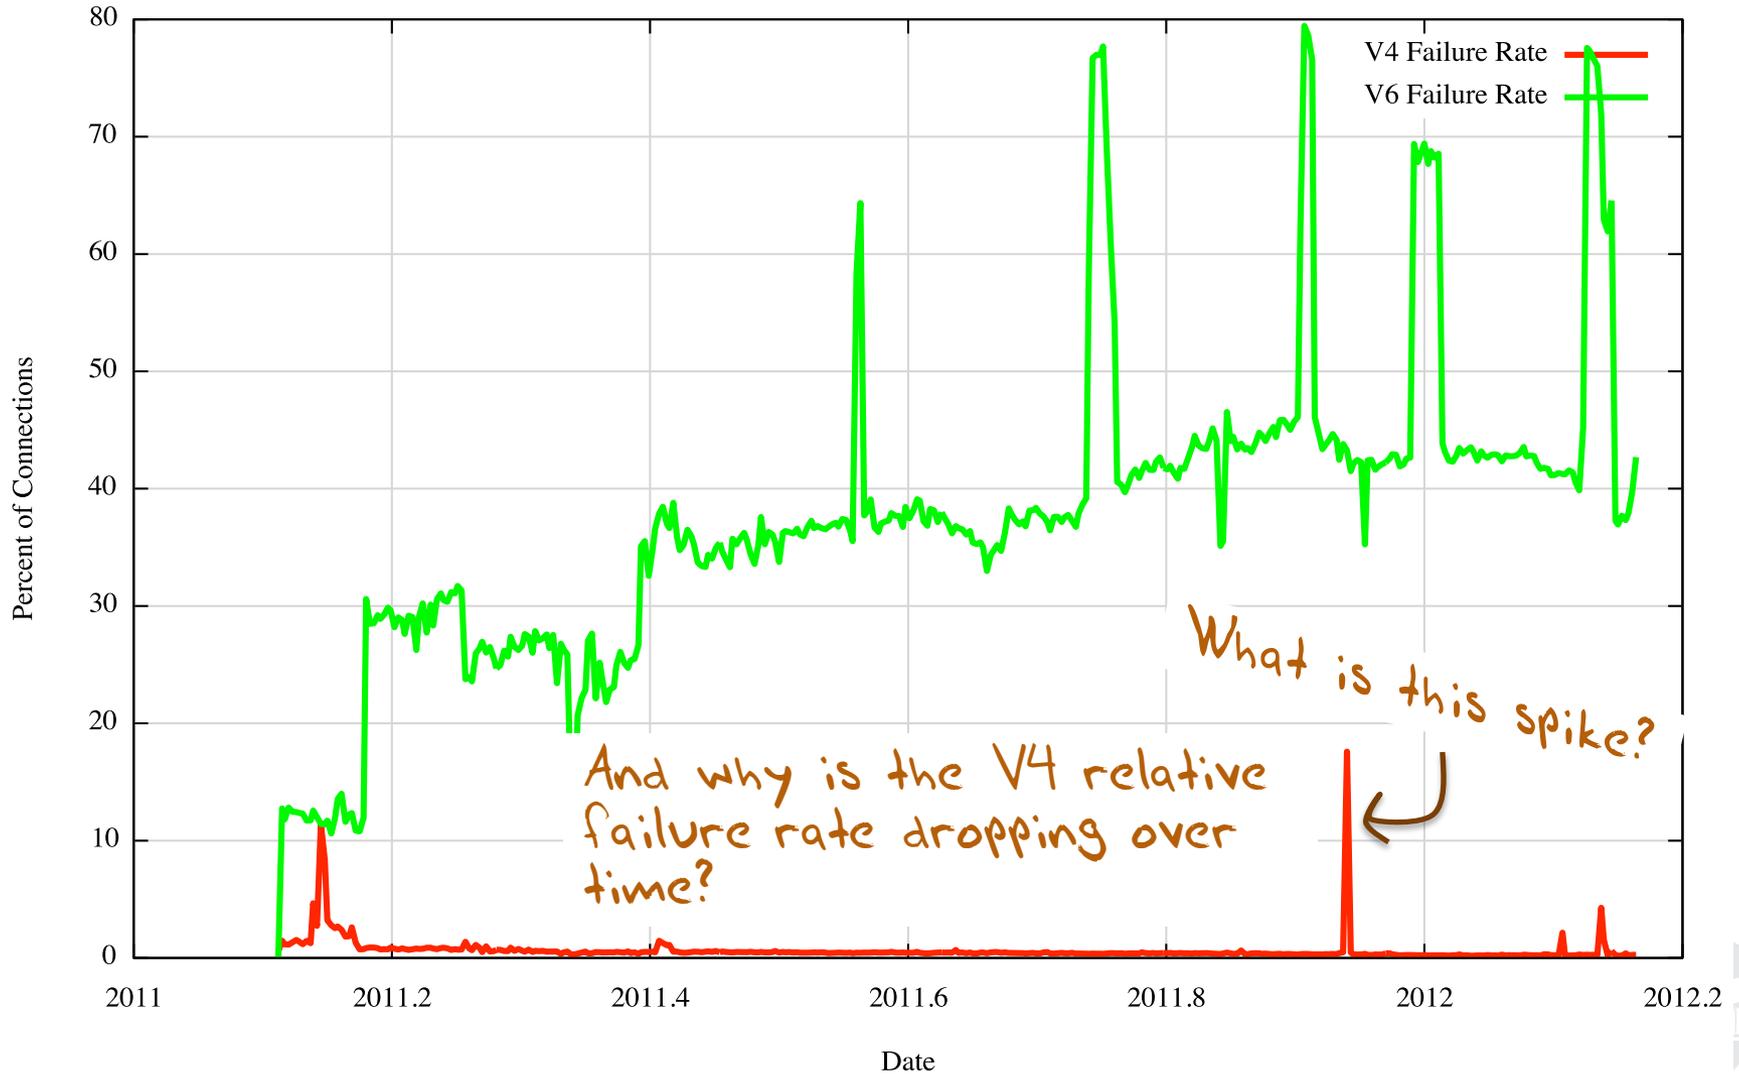
# 2011 - Relative Connection Failure Rates

Connection Failure Rate



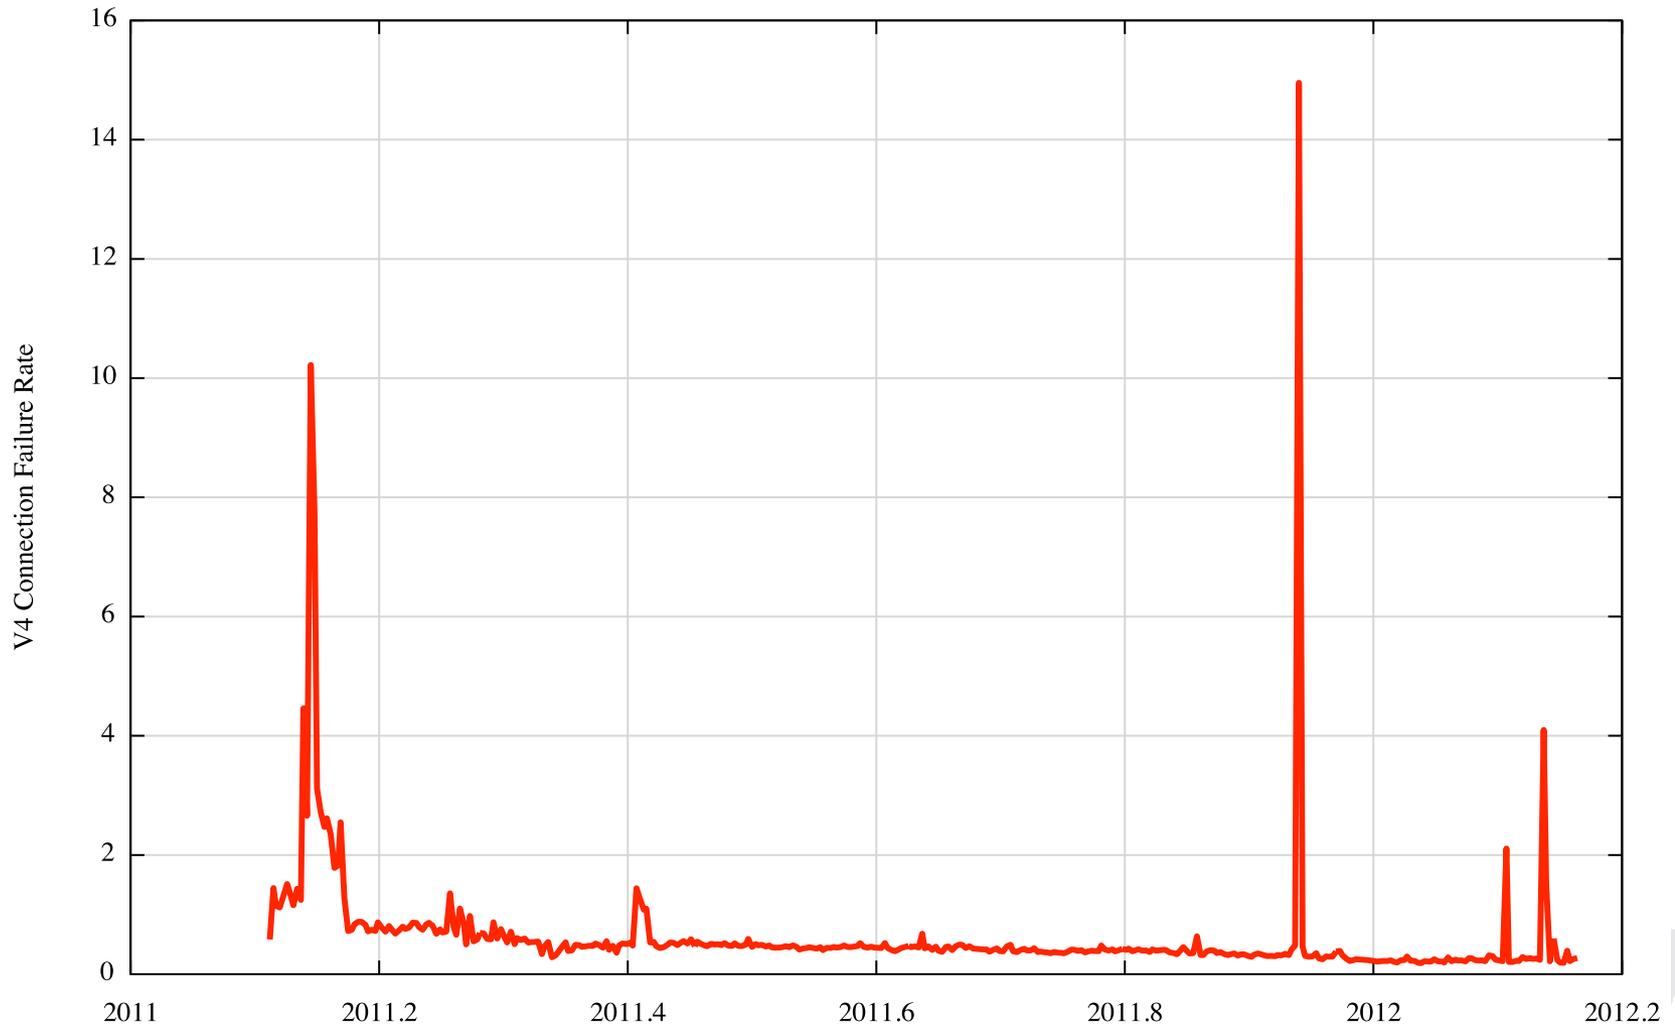
# 2011 - Relative Connection Failure Rates

Connection Failure Rate



# What is going on with IPv4?

Connection Failures - IPv4



# What is going on with IPv4?

The failure rate for V4 decreases as the volume of experiments increases – which implies that the number of “naked SYNs” being sent to the servers is not related to the number of tests being performed.

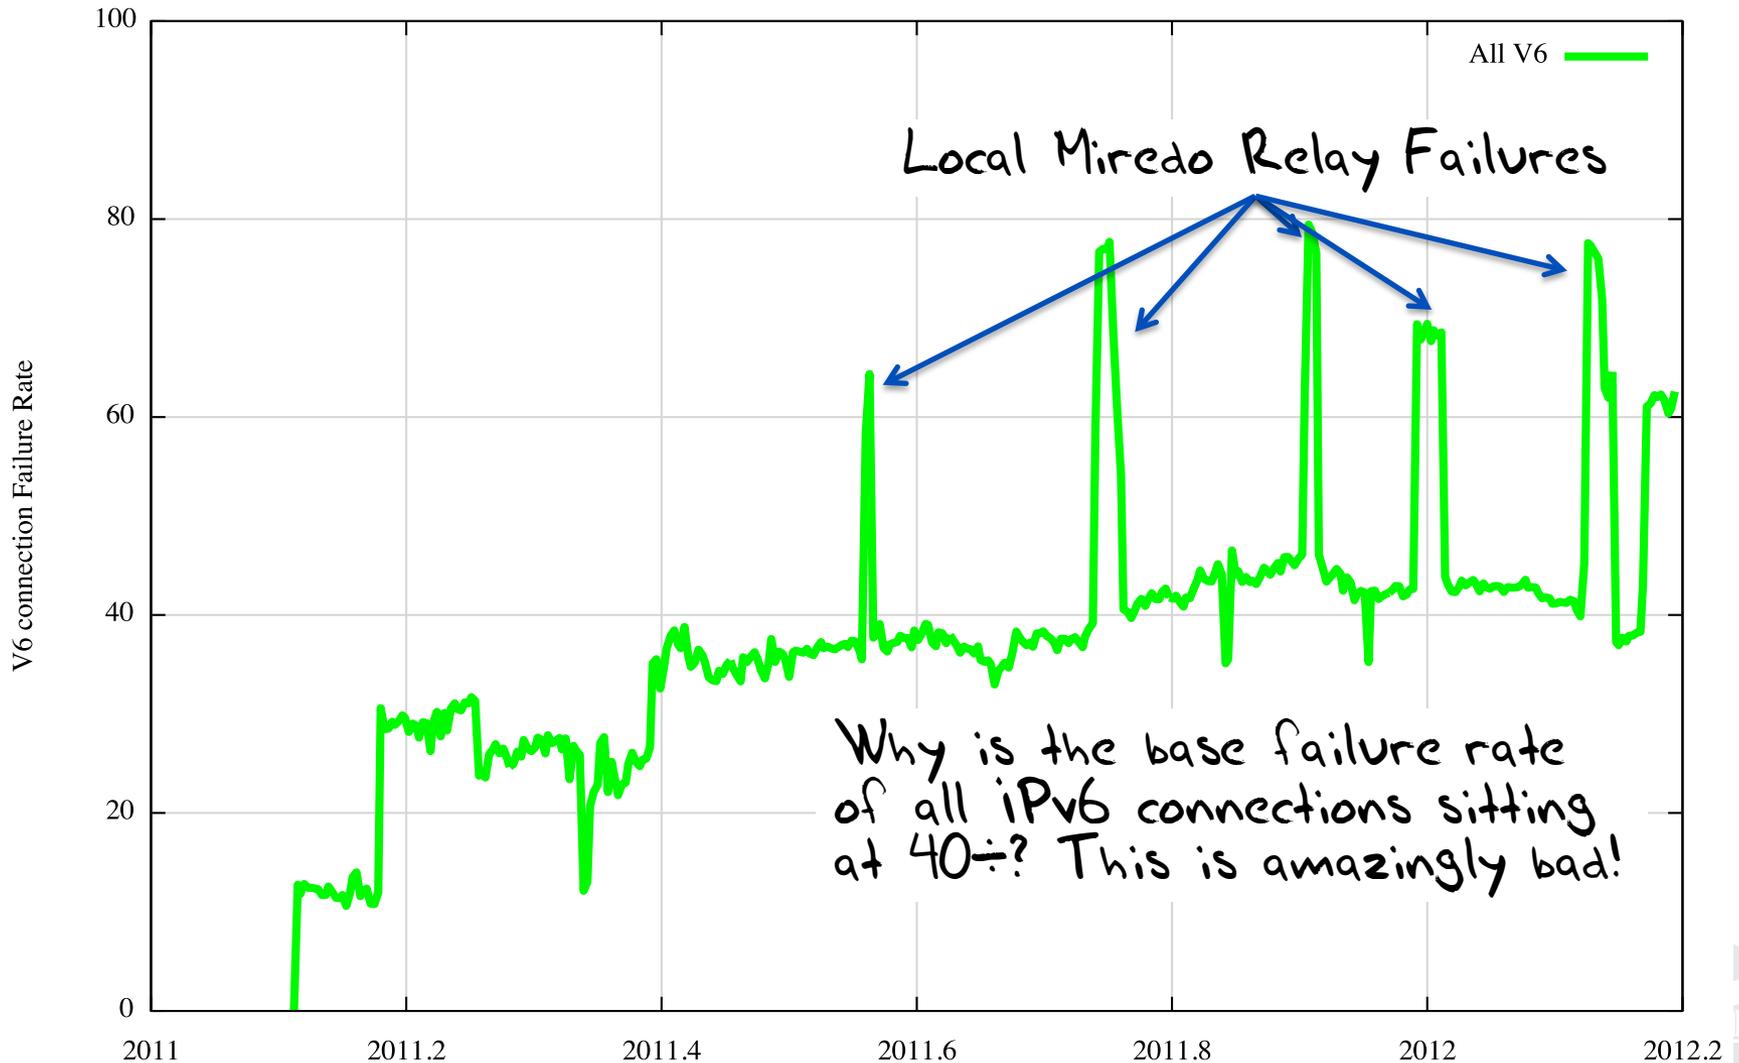
Aside from residual IPv4 failures in the image fetch due to device resets, connection dropouts, etc, the bulk of the recorded failures here is probably attributable to bots doing all-of-address scanning on port 80 – or maybe not!





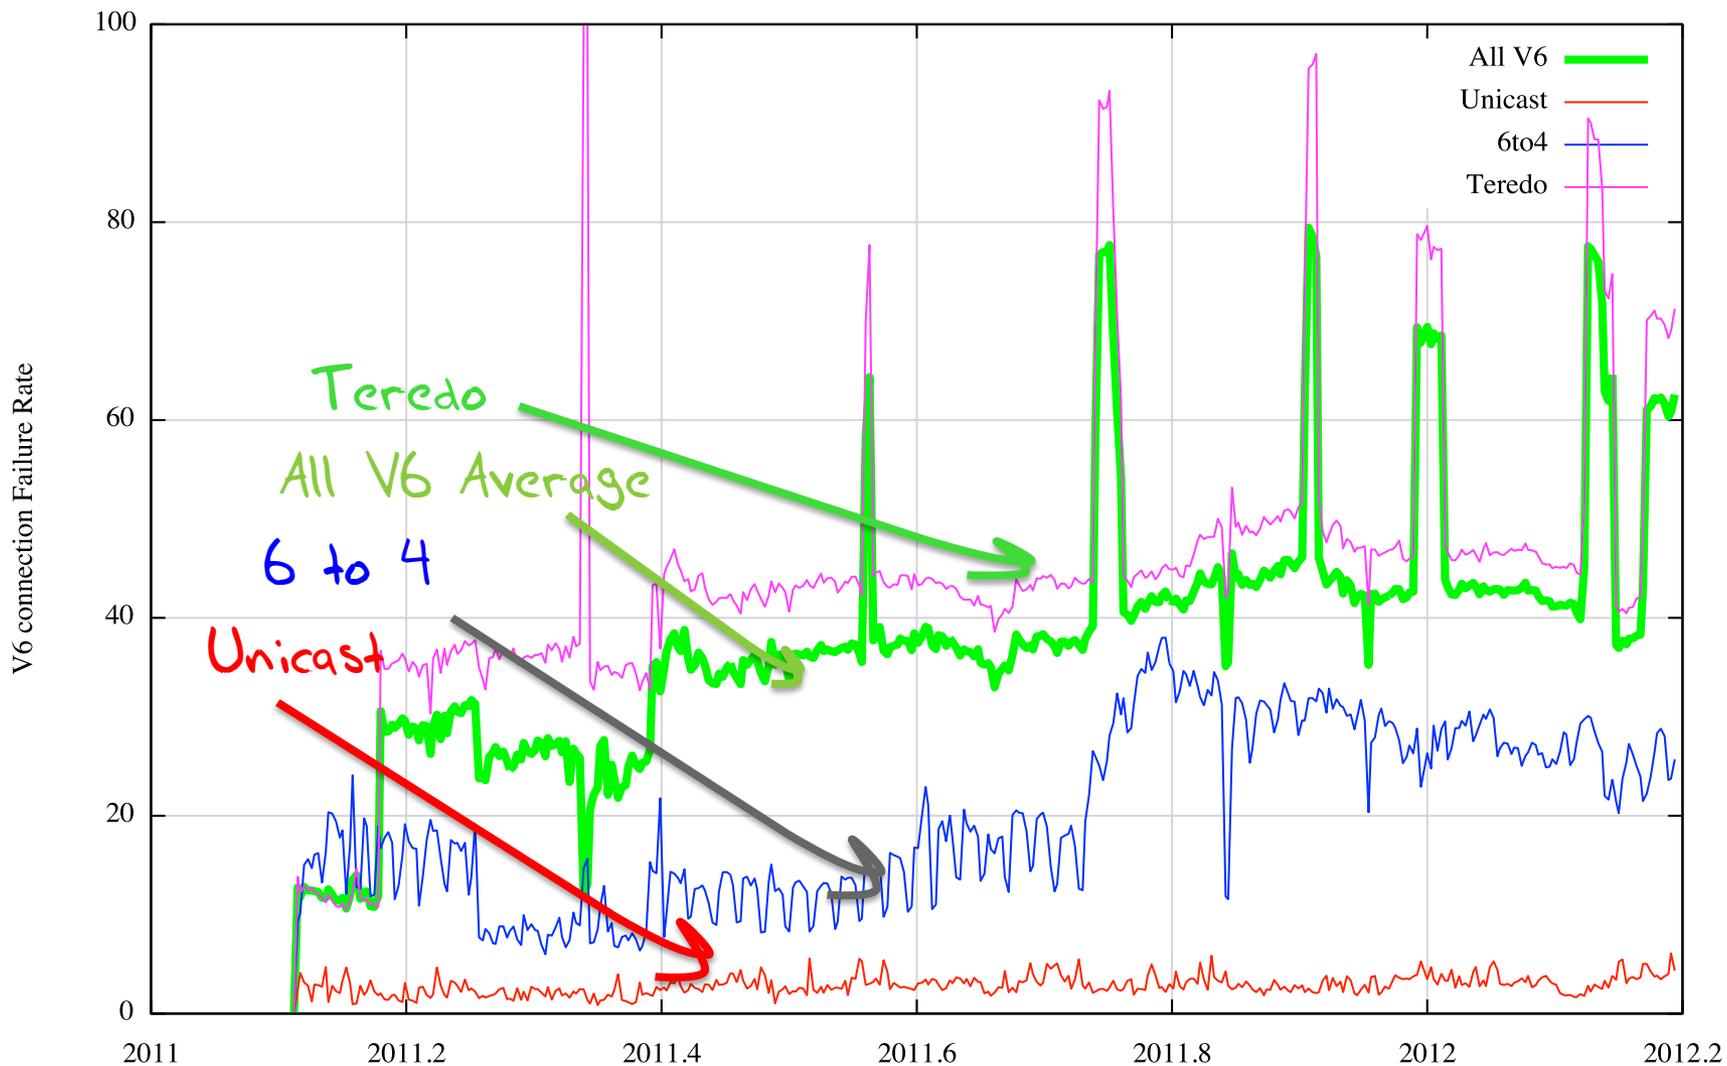
# What about IPv6?

Connection Failure Rate - V6



# V6 Failure Rate by Address Type

Connection Failure Rate - V6



# 6to4 Failure is Local Failure

6to4 failure appears to be related to two factors:

1. The client's site has a protocol 41 firewall filter rule for incoming traffic (this is possibly more prevalent in AsiaPac than in Europe)
2. Load / delay / reliability issues in the server's chosen outbound 6to4 relay (noted in the data gathered at the US server)

Even so, the 10% to 20% connection failure rate for 6to4 is unacceptably high!



# V6 Unicast Failures

January – March 2012:

110,761 successful V6 connecting endpoints

6,227 failures

That's a failure rate of 5.3%!

7 clients used fe80:: link local addresses

7 clients used fc00:/7 ULA source addresses

2 clients used fec0::/16 deprecated site local addresses

16 clients used 1f02:d9fc::/16

Nobody used 3ffe::/16 prefixes!



# Data Set 2: Connection Failure in 2015/2016

January 2015– January 2016

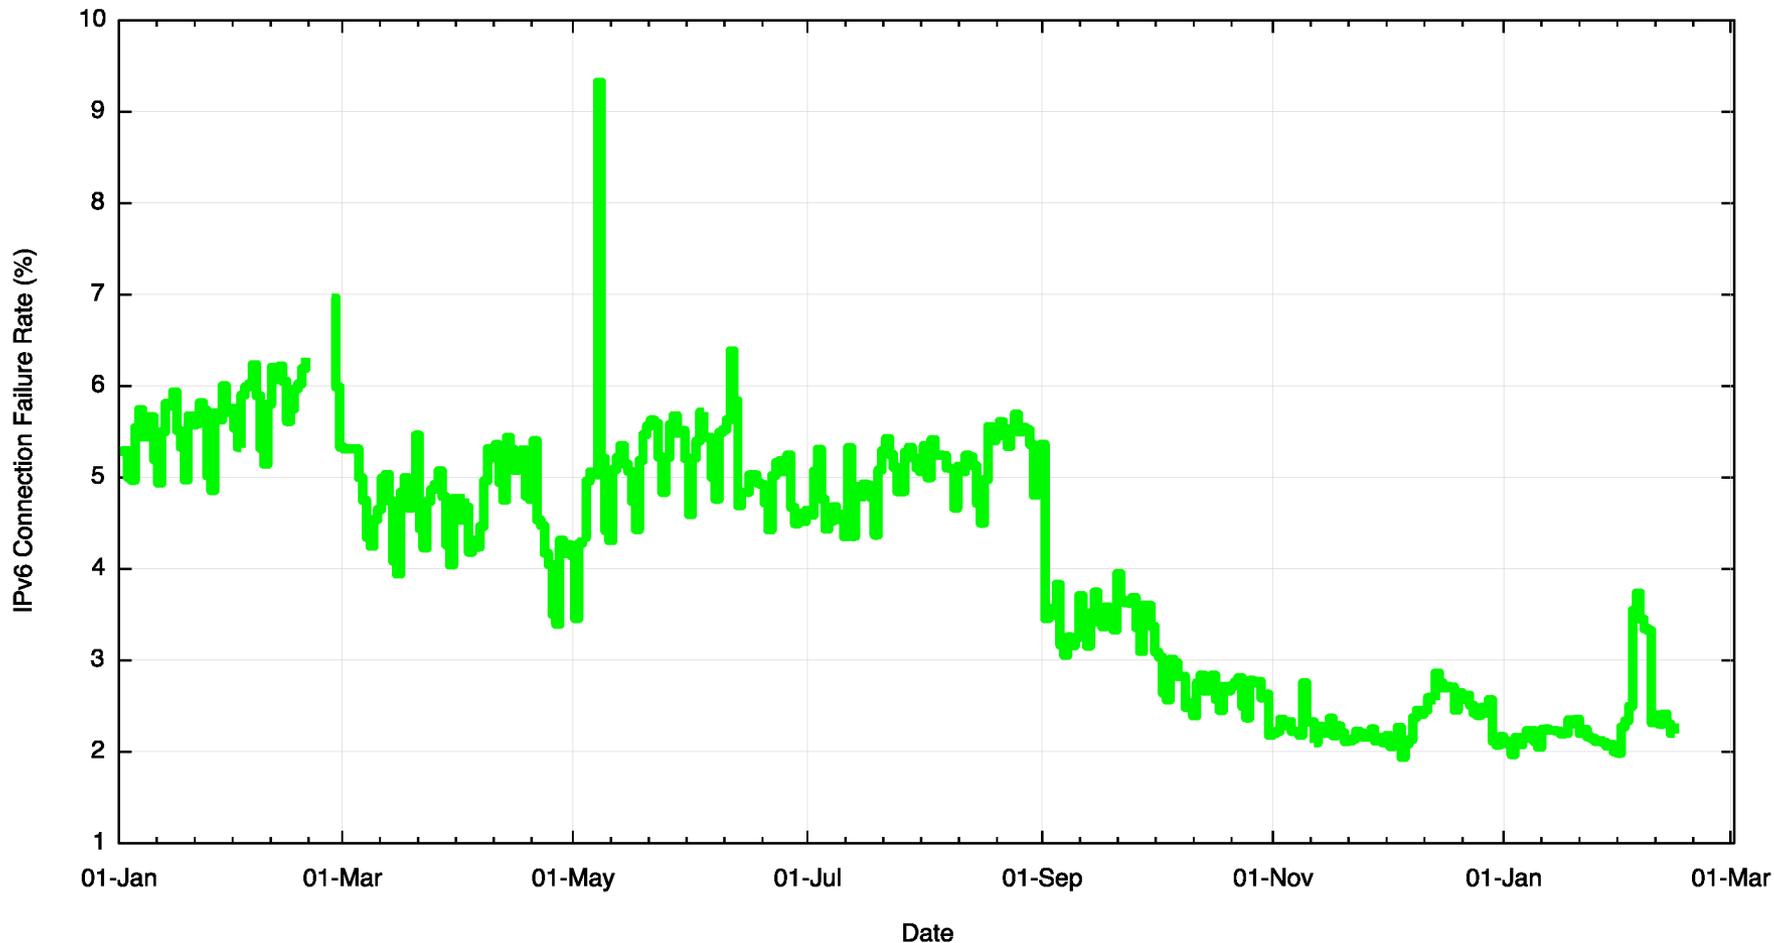
40,359,805 IPv6 endpoints

1,361,256 Failure rate (**3.37%**)



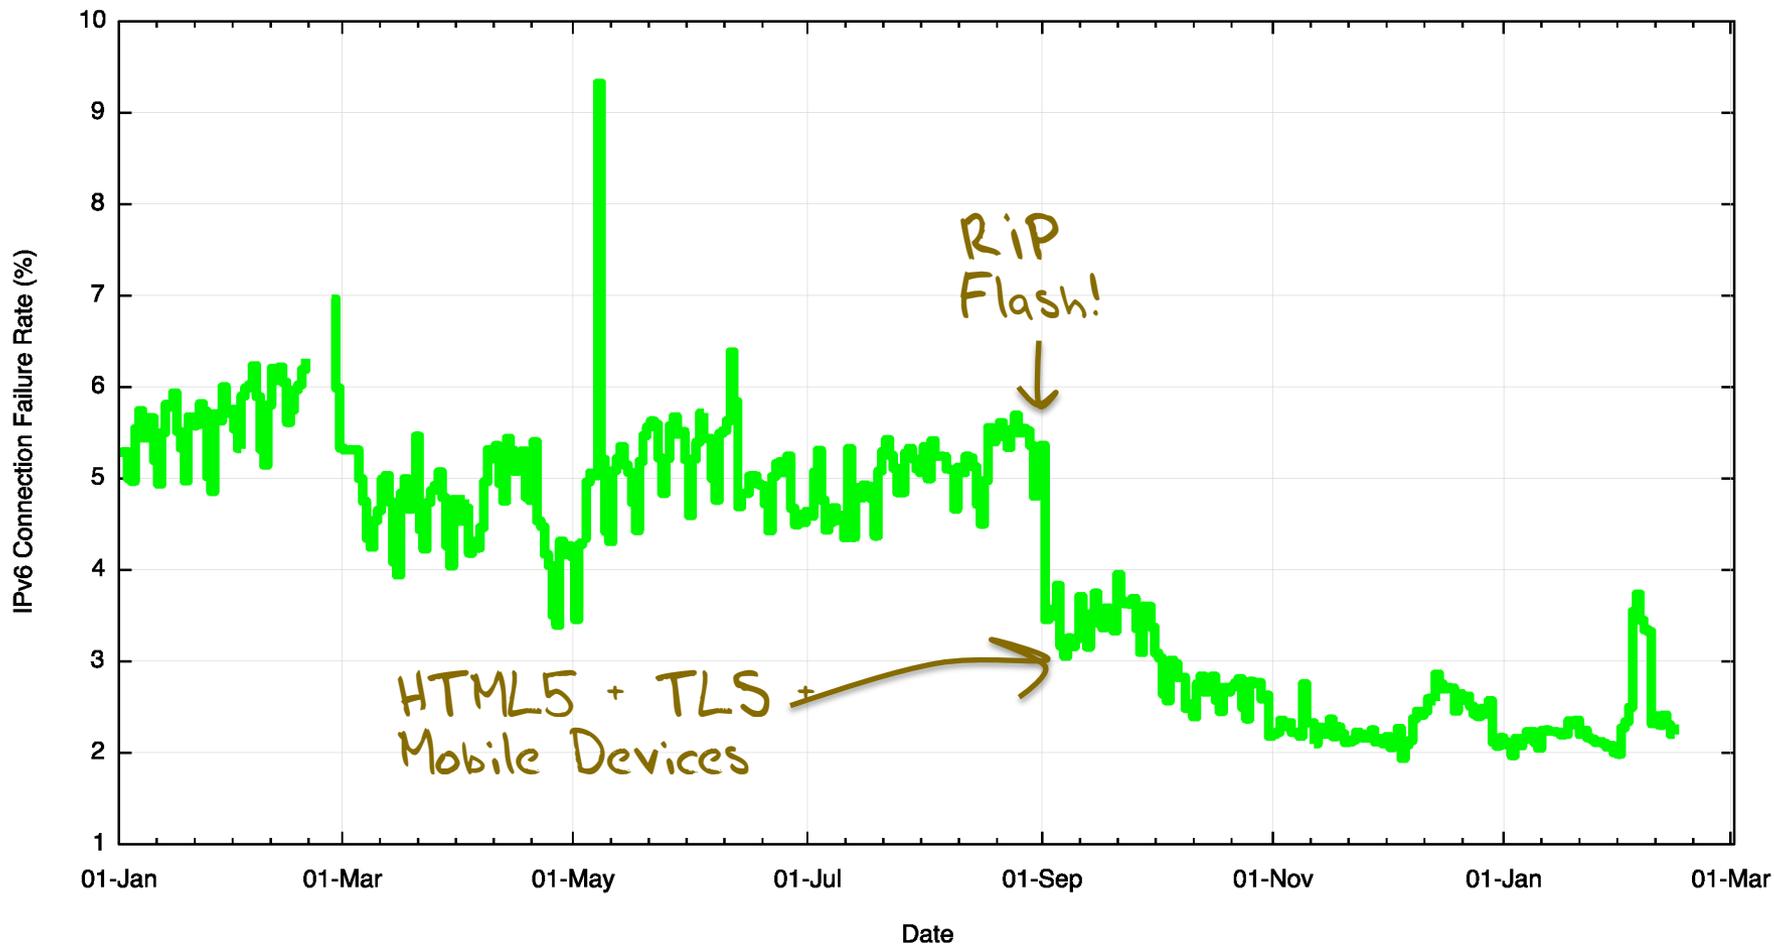
# Daily IPv6 Failures

IPv6 Daily Connection Failure Rate - 2015



# Daily IPv6 Failures

IPv6 Daily Connection Failure Rate - 2015



# 6to4

7,693,849 6to4 endpoints

- 19% of all IPv6 used 6to4

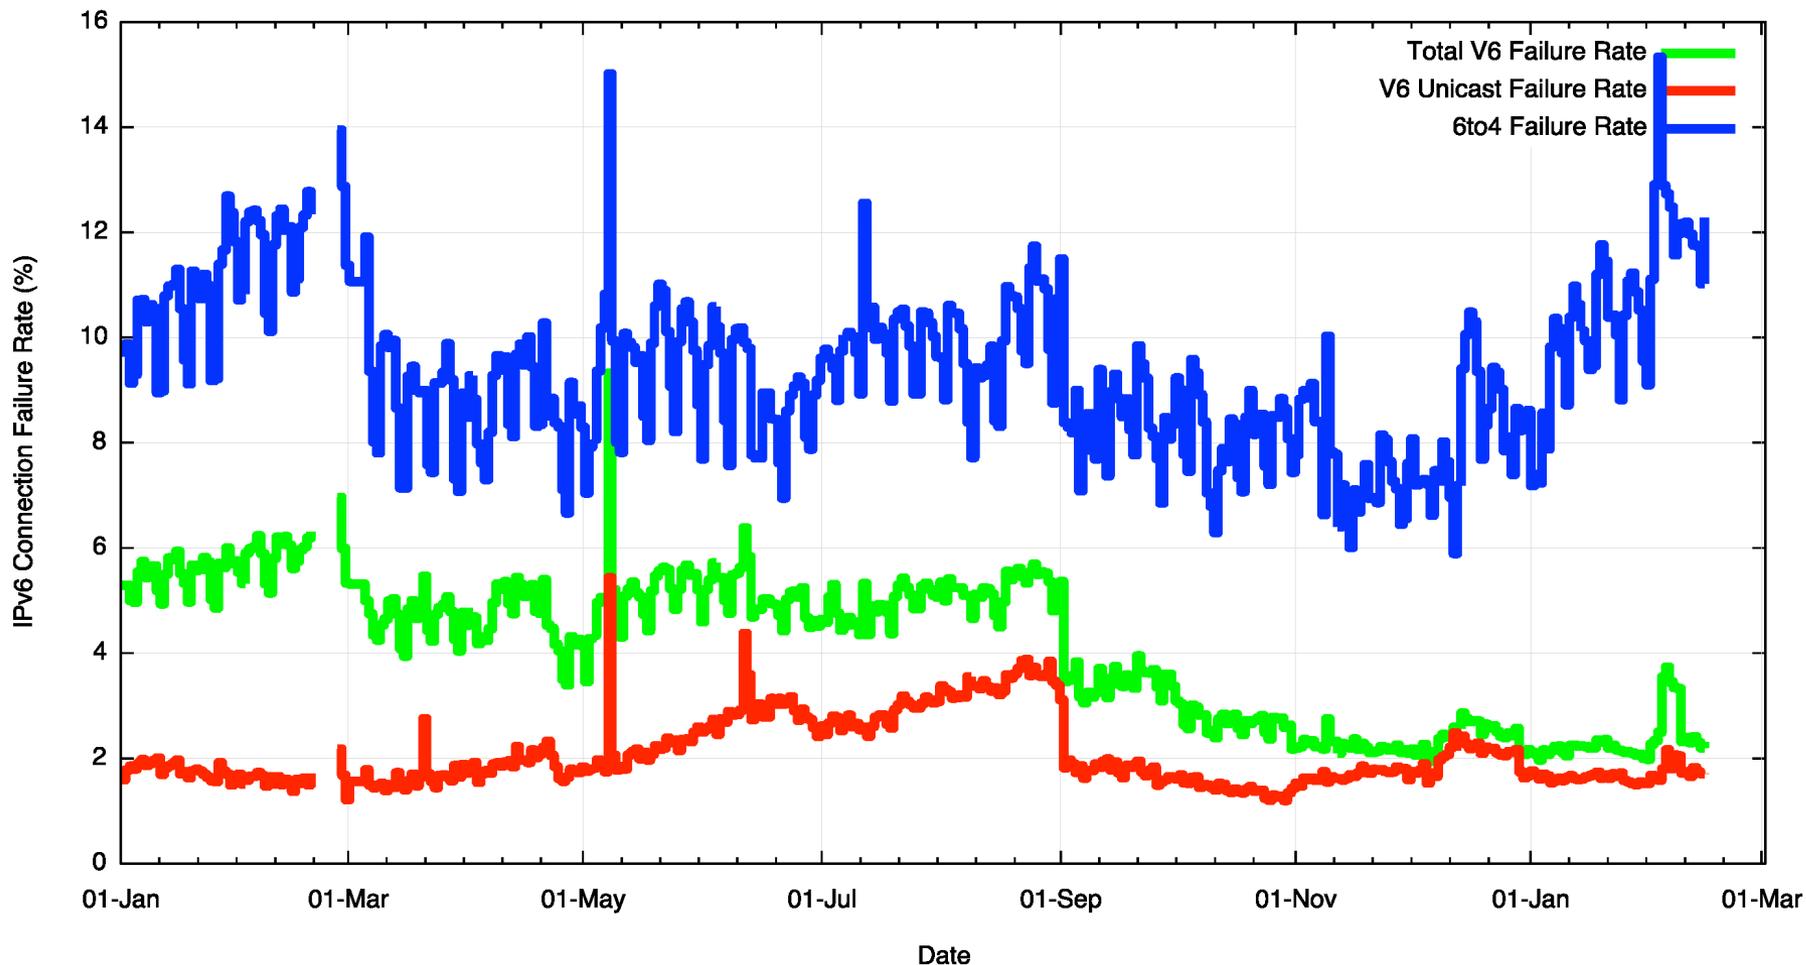
- 9% failure rate within the set of 6to4 connections

*This is still very high!*



# Daily IPv6 Failures

IPv6 Daily Connection Failure Rate - 2015

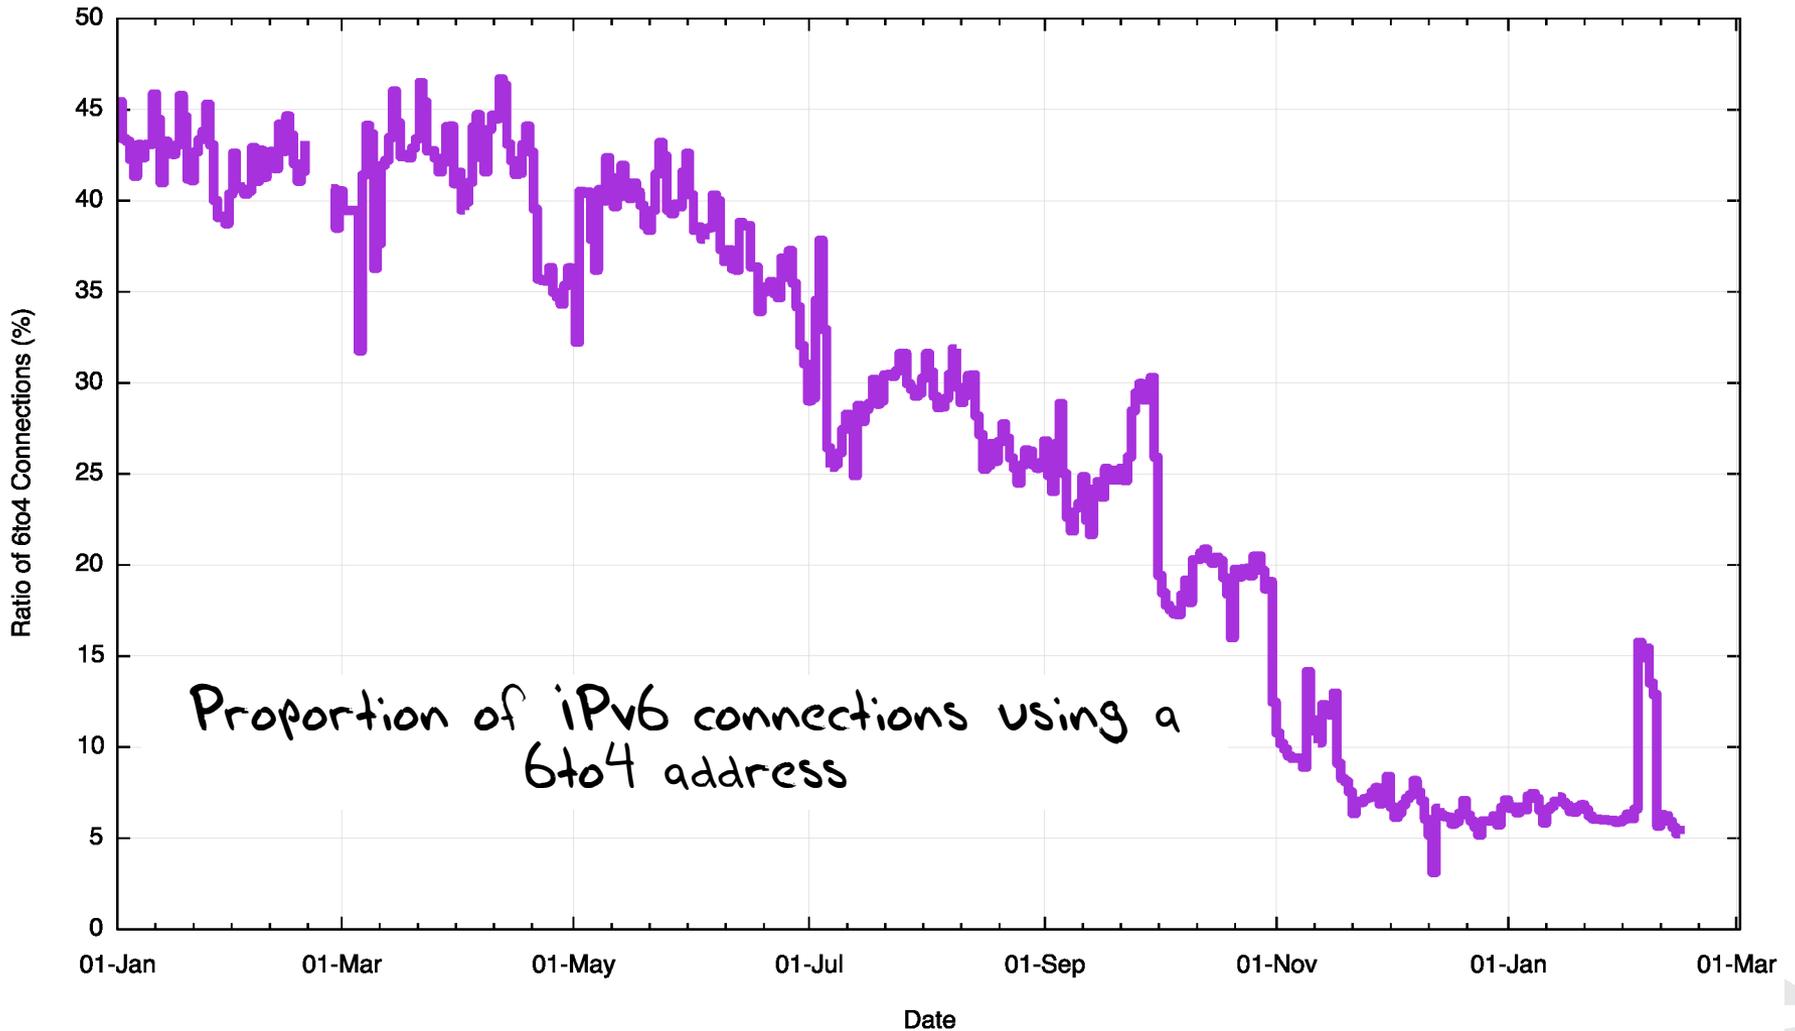


# Daily IPv6 Failures

- 6to4 failure rate has improved from 15%-20% in 2011 to 9% in 2015
- Teredo has all but disappeared
- Unicast failure rate is between 1.5% and 4% in 2015
  - Current unicast failure rate is 2%



# Killing off 6to4



# IPv6 Failures - Sep 2015 - Jan 2016

20,872,173 unique IPv6 Addresses

464,344 failing IPv6 addresses

142,362 6to4 addresses

138 teredo addresses

68 fe80:: local scope addresses

834 unallocated addresses

1,244 unannounced addresses

319,698 addresses from unicast allocated routed space

216,620 unique /64s



# Origin AS's with High IPv6 Failure Rates

AS	Failure Rate	Samples	AS Name
AS13679	97.33%	374	Centros Culturales de Mexico, A.C.,MX
AS201986	93.69%	222	ARPINET Arpinet LLC,AM
AS17660	65.14%	1,374	DRUKNET-AS DrukNet ISP,BT
AS10349	60.29%	763	TULANE - Tulane University,US
AS21107	46.97%	692	BLICNET-AS Blicnet d.o.o.,BA
AS20880	42.65%	762	TELECOLUMBUS Tele Columbus AG,DE
AS12779	36.70%	109	ITGATE IT.Gate S.p.A.,IT
AS46261	35.64%	101	QUICKPACKET - QuickPacket, LLC,US
AS9329	35.29%	119	SLTINT-AS-AP Sri Lanka Telecom Internet,LK
AS52888	27.92%	265	UNIVERSIDADE FEDERAL DE SAO CARLOS,BR
AS30036	27.55%	60,228	Mediacom Communications Corp,US
AS45920	25.77%	163	SKYMESH-AS-AP SkyMesh Pty Ltd,AU
AS210	25.04%	571	WEST-NET-WEST - Utah Education Network,US
AS28343	24.57%	985	TPA TELECOMUNICACOES LTDA,BR
AS7477	21.72%	488	TEREDONN-AS-AP SkyMesh Pty Ltd,AU
AS24173	21.48%	256	NETNAM-AS-AP Netnam Company,VN
AS28580	21.48%	1,341	CILNET Comunicacao e Informatica LTDA.,BR
AS32329	20.63%	126	MONKEYBRAINS - Monkey Brains,US
AS17451	19.35%	248	BIZNET-AS-AP BIZNET NETWORKS,ID
AS5707	19.35%	155	UTHSC-H - The University of Texas Health



# Origin AS's with Zero Failure Rates

AS	Fail Rate	Samples	AS Name
AS3676	0.00%	2,149	UIOWA-AS - University of Iowa,US
AS55536	0.00%	1,548	PSWITCH-HK PACSWITCH GLOBAL IP NETWORK,HK
AS57026	0.00%	1,188	CHEB-AS JSC "ER-Telecom Holding",RU
AS133414	0.00%	1,179	FOXTEL-AS-AP Foxtel Management Pty Ltd,AU
AS18144	0.00%	1,179	AS-ENECOM Energia Communications,Inc.,JP
AS196705	0.00%	936	ARDINVEST Ardinvest LTD,UA
AS21191	0.00%	816	ASN-SEVERTTK TransTeleCom,RU
AS1239	0.00%	734	SPRINTLINK - Sprint,US
AS56420	0.00%	717	RYAZAN-AS JSC "ER-Telecom Holding",RU
AS33070	0.00%	656	RMH-14 - Rackspace Hosting,US
AS51819	0.00%	651	YAR-AS JSC "ER-Telecom Holding",RU
AS27357	0.00%	625	RACKSPACE - Rackspace Hosting,US
AS7233	0.00%	623	YAHOO-US - Yahoo,US
AS20130	0.00%	606	DEPAUL - Depaul University,US
AS49048	0.00%	604	TVER-AS JSC "ER-Telecom Holding",RU
AS25513	0.00%	481	ASN-MGTS-USPD OJS Moscow city telephone ne
AS53264	0.00%	426	CDC-LMB1 - Continuum Data Centers, LLC.,US
AS29854	0.00%	392	WESTHOST - WestHost, Inc.,US
AS13238	0.00%	391	YANDEX Yandex LLC,RU
AS10359	0.00%	372	EPICSYS - Epic Systems Corporation,US

Ranked by IPv6 measurement count



# What about IPv4 Connection Failures?

2011: failure rate 0.2%



# What about IPv4 Connection Failures?

2011: failure rate 0.2%

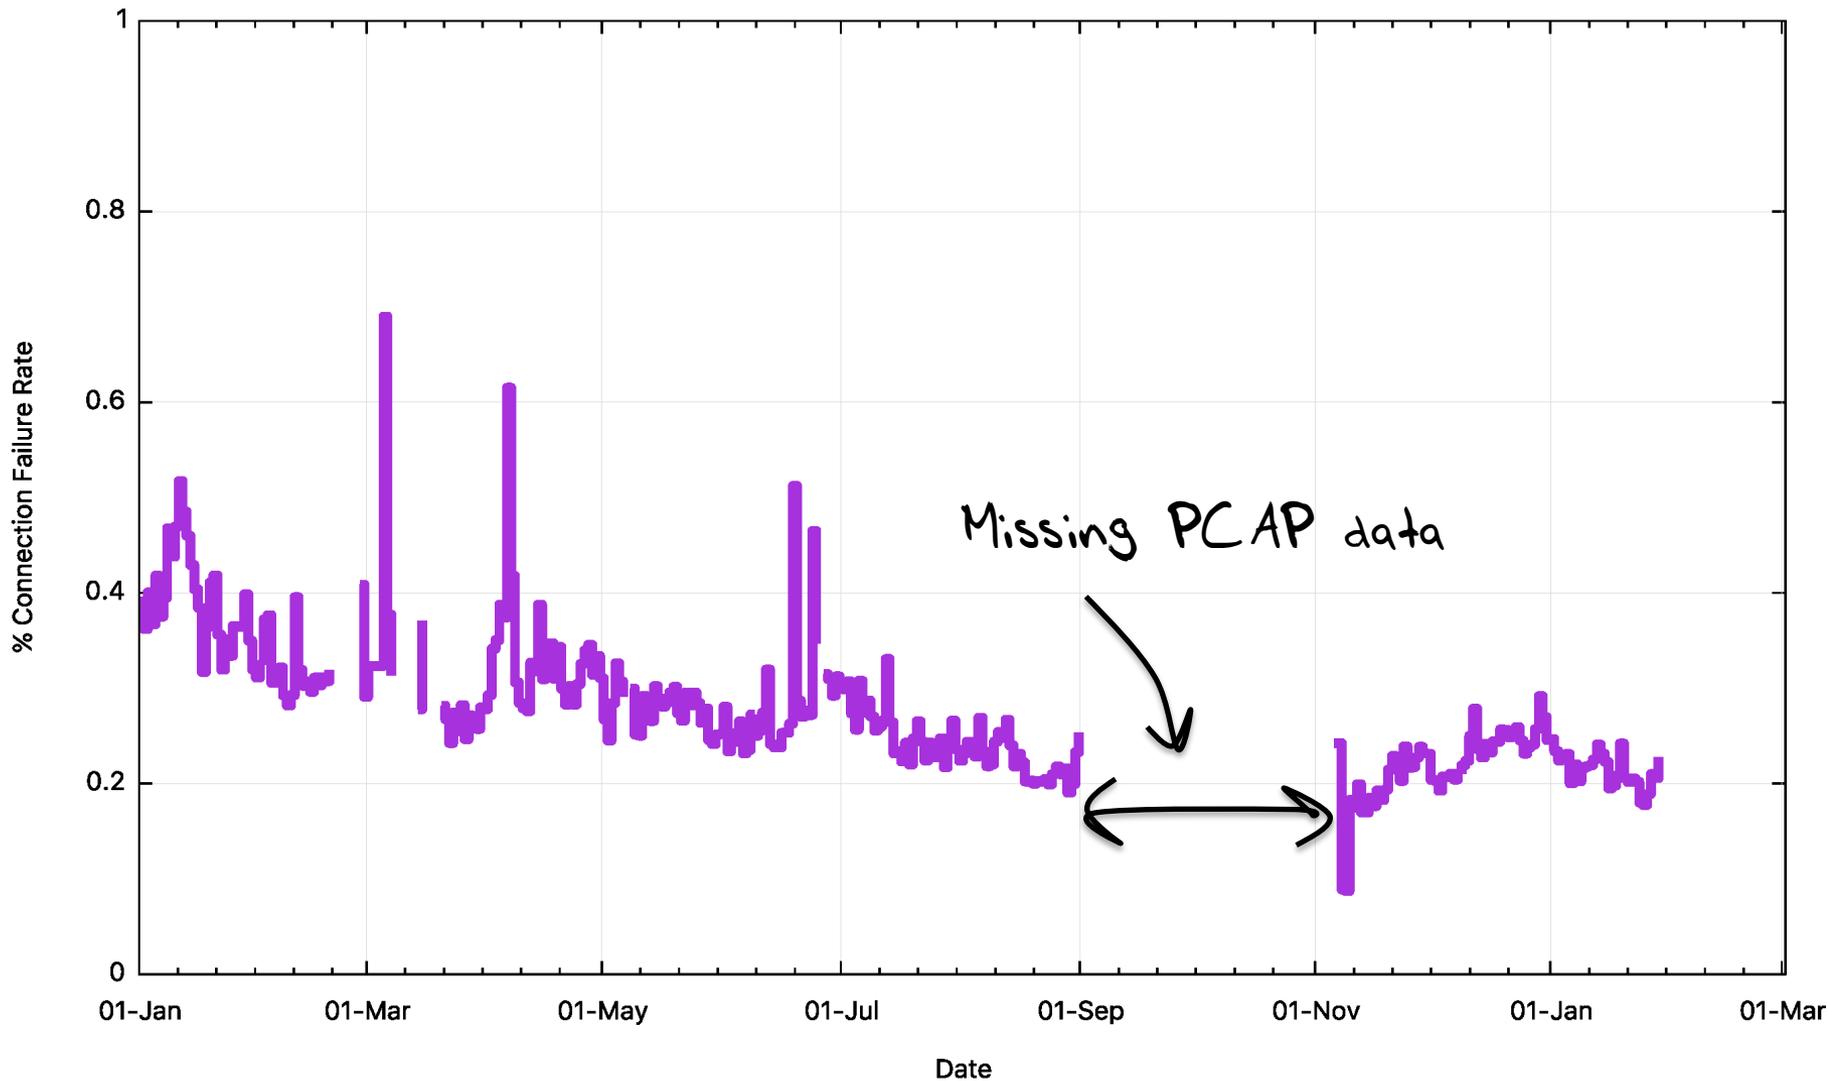
2015:

446,414,857 IPv4 endpoints

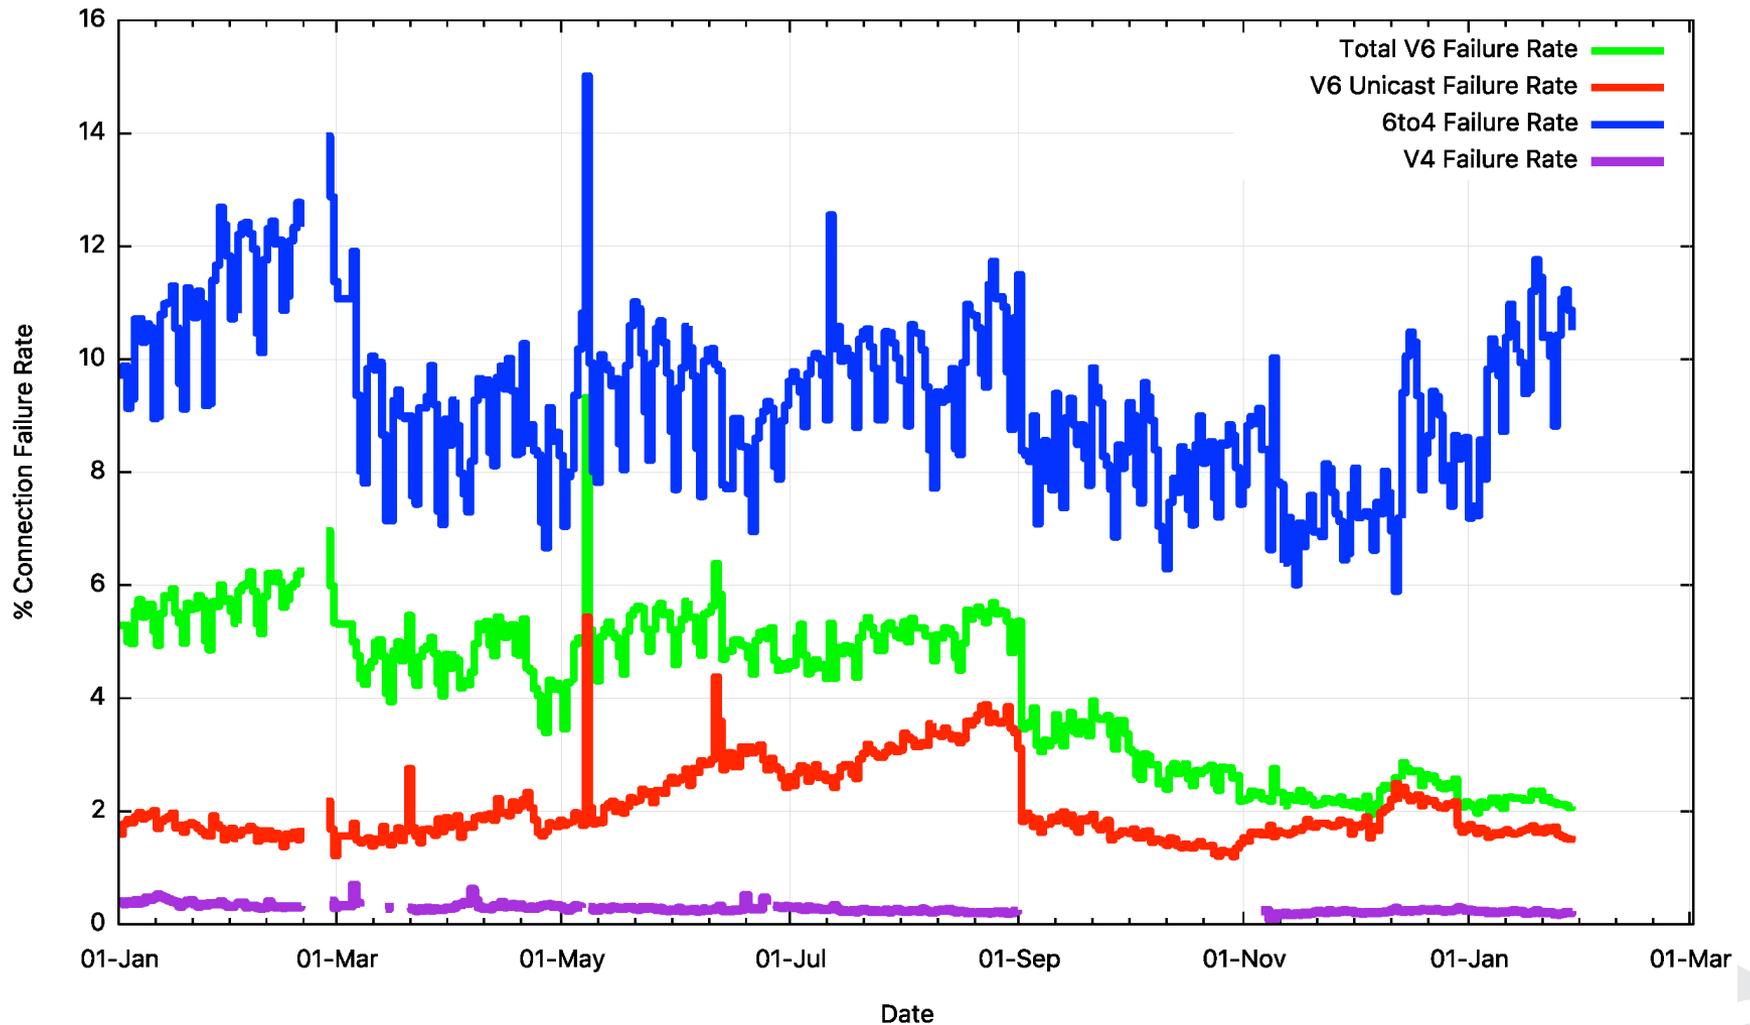
1,166,332 Connection Failures (0.26%)



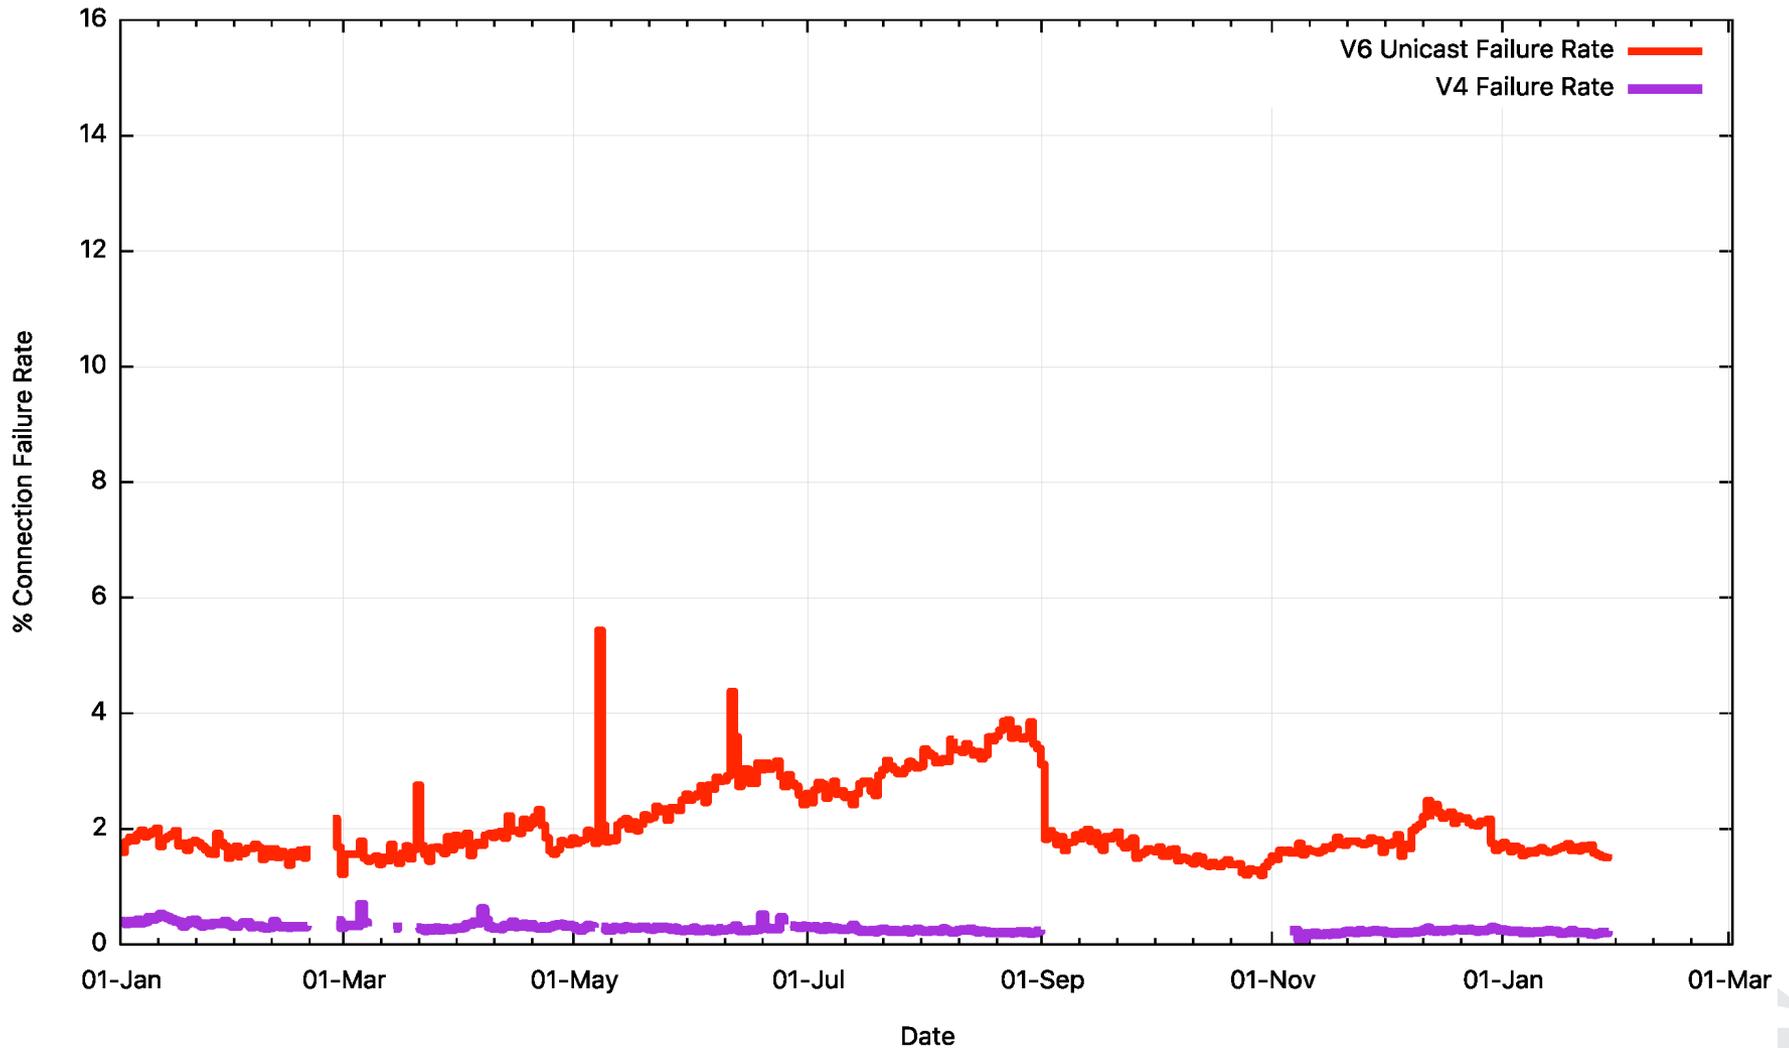
# IPv4 Connection Failure



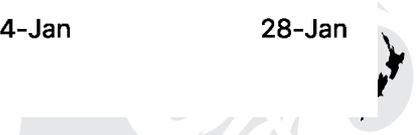
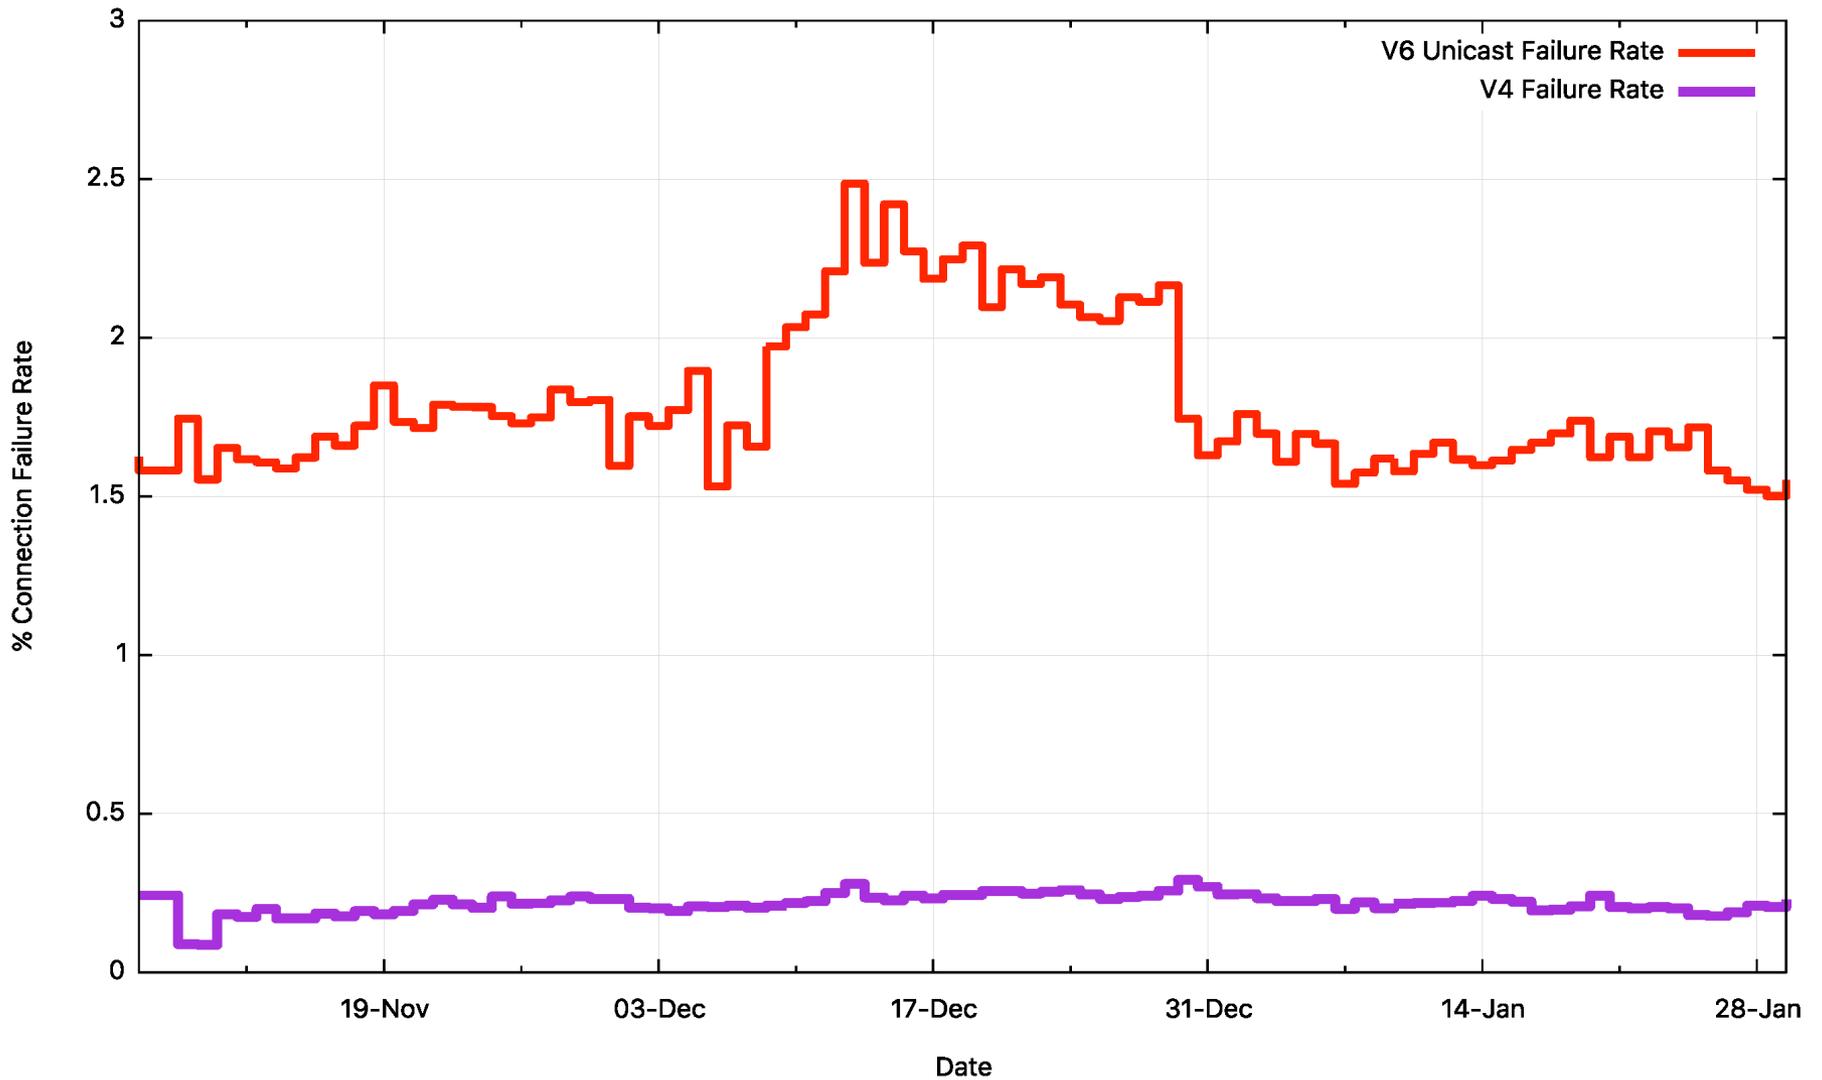
# Comparison



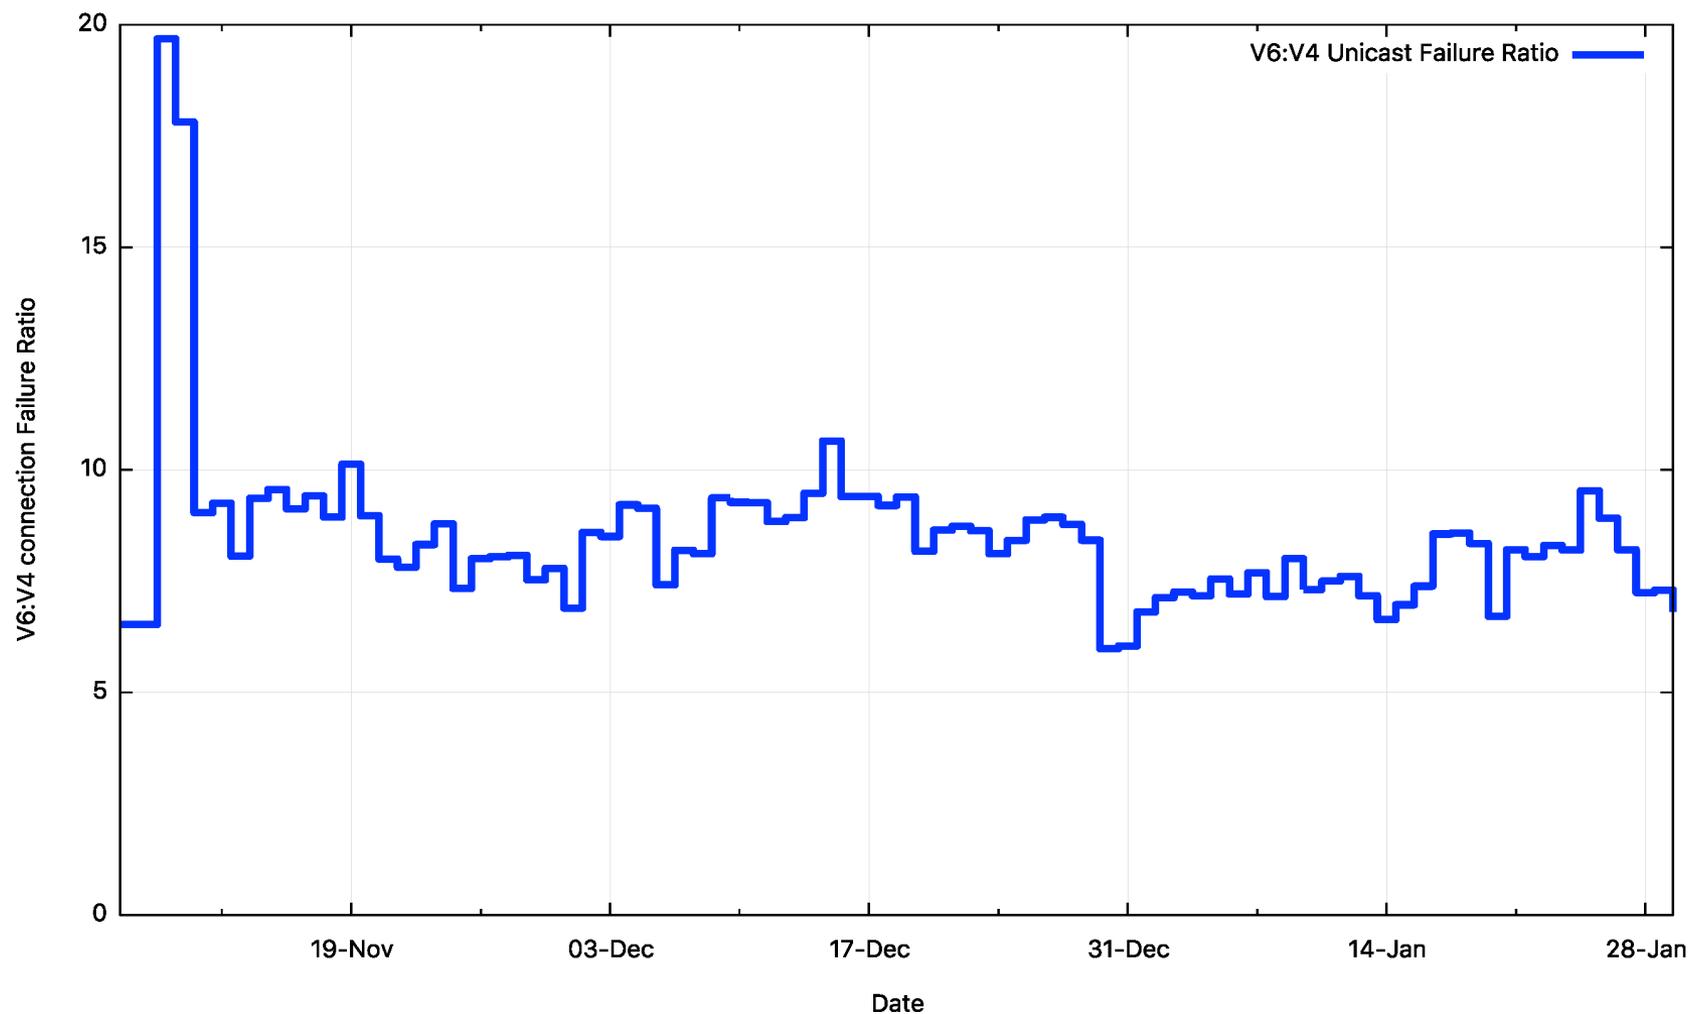
# Comparison: Unicast



# Comparison: Unicast



# Comparison Ratio: Unicast



# It's still not good!

IPv6 Unicast Failure rate: 1.6% (falling)

IPv4 Failure rate: 0.2% (steady)



# What are we looking at:

- How “reliable” are IPv6 connections?

*Do all TCP connection attempts succeed?*

- How “fast” are IPv6 connections?

*is V6 slower than V4?*



# Why SYNs?

- Every TCP session starts with a SYN handshake
- It's typically a kernel level operation, which means that there is little in the way of application level interaction with the SYN exchange
- On the downside there is only a single sample point per measurement



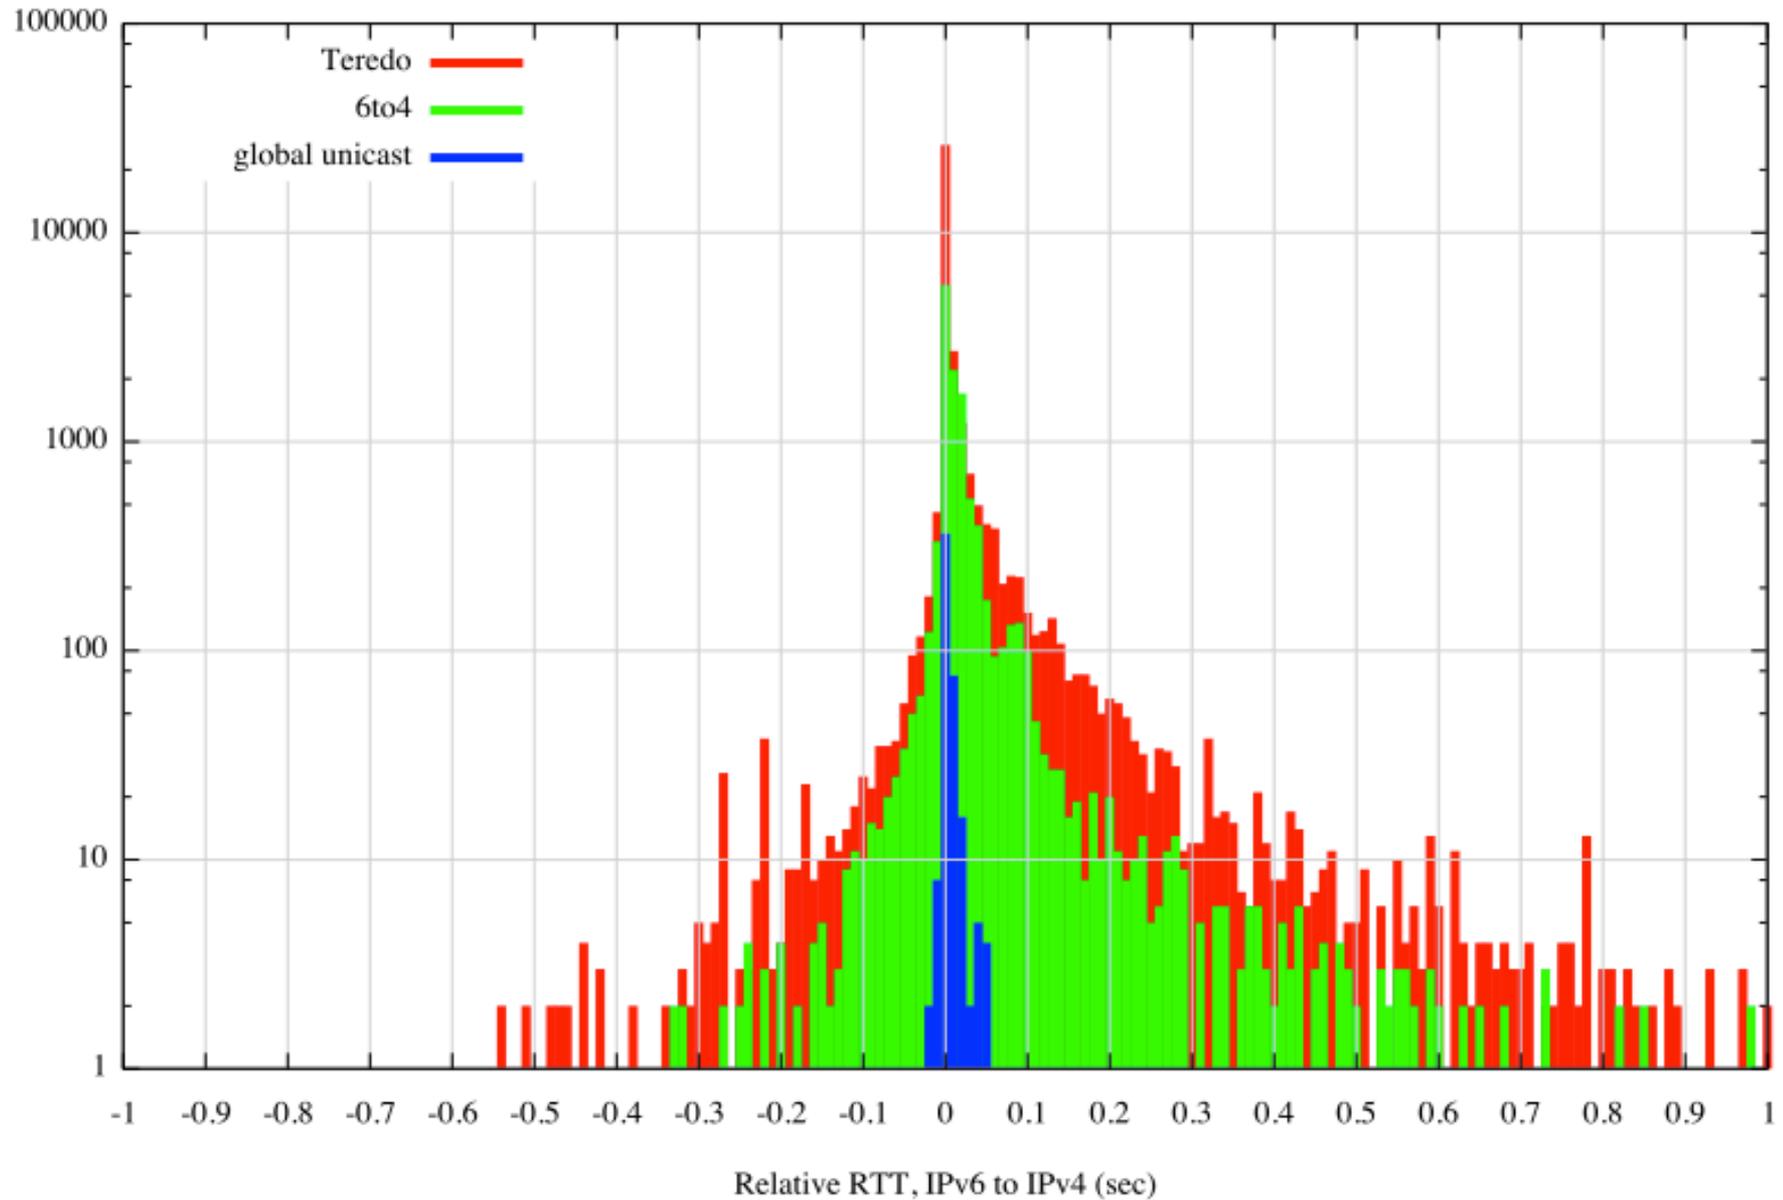
# Generating a comparative RTT profile

- For each successful connection couplet (IPv4 and IPv4) from the same endpoint, gather the pair of RTT measurements from the SYN-ACK exchanges
- Use the server's web logs to associate a couplet of IPv4 and IPv6 addresses
- Use the packet dumps to collect RTT information from the SYN-ACK Exchange
- Plot the difference in RTT in buckets



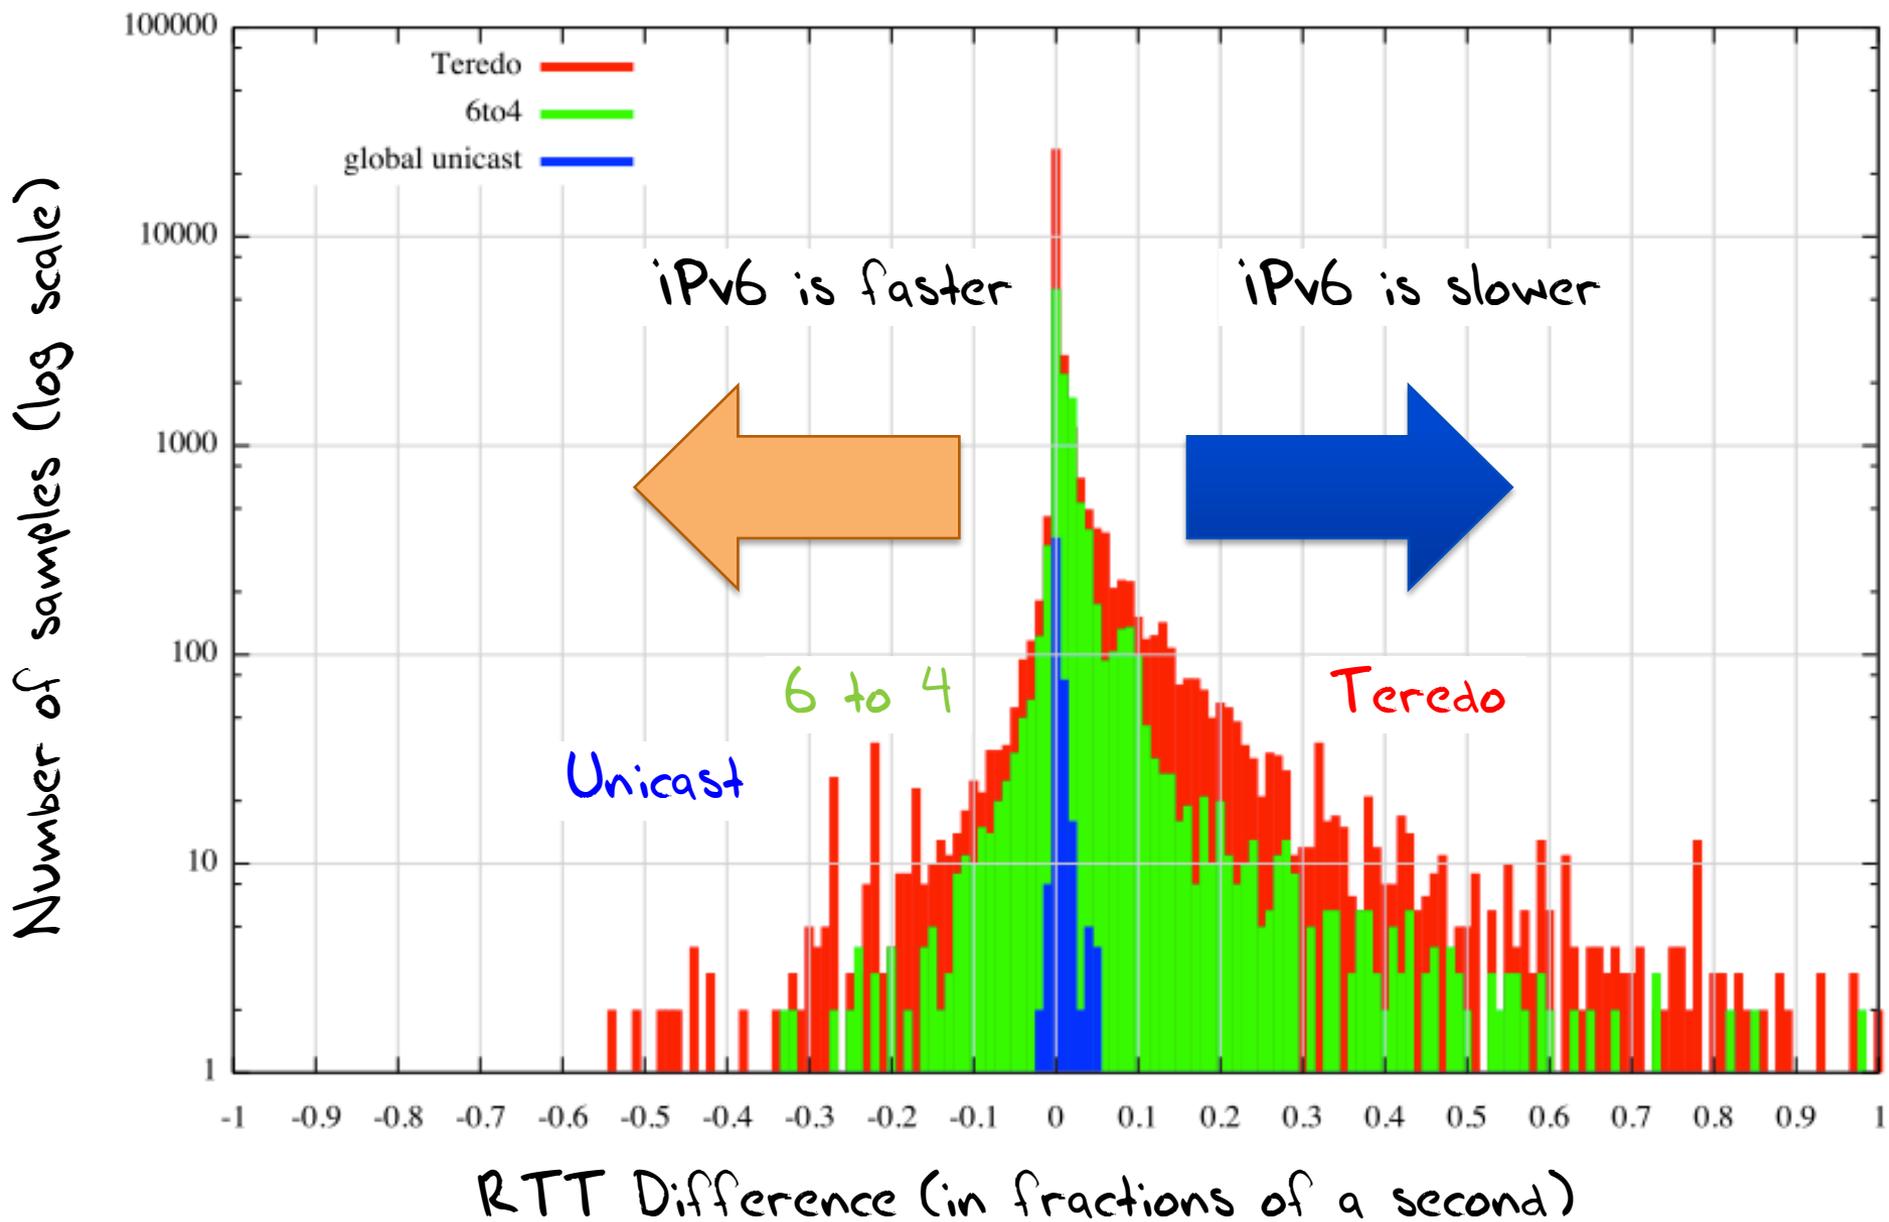
# 2012 Data

Relative RTT, IPv6 to IPv4 (sec) for bilby on 2012/03/01



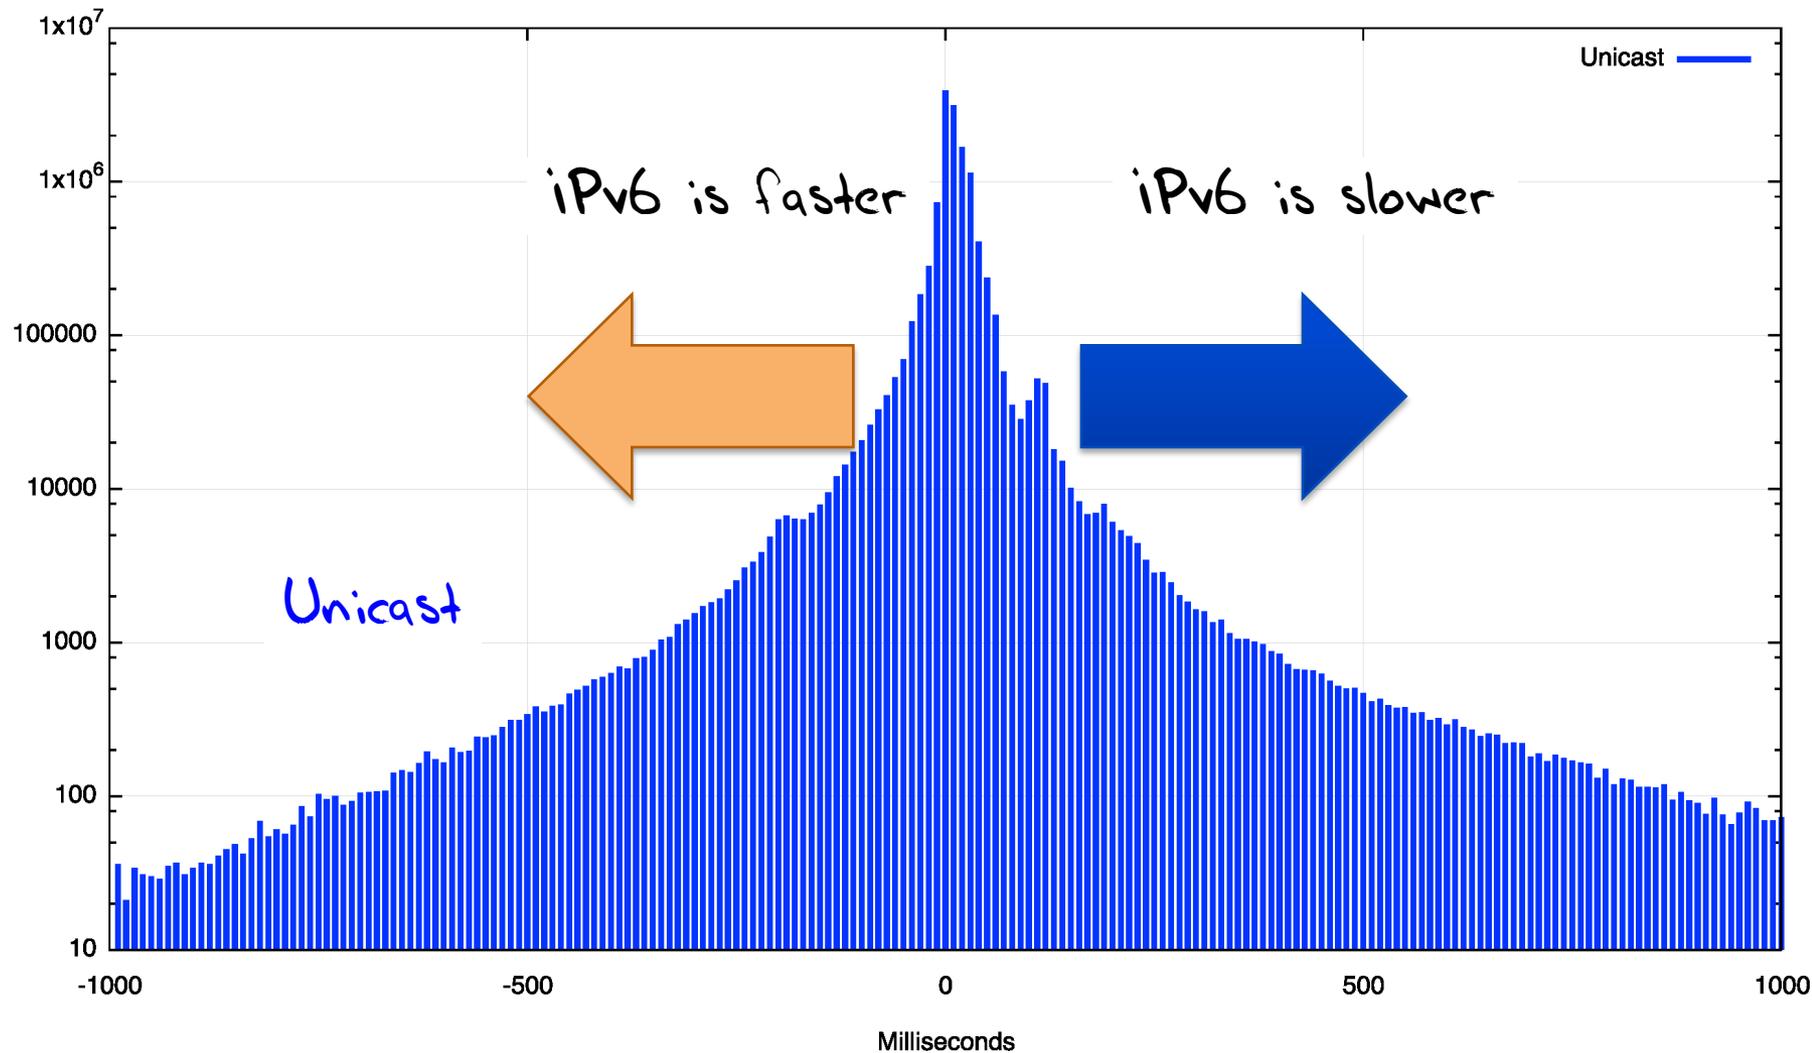
# 2012 Data

Relative RTT, IPv6 to IPv4 (sec) for bilby on 2012/03/01

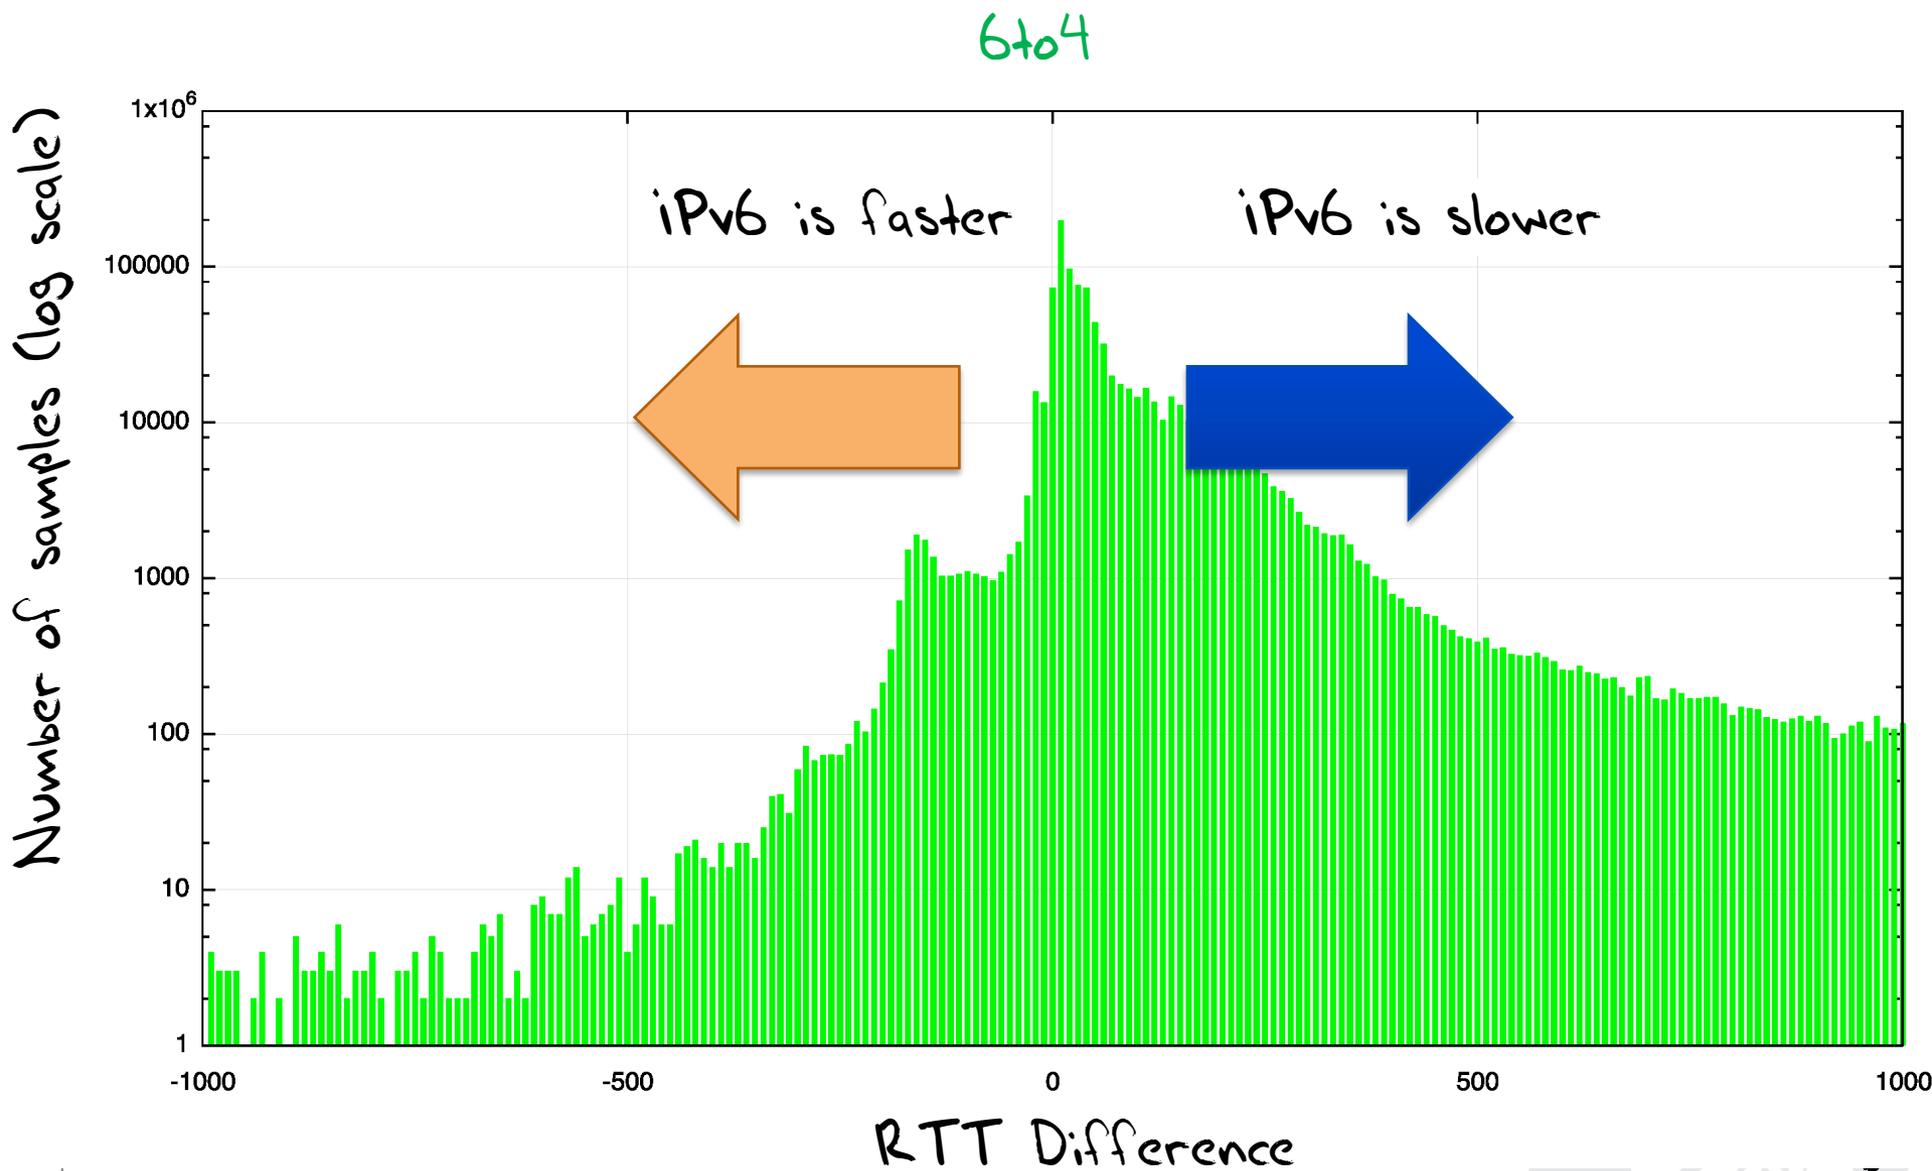


# December 2015/January 2016

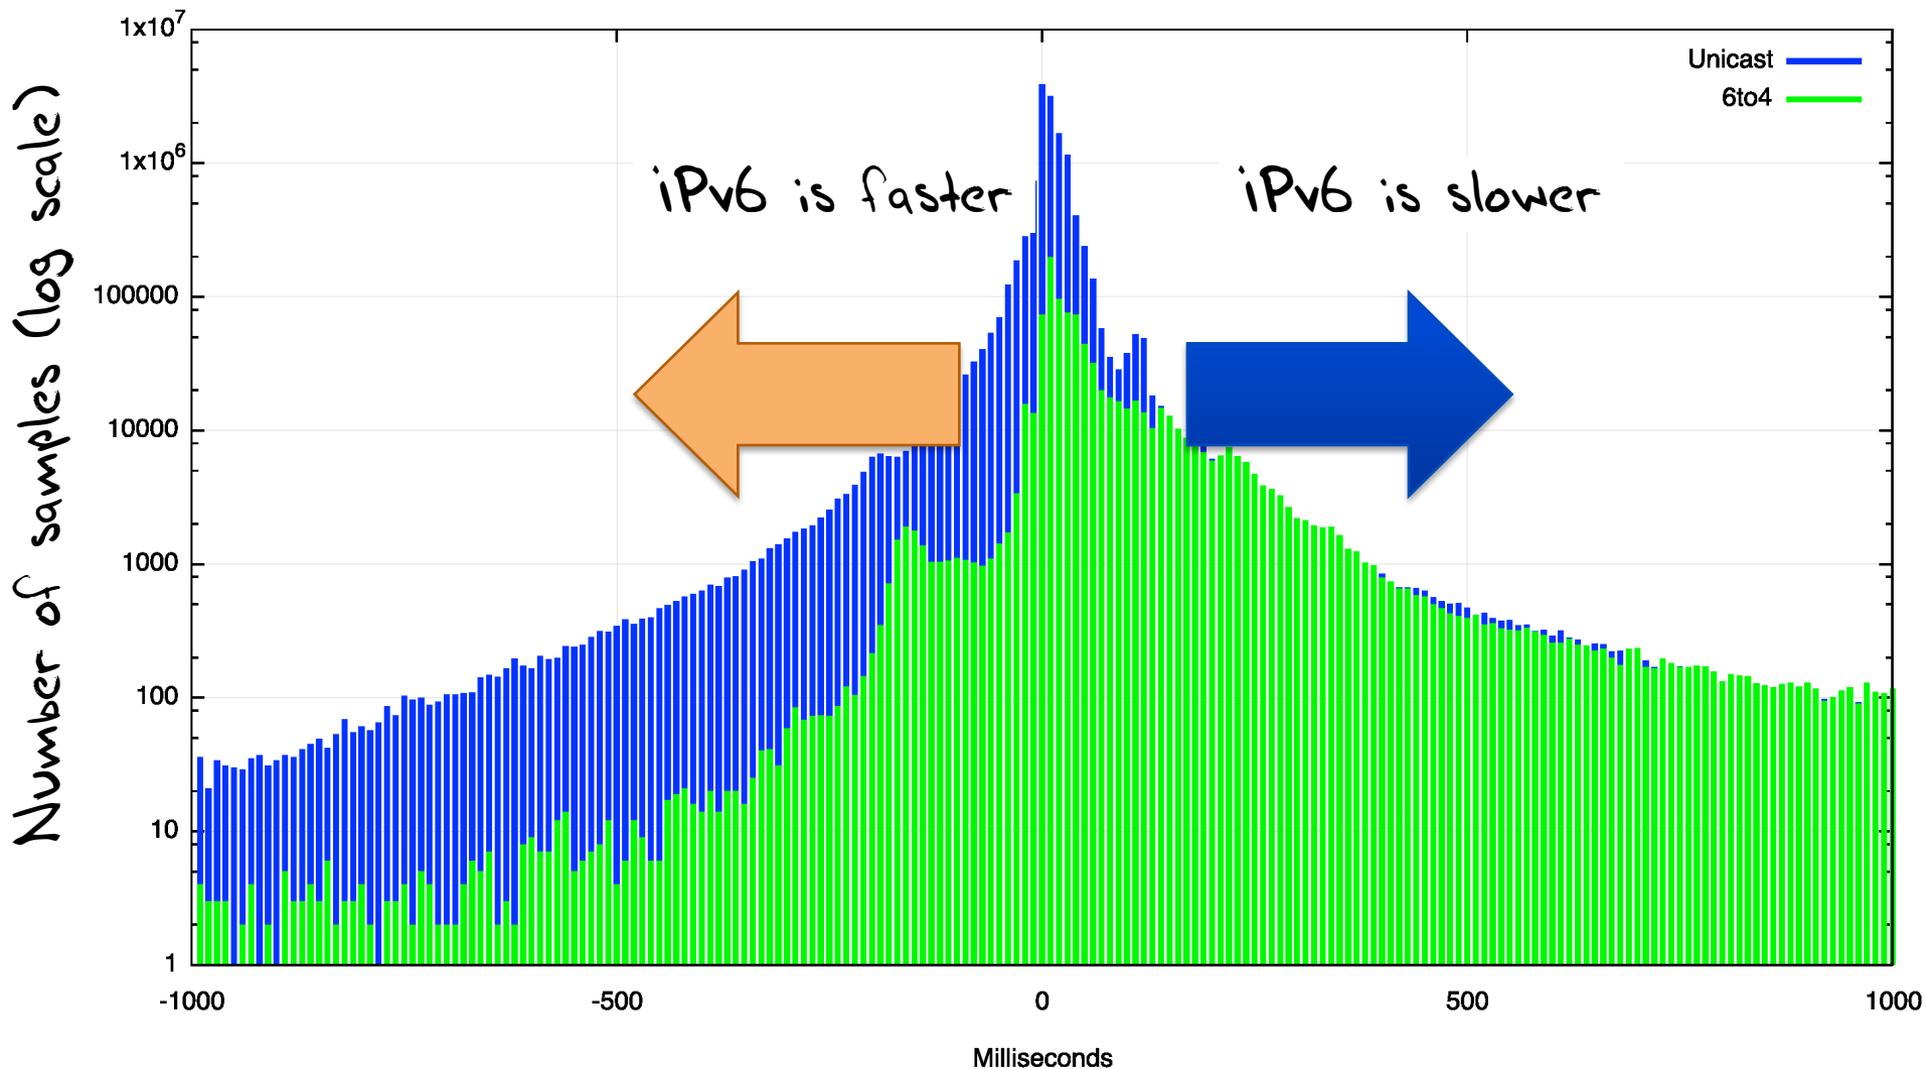
Number of samples (log scale)



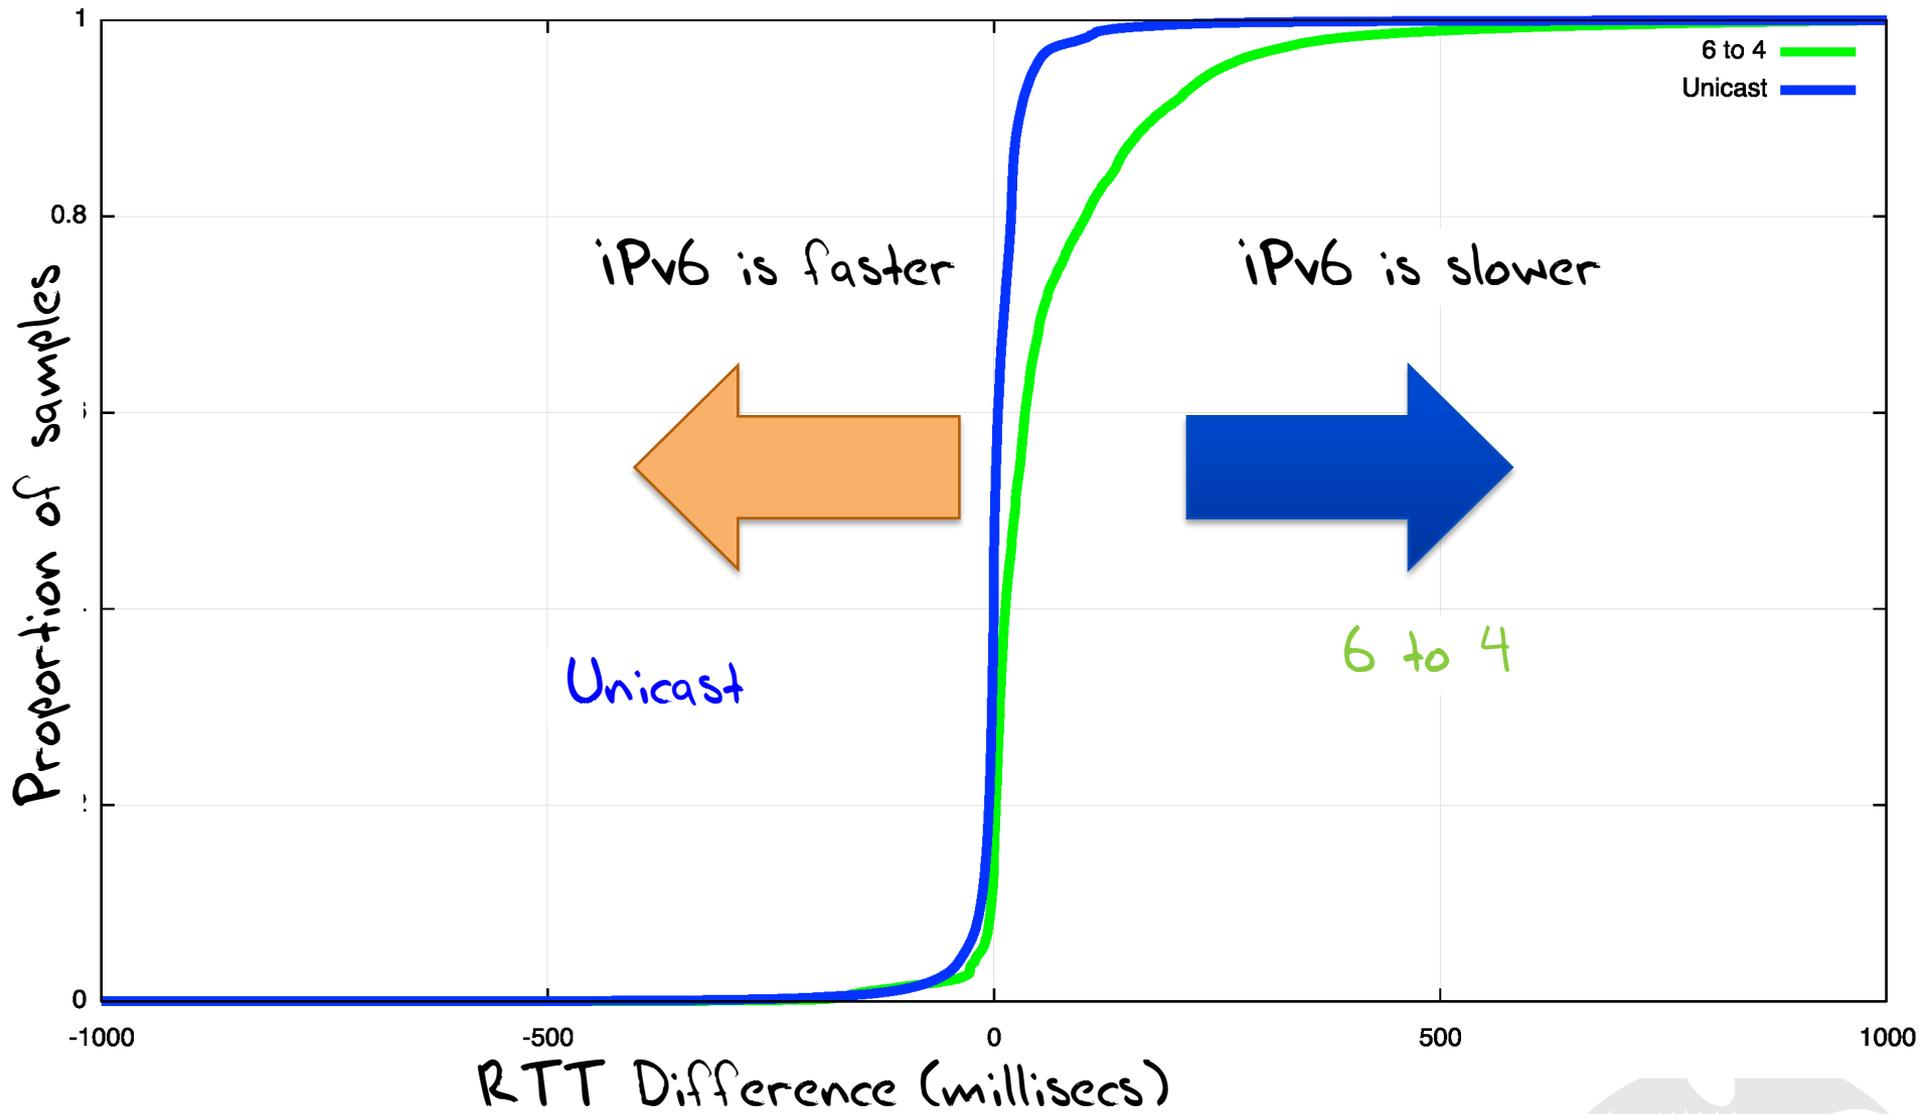
# December 2015/January 2016



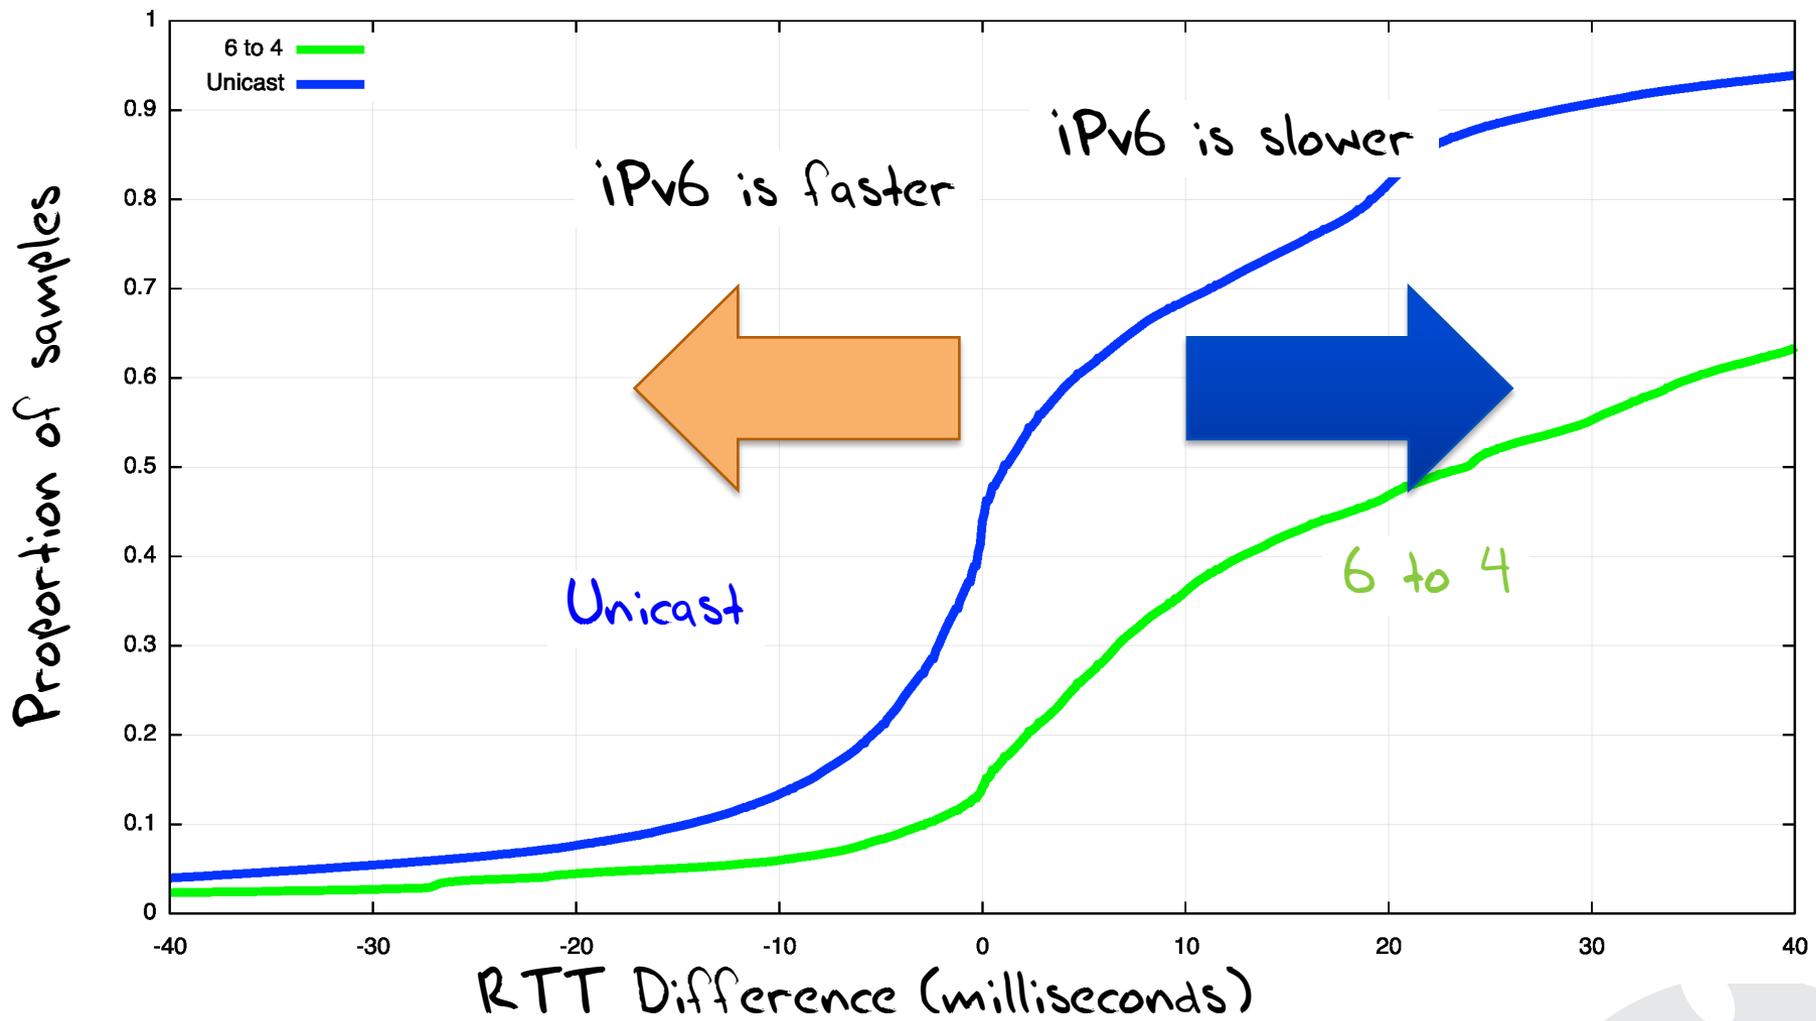
# December 2015/January 2016



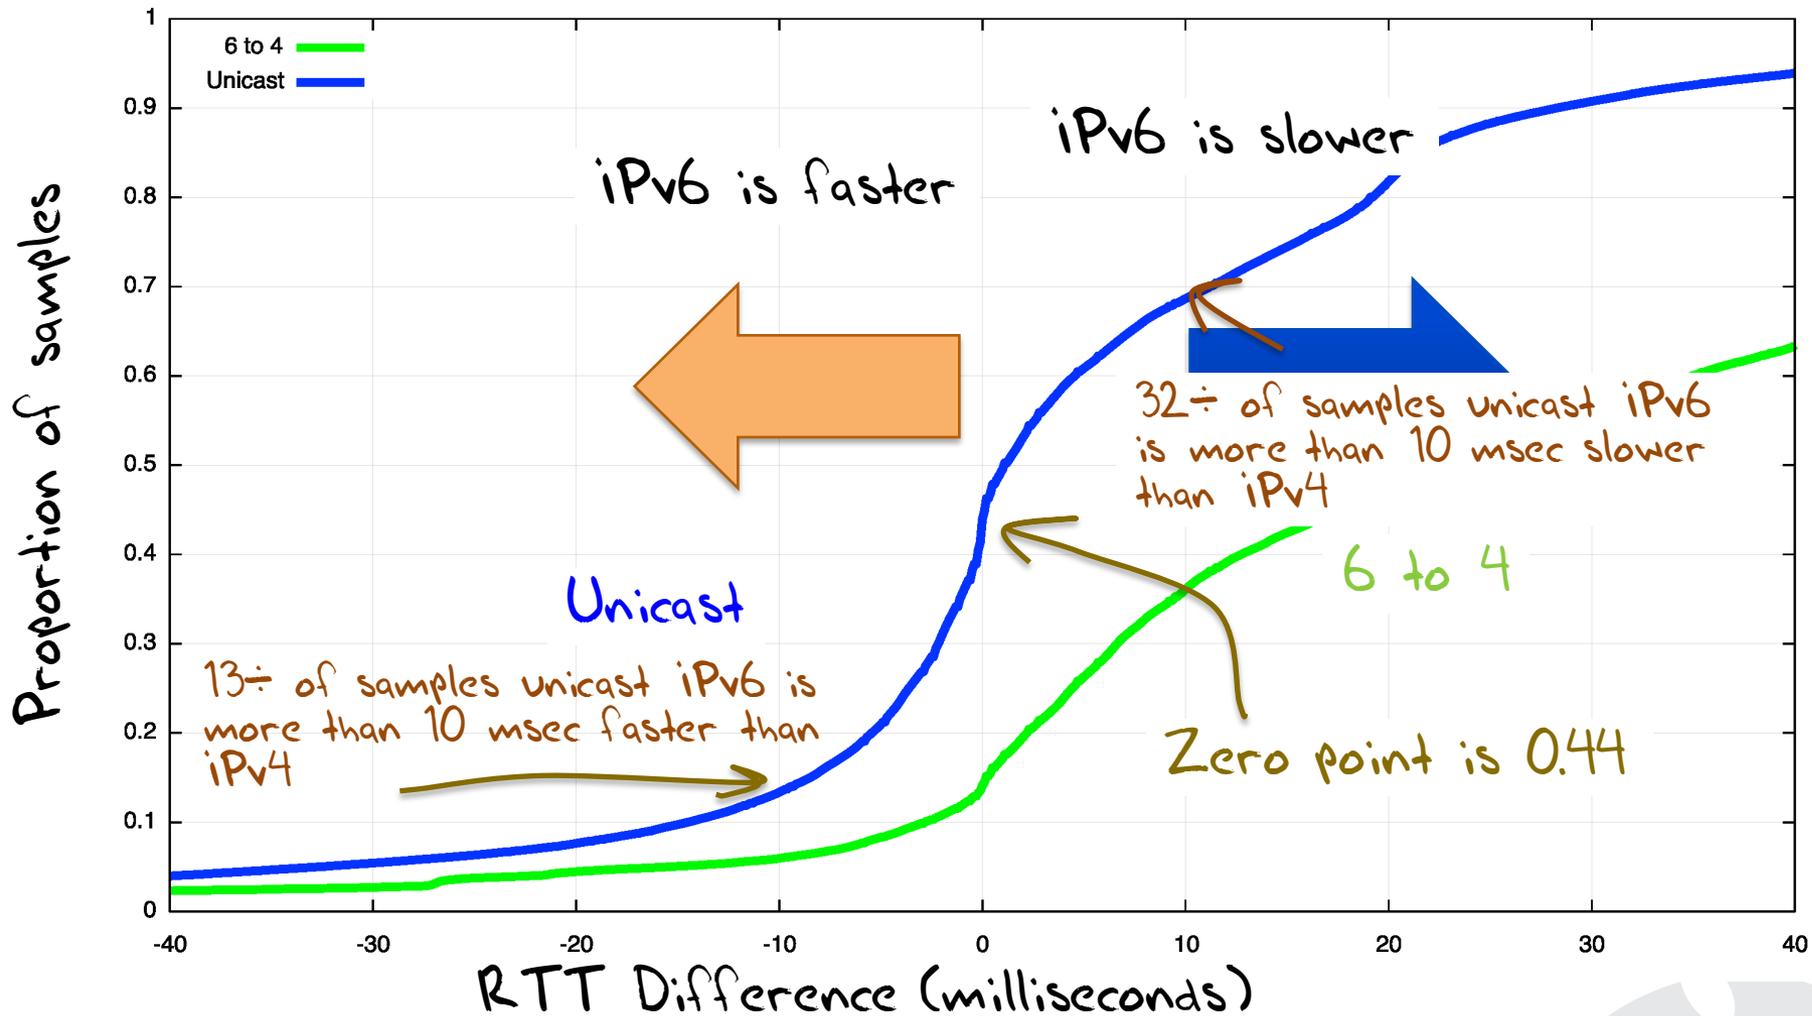
# 2015/6 RTT Data CDF



# 2015/6 RTT Data CDF



# 2015/6 RTT Data CDF



Is IPv6 as "good" as IPv4?



# Is IPv6 as "good" as IPv4?

Is IPv6 as **fast** as IPv4?

Basically, yes

IPv6 is faster about half of the time

For 65% of unicast cases, IPv6 is within 10ms RTT of IPv4

So they perform at much the same rate

But that's just for unicast IPv6

The use of 6to4 makes this a whole lot worse!



# Is IPv6 as "good" as IPv4?

## Is IPv6 as **robust** as IPv4?

IPv4 connection reliability currently sits at 0.2%

The base failure rate of Unicast V6 connection attempts at 1.8% of the total V6 unicast connections is not brilliant.

6to4 is still terrible!

It could be better.

It could be a whole lot better!



# Is IPv6 as "good" as IPv4?

If you can establish a connection, then IPv4 and IPv6 appear to have comparable RTT measurements across most of the Internet

But the odds of establishing that connection are still weighted in favour of IPv4!



That's it!

Questions?

