Temporal & Spatial Classification of Active IPv6 Addresses

v6ops Working Group, IETF 94
November 2, 2015

David Plonka and Arthur Berger {plonka,arthur}@akamai.com

Premise: IPv6 is not merely a bigger IPv4

To understand the structure of the ***active*** *IPv6 Internet, we must observe how addresses are assigned and used.*

Current levels of native IPv6 activity raise urgent operational issues:

- Selecting **targets** for active measurements
- Measuring IPv6 in operation by **census** or survey
- Informing **logging** re: data retention and resource exhaustion
- Informing **host reputation or *prefix intelligence***,
  e.g., DOS mitigation

Challenges

- Sheer size of IPv6's **sparse address space**
  "*Just tell me where the action is.*"
- IPv6 addresses and prefixes **differ in *kind*** from those of IPv4.
- **Billions** of active IPv6 WWW client addresses per week

Counting IPv6? What do the counts mean?

"So far, only about 3% to 4% of Internet users have IPv6, Cerf said." (May 4, 2015)

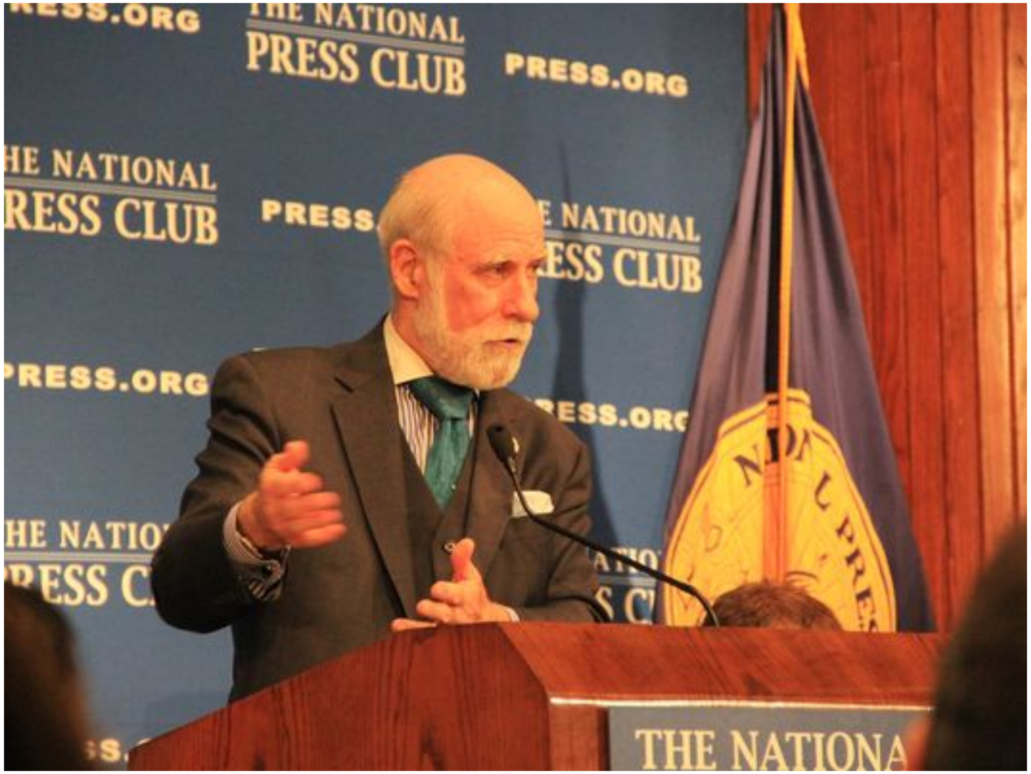

Thus, ~87.6–117 million v6 users...

based on ~2.92 billion users worldwide (estimate for July 1, 2014)

Sources: Internet Live Stats (elaboration of data ITU and United Nations Population Division)
Quote and Photo: Holly LaFon, Medill News Service

Counting IPv6? What do the counts mean?

"So far, only about 3% to 4% of Internet users have IPv6, Cerf said." (May 4, 2015)

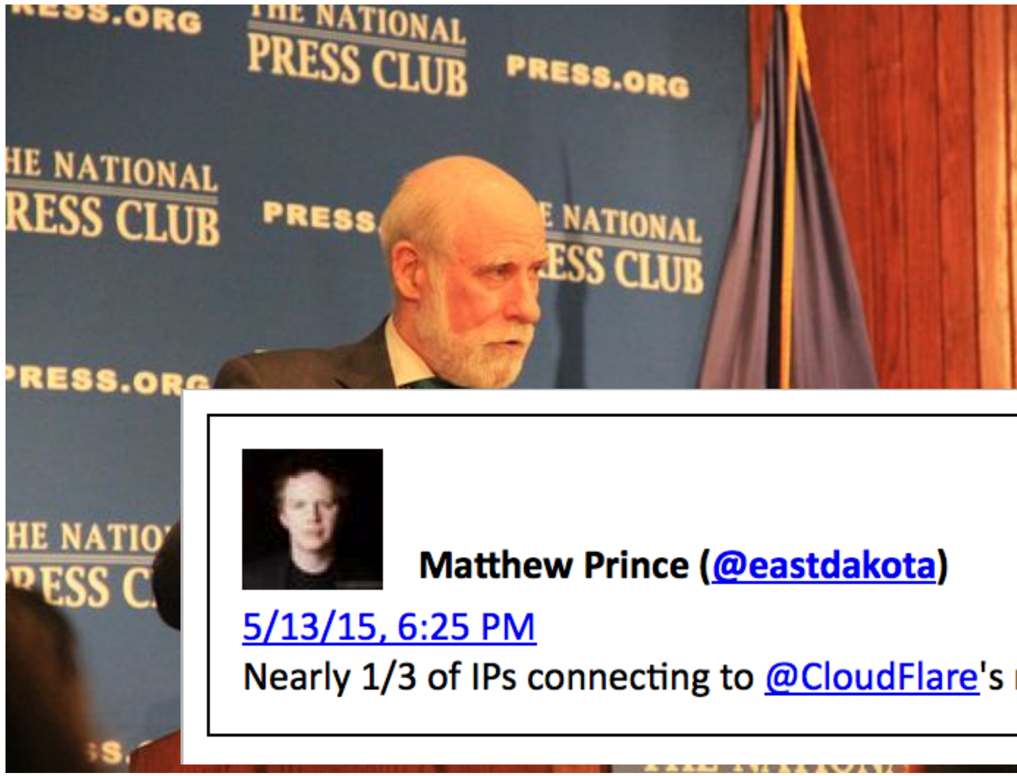

Matthew Prince (@eastdakota)

5/13/15, 6:25 PM

Nearly 1/3 of IPs connecting to @CloudFlare's network are now IPv6.

Thus, ~87.6–117 million v6 users worldwide...

based on ~2.92 billion users worldwide (estimate for July 1, 2014)

Sources: Internet Live Stats (elaboration of data ITU and United Nations Population Division)
Quote and Photo: Holly LaFon, Medill News Service

Background: Sample IPv6 Addresses

IPv6 addresses in presentation format:

```
2001:db8:0:1cdf:21e:c2ff:fec0:11db
2001:db8:10:1::103
2001:db8:167:1109::10:901
2001:db8:4137:9e76:3031:f3fd:bbdd:2c2a
```

Sample IPv6 Addresses

IPv6 addresses in presentation format:

2001:db8:0:1cdf**:21e:c2ff:fec0:11db**

2001:db8:10:1:**:103**

2001:db8:167:1109:**:10:901**

2001:db8:4137:9e76**:3031:f3fd:bbdd:2c2a**

Sample IPv6 Addresses

IPv6 addresses in presentation format:

2001:db8:0:1cdf:**21e:c2ff:fec0:11db**
2001:db8:10:1::**103**
2001:db8:167:1109::**10:901**
2001:db8:4137:9e76:**3031:f3fd:bbdd:2c2a**

Consider 16-bit (4 character) and 4-bit (1 character) segments:

2001:0db8:0000:1cdf:**021e:c2ff:fec0:11db**
2001:0db8:0010:0001:**0000:0000:0000:0103**
2001:0db8:0167:1109:**0000:0000:0010:0901**
2001:0db8:4137:9e76:**3031:f3fd:bbdd:2c2a**

Data sets

- **Active IPv6 WWW Client Addresses - passive discovery**
  - *Tens of billions* observed with successful hits on ~55 K CDN servers
    Full month’s addresses, semi-annually, Mar & Sep, 2014 & 2015
  - e.g., September 2015:
    8,844 BGP prefixes
    4,923 ASNs (47% of those advertising IPv6 prefixes)

    March 2014:
    5,531 BGP prefixes
    3,842 ASNs (40%)
- **IPv6 Router Addresses - active discovery**
  - February, 2015:
    3.2 million, discovered by TTL-limited probes, like `traceroute`

WWW Client Data Characteristics: Daily (Table 1a)

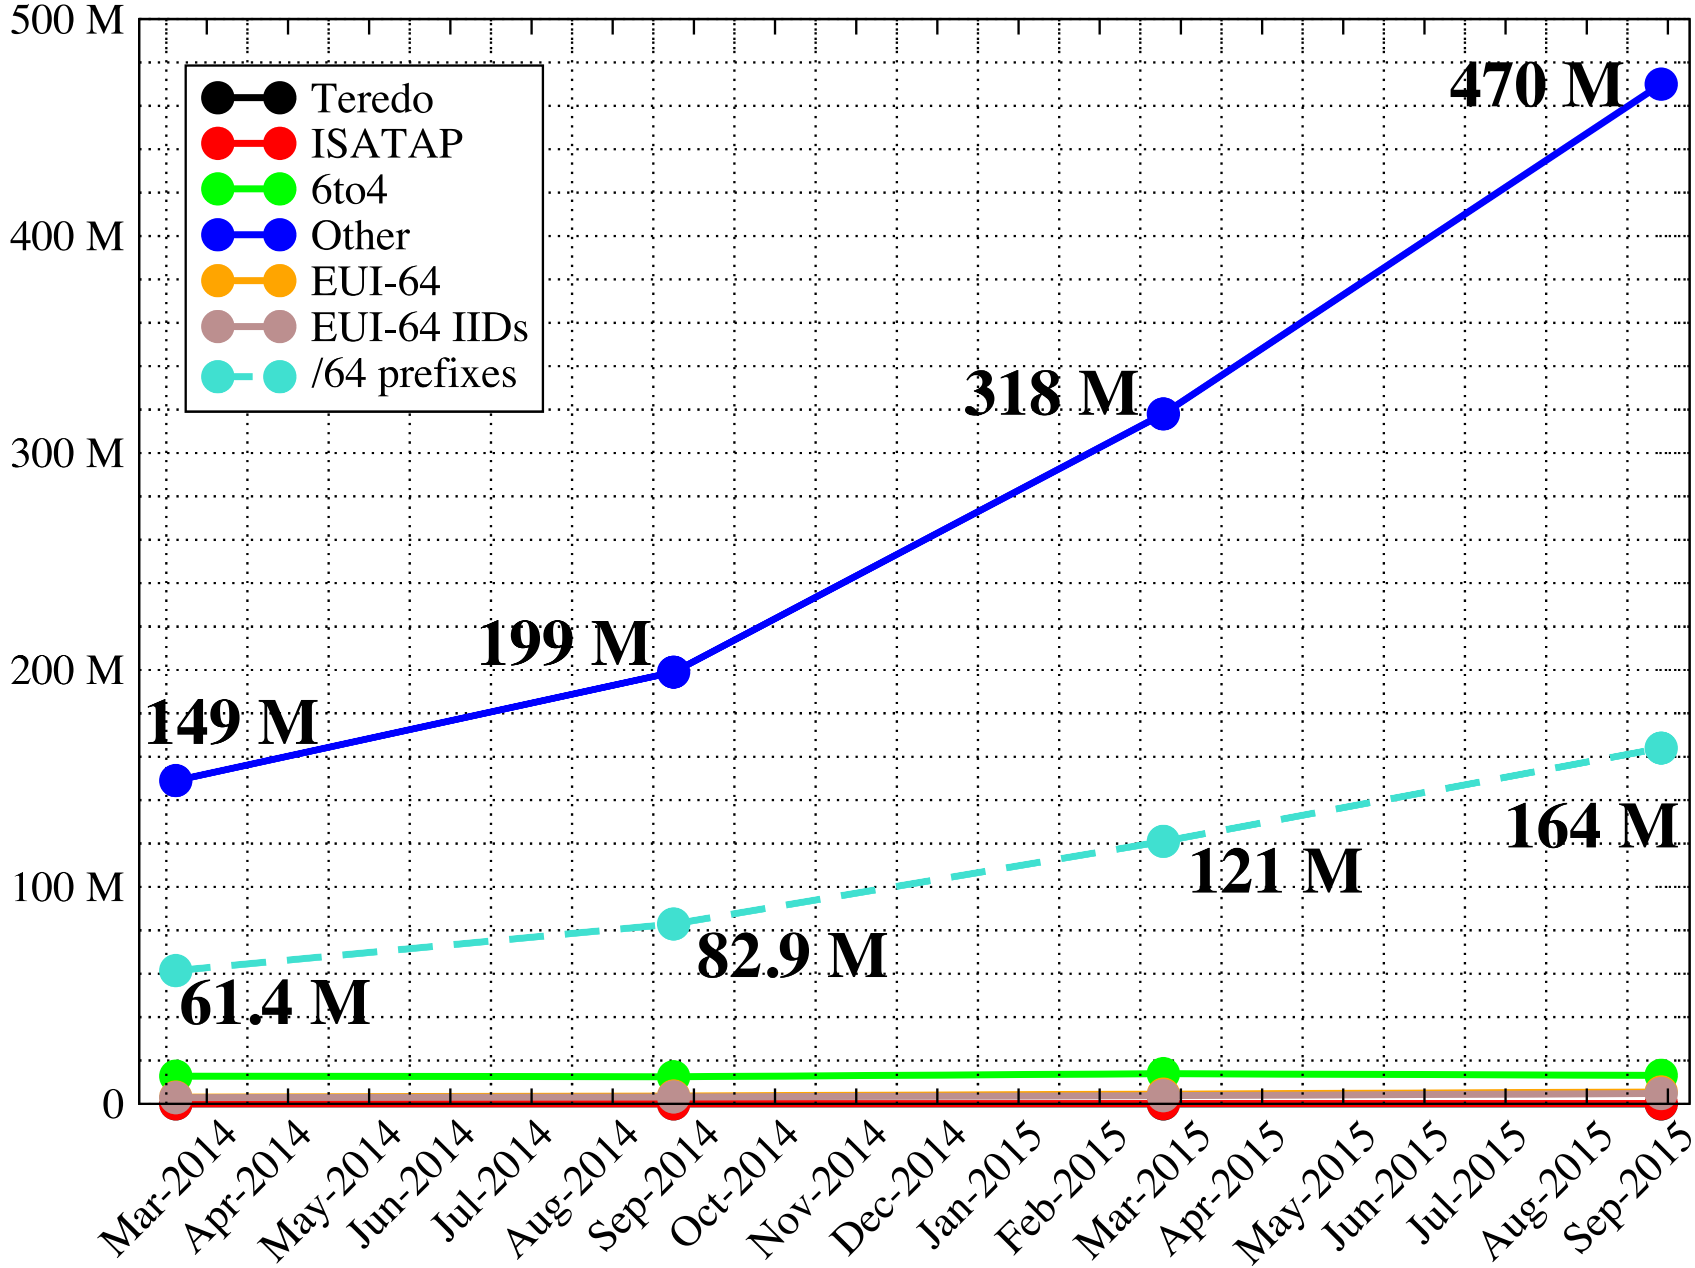

Visualizing the Active IPv4 Internet: Hilbert Curve Heatmap

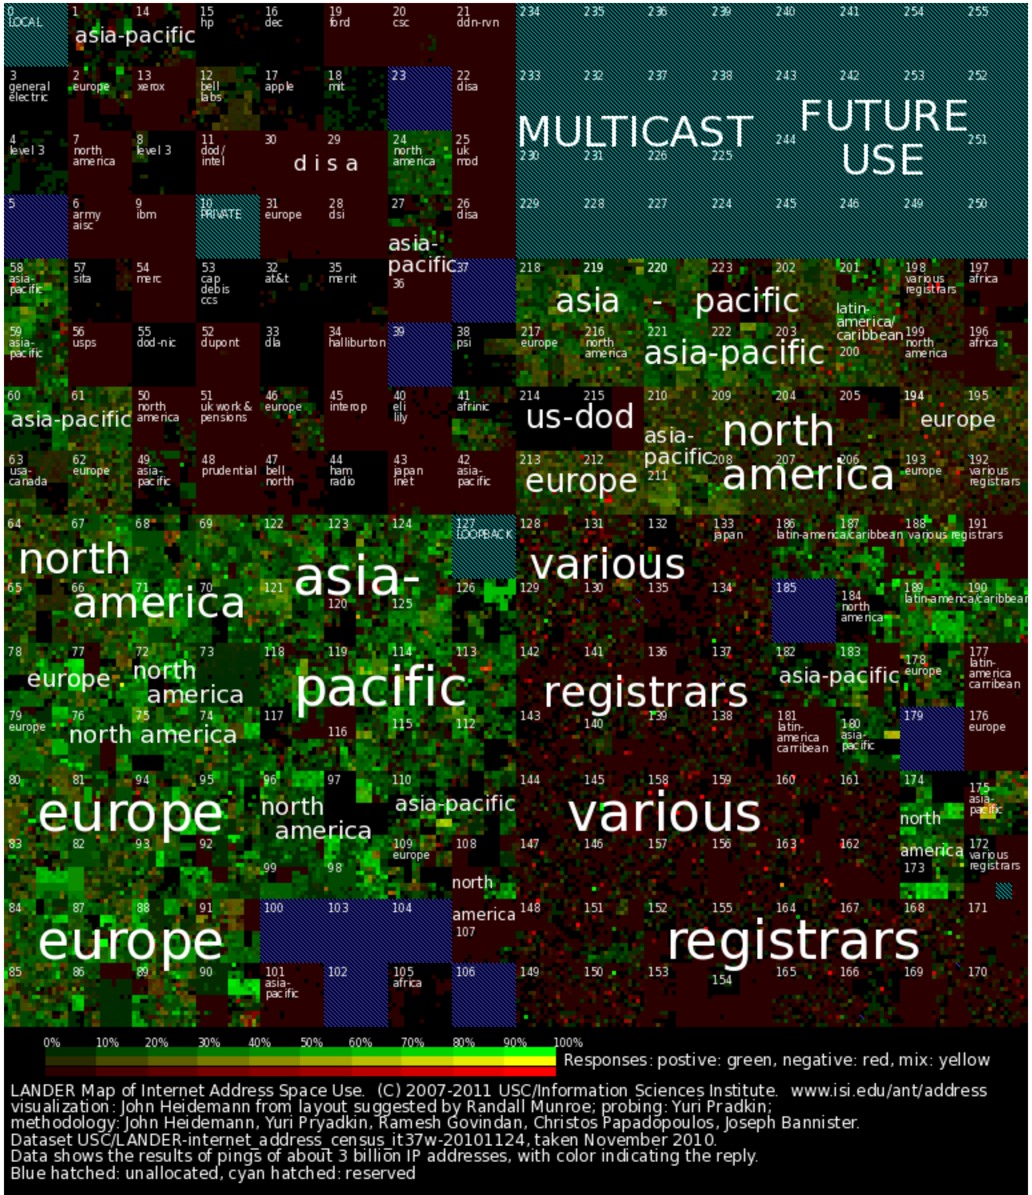

Image courtesy of John Heidemann, `https://ant.isi.edu/address/`

Visualizing the Active IPv6 Internet: Classified Treemap

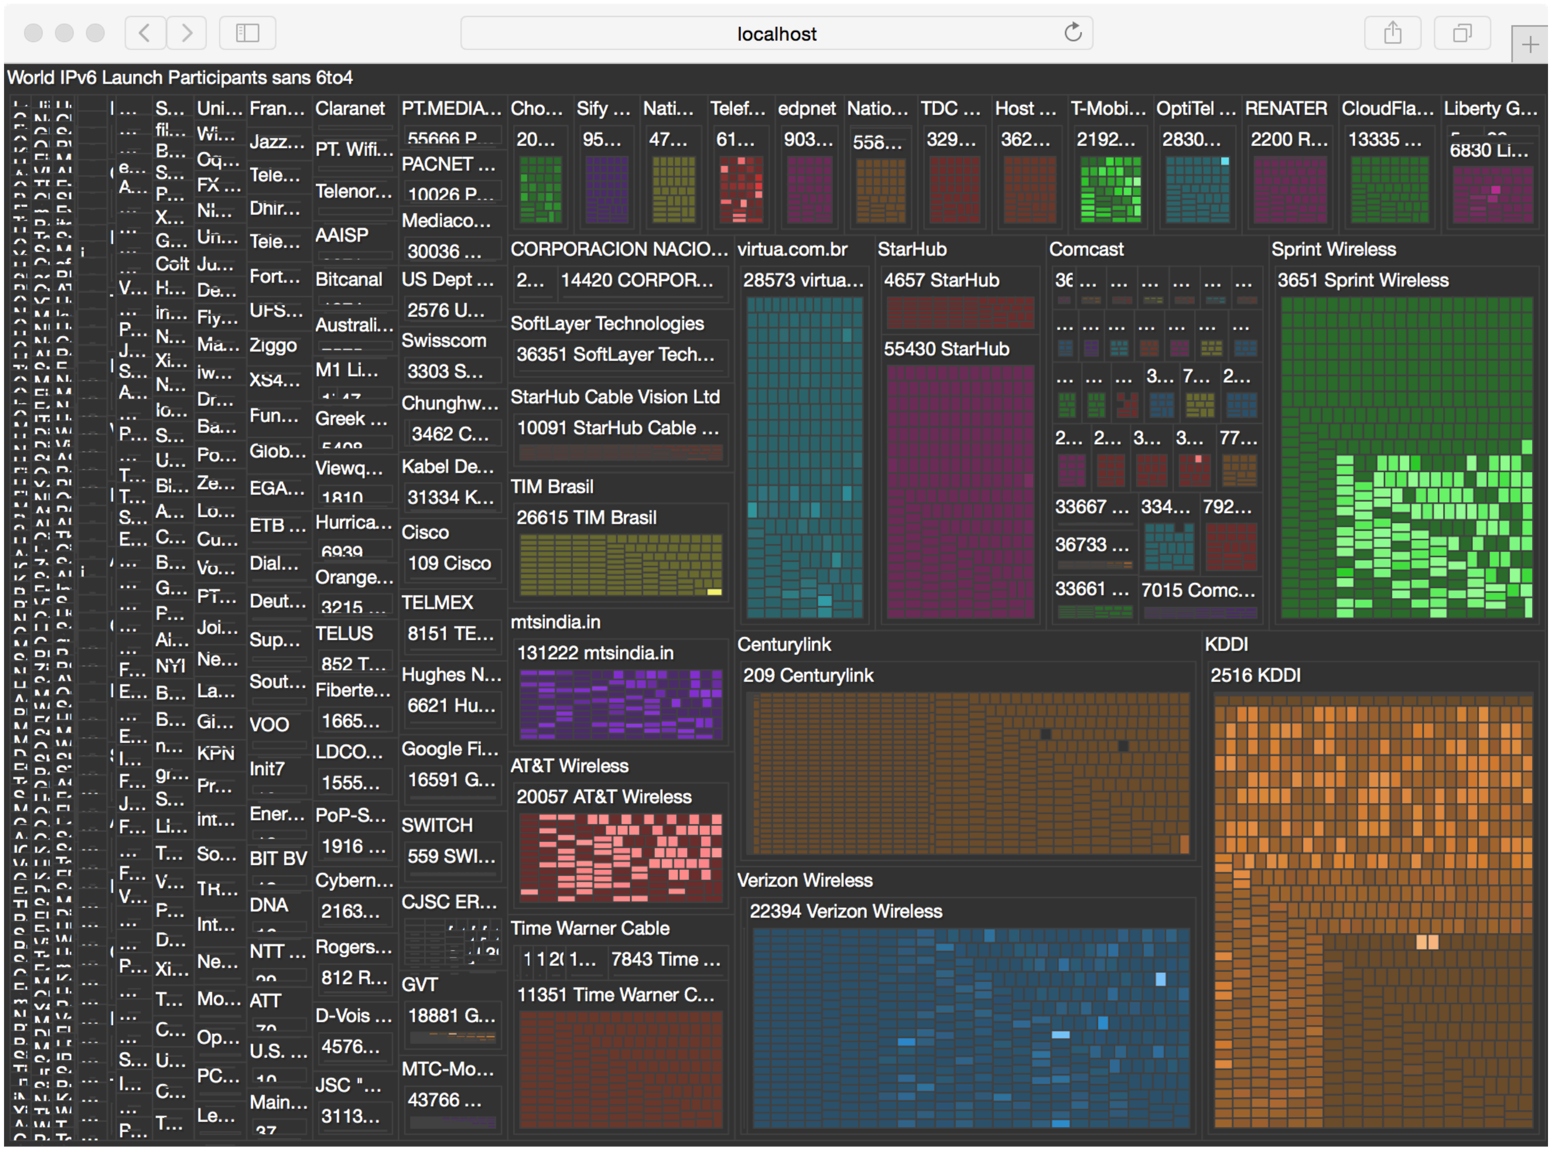

BGP prefix size, log2 scale

WWW Client Data Characteristics: Weekly (Table 1b)

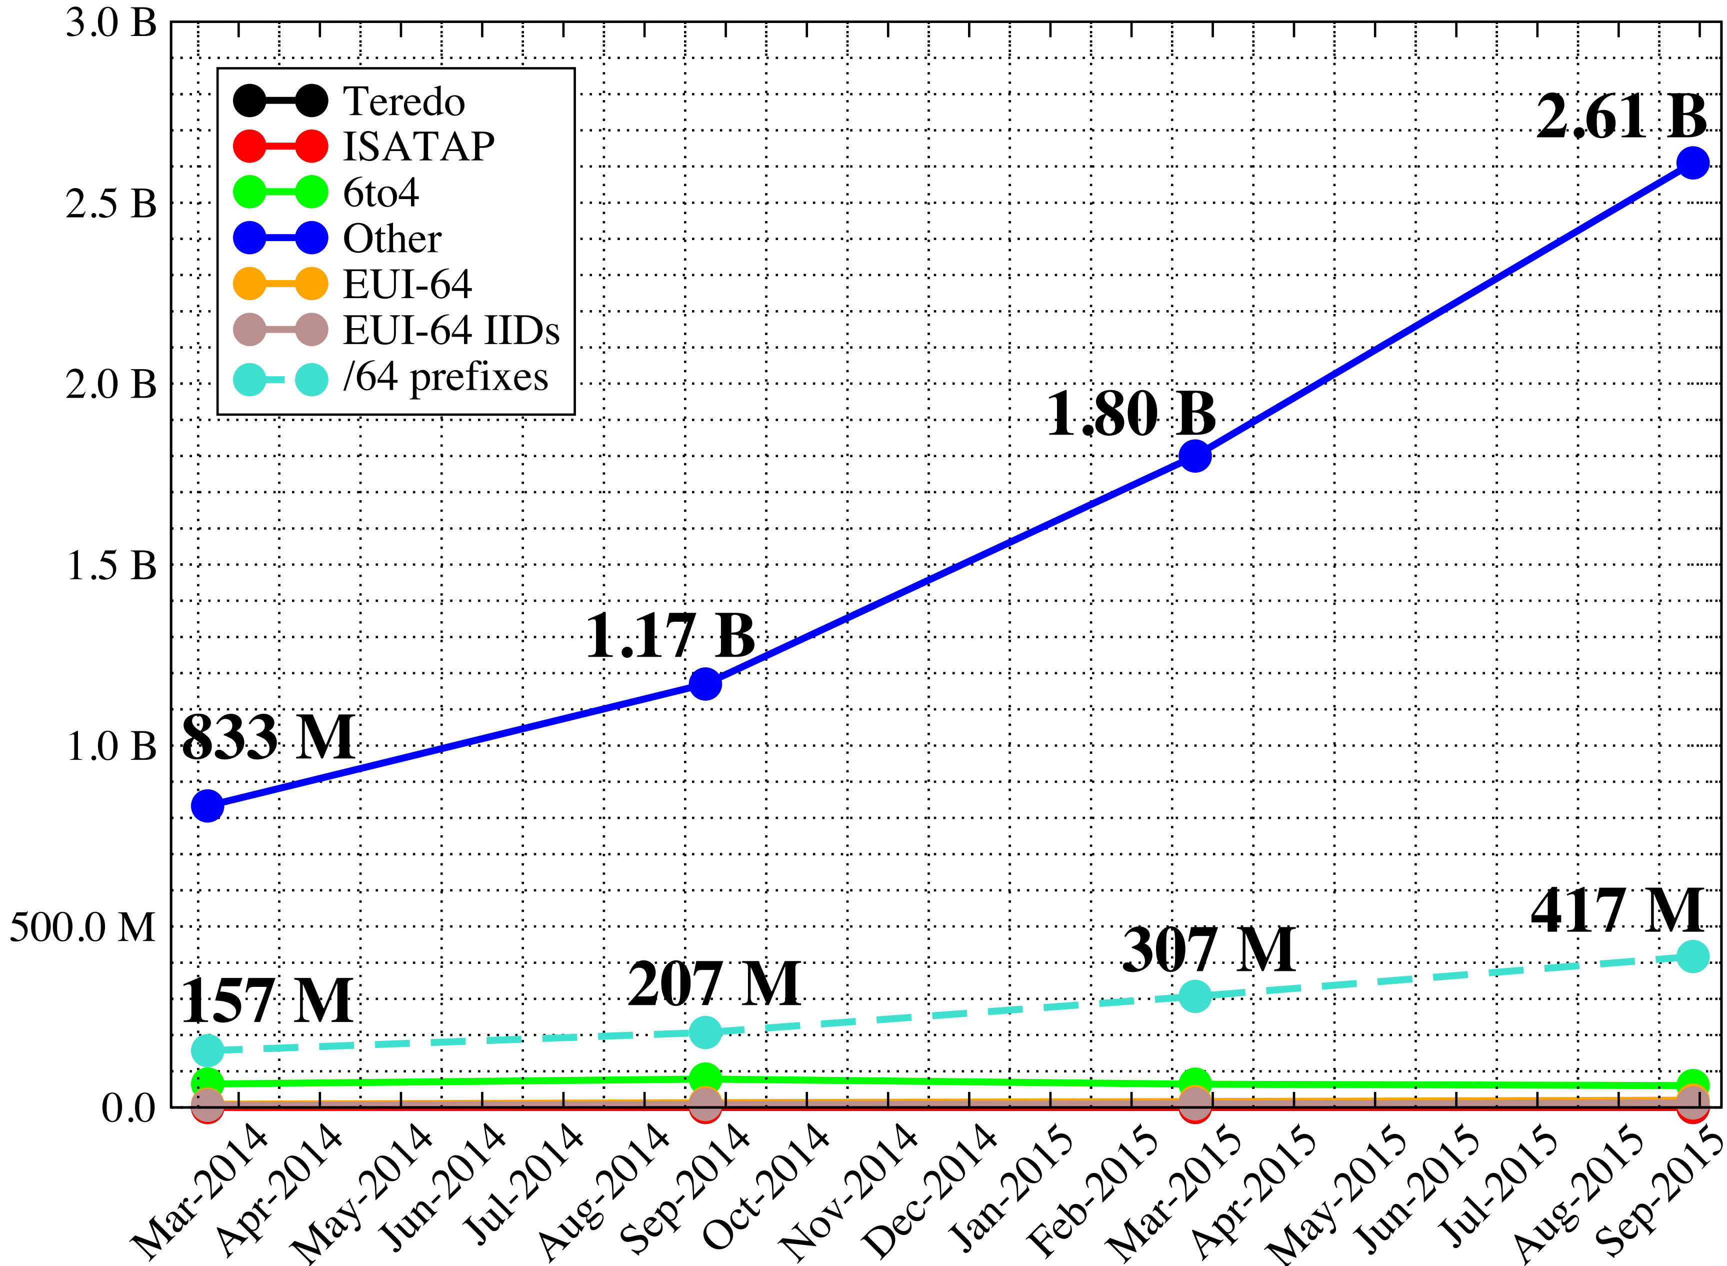

WWW Client Data Characteristics: Weekly (Table 1b)

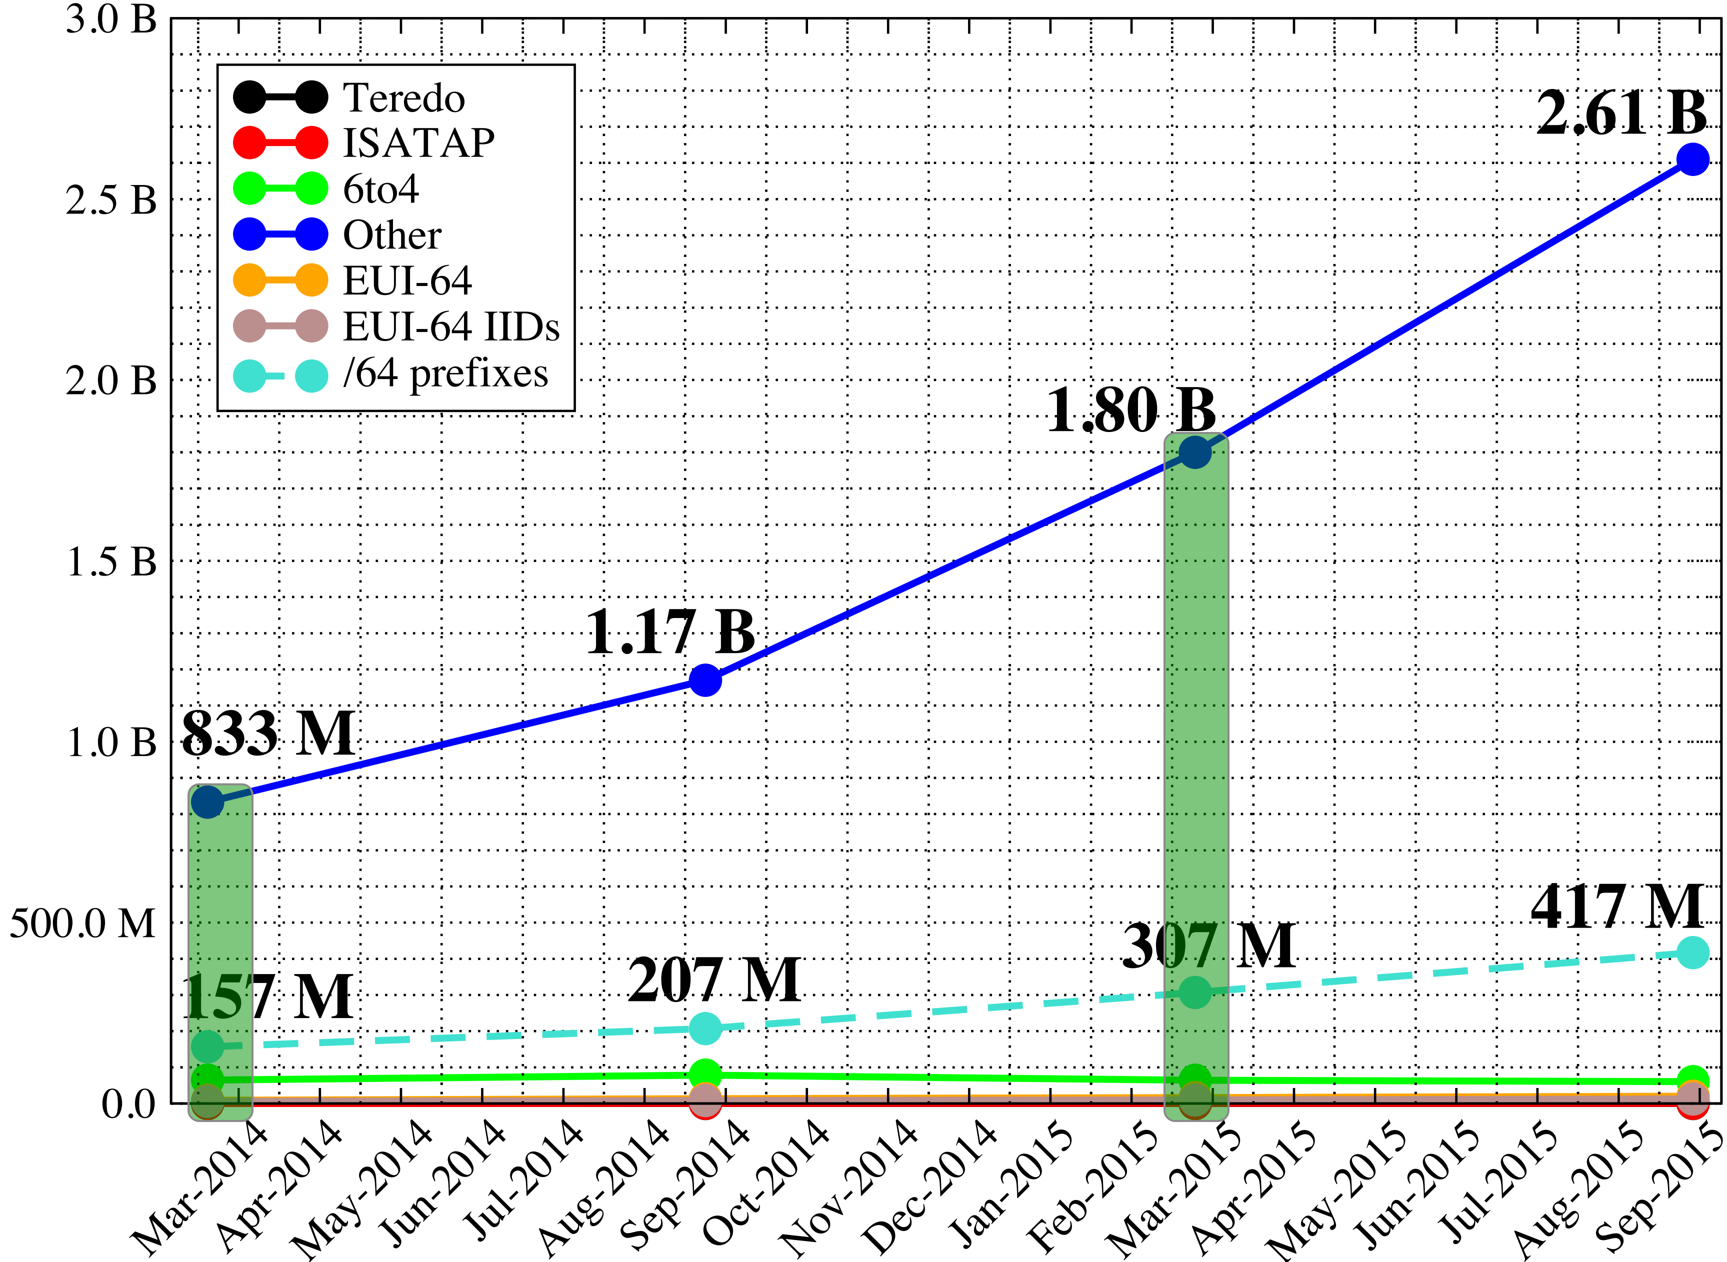

WWW Client Data Characteristics (log scale): Weekly (Table 1b)

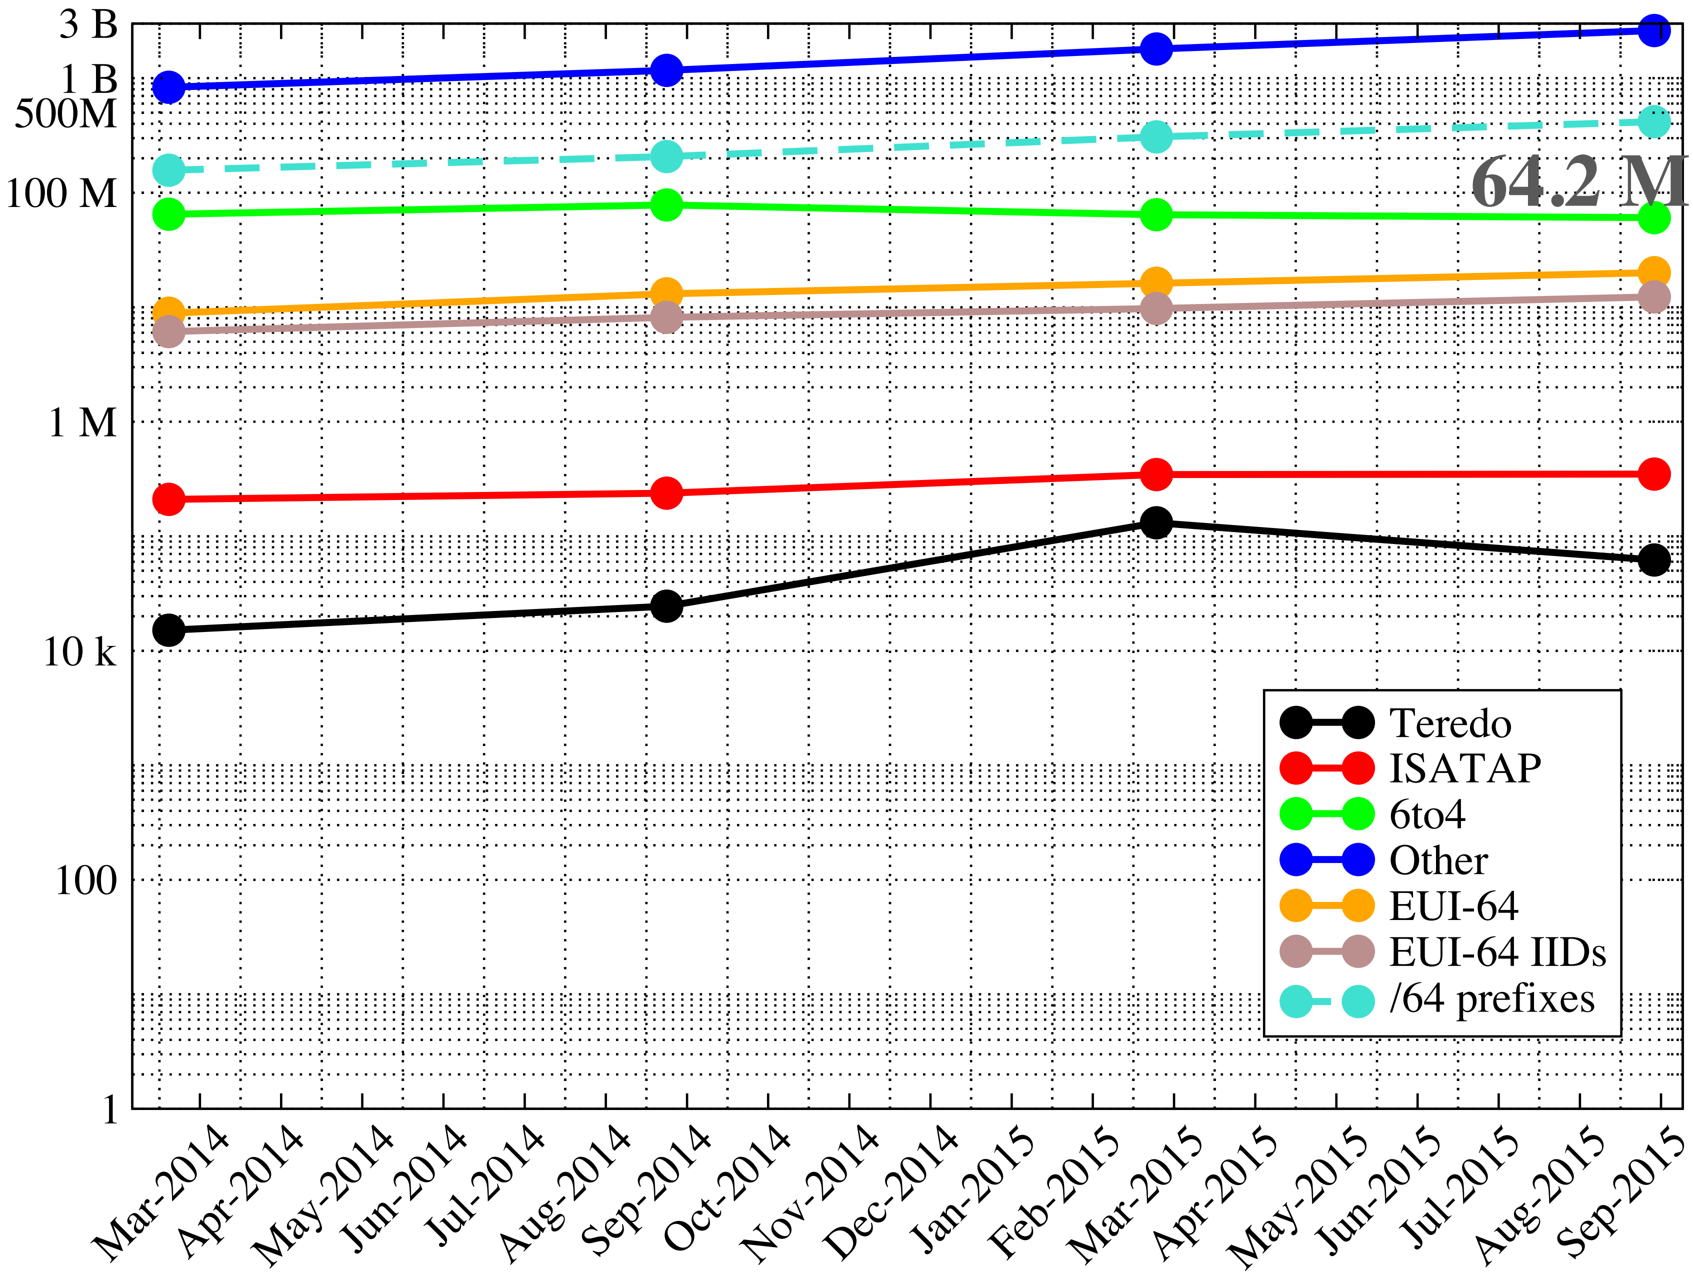

Temporal Address Classification: Stability Analysis

- Determine address stability or transience by set membership tests across *n* days, *e.g.*, “3d-stable” addresses:

  “Is this address known to have been active across a 3-day span?”

Address Stability: March 17-23, 2015 (1.80B v6 addresses)

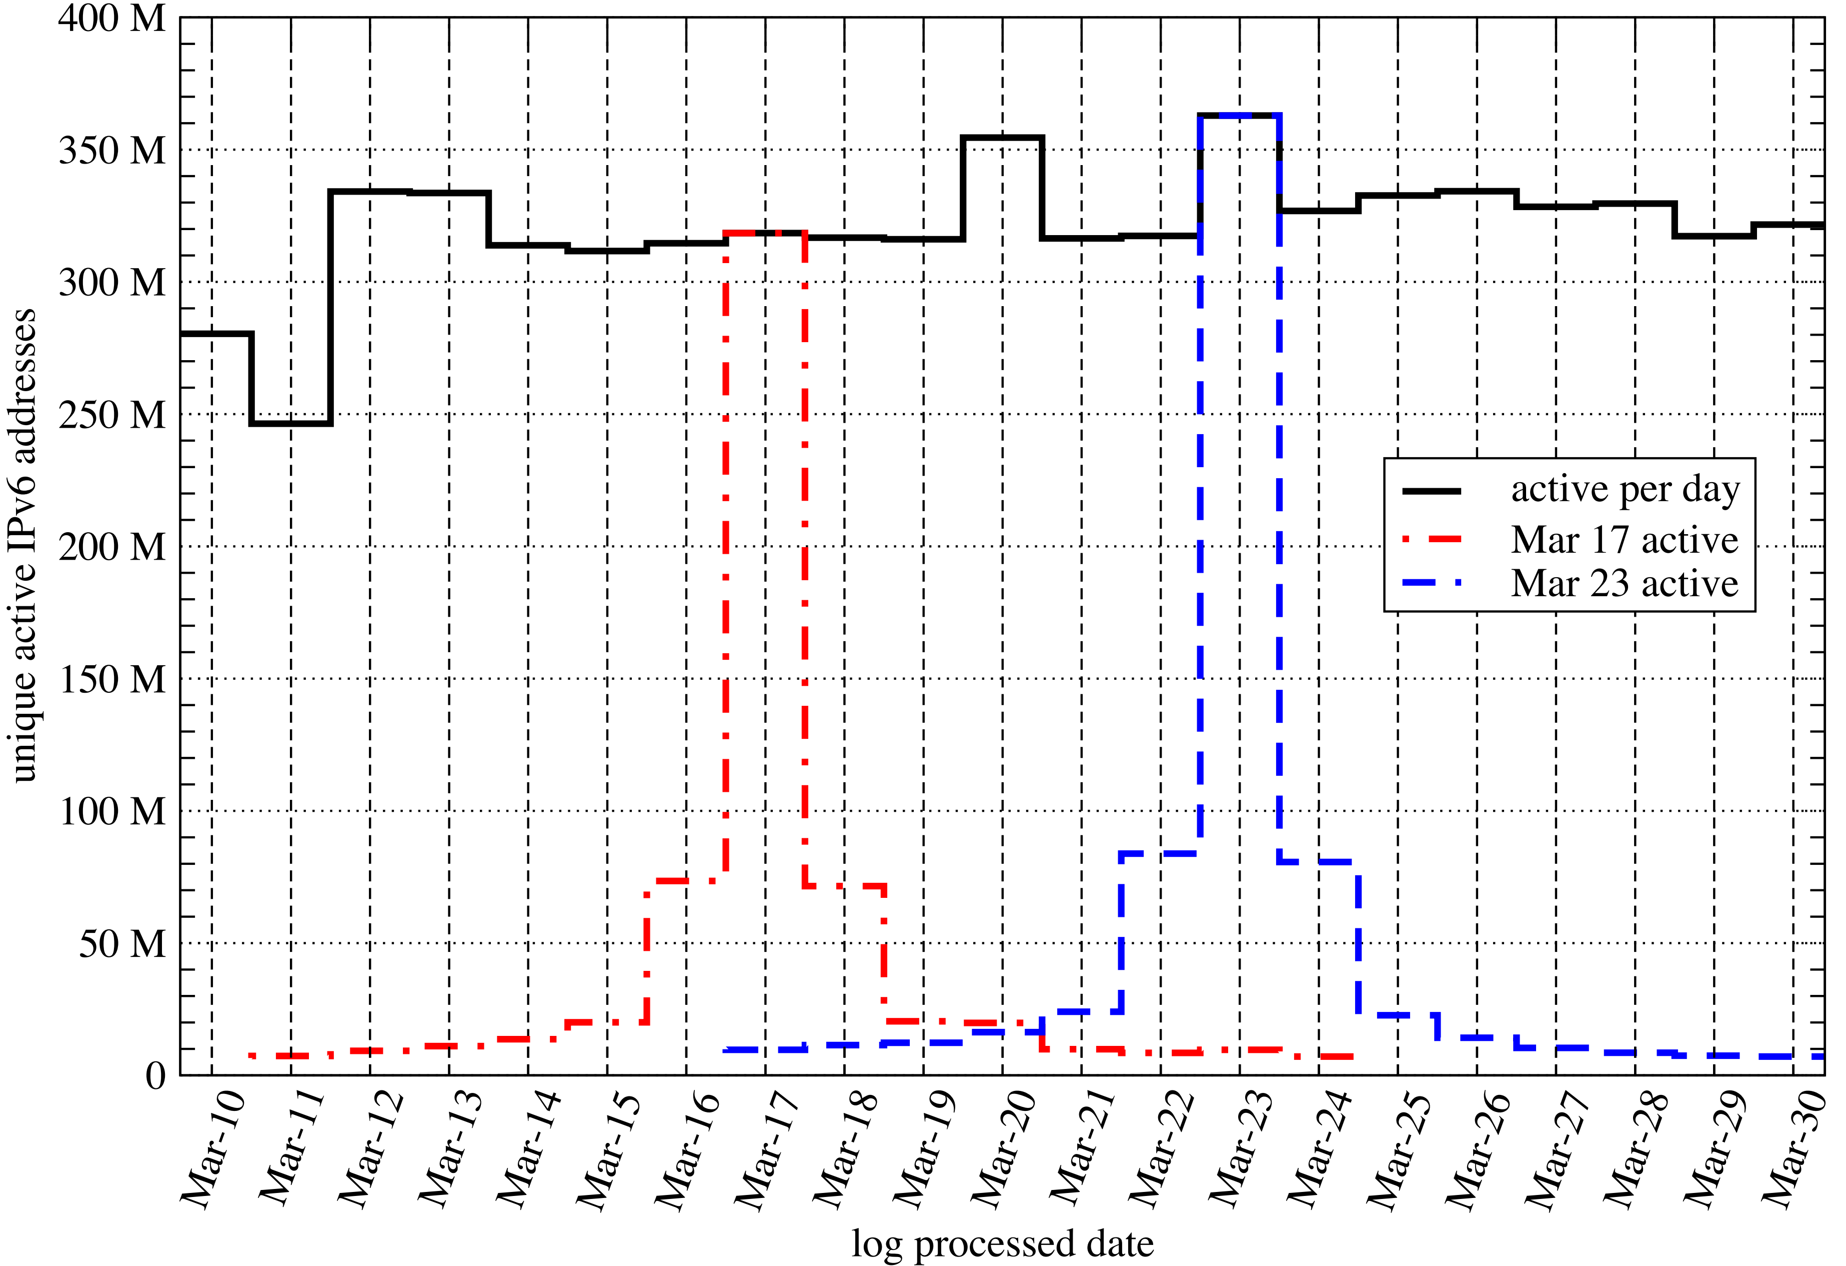

Address Stability: Daily (Table 2a)

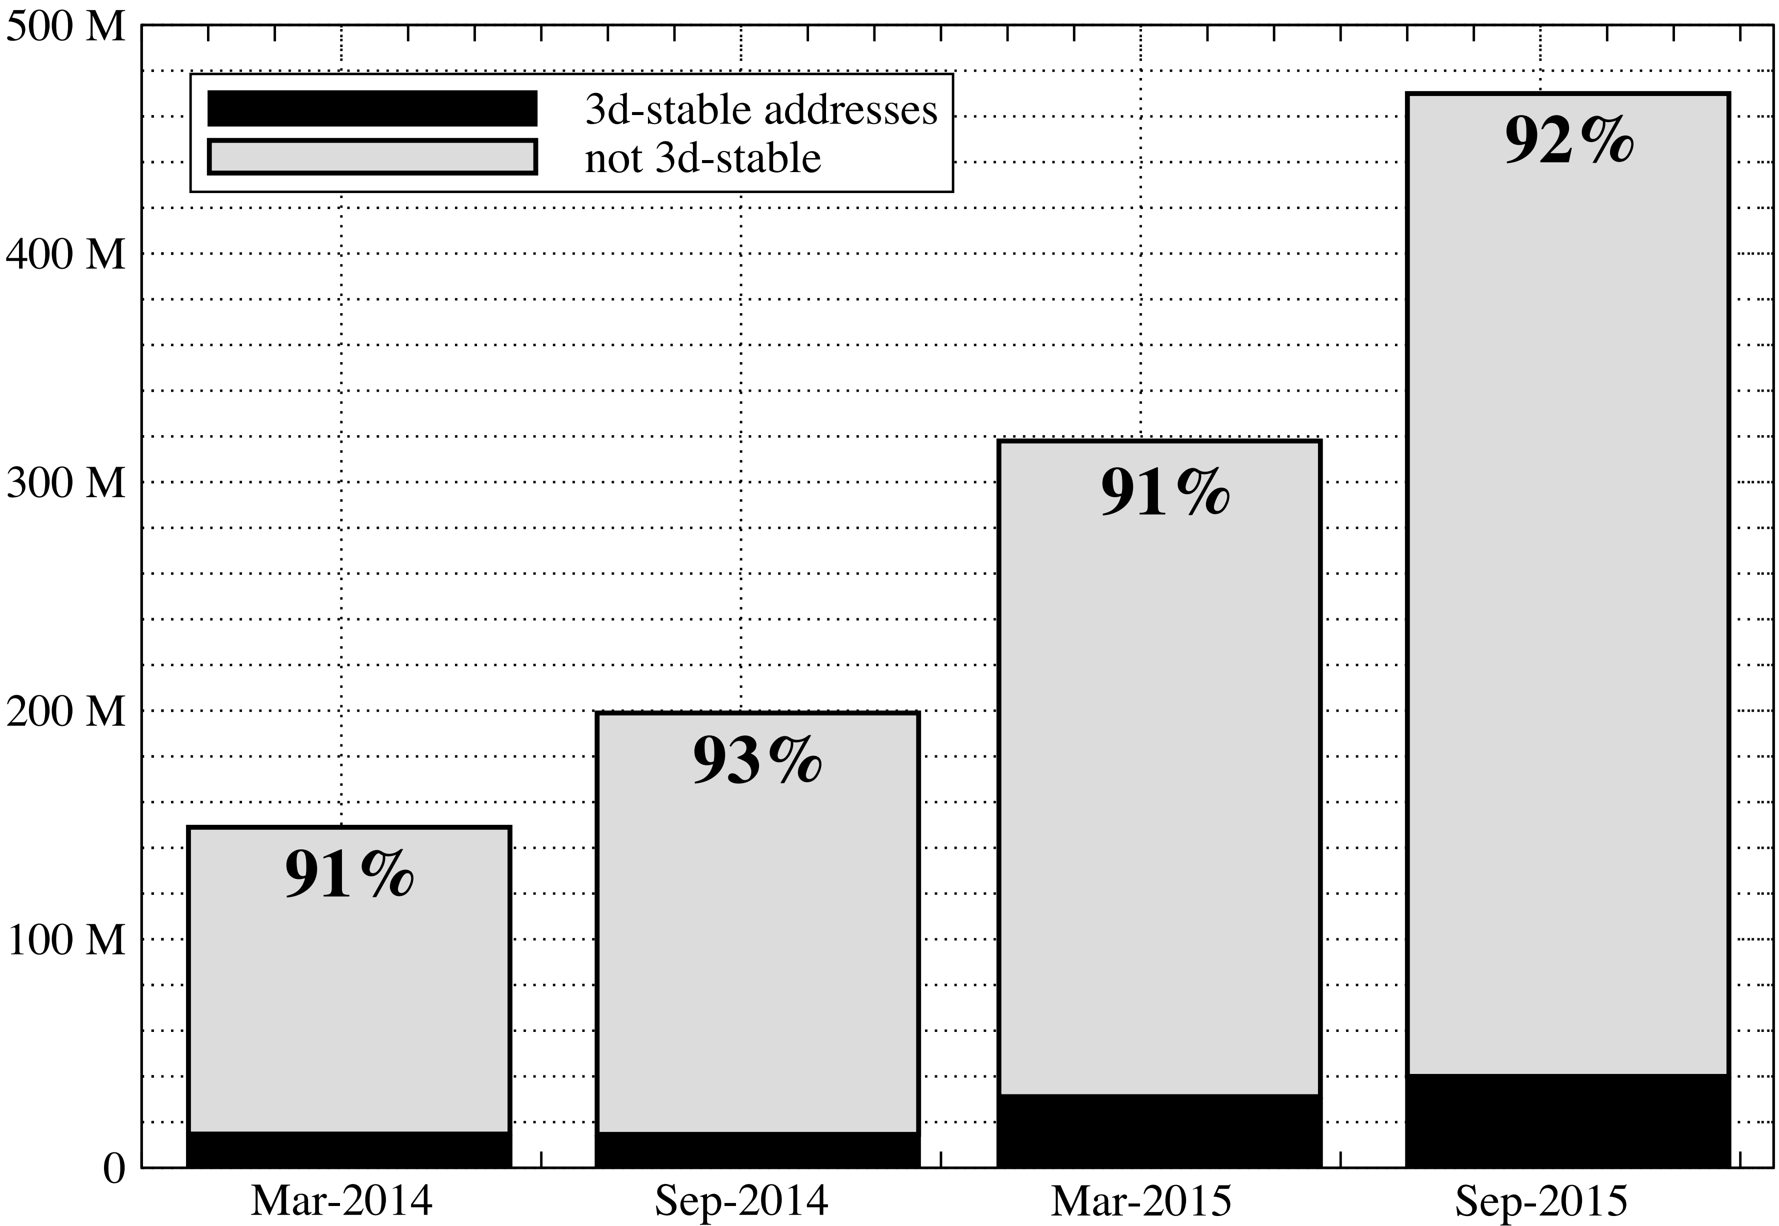

Address Stability: Weekly (Table 2c)

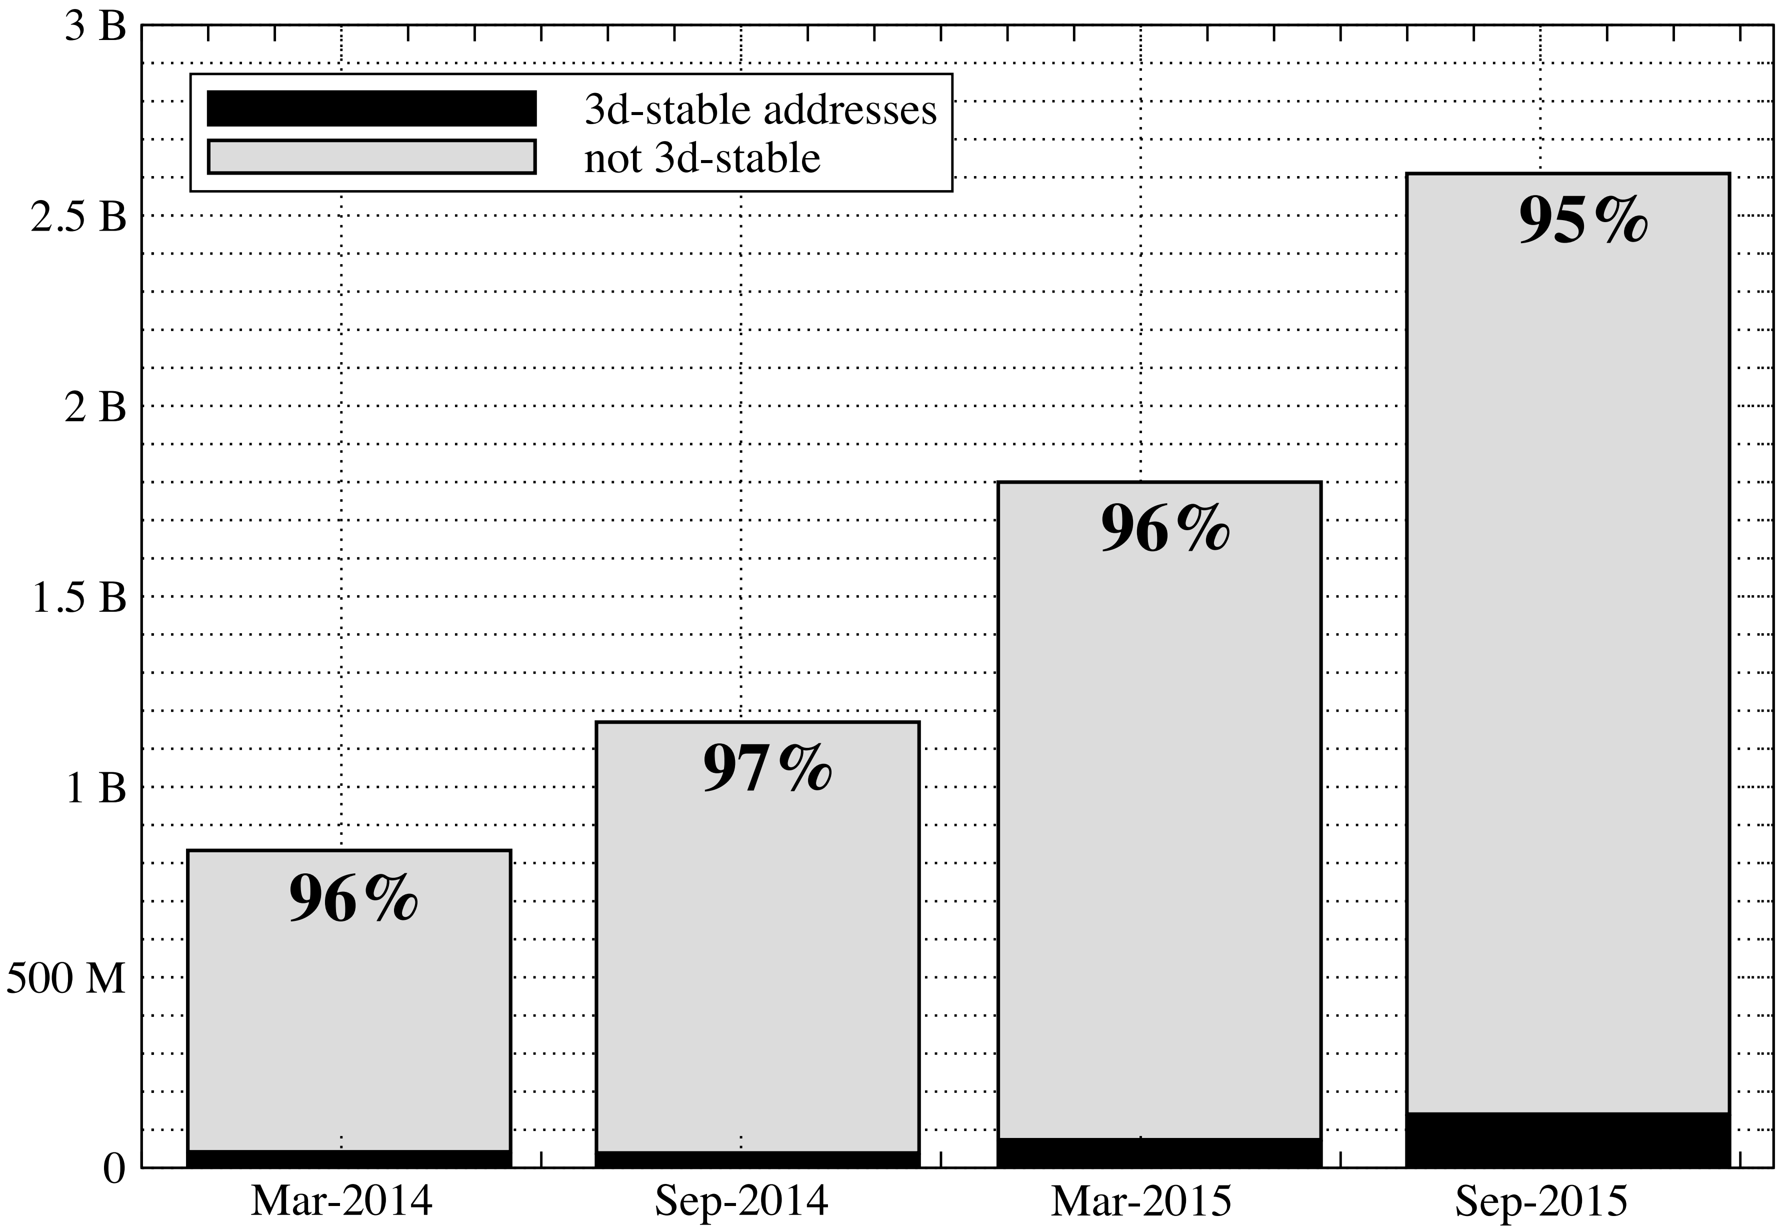

/64 Prefix Stability: March 17-23, 2015 (307M /64 prefixes)

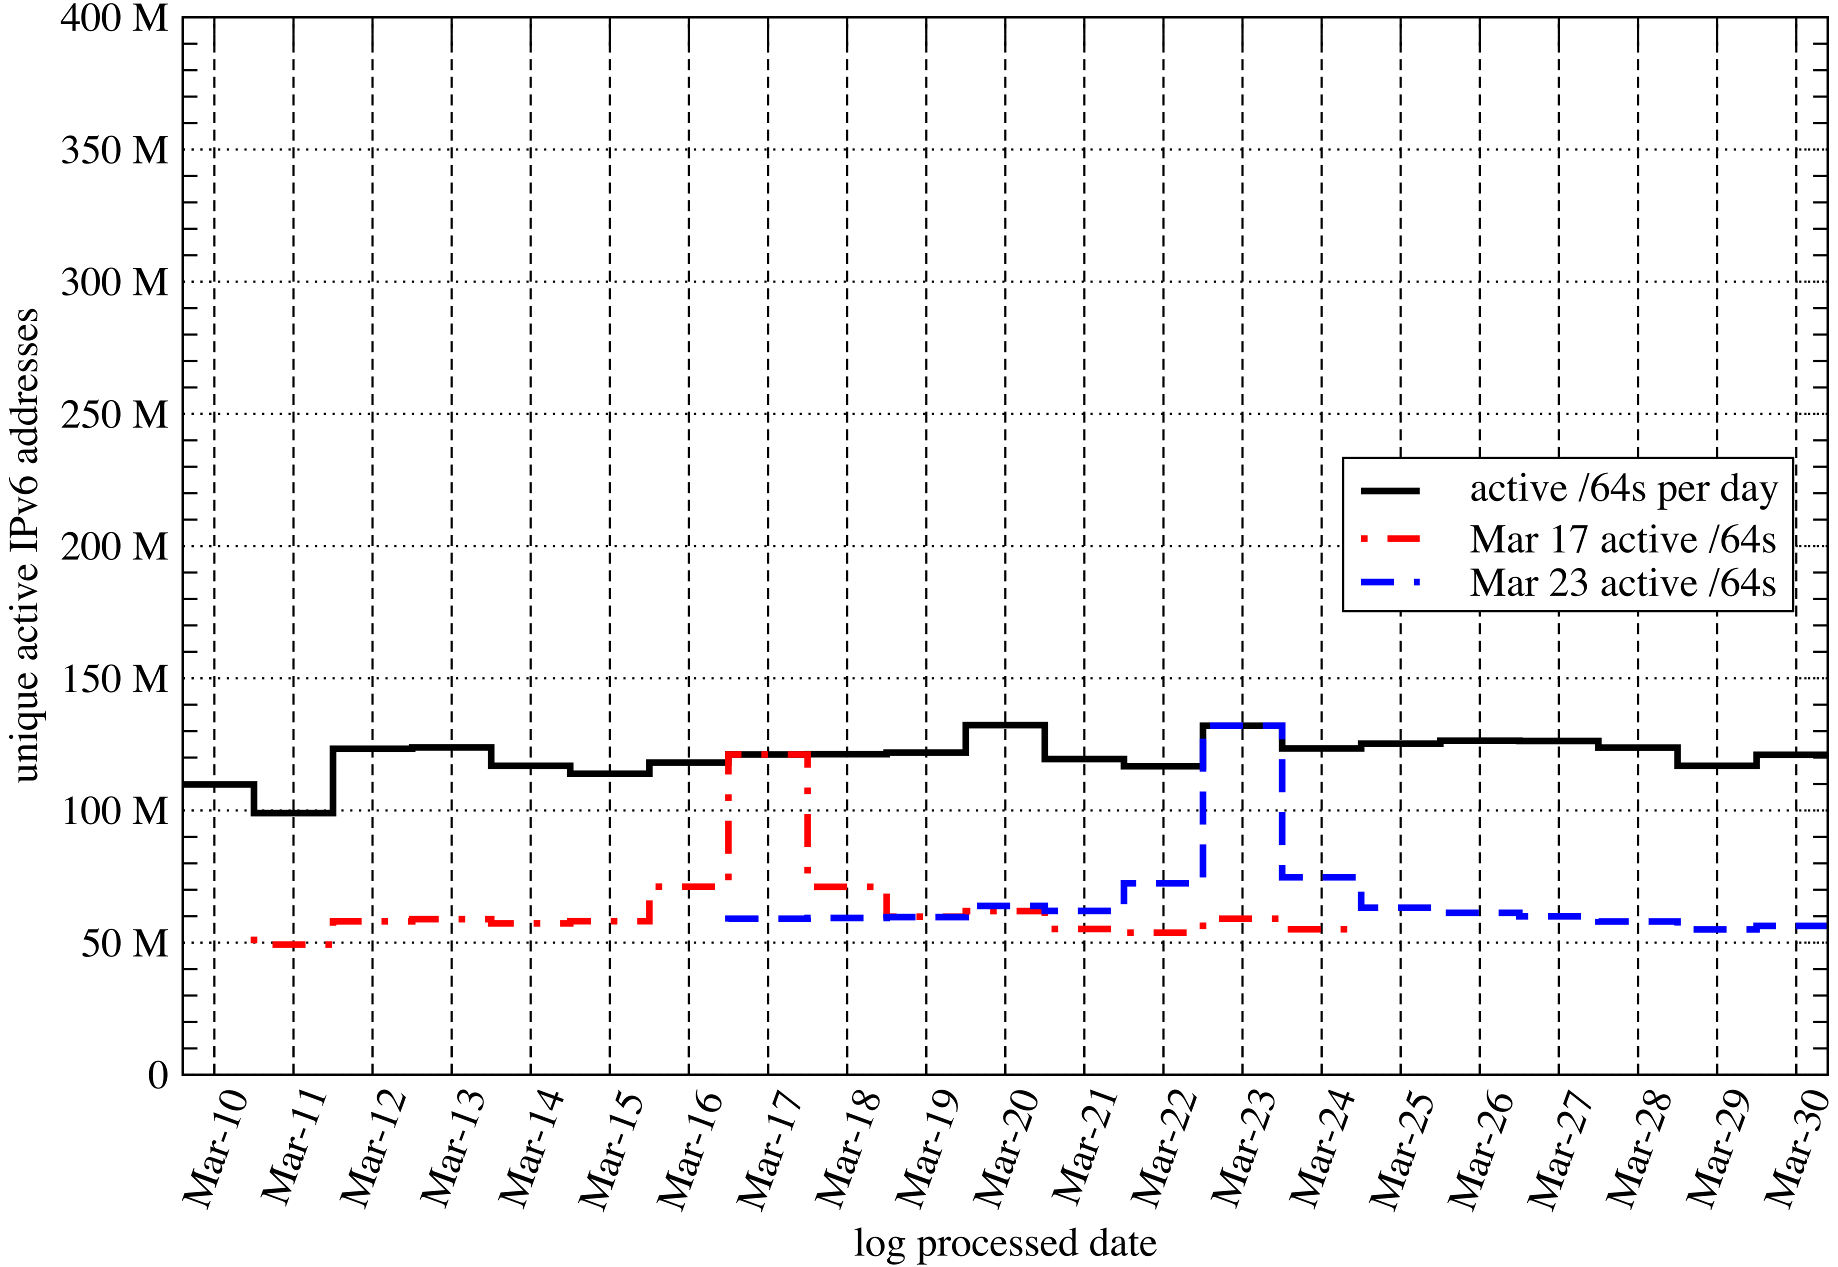

/64 Prefix Stability: Daily (Table 2b)

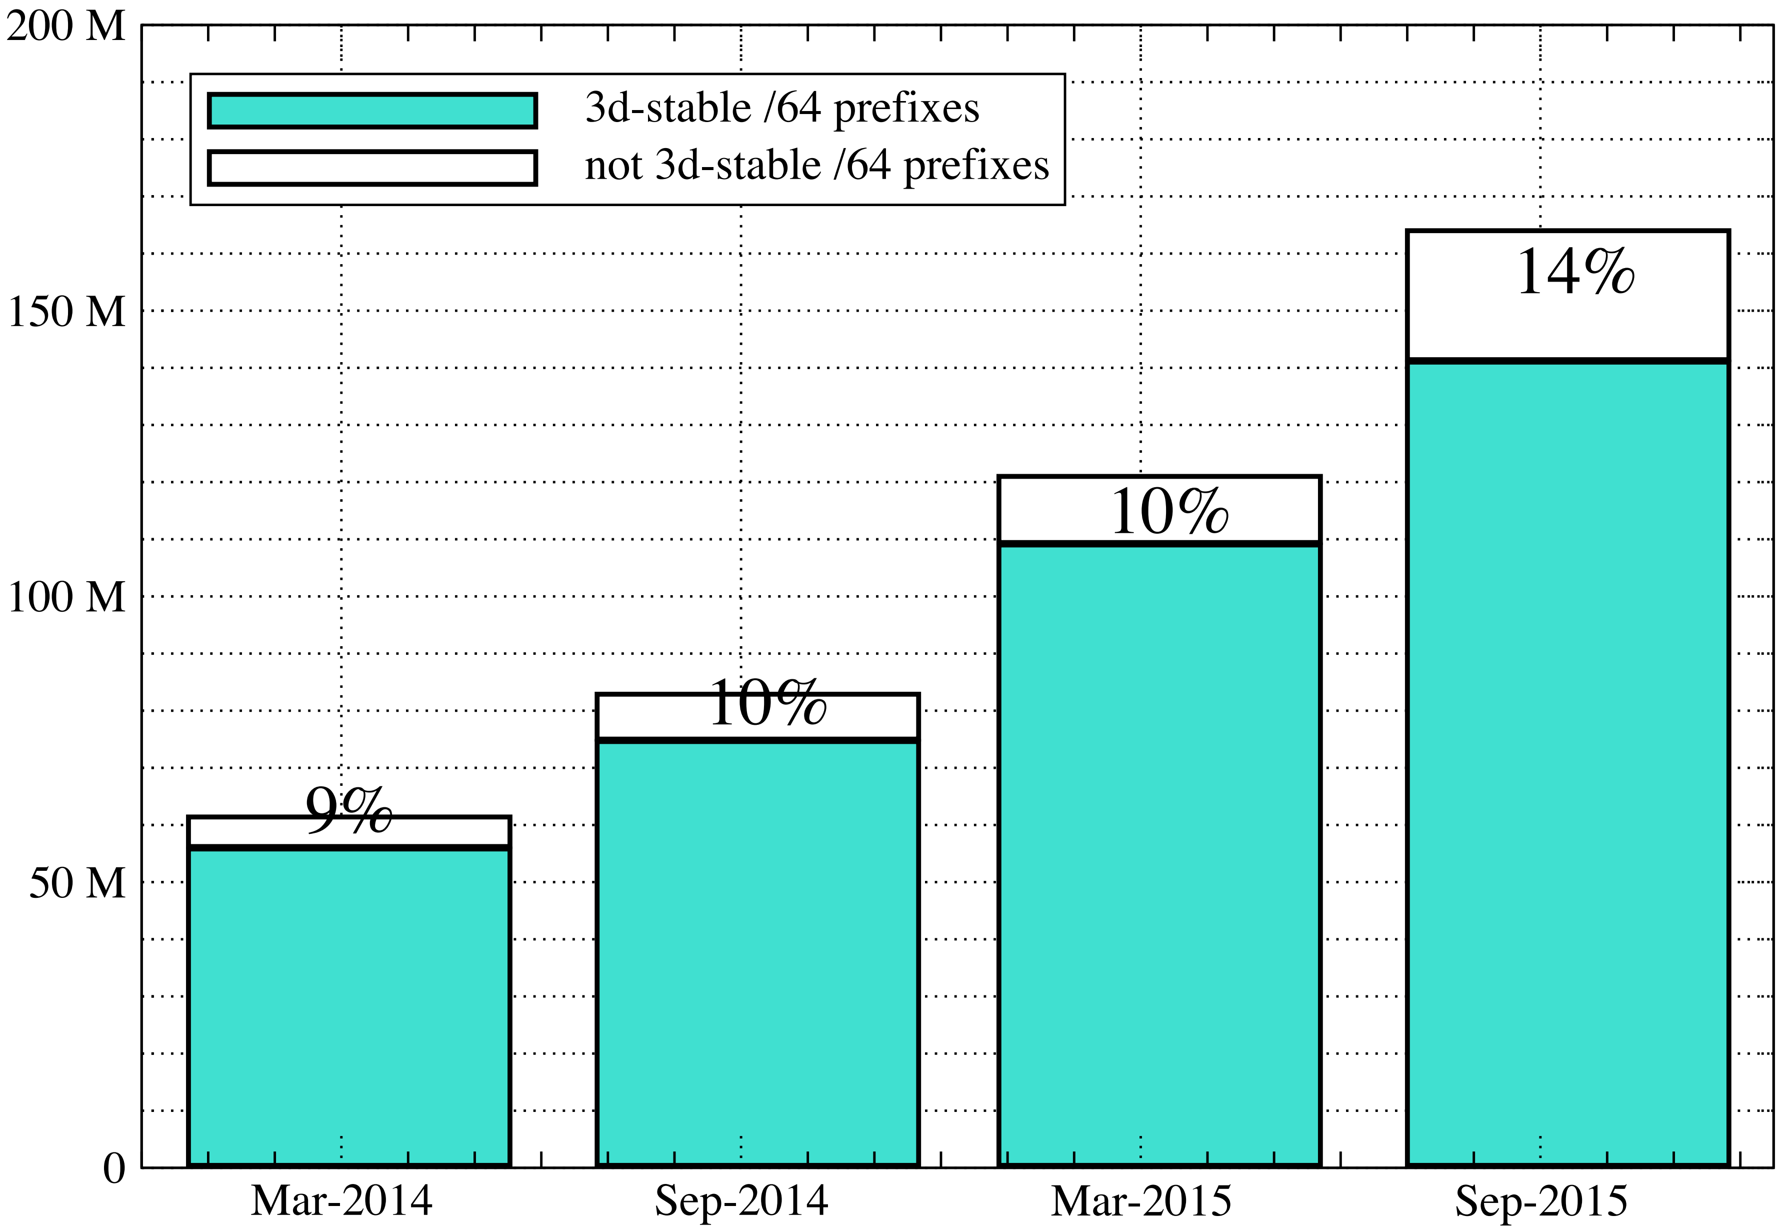

/64 Prefix Stability: Weekly (Table 2d)

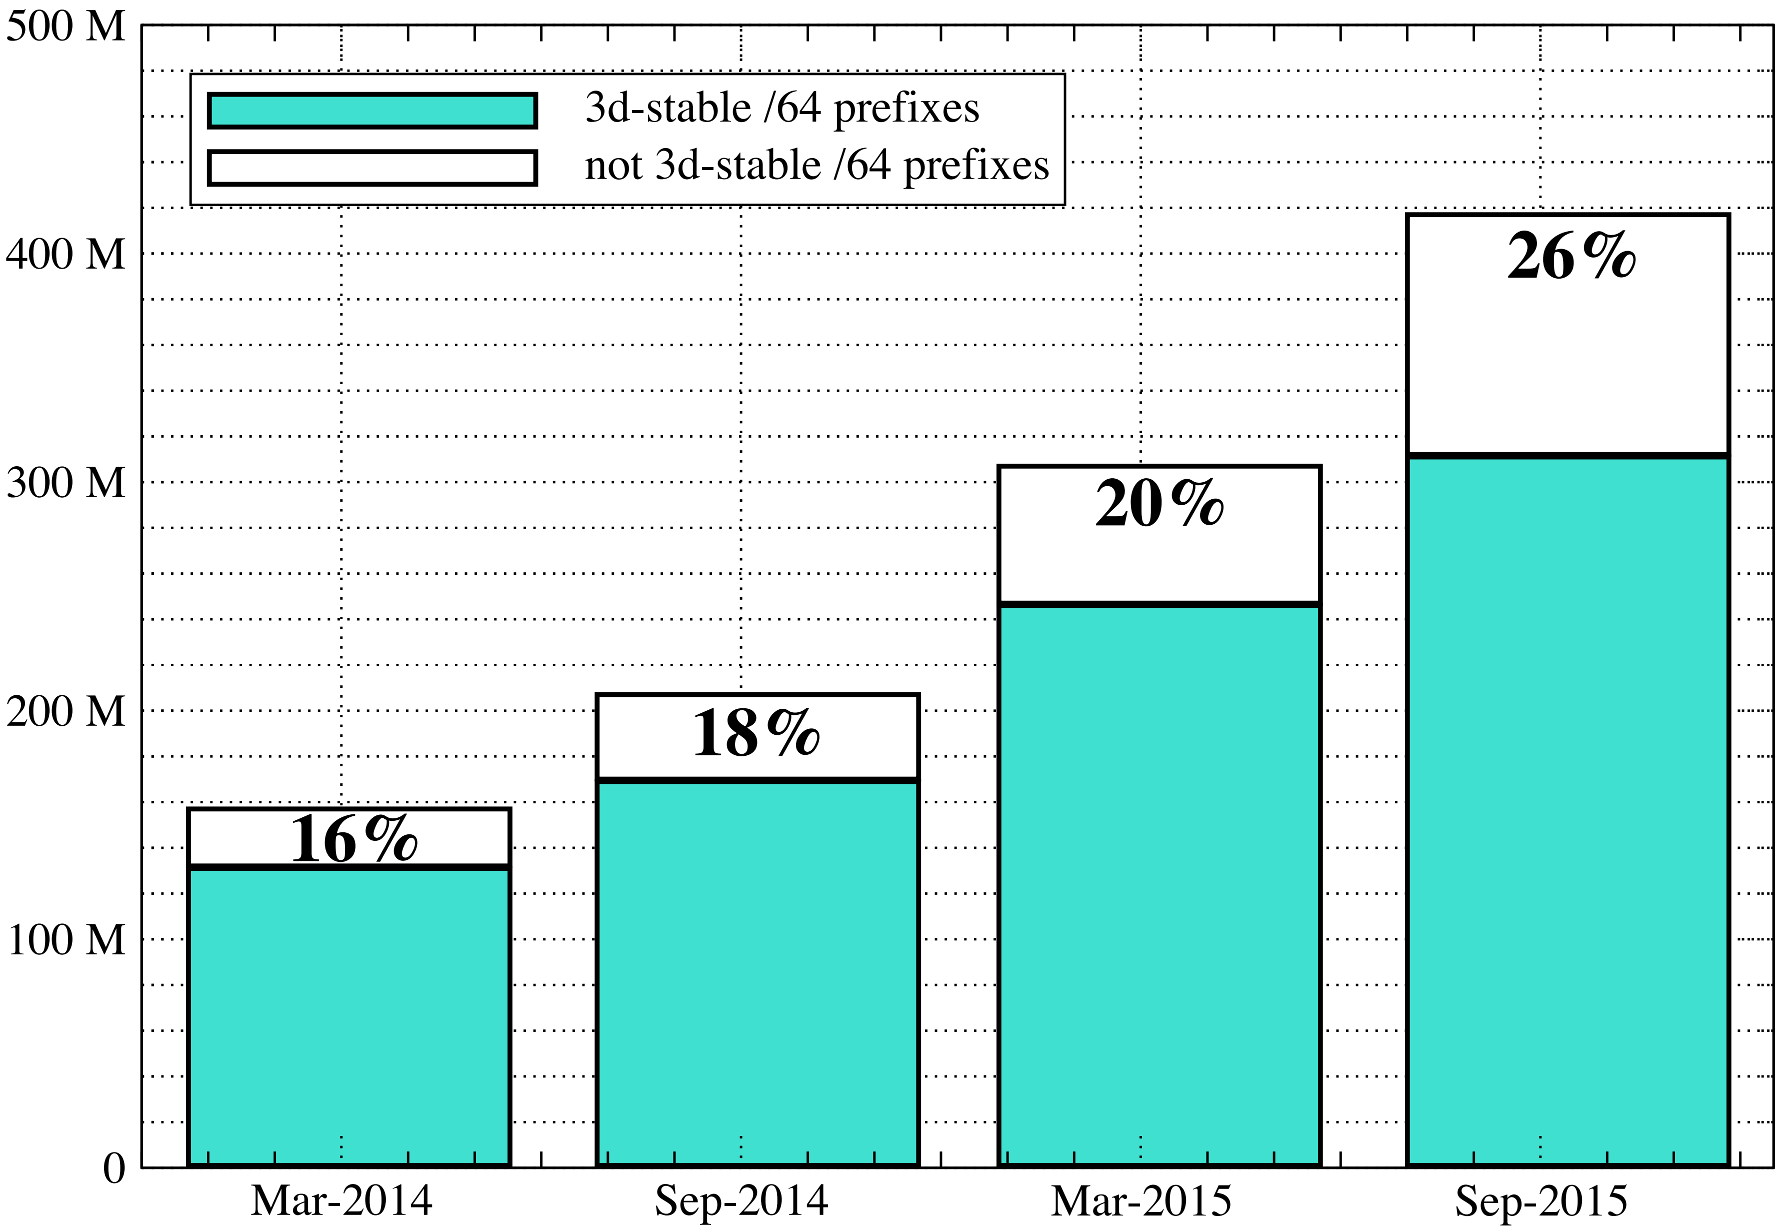

Binary PATRICIA Trie: JP TelCo /32

```
2001:db8:10:8::17f
2001:db8:10:9::68
2001:db8:10:c::109d
2001:db8:10:e::92d
2001:db8:20:c000:568:acf2:32ba:c6bf
```

Binary PATRICIA Trie: JP TelCo /32

```
2001:db8:10:8::17f
2001:db8:10:9::68
2001:db8:10:c::109d
2001:db8:10:e::92d
2001:db8:20:c000:568:acf2:32ba:c6bf
```

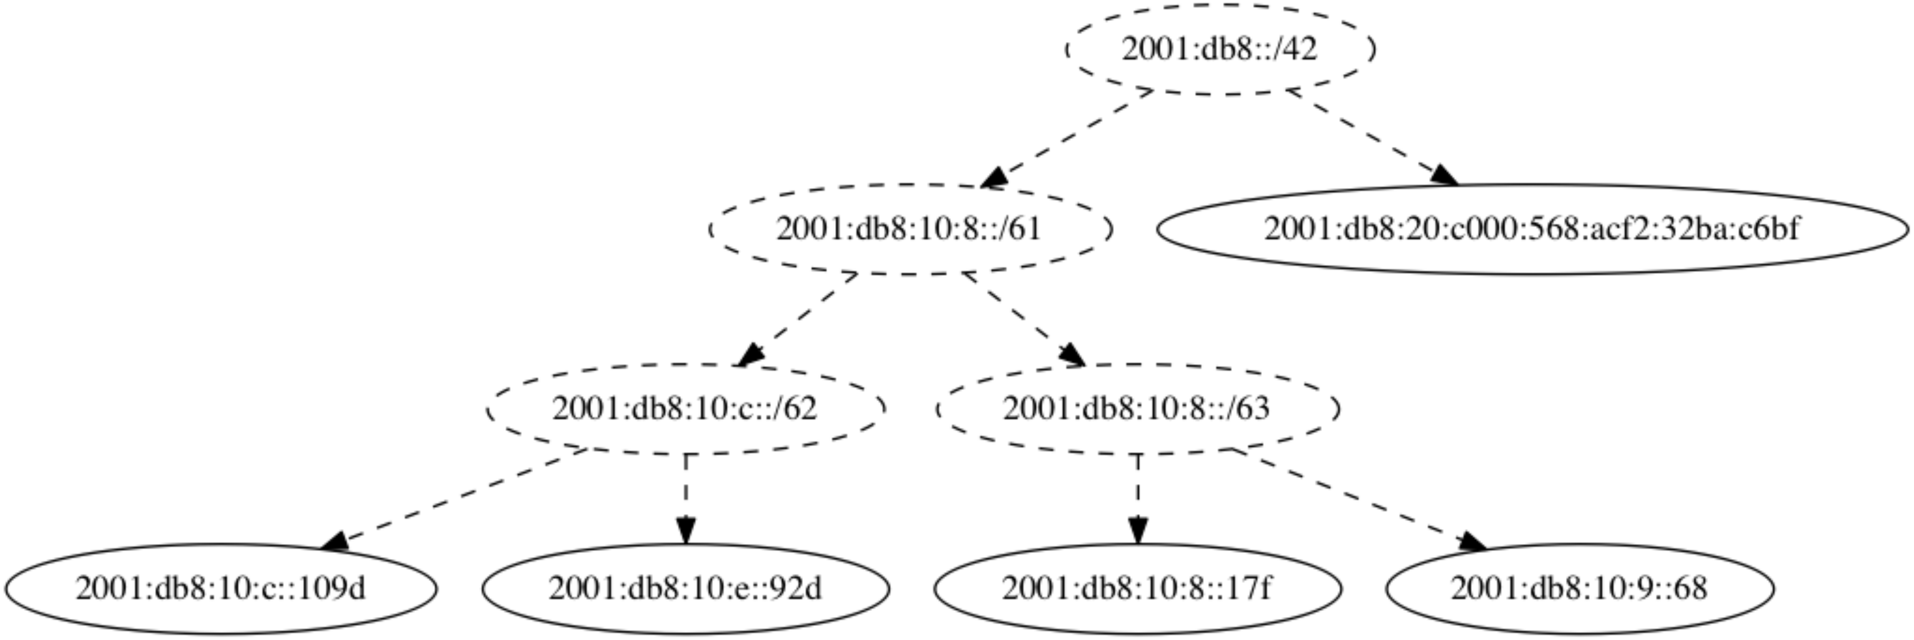

Binary PATRICIA Trie: JP TelCo /32

```
2001:db8:10:8::17f
2001:db8:10:9::68
2001:db8:10:c::109d
2001:db8:10:e::92d
2001:db8:20:c000:568:acf2:32ba:c6bf
```

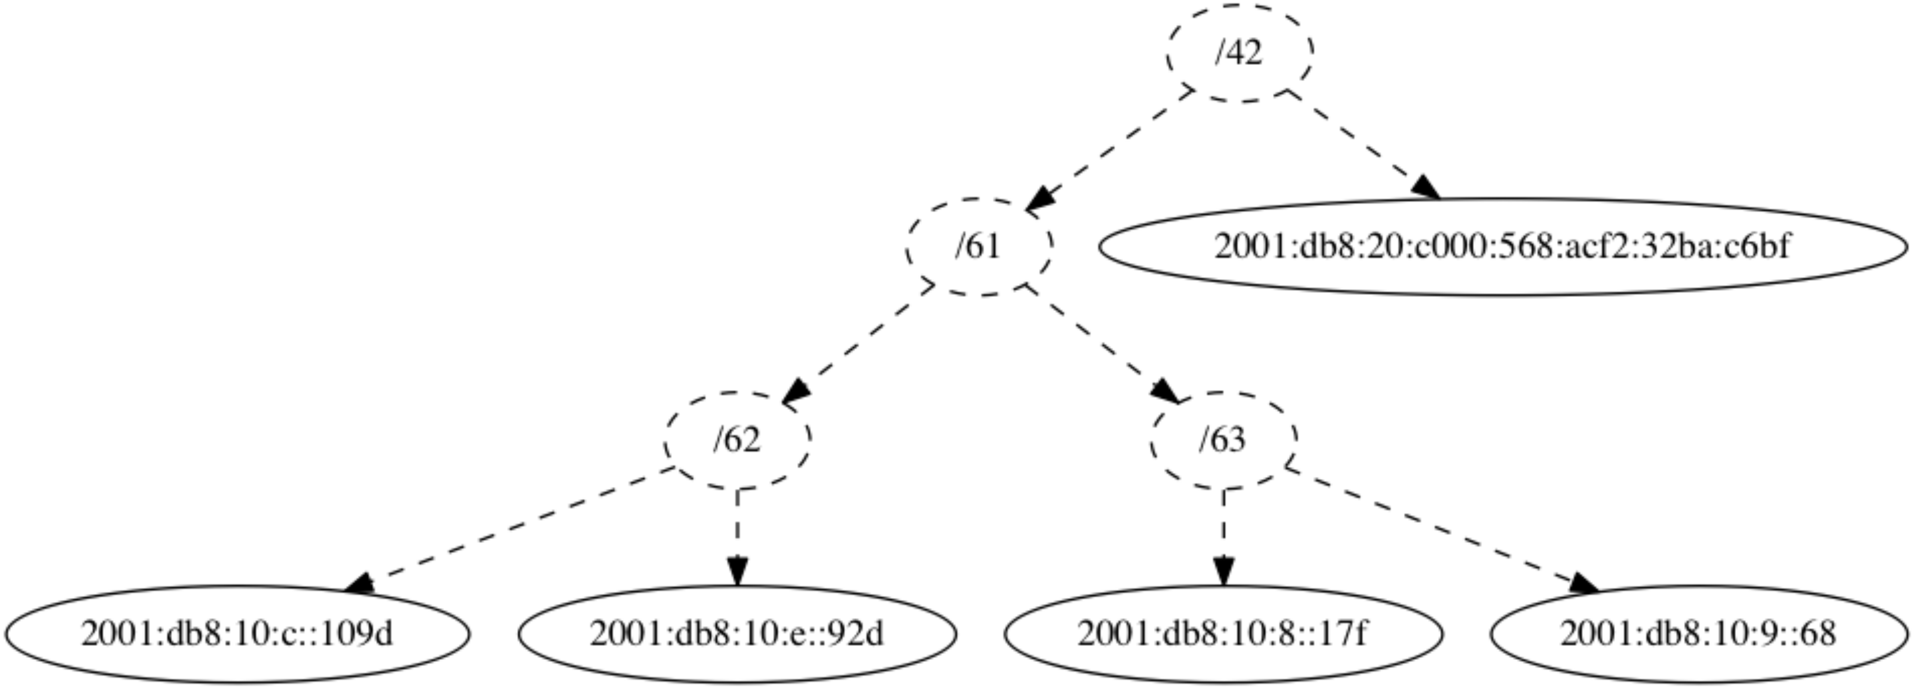

Binary PATRICIA Trie: JP TelCo /32

```
2001:db8:10:8::17f
2001:db8:10:9::68
2001:db8:10:c::109d
2001:db8:10:e::92d
2001:db8:20:c000:568:acf2:32ba:c6bf
```

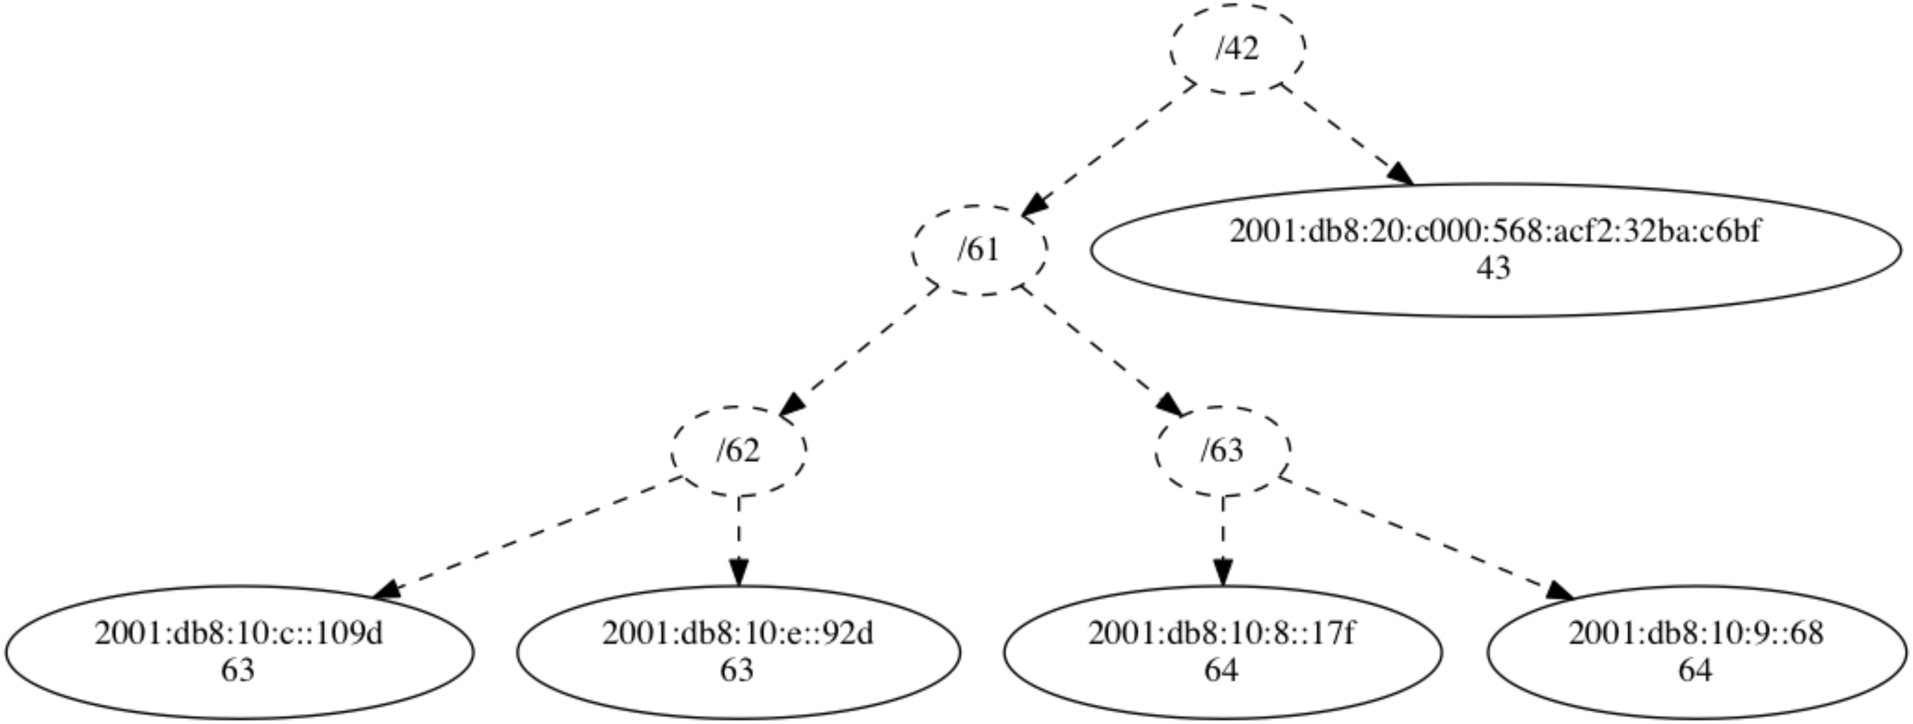

A Tokyo-inspired explanation of the MRA Plot

A Japan-inspired explanation of the MRA Plot

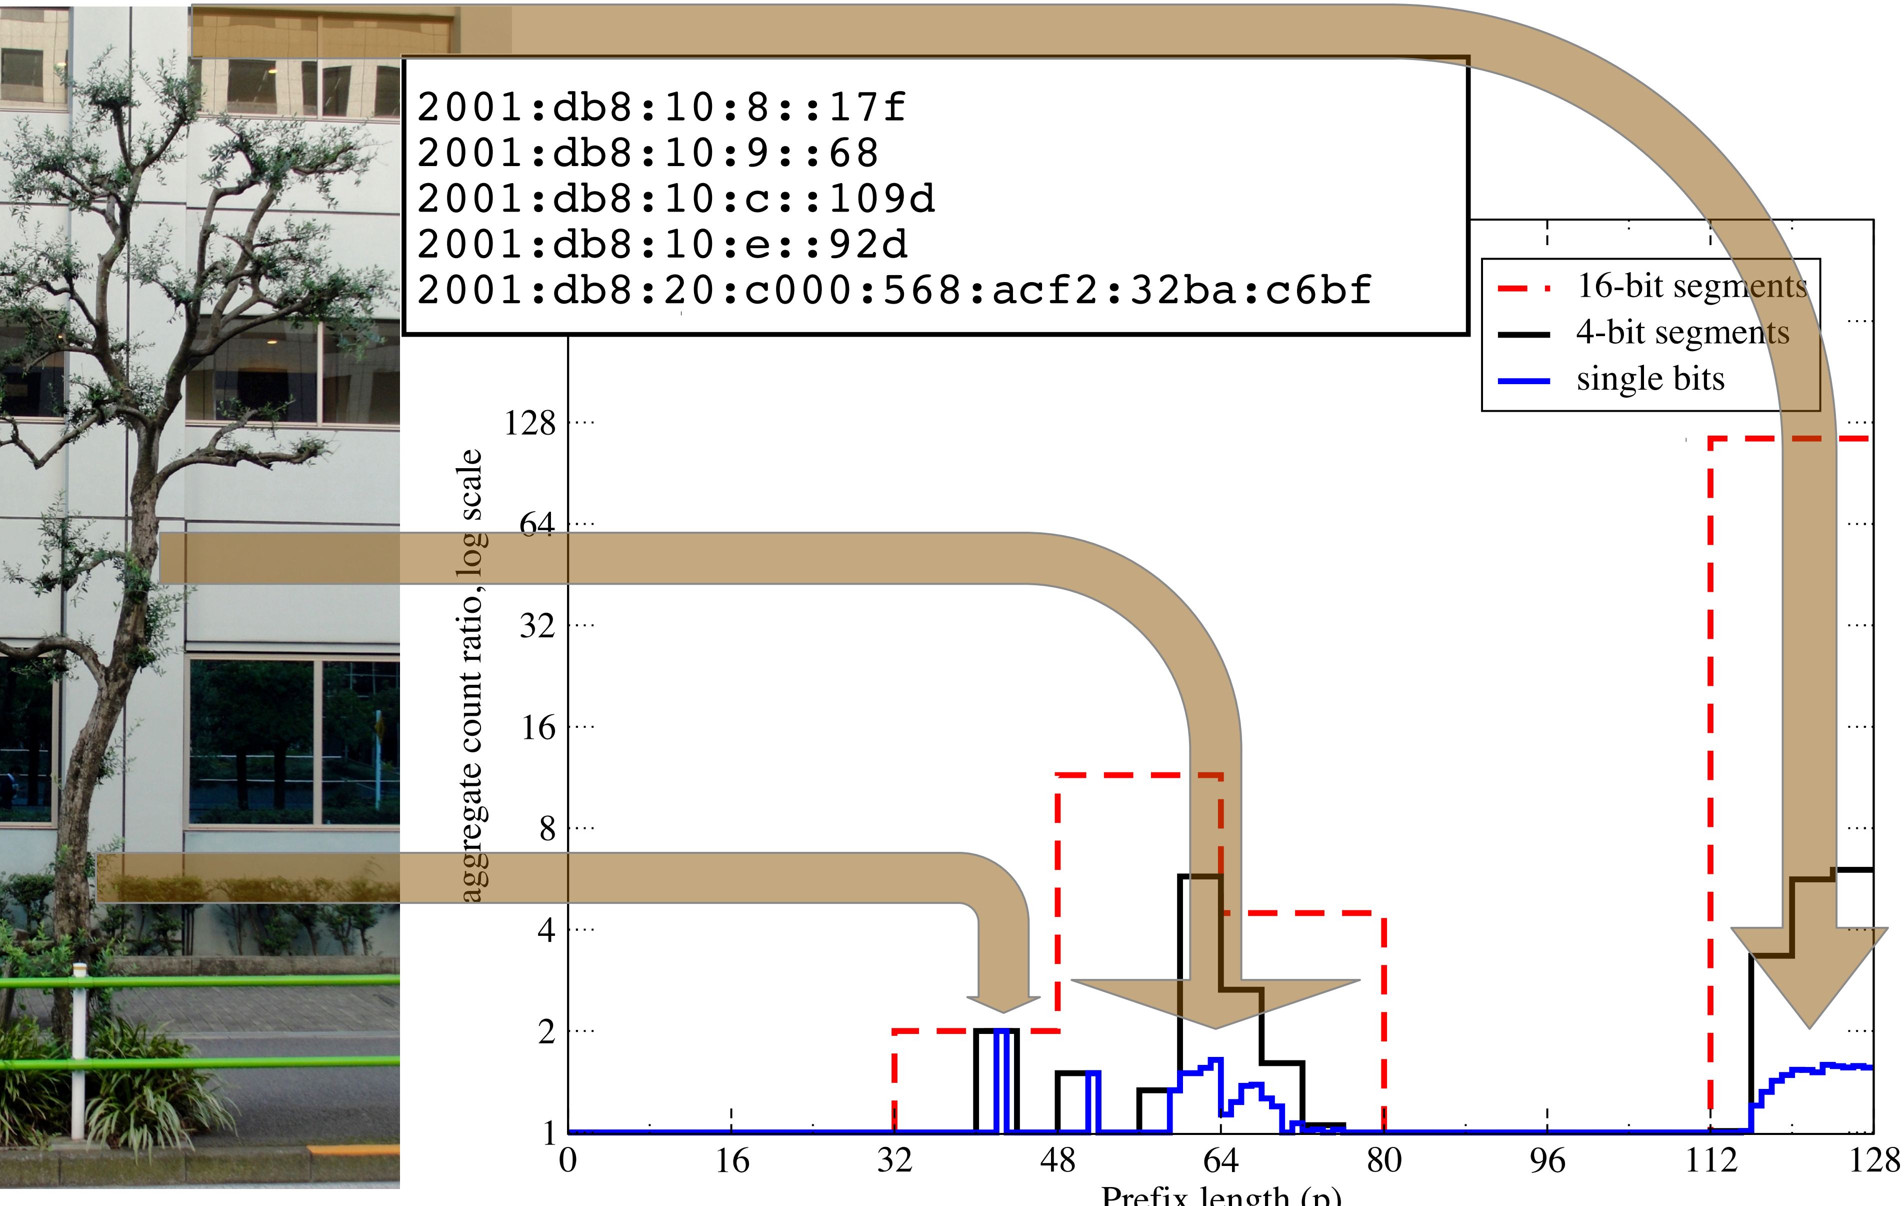

IPv4 Aggregate Count Ratio [Kohler *et al.* 2002]

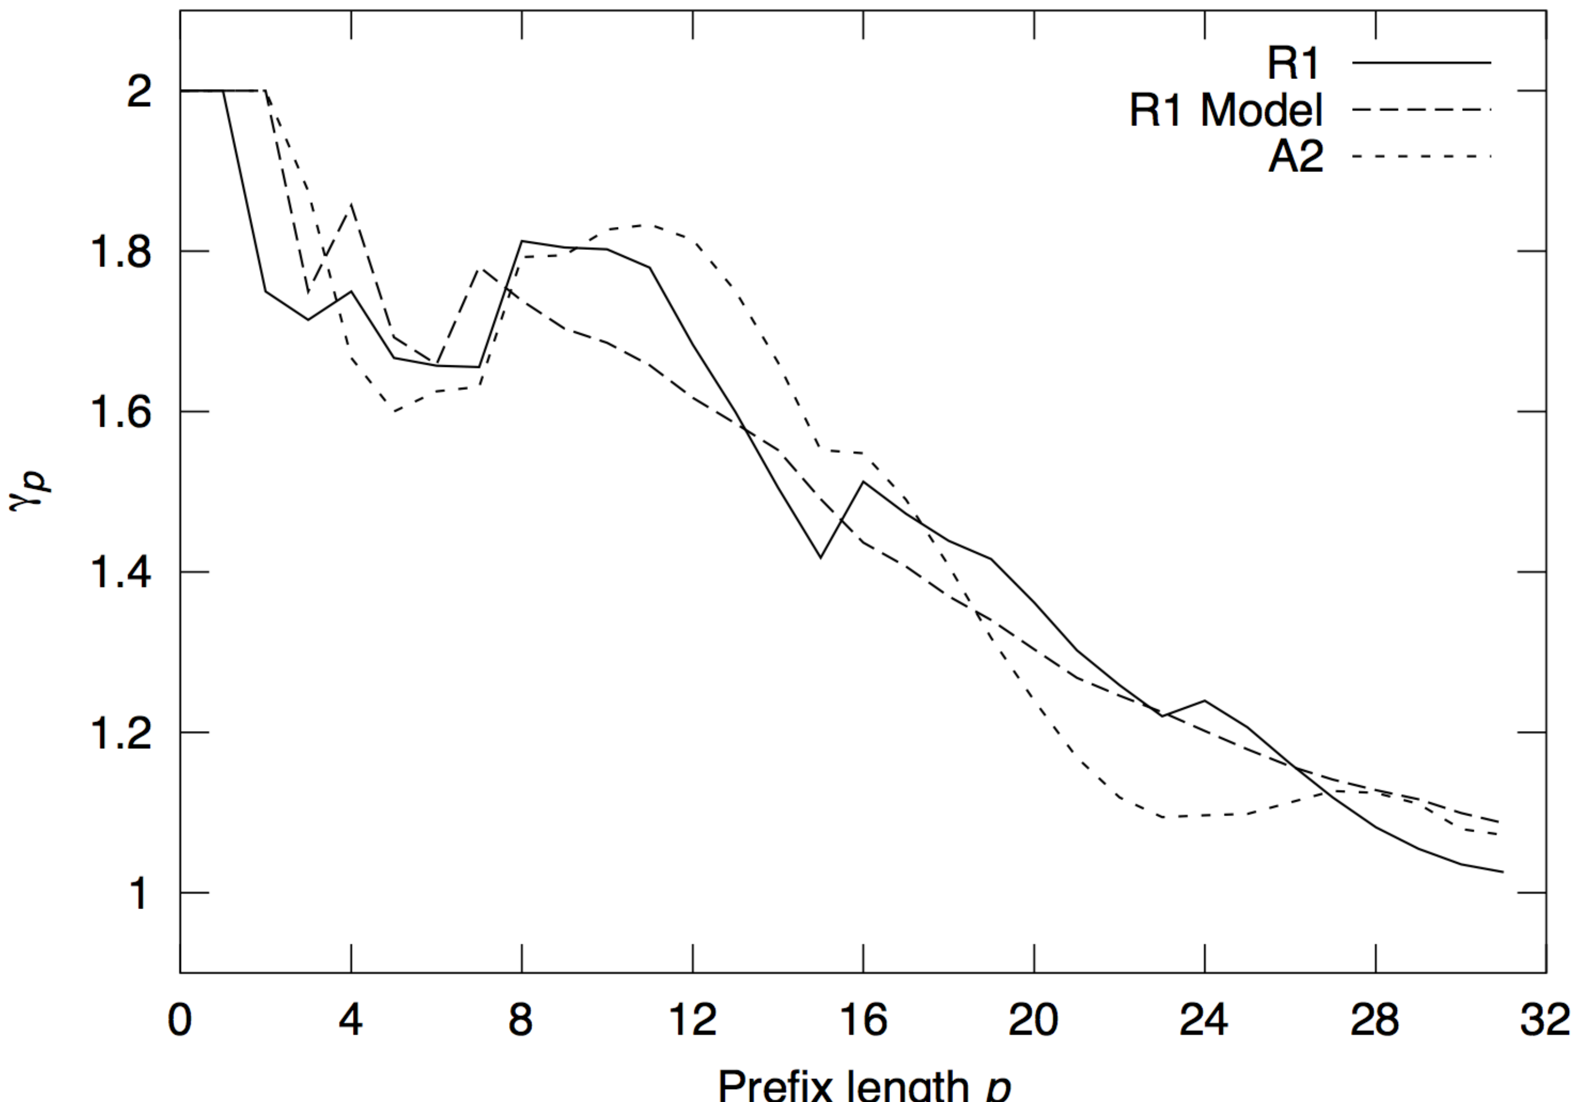

IPv6 Aggregate Counts: JP TelCo (~12K active client addrs)

```
2001:db8:10:8::17f
2001:db8:10:9::68
2001:db8:10:c::109d
2001:db8:10:e::92d
2001:db8:20:c000:568:acf2:32ba:c6bf
```

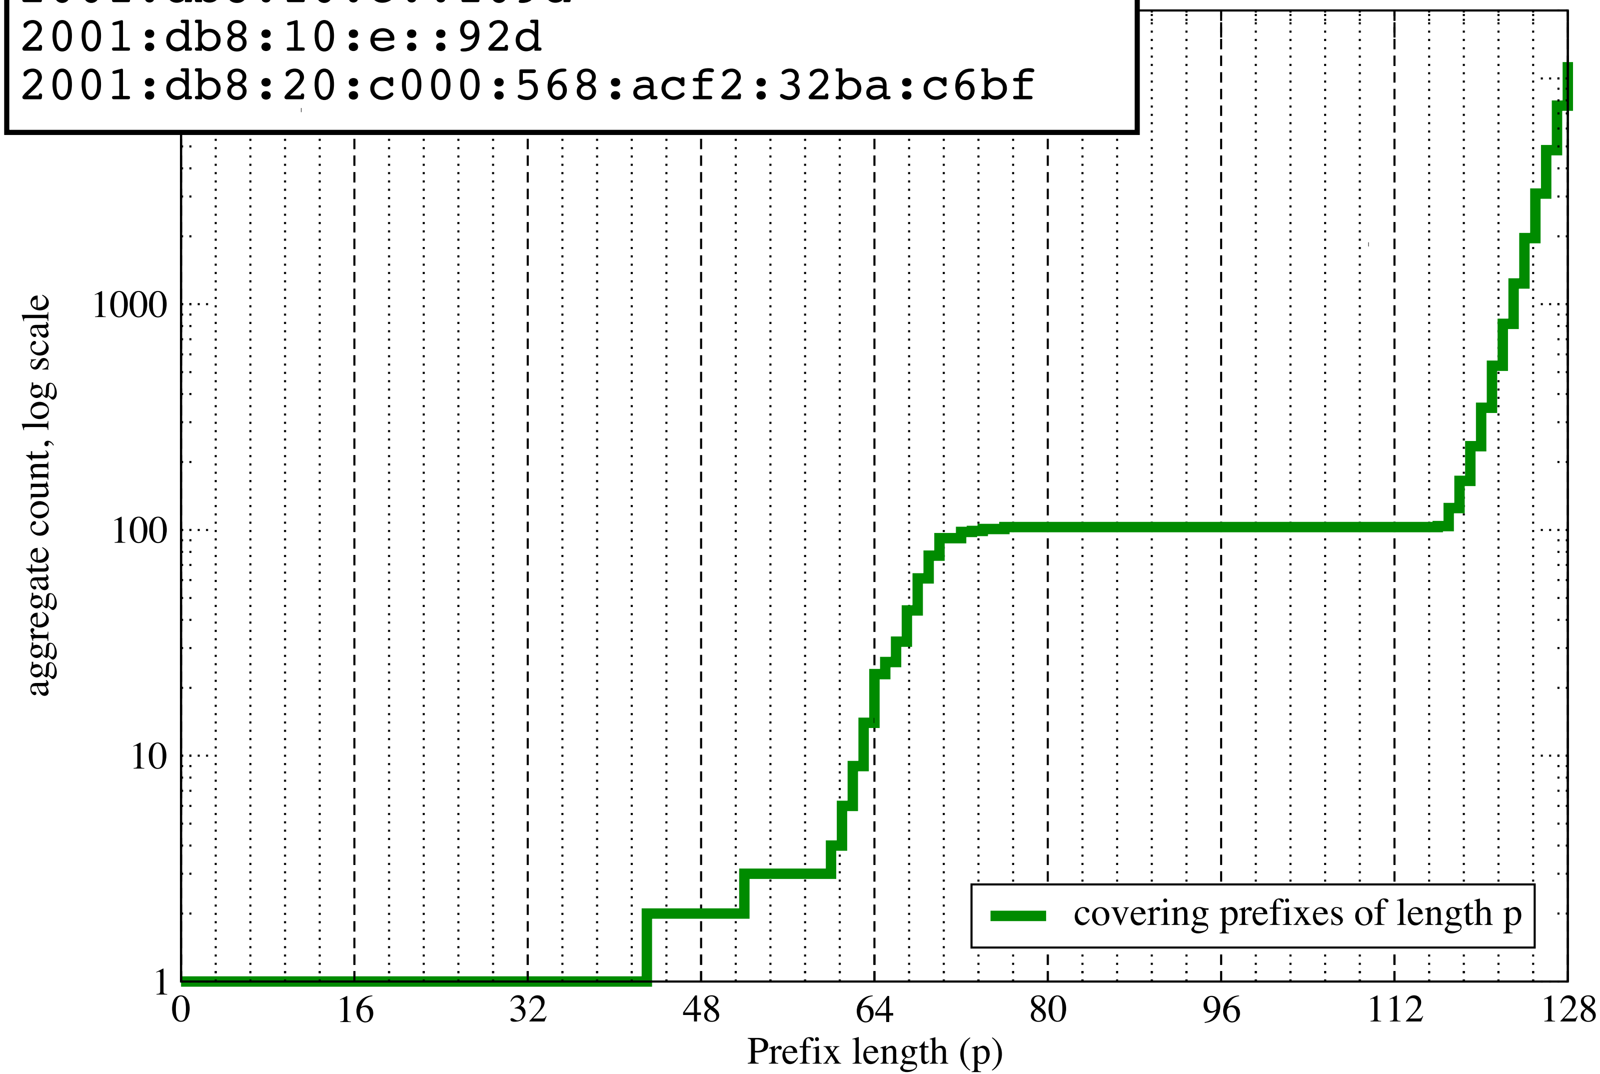

IPv6 Aggregate Counts: JP TelCo (~12K active client addrs)

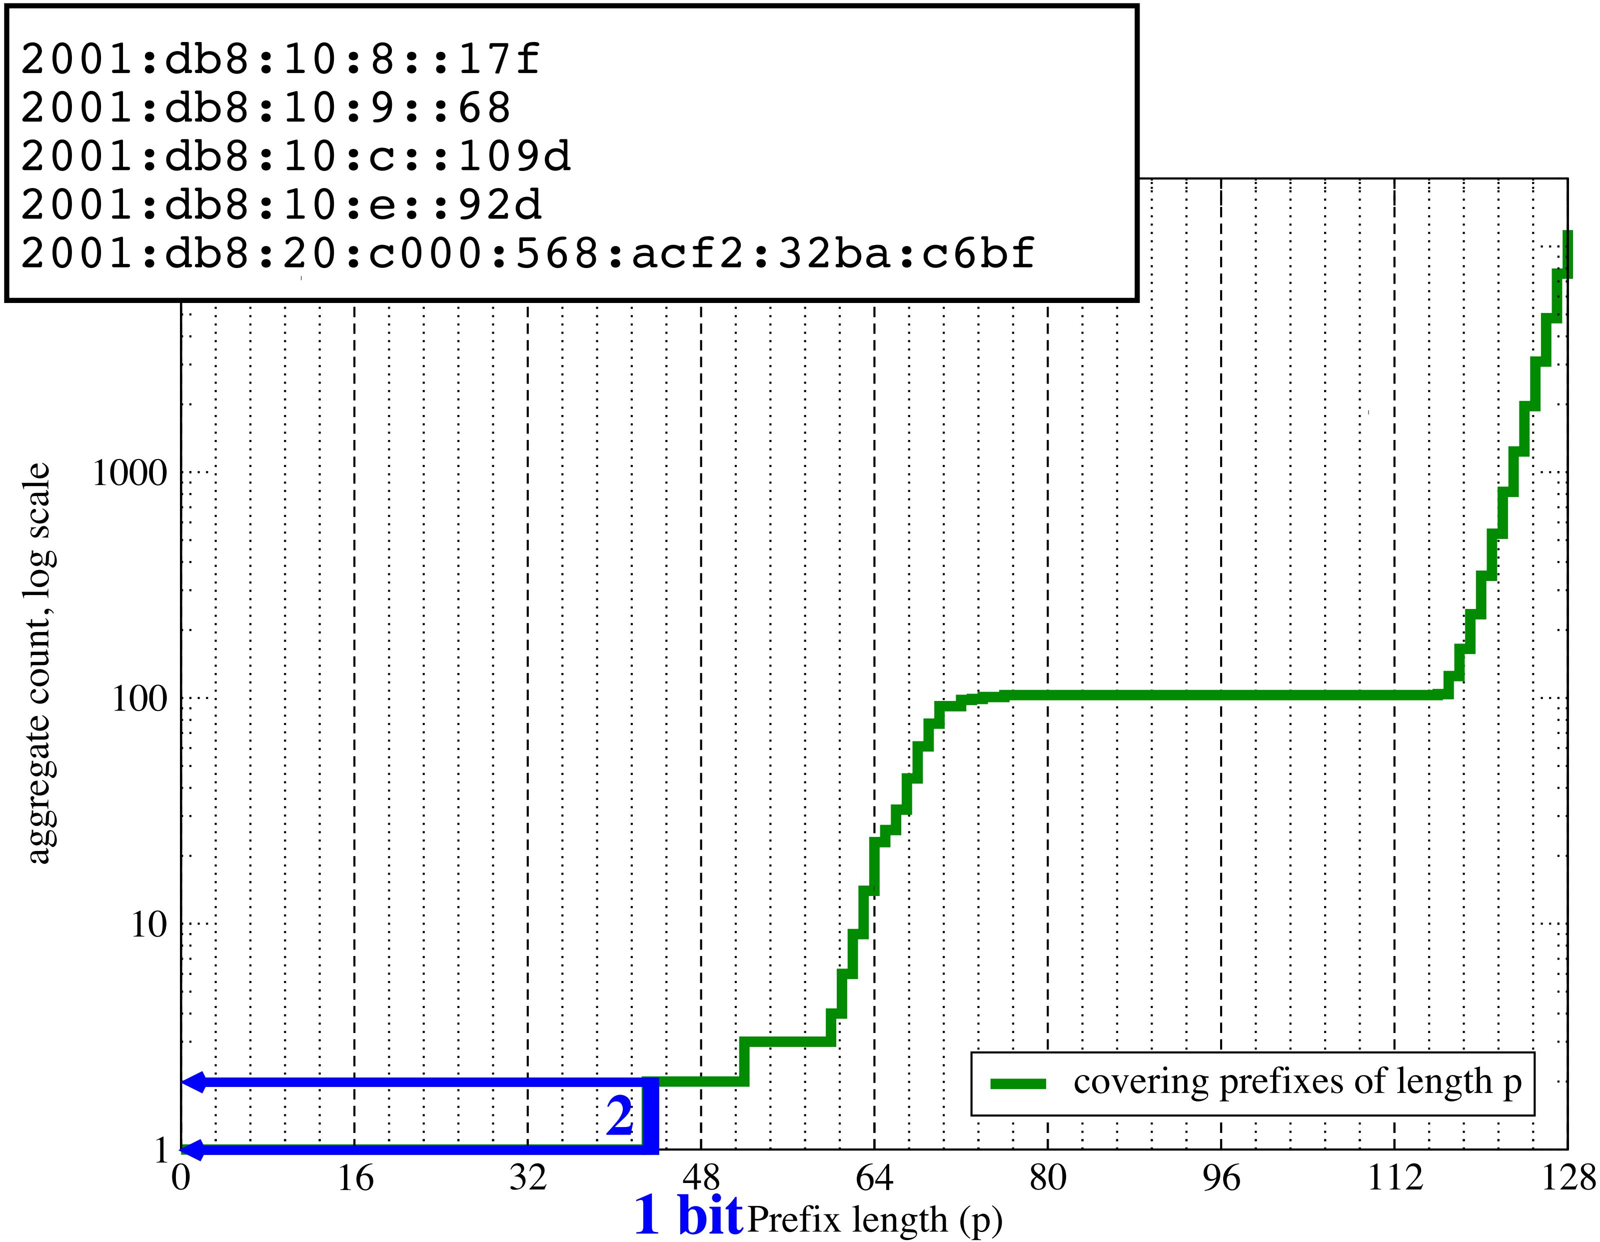

IPv6 Aggregate Counts: JP TelCo (~12K active client addrs)

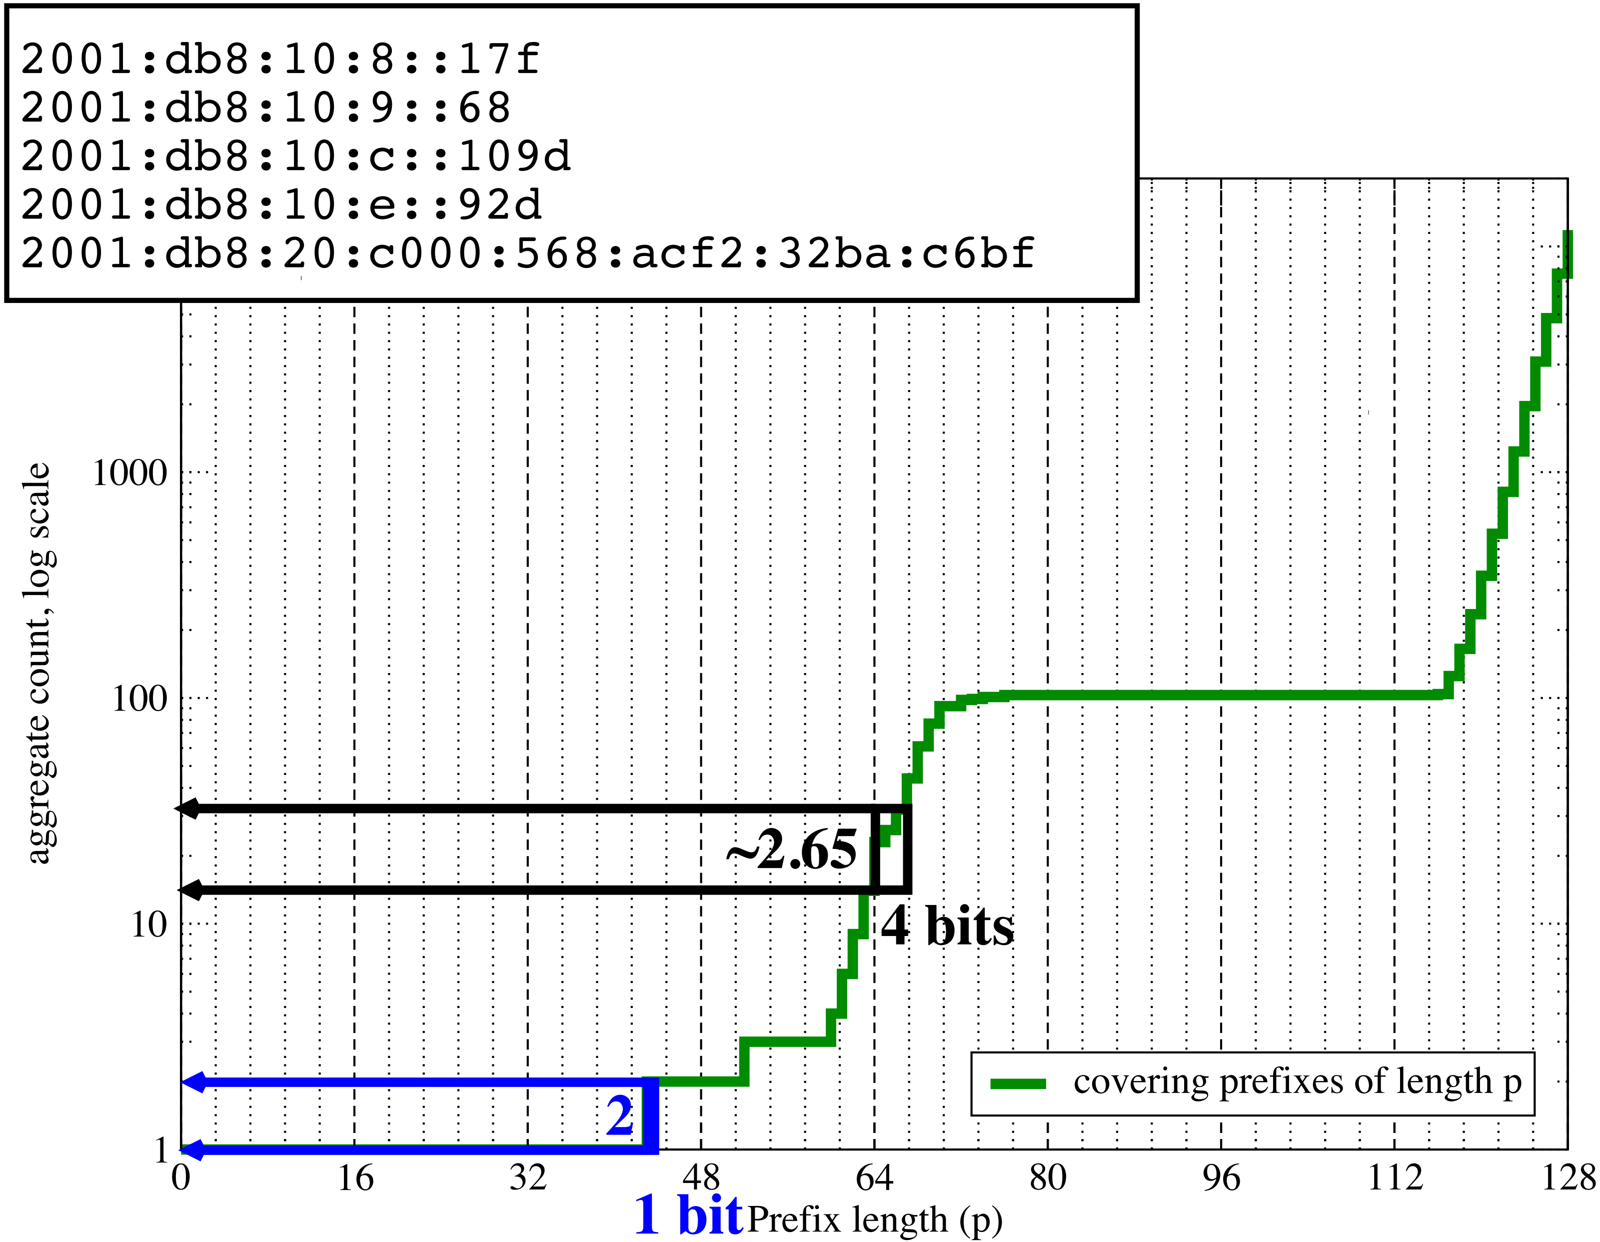

IPv6 Aggregate Counts: JP TelCo (~12K active client addrs)

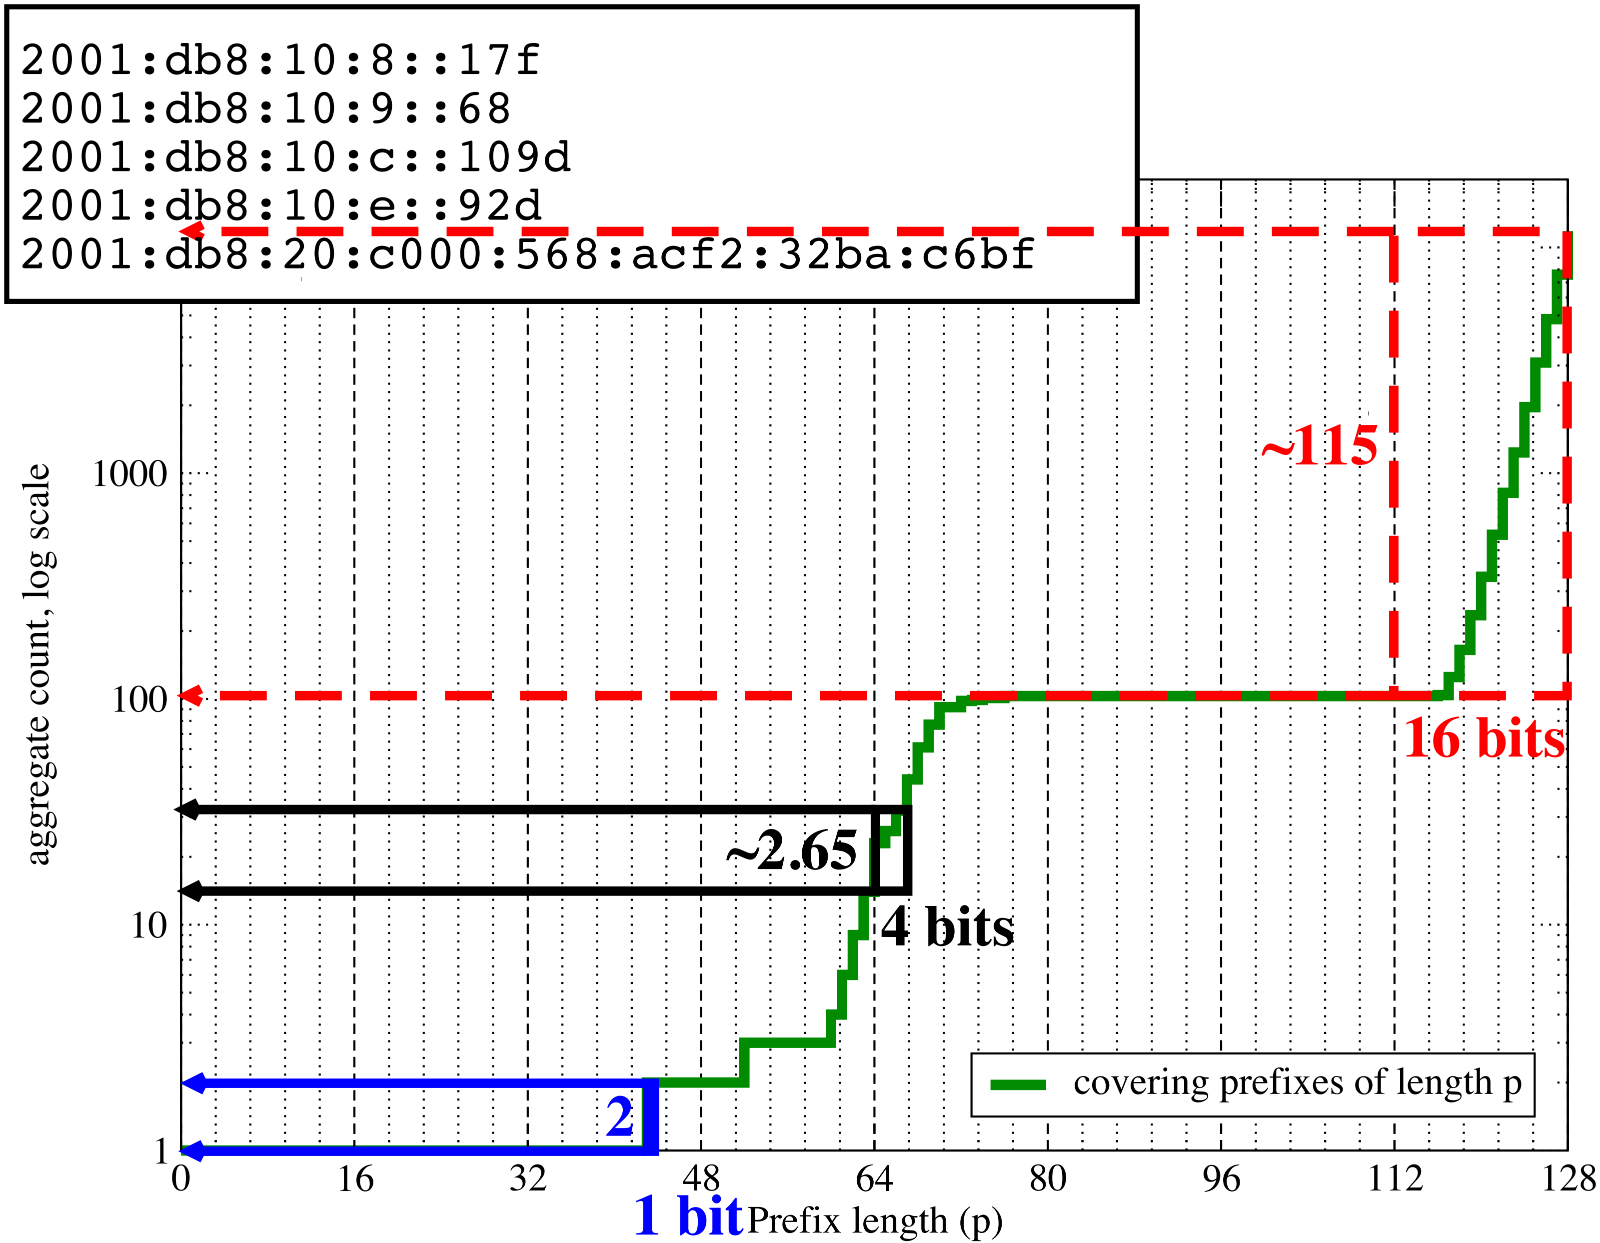

MRA Plot: JP TelCo /32 (~12K active WWW client addrs)

```
2001:db8:10:8::17f
2001:db8:10:9::68
2001:db8:10:c::109d
2001:db8:10:e::92d
2001:db8:20:c000:568:acf2:32ba:c6bf
```

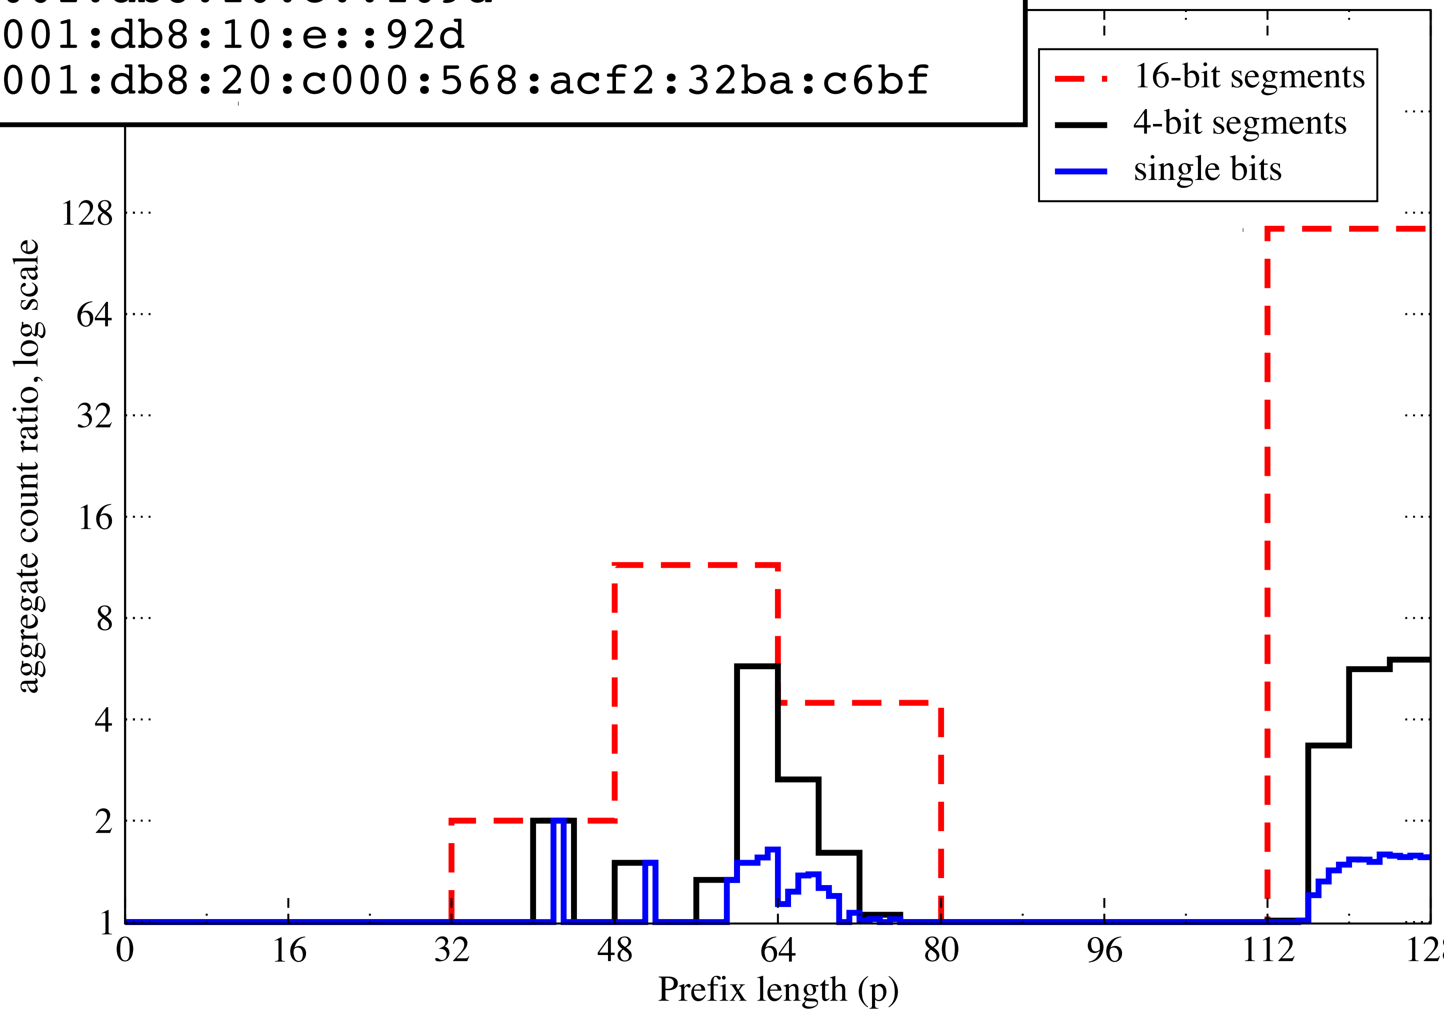

MRA Plot: JP TelCo /32 (~12K active WWW client addrs)

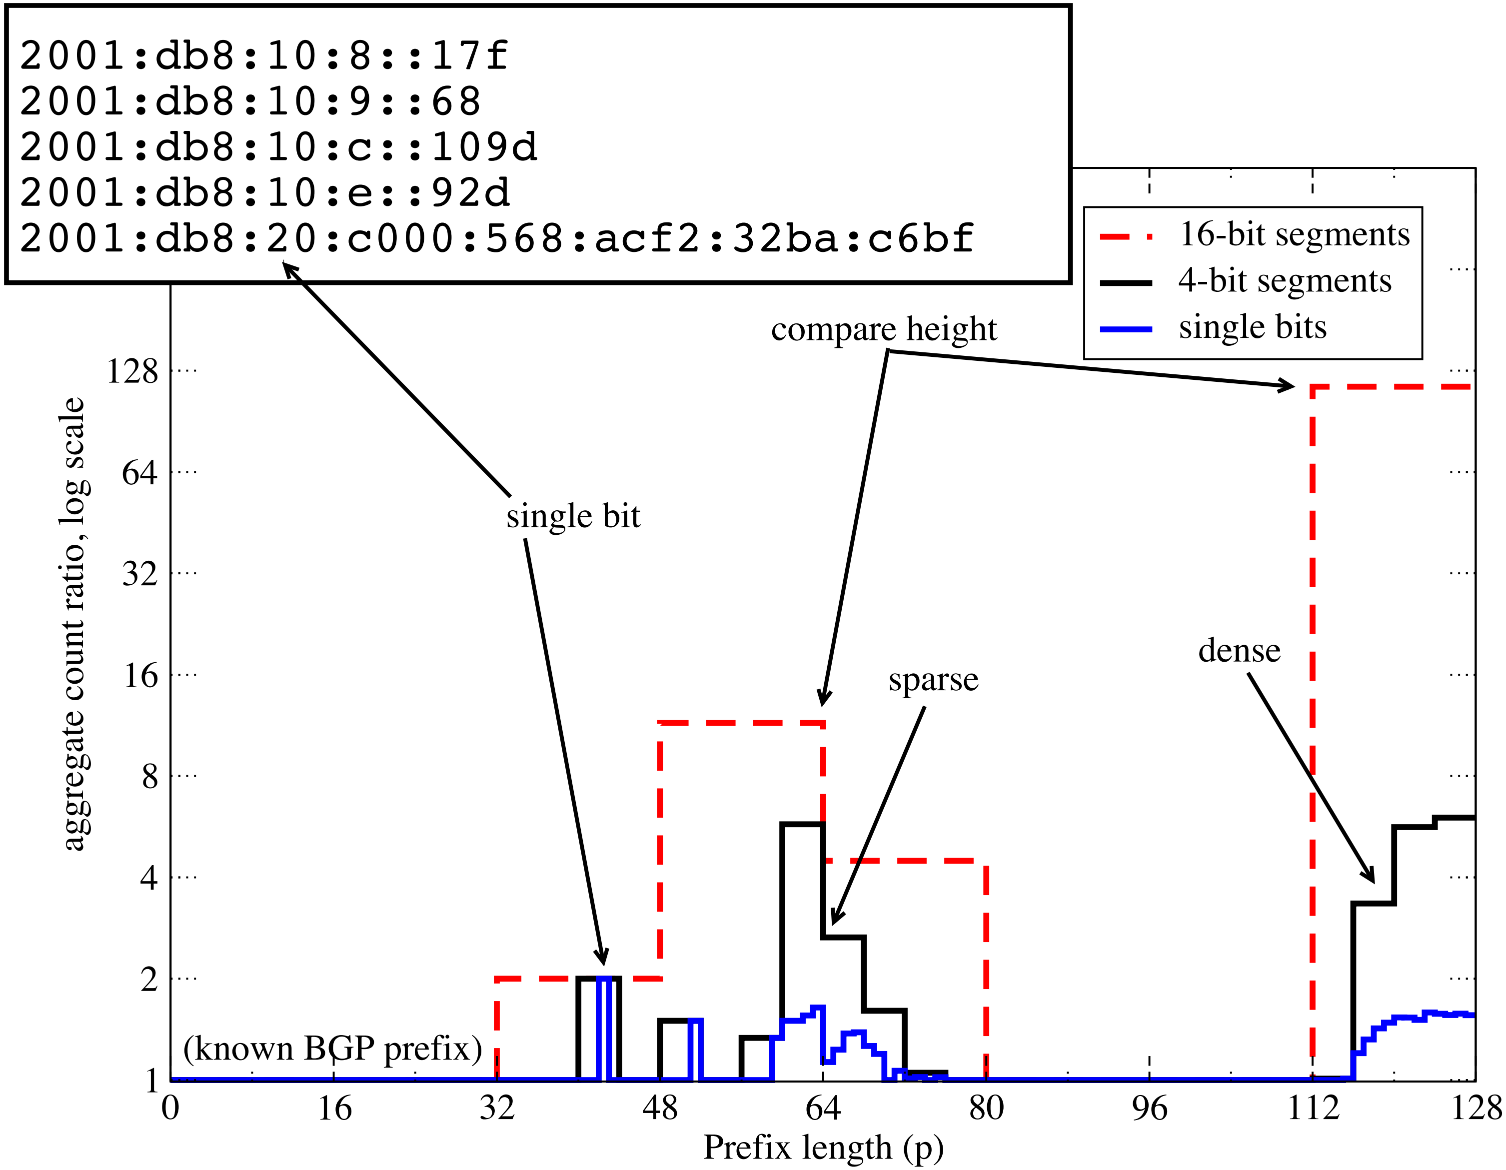

MRA Plot: US Uni /32 (~7K active WWW client addrs)

```
2001:db8:e:0:e174:5522:1ada:1e5b
2001:db8:1082:fff8:ab:ebfd:9b16:6095
2001:db8:1082:fff8:9185:20eb:4349:816b
2001:db8:1082:fffa:245d:21cc:69a5:ac39
2001:db8:1082:ffff:d4b8:7d56:ad92:252f
```

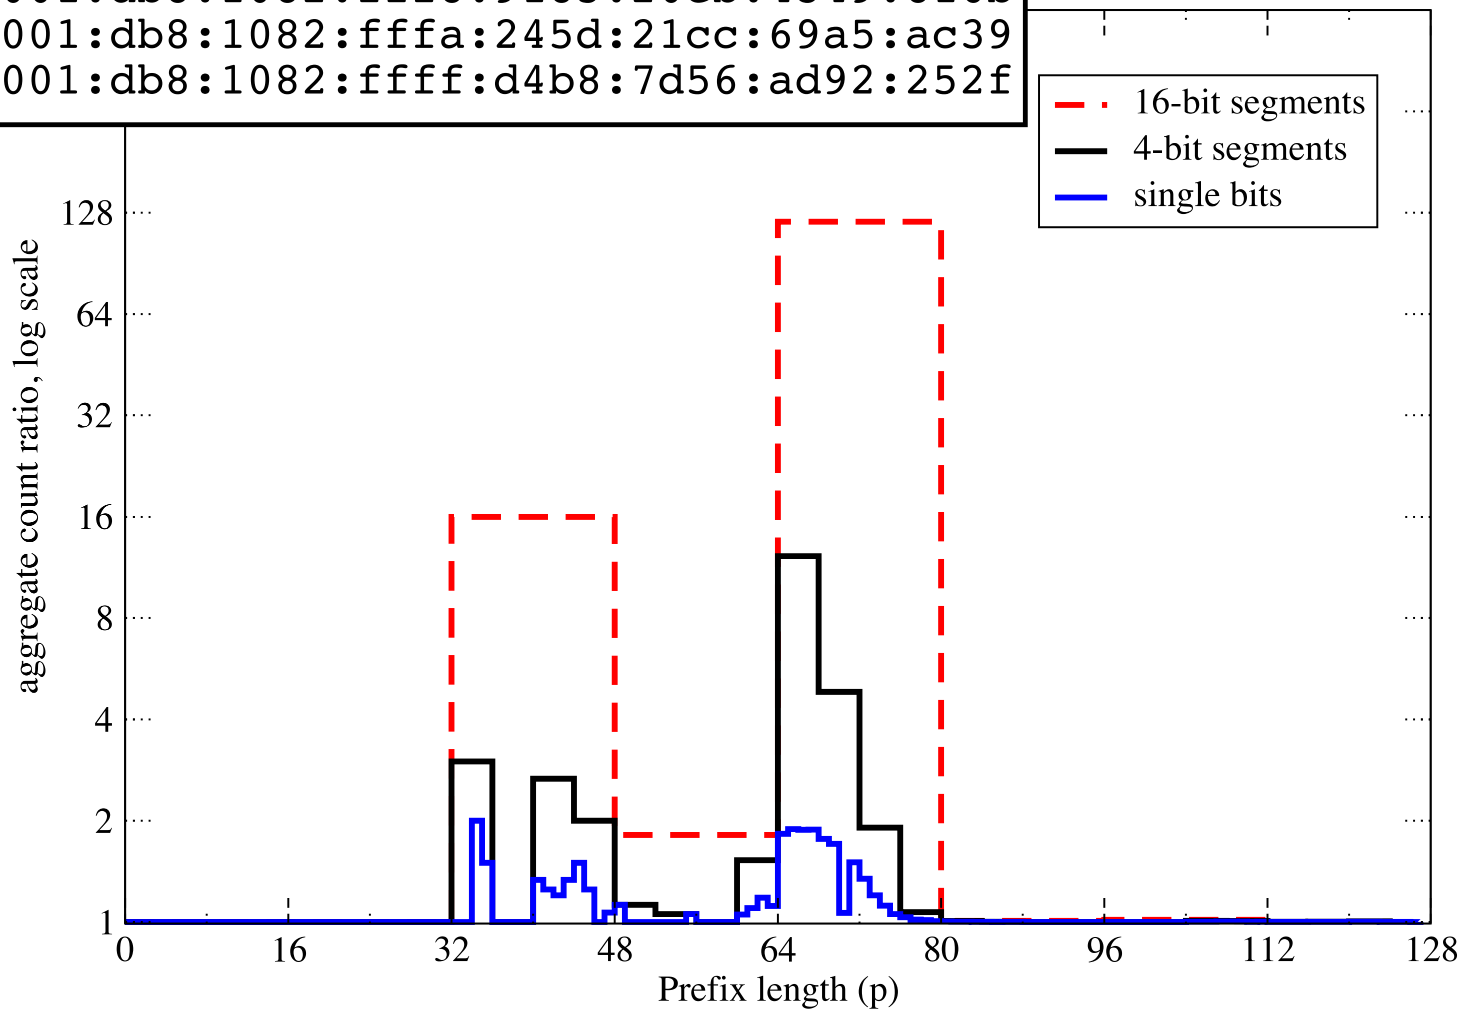

MRA Plot: US Uni /32 (~7K active WWW client addrs)

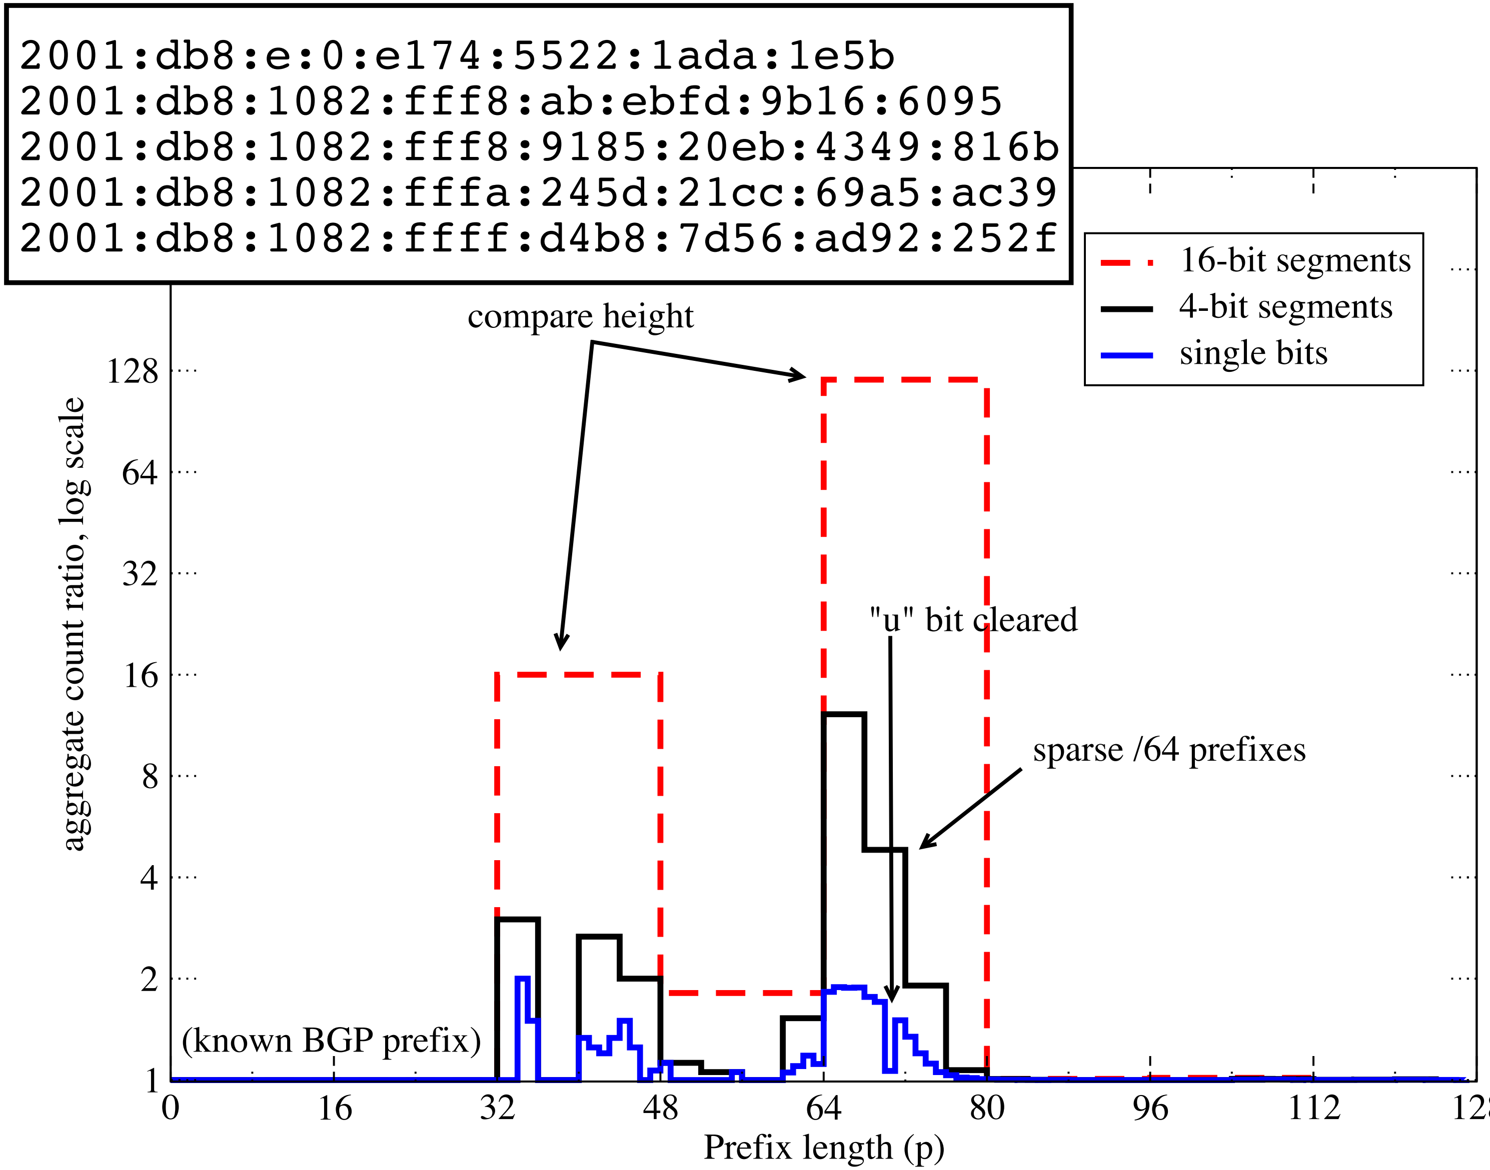

A Tour of the Active IPv6 Address Space, Mar 17-23, 2015

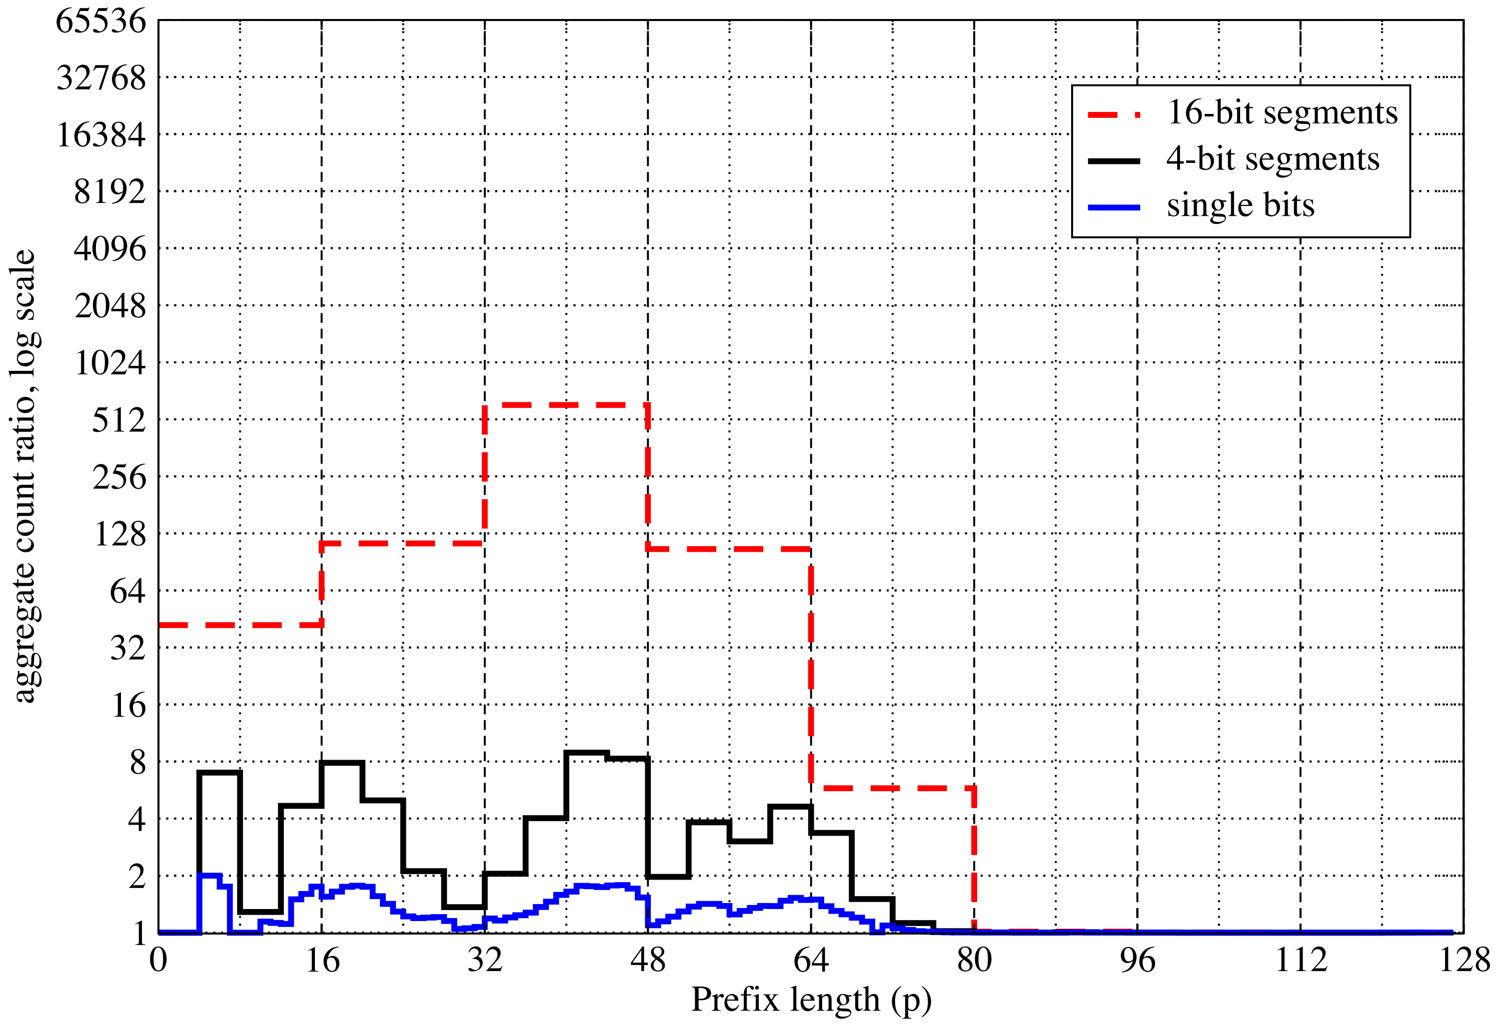

MRA Plot: 1.81B active WWW client addrs (not 6to4, Teredo)

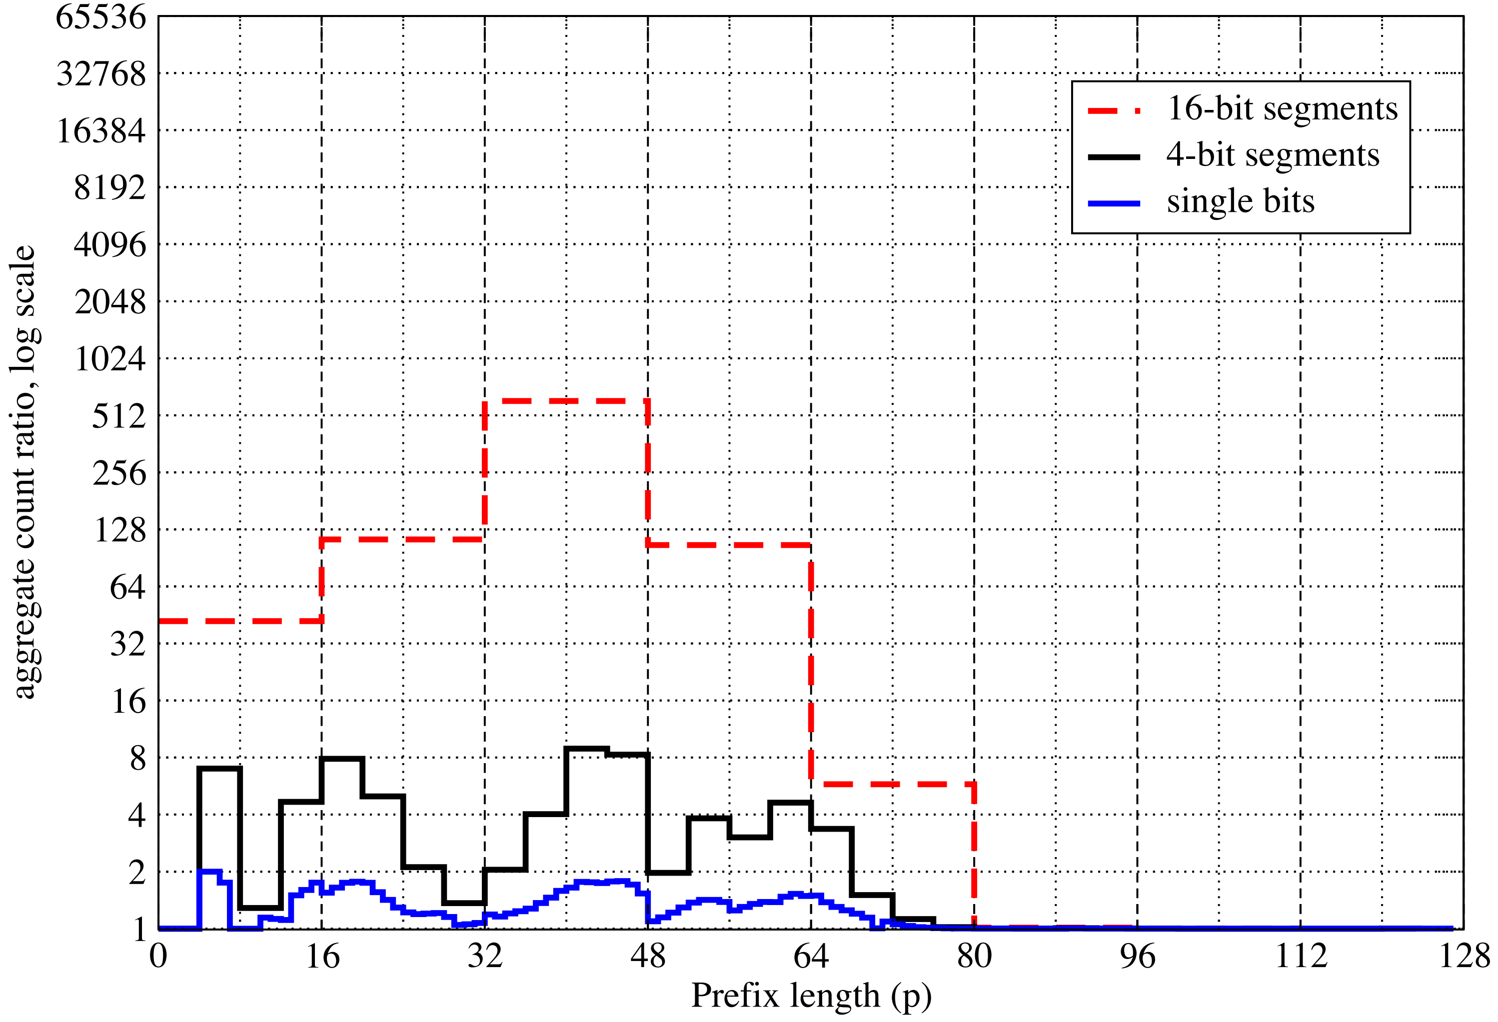

MRA Plot: 6to4: 64.2M active client addrs (49.3M IPv4 addrs)

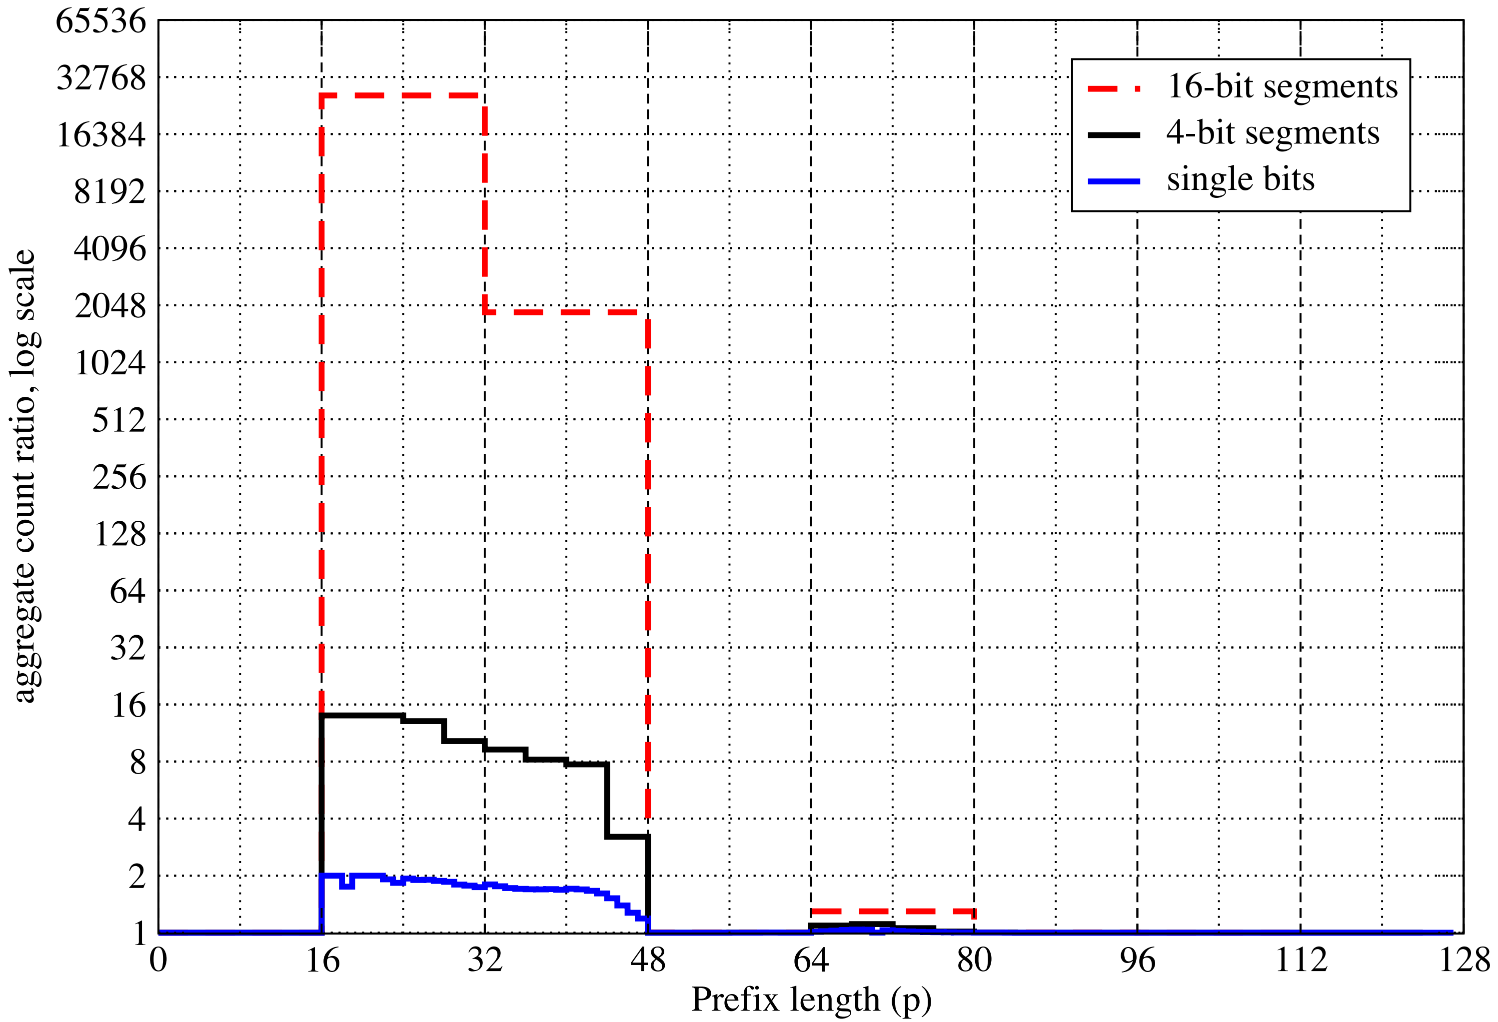

MRA Plot: US mobile: 510M active client addrs, 167M /64s

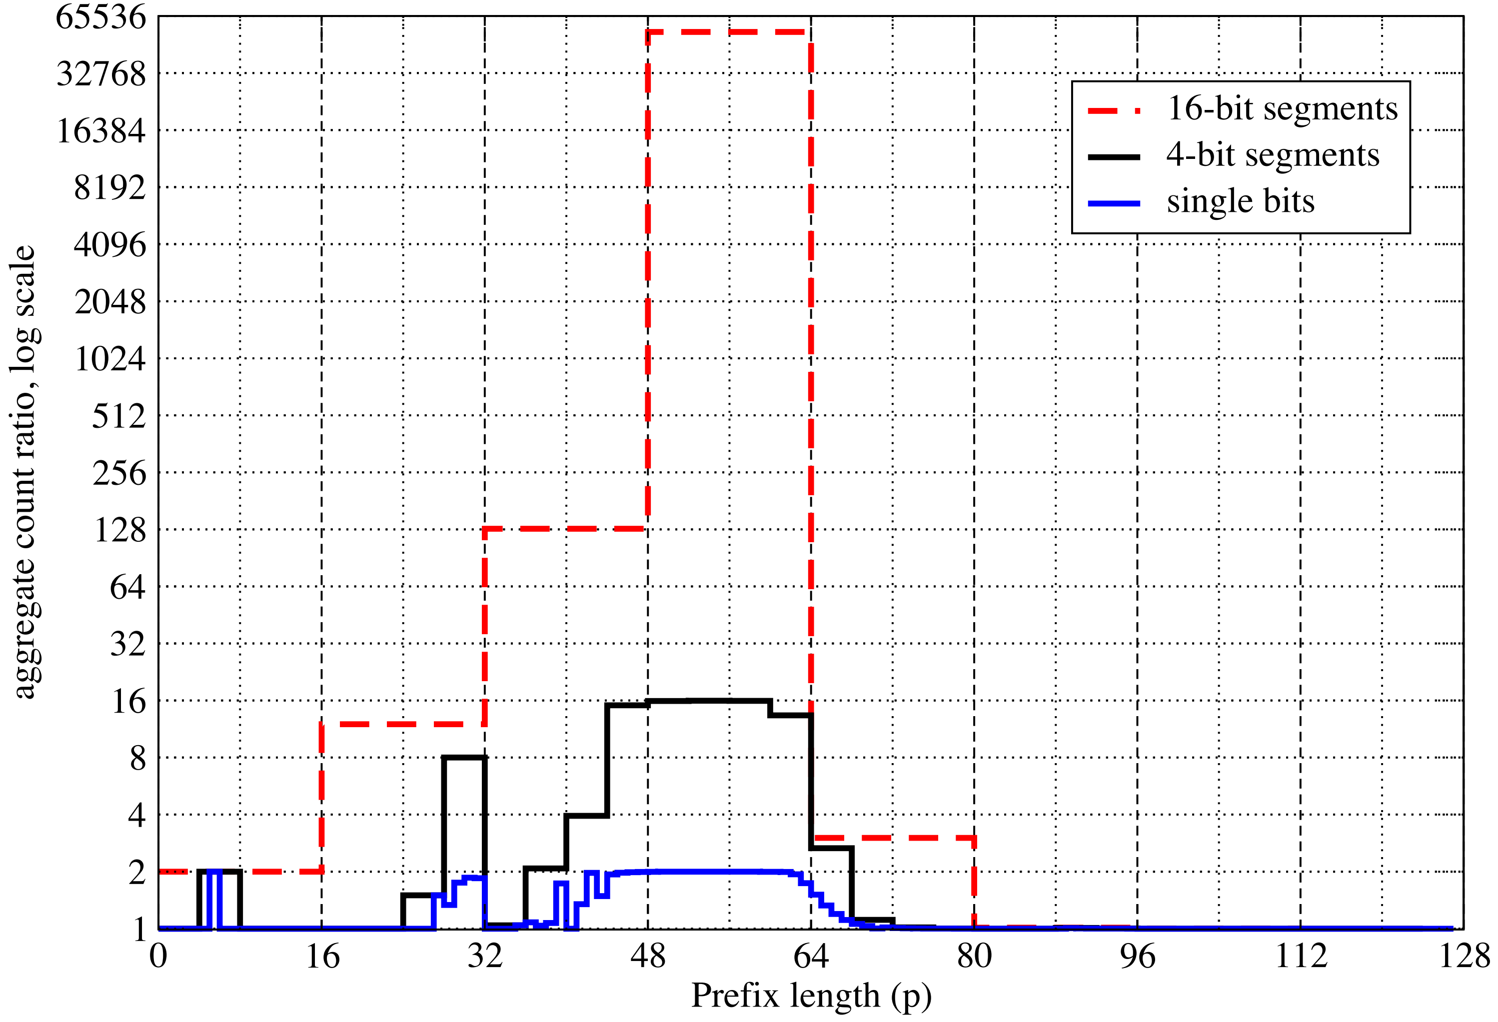

MRA Plot: JP ISP: 57.0M active client addrs, 2.18M /64s

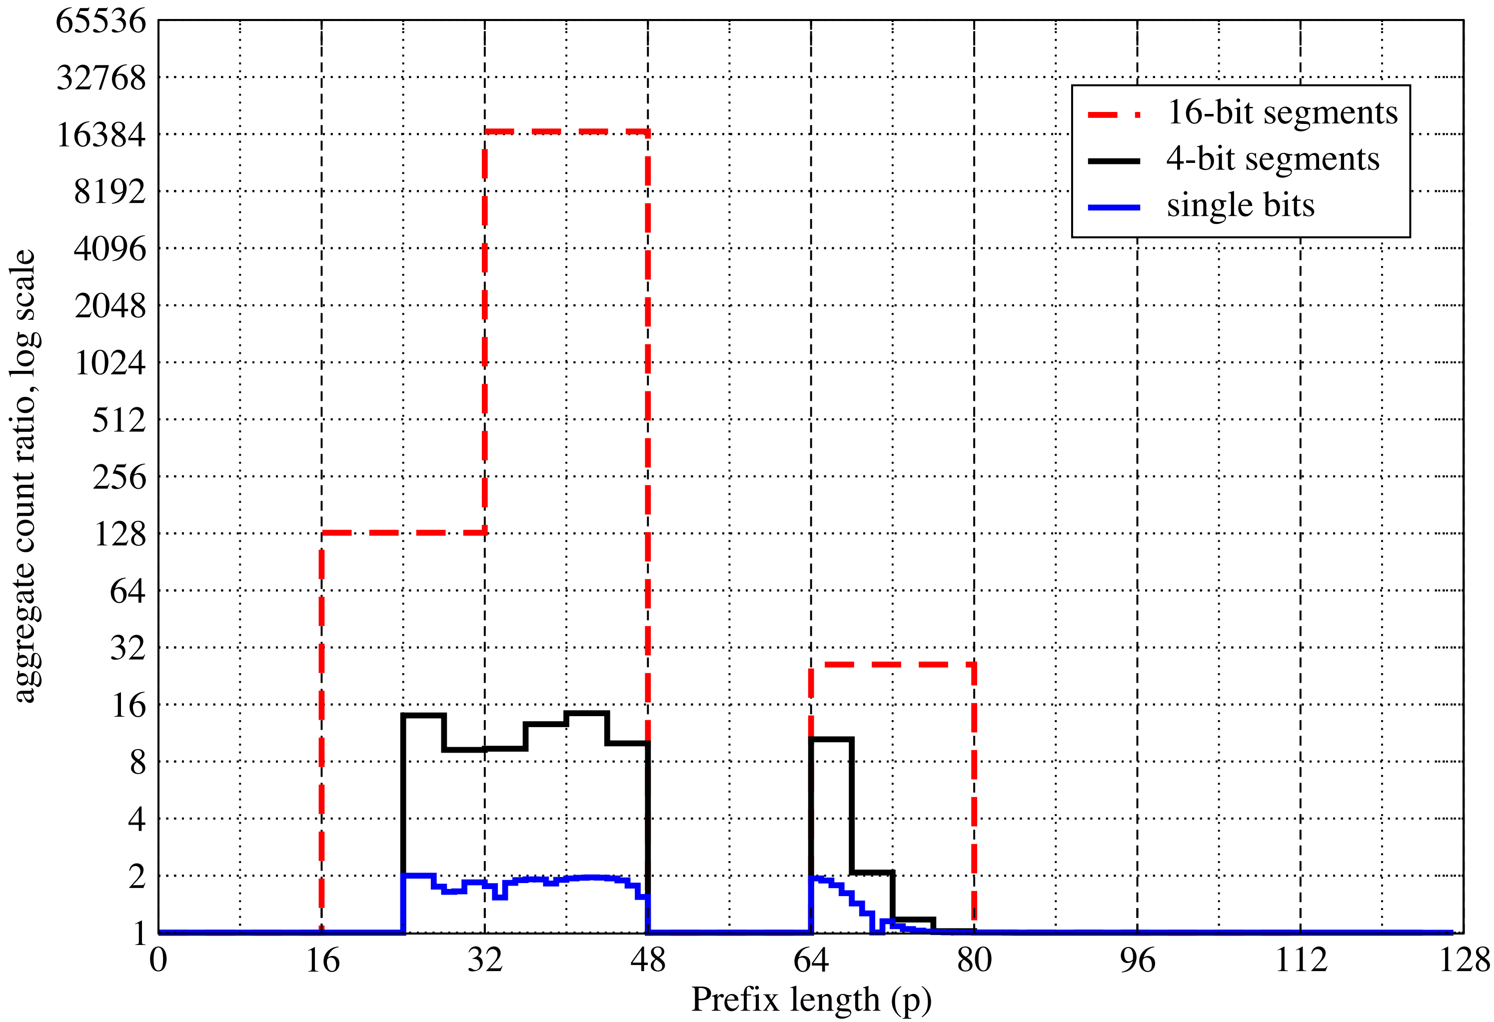

MRA Plot: EU ISP: 86.2M active client addrs, 15.5M /64s

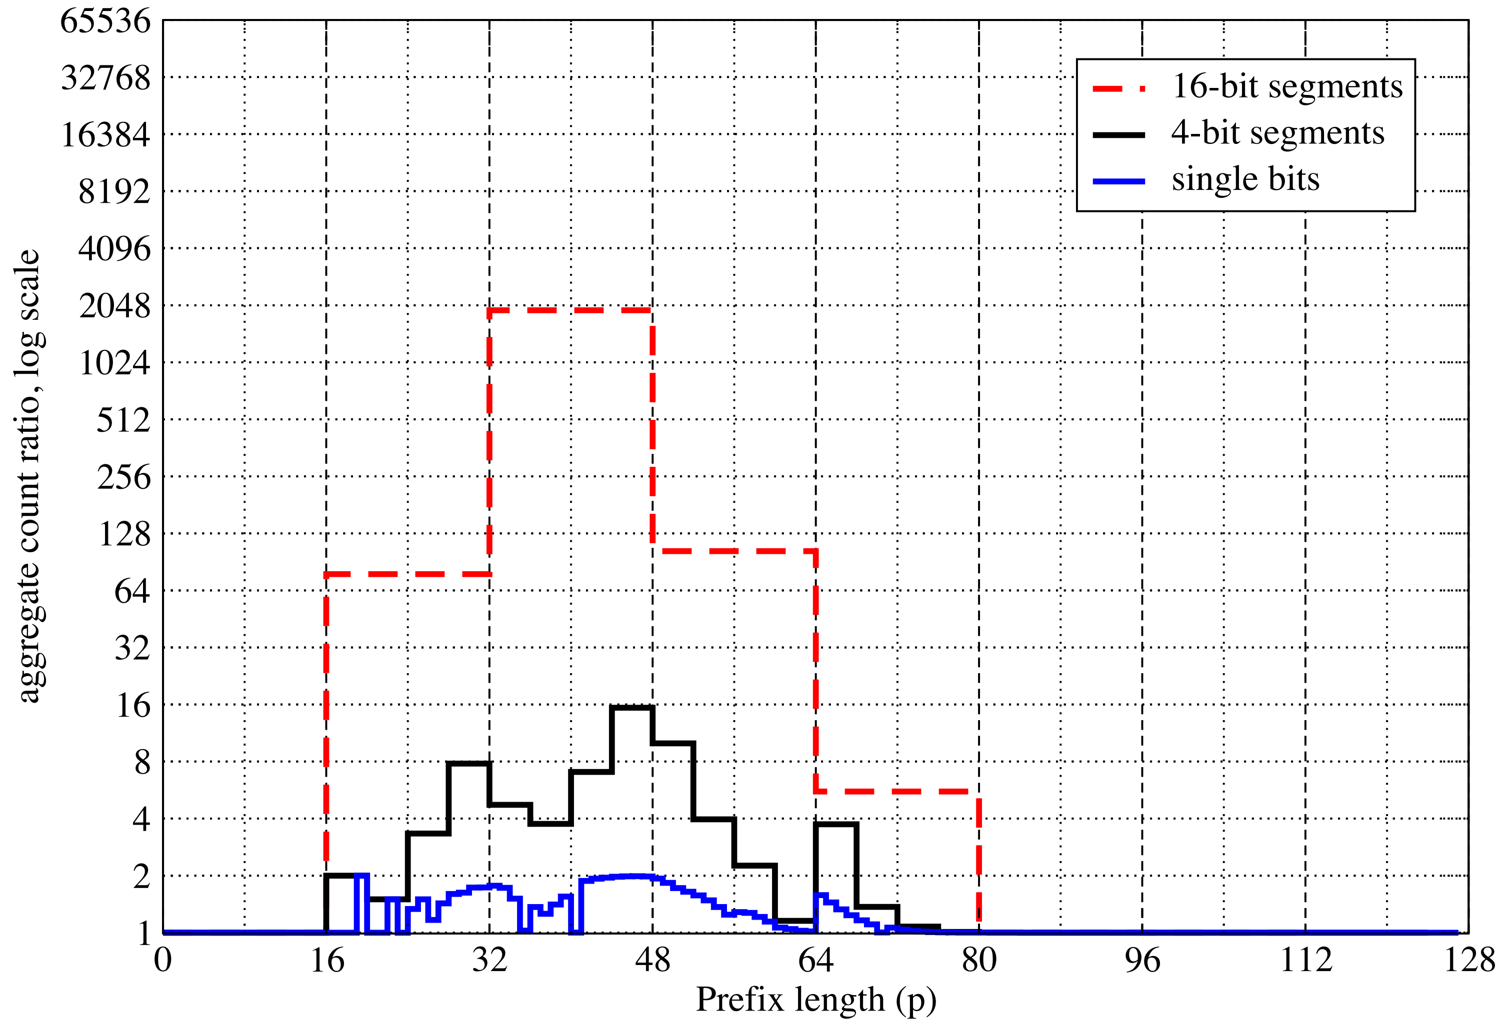

MRA Plot: EU Univ. dept. prefix: 94 active client addrs, 1 /64

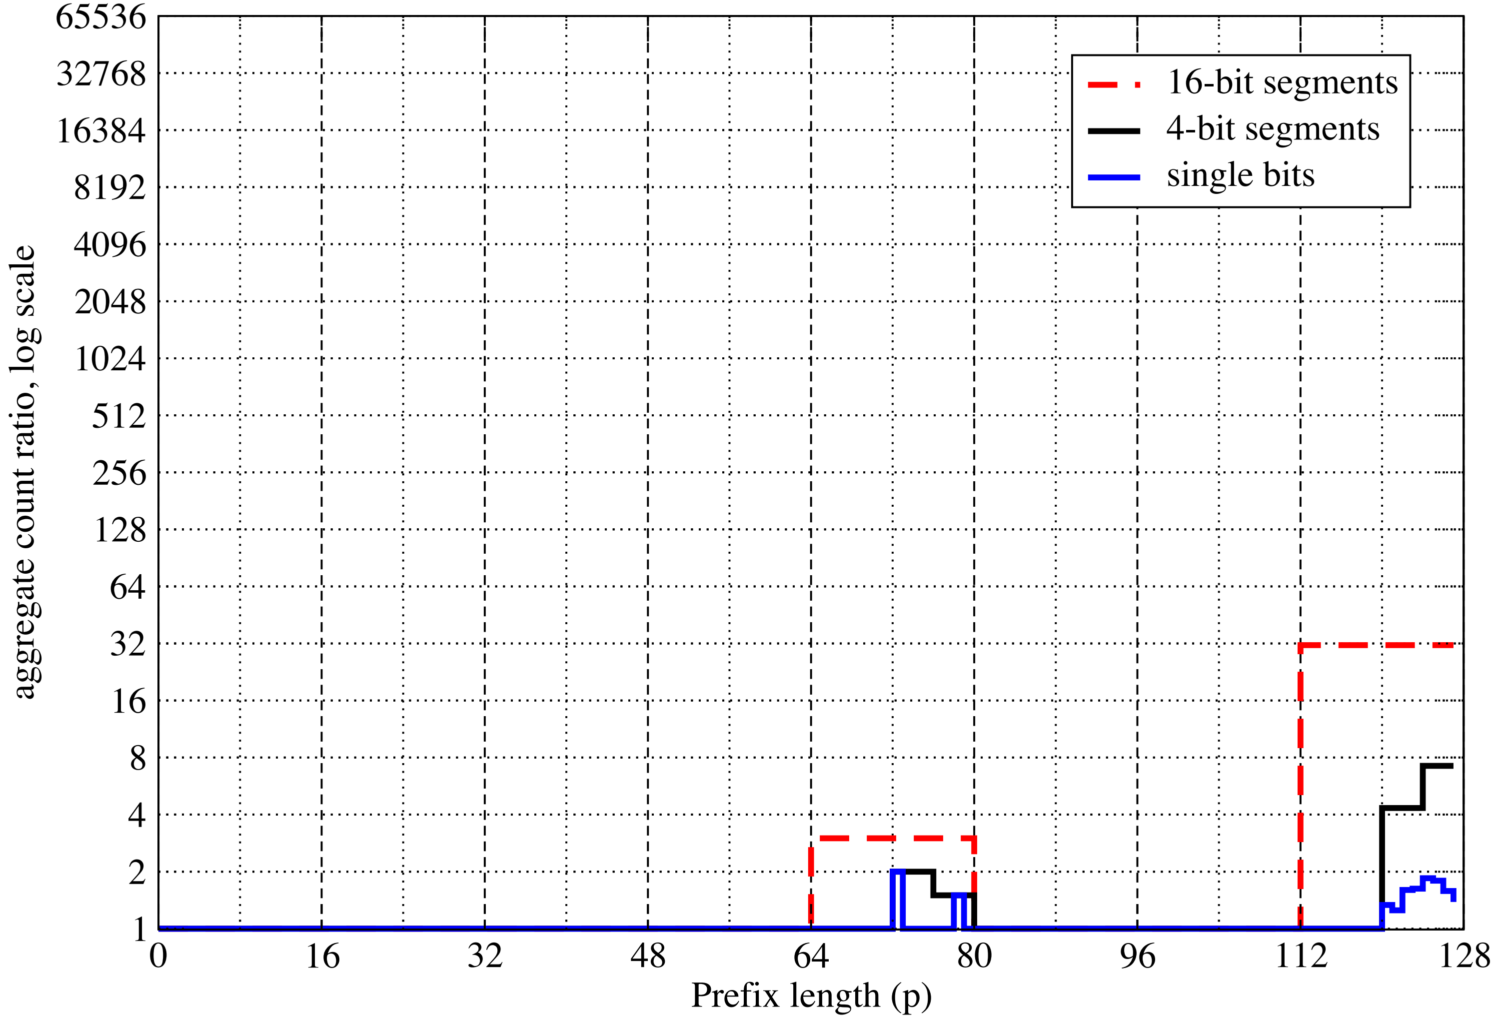

Spatial Address Classification: Density Analysis

- Count the aggregate prefixes for active hosts (at multiple resolutions) and identify “small” dense prefixes that are feasible to probe, scan, or to resolve via `ip6.arpa PTR` queries.

  “How densely-packed are active IPv6 addresses in a given prefix?”

- “What are the set of dense prefixes? Which are `*n*@/*p*-dense’?”

Spatial Address Classification: Density Analysis

- Count the aggregate prefixes for active hosts (at multiple resolutions) and identify “small” dense prefixes that are feasible to probe, scan, or to resolve via `ip6.arpa PTR` queries.

  “How densely-packed are active IPv6 addresses in a given prefix?”

- “What are the set of dense prefixes? Which are `*n@/p*-dense’?”

- **We’ll skip this to save time; see Figure 4 in the paper.**
  I’m happy to talk with you about it offline.

Summary and Current Work

- *Temporal Classification*
  - Feasible for online analysis, effective in **target selection for active probes,** traceroutes
- *Spatial Classification*
  - Feasible for online analysis, effective in **target selection** for DNS lookups `ip6.arpa` PTRs; scans?
- *Current work:* Longest Stable Prefixes and "Prefix Intelligence"
  - Combines aspects of both temporal and spatial classification
- *Current work:* Counting and Visualizing IPv6…
  - Comprehensive survey of the IPv6 Internet *likely must* be informed by address classification.
-

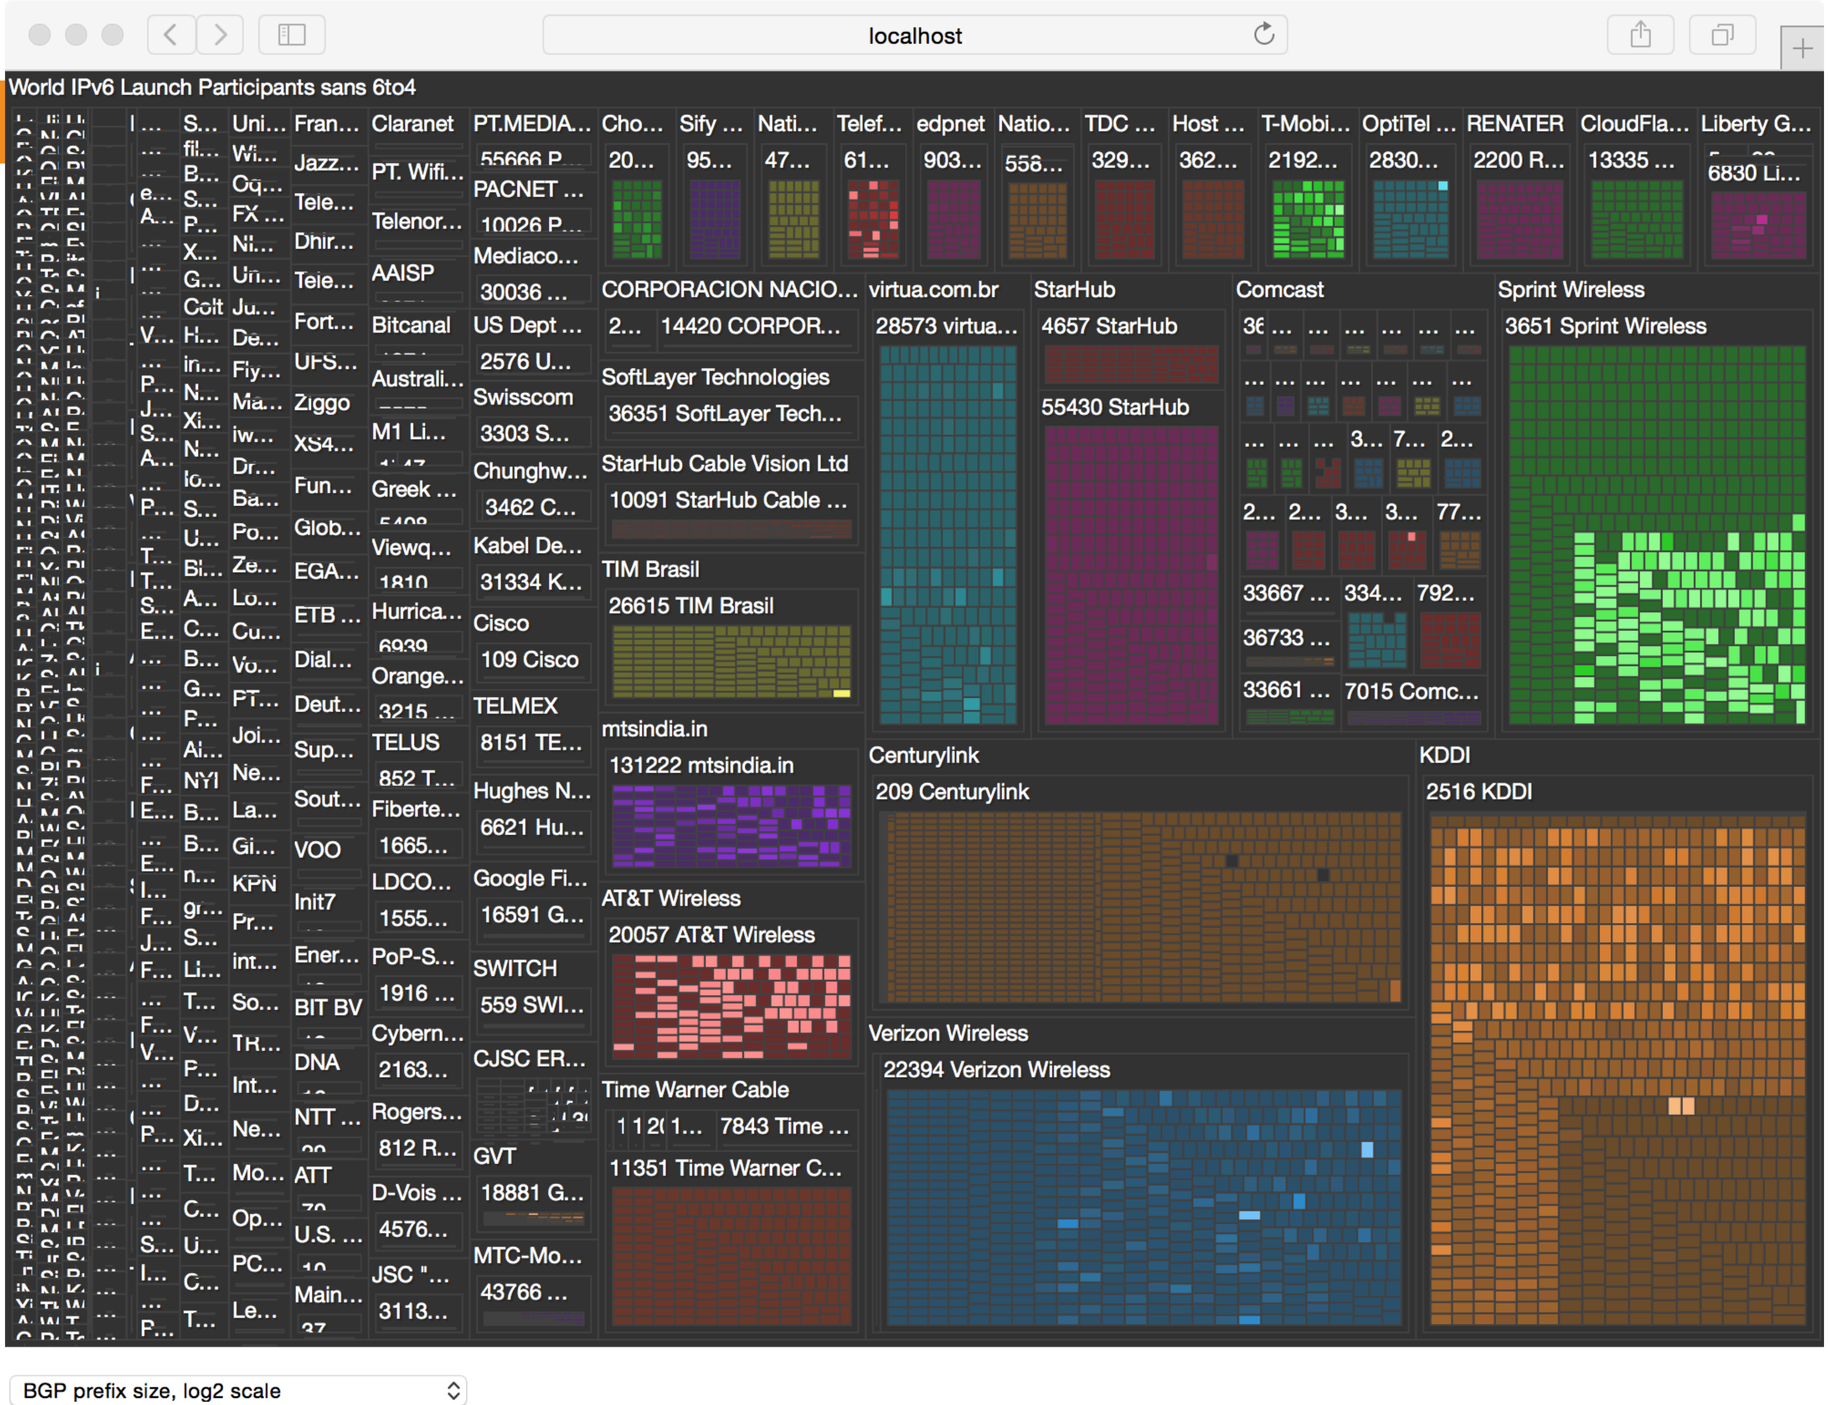

localhost
World IPv6 Launch Participants sans 6to4
Claranet
PT.MEDIA...
PACNET ...
Mediaco...
AAISP
Bitcanal
US Dept ...
Swisscom
Chunghw...
Kabel De...
Cisco
109 Cisco
TELMEX
TELUS
Hughes N...
Google Fi...
SWITCH
CJSC ER...
GVT
MTC-Mo...
edpnet
RENATER
CloudFla...
Liberty G...
CORPORACION NACIO...
14420 CORPOR...
SoftLayer Technologies
36351 SoftLayer Tech...
StarHub Cable Vision Ltd
10091 StarHub Cable ...
TIM Brasil
26615 TIM Brasil
mtsindia.in
131222 mtsindia.in
AT&T Wireless
20057 AT&T Wireless
Time Warner Cable
7843 Time ...
11351 Time Warner C...
virtua.com.br
28573 virtua...
StarHub
4657 StarHub
55430 StarHub
Comcast
7015 Comc...
Sprint Wireless
3651 Sprint Wireless
Centurylink
209 Centurylink
KDDI
2516 KDDI
Verizon Wireless
22394 Verizon Wireless
BGP prefix size, log2 scale

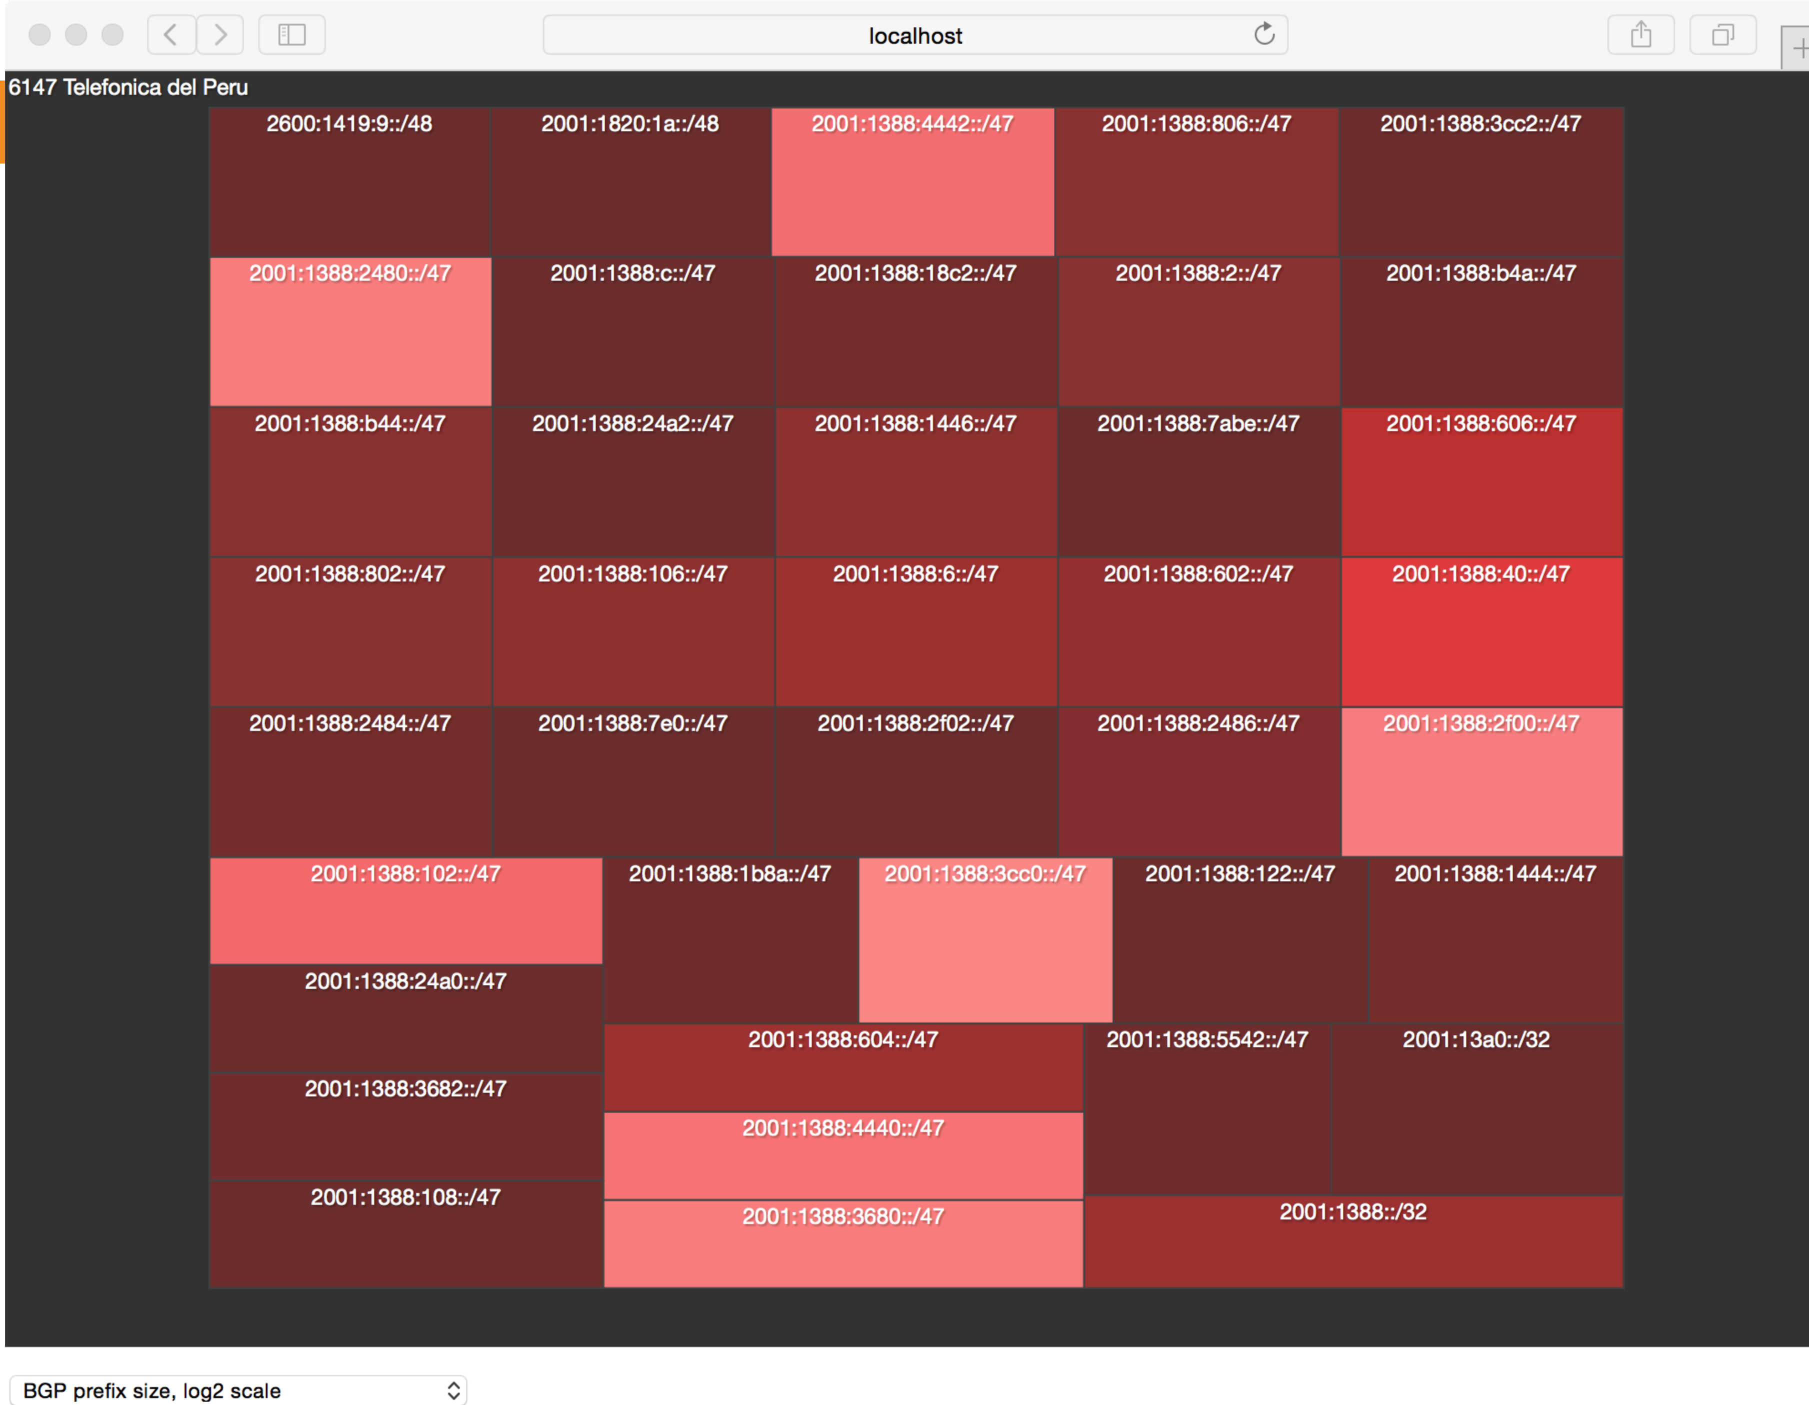
localhost
6147 Telefonica del Peru
2600:1419:9::/48
2001:1820:1a::/48
2001:1388:4442::/47
2001:1388:806::/47
2001:1388:3cc2::/47
2001:1388:2480::/47
2001:1388:c::/47
2001:1388:18c2::/47
2001:1388:2::/47
2001:1388:b4a::/47
2001:1388:b44::/47
2001:1388:24a2::/47
2001:1388:1446::/47
2001:1388:7abe::/47
2001:1388:606::/47
2001:1388:802::/47
2001:1388:106::/47
2001:1388:6::/47
2001:1388:602::/47
2001:1388:40::/47
2001:1388:2484::/47
2001:1388:7e0::/47
2001:1388:2f02::/47
2001:1388:2486::/47
2001:1388:2f00::/47
2001:1388:102::/47
2001:1388:1b8a::/47
2001:1388:3cc0::/47
2001:1388:122::/47
2001:1388:1444::/47
2001:1388:24a0::/47
2001:1388:604::/47
2001:1388:5542::/47
2001:13a0::/32
2001:1388:3682::/47
2001:1388:4440::/47
2001:1388:108::/47
2001:1388:3680::/47
2001:1388::/32
BGP prefix size, log2 scale

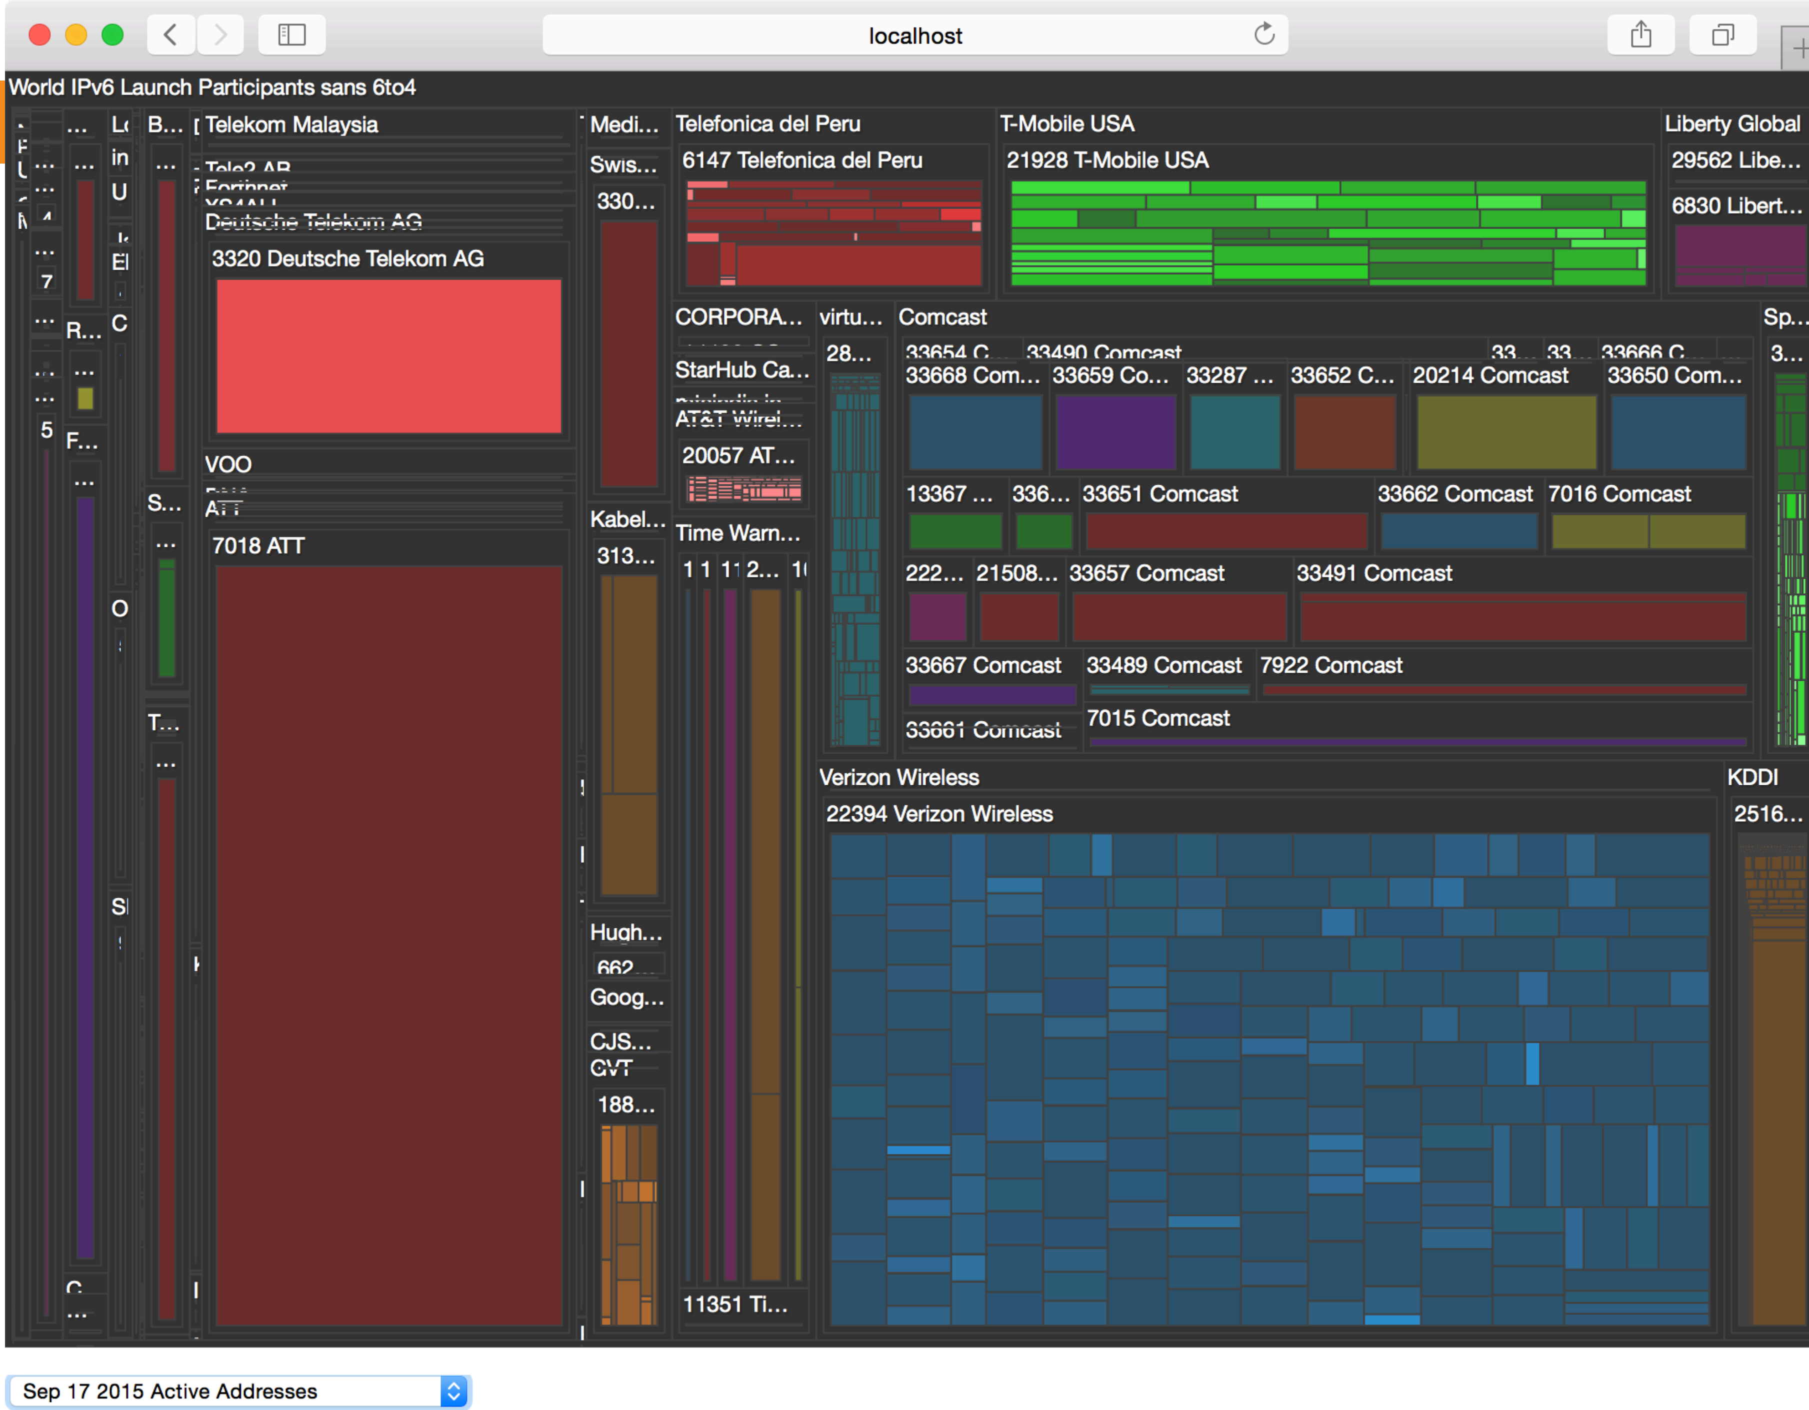

localhost
World IPv6 Launch Participants sans 6to4
Telekom Malaysia
Deutsche Telekom AG
3320 Deutsche Telekom AG
VOO
7018 ATT
Medi…
Swis…
Kabel…
Hugh…
Goog…
CJS…
GVT
Telefonica del Peru
6147 Telefonica del Peru
T-Mobile USA
21928 T-Mobile USA
Liberty Global
29562 Libe…
6830 Libert…
CORPORA…
StarHub Ca…
20057 AT…
Time Warn…
11351 Ti…
Comcast
33668 Com…
33659 Co…
33287 …
33652 C…
20214 Comcast
33650 Com…
33651 Comcast
33662 Comcast
7016 Comcast
33657 Comcast
33491 Comcast
33667 Comcast
33489 Comcast
7922 Comcast
33661 Comcast
7015 Comcast
Verizon Wireless
22394 Verizon Wireless
KDDI
2516…
Sep 17 2015 Active Addresses

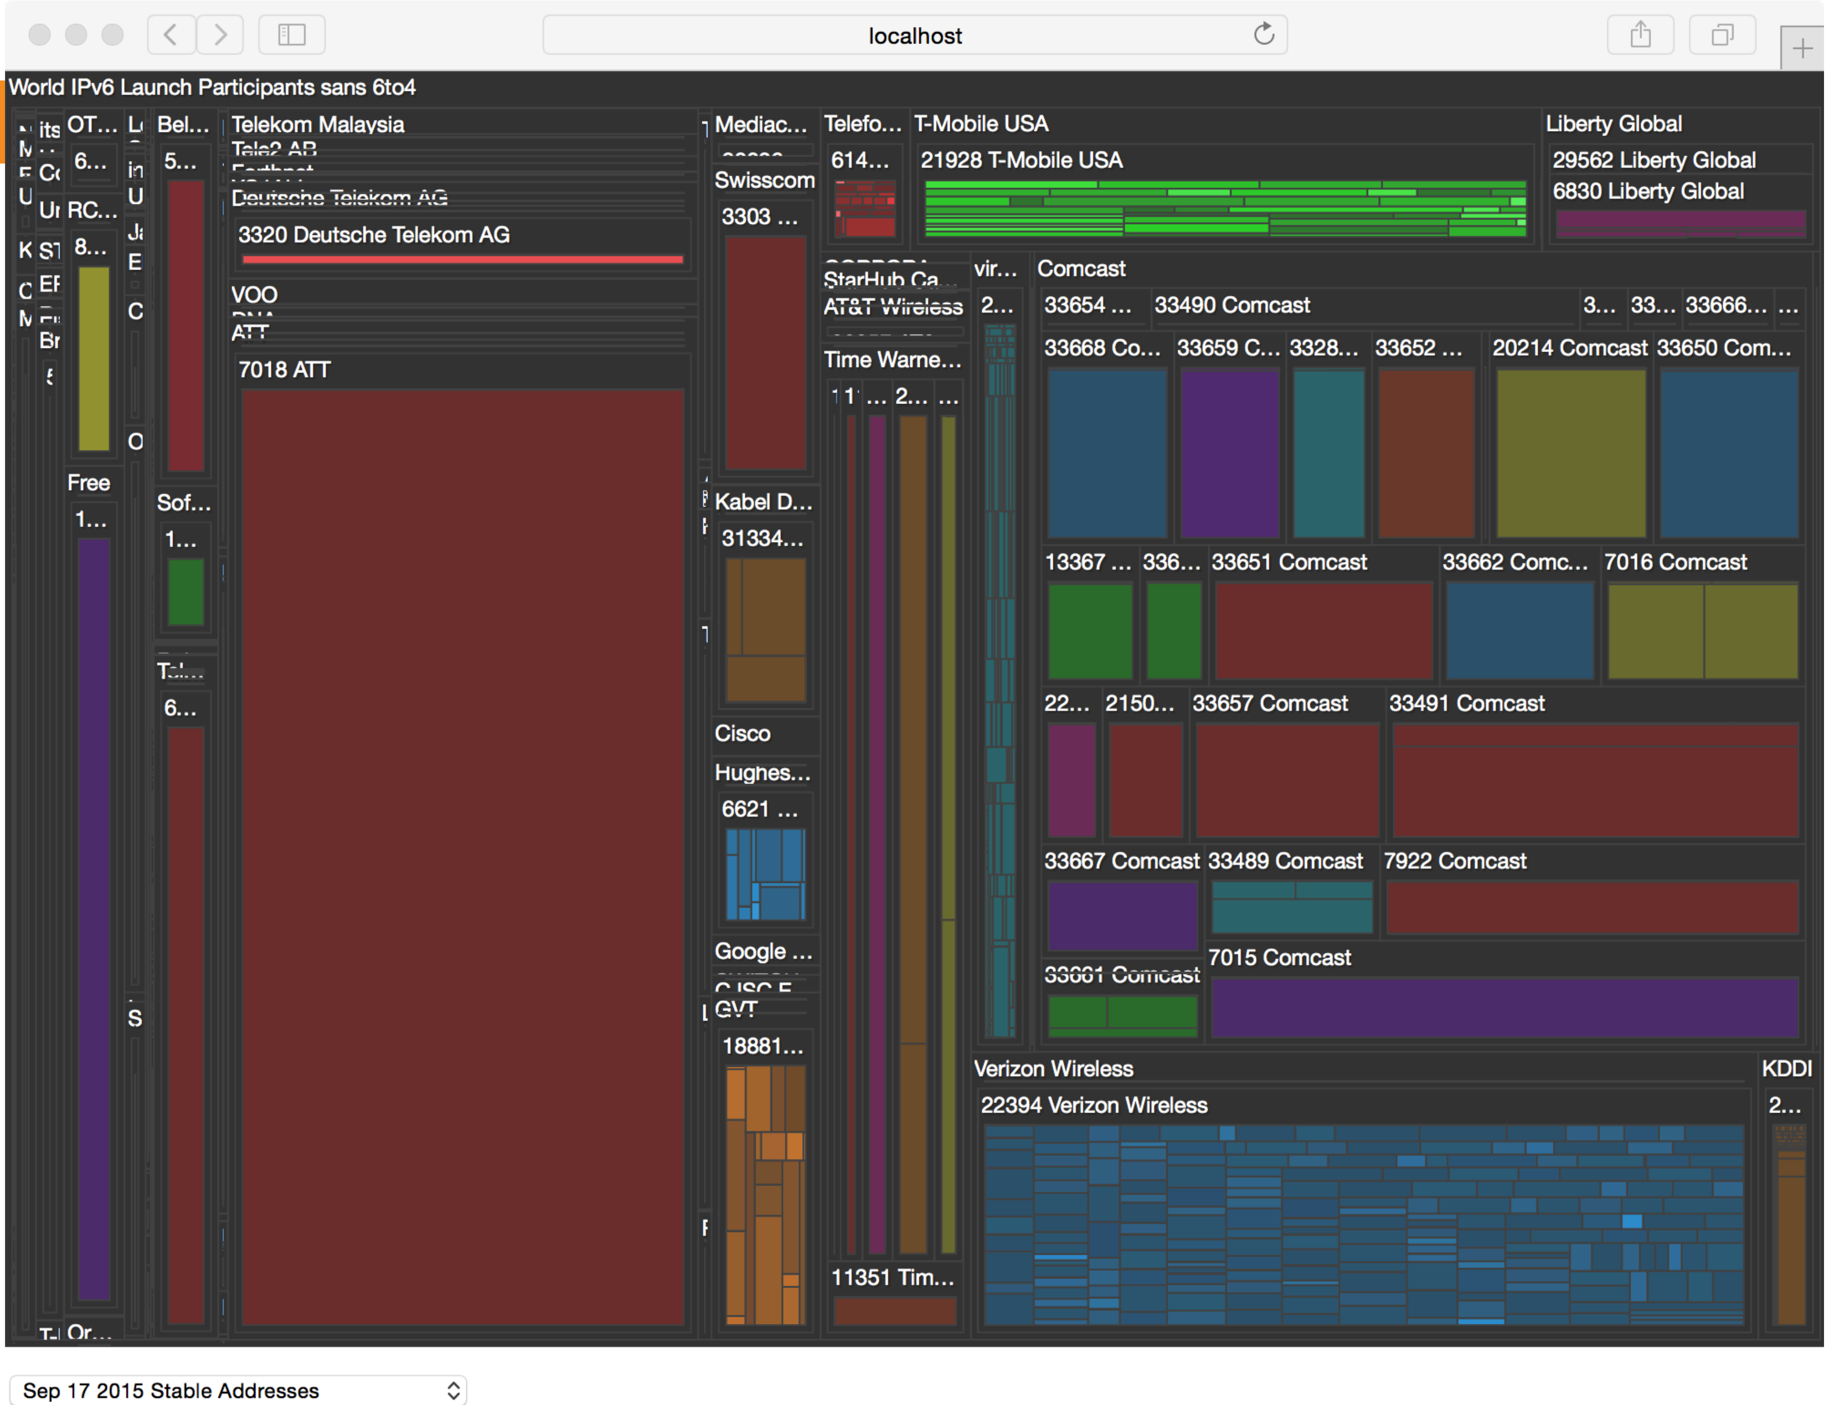
localhost
World IPv6 Launch Participants sans 6to4
Telekom Malaysia
Tele2 AB
Deutsche Telekom AG
3320 Deutsche Telekom AG
VOO
ATT
7018 ATT
Free
Sof...
Mediac...
Swisscom
3303 ...
Kabel D...
31334...
Cisco
Hughes...
6621 ...
Google ...
GVT
18881...
Telefo...
614...
T-Mobile USA
21928 T-Mobile USA
Liberty Global
29562 Liberty Global
6830 Liberty Global
StarHub Ca...
AT&T Wireless
Time Warne...
11351 Tim...
vir...
Comcast
33654 ...
33490 Comcast
33668 Co...
33659 C...
33652 ...
20214 Comcast
33650 Com...
13367 ...
33651 Comcast
33662 Comc...
7016 Comcast
33657 Comcast
33491 Comcast
33667 Comcast
33489 Comcast
7922 Comcast
33661 Comcast
7015 Comcast
Verizon Wireless
22394 Verizon Wireless
KDDI
Sep 17 2015 Stable Addresses


localhost
World IPv6 Launch Participants sans 6to4
CloudFlare, Inc.
13335 CloudFlare, Inc.
StarHub
55430 StarHub
Comcast
7015 Comcast
Centurylink
Verizon Wireless
KDDI
Cisco
109 Cisco
Swisscom
TDC Sv...
Host Virt...
SoftLayer Tech...
StarHub Cable...
AT&T Wireless
20057 AT&T ...
Time Warner C...
TELMEX
SWITCH
Telekom...
Claranet
FranTec...
Sep 17 2015 % Stable Addresses

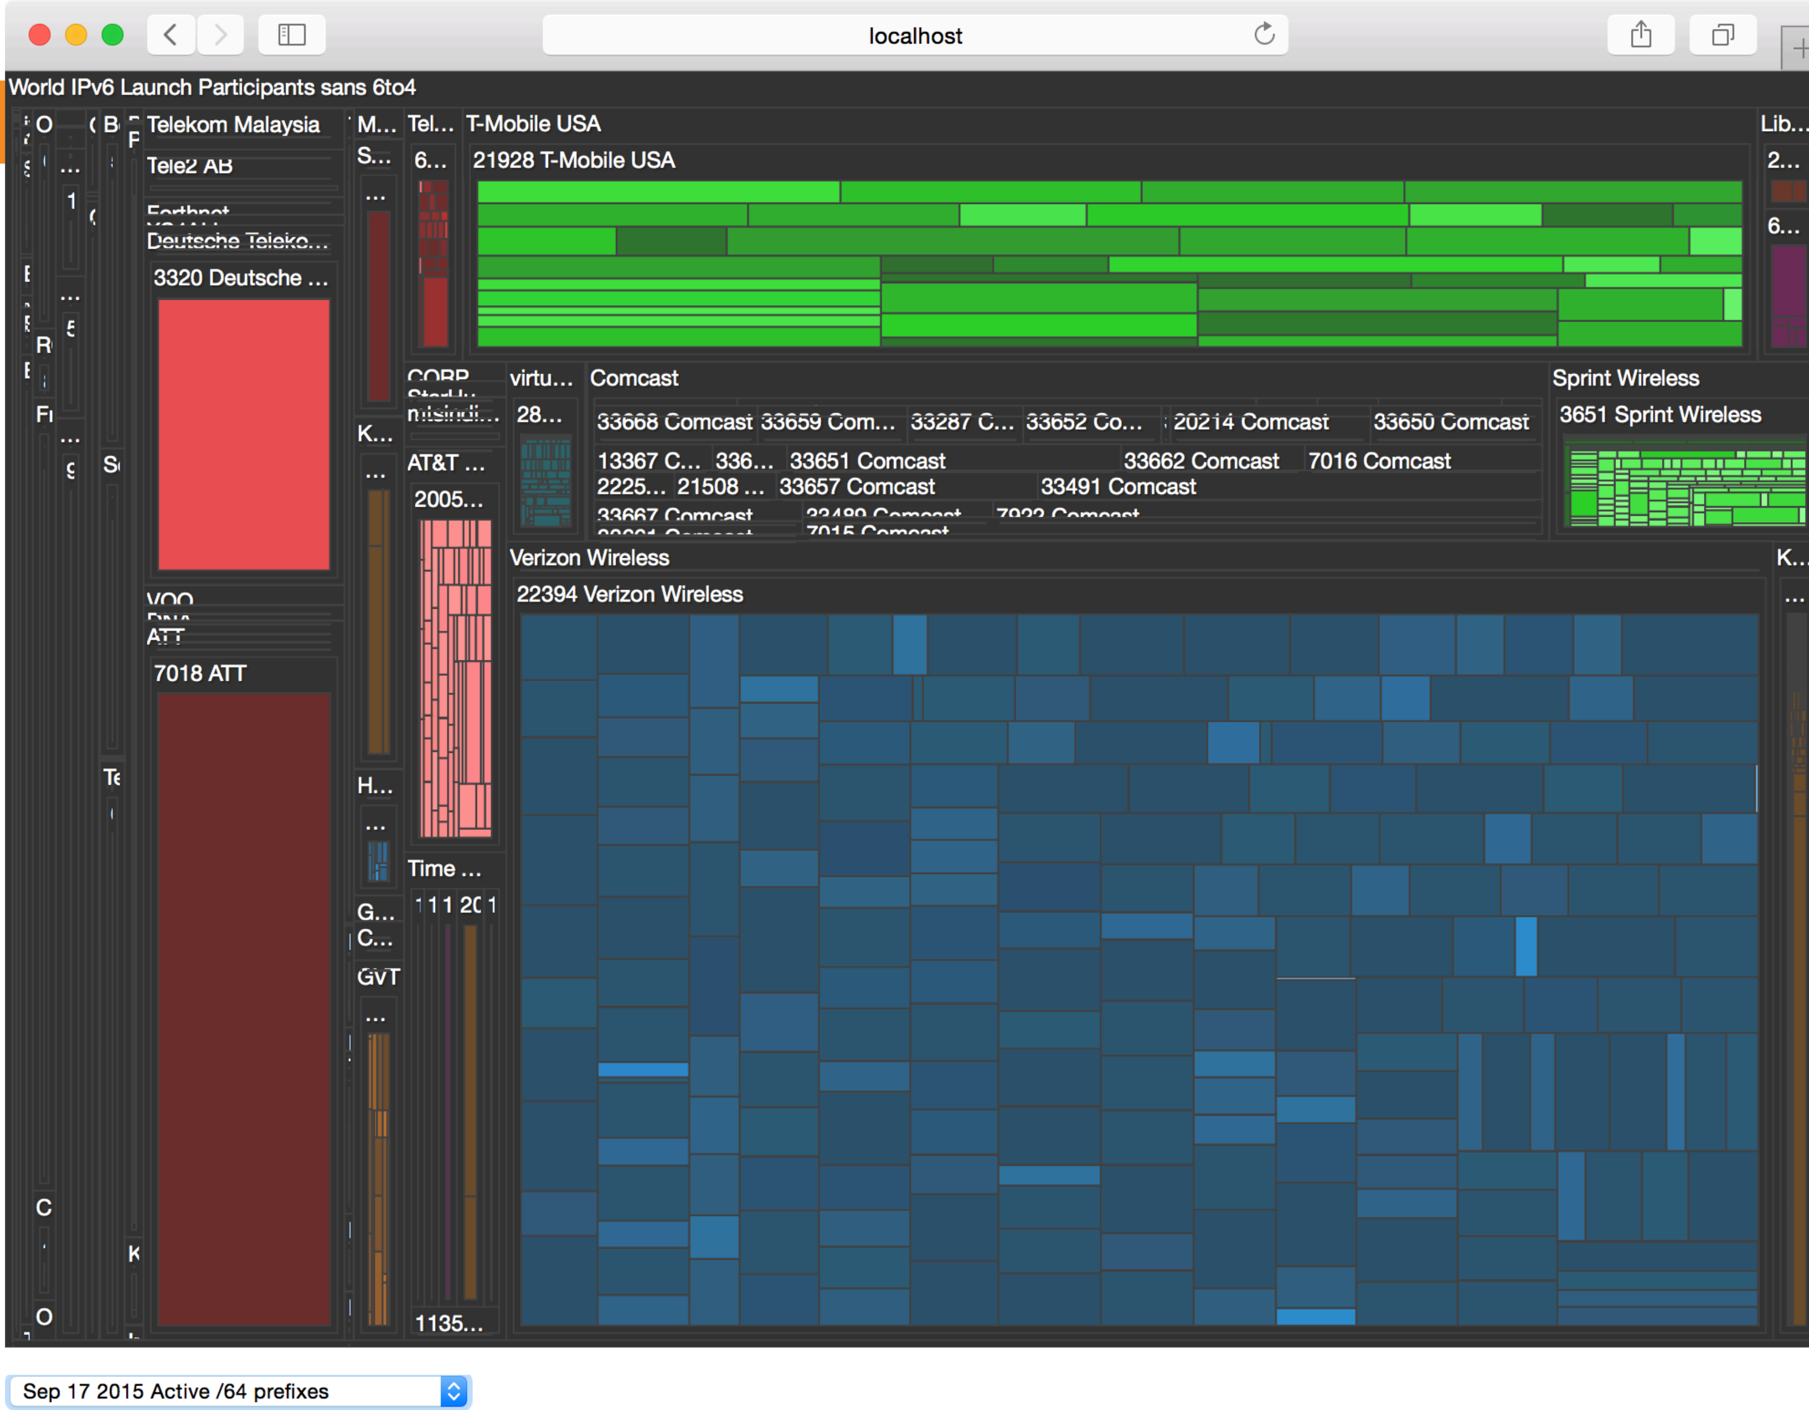
localhost
World IPv6 Launch Participants sans 6to4
Telekom Malaysia
Tele2 AB
Deutsche Teleko…
3320 Deutsche …
VOO
ATT
7018 ATT
M…
S…
K…
H…
G…
C…
GVT
Tel…
6…
AT&T …
2005…
Time …
1135…
T-Mobile USA
21928 T-Mobile USA
virtu…
28…
Comcast
33668 Comcast
33659 Com…
33287 C…
33652 Co…
20214 Comcast
33650 Comcast
13367 C…
336…
33651 Comcast
33662 Comcast
7016 Comcast
2225…
21508 …
33657 Comcast
33491 Comcast
33667 Comcast
Sprint Wireless
3651 Sprint Wireless
Verizon Wireless
22394 Verizon Wireless
Lib…
2…
6…
K…
Sep 17 2015 Active /64 prefixes


localhost
World IPv6 Launch Participants sans 6to4
Chin...
41...
Telekom Malaysia
4788 Telekom Malaysia
Deutsche Telekom AG
3320 Deutsche Telekom AG
ATT
T-Mobile USA
21928 T-Mobile USA
mtsindia.in
AT&T Wireless
20057 AT&T Wireless
Sprint Wireless
3651 Sprint Wireless
Verizon Wireless
22394 Verizon Wireless
Sep 17 2015 Transient /64 prefixes

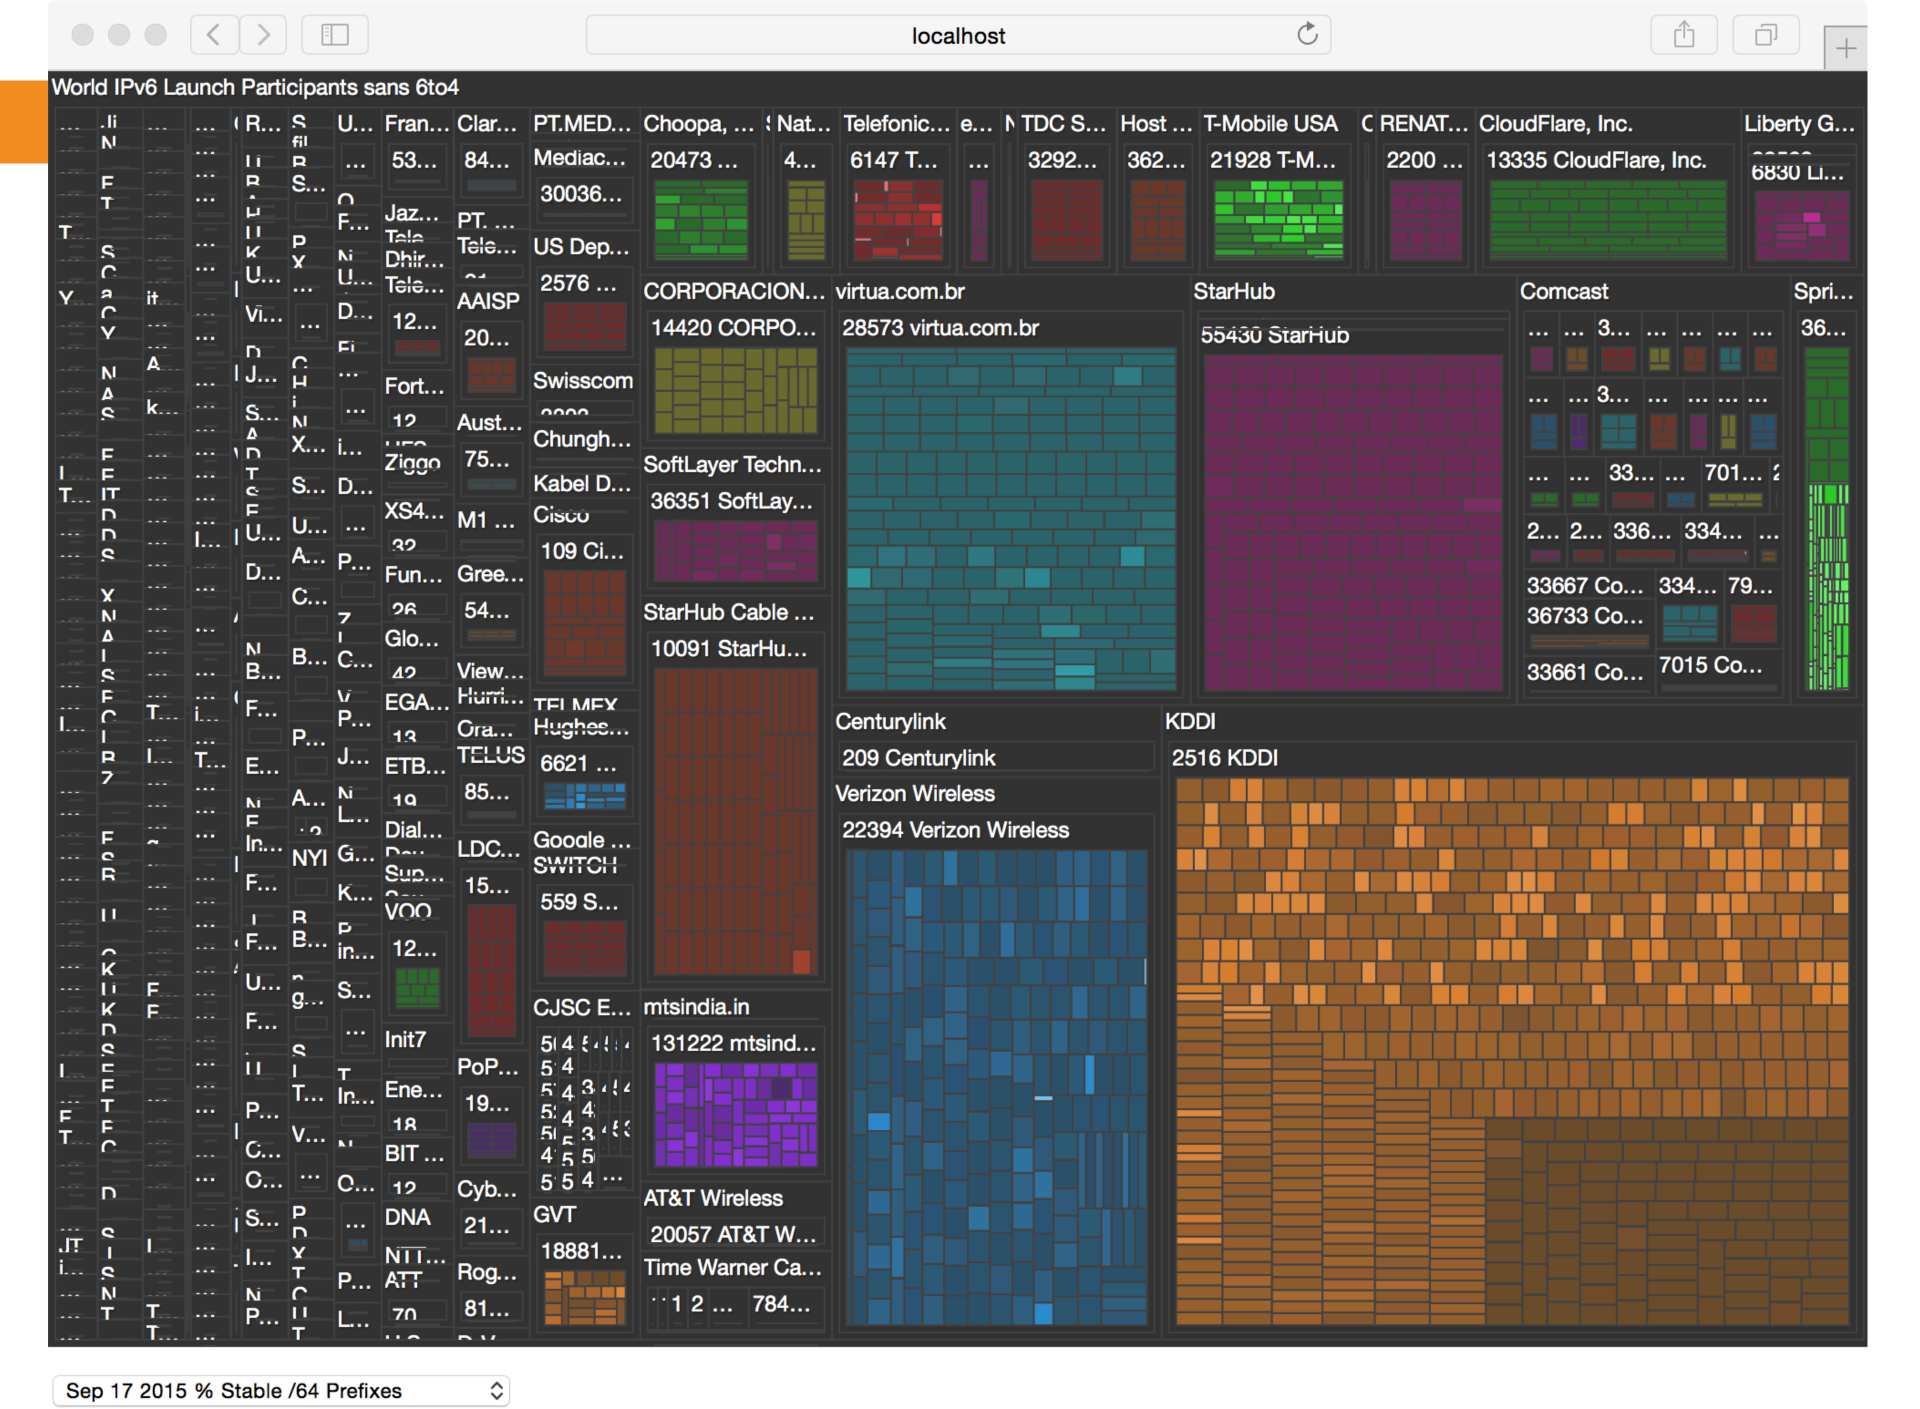

localhost
World IPv6 Launch Participants sans 6to4
PT.MED...
Mediac...
30036...
US Dep...
2576 ...
Swisscom
Chungh...
Kabel D...
Cisco
109 Ci...
Choopa, ...
20473 ...
Telefonic...
6147 T...
TDC S...
3292...
Host ...
362...
T-Mobile USA
21928 T-M...
RENAT...
2200 ...
CloudFlare, Inc.
13335 CloudFlare, Inc.
Liberty G...
6830 Li...
CORPORACION...
14420 CORPO...
SoftLayer Techn...
36351 SoftLay...
StarHub Cable ...
10091 StarHu...
virtua.com.br
28573 virtua.com.br
StarHub
55430 StarHub
Comcast
33667 Co...
36733 Co...
33661 Co...
7015 Co...
Spri...
AAISP
Ziggo
XS4...
Fun...
Gree...
Glo...
View...
EGA...
TELUS
ETB...
Hughes...
6621 ...
Google ...
SWITCH
559 S...
LDC...
VOO
Init7
CJSC E...
mtsindia.in
131222 mtsind...
AT&T Wireless
20057 AT&T W...
Time Warner Ca...
GVT
18881...
PoP...
Ene...
BIT ...
Cyb...
DNA
ATT
Rog...
Centurylink
209 Centurylink
Verizon Wireless
22394 Verizon Wireless
KDDI
2516 KDDI
Sep 17 2015 % Stable /64 Prefixes

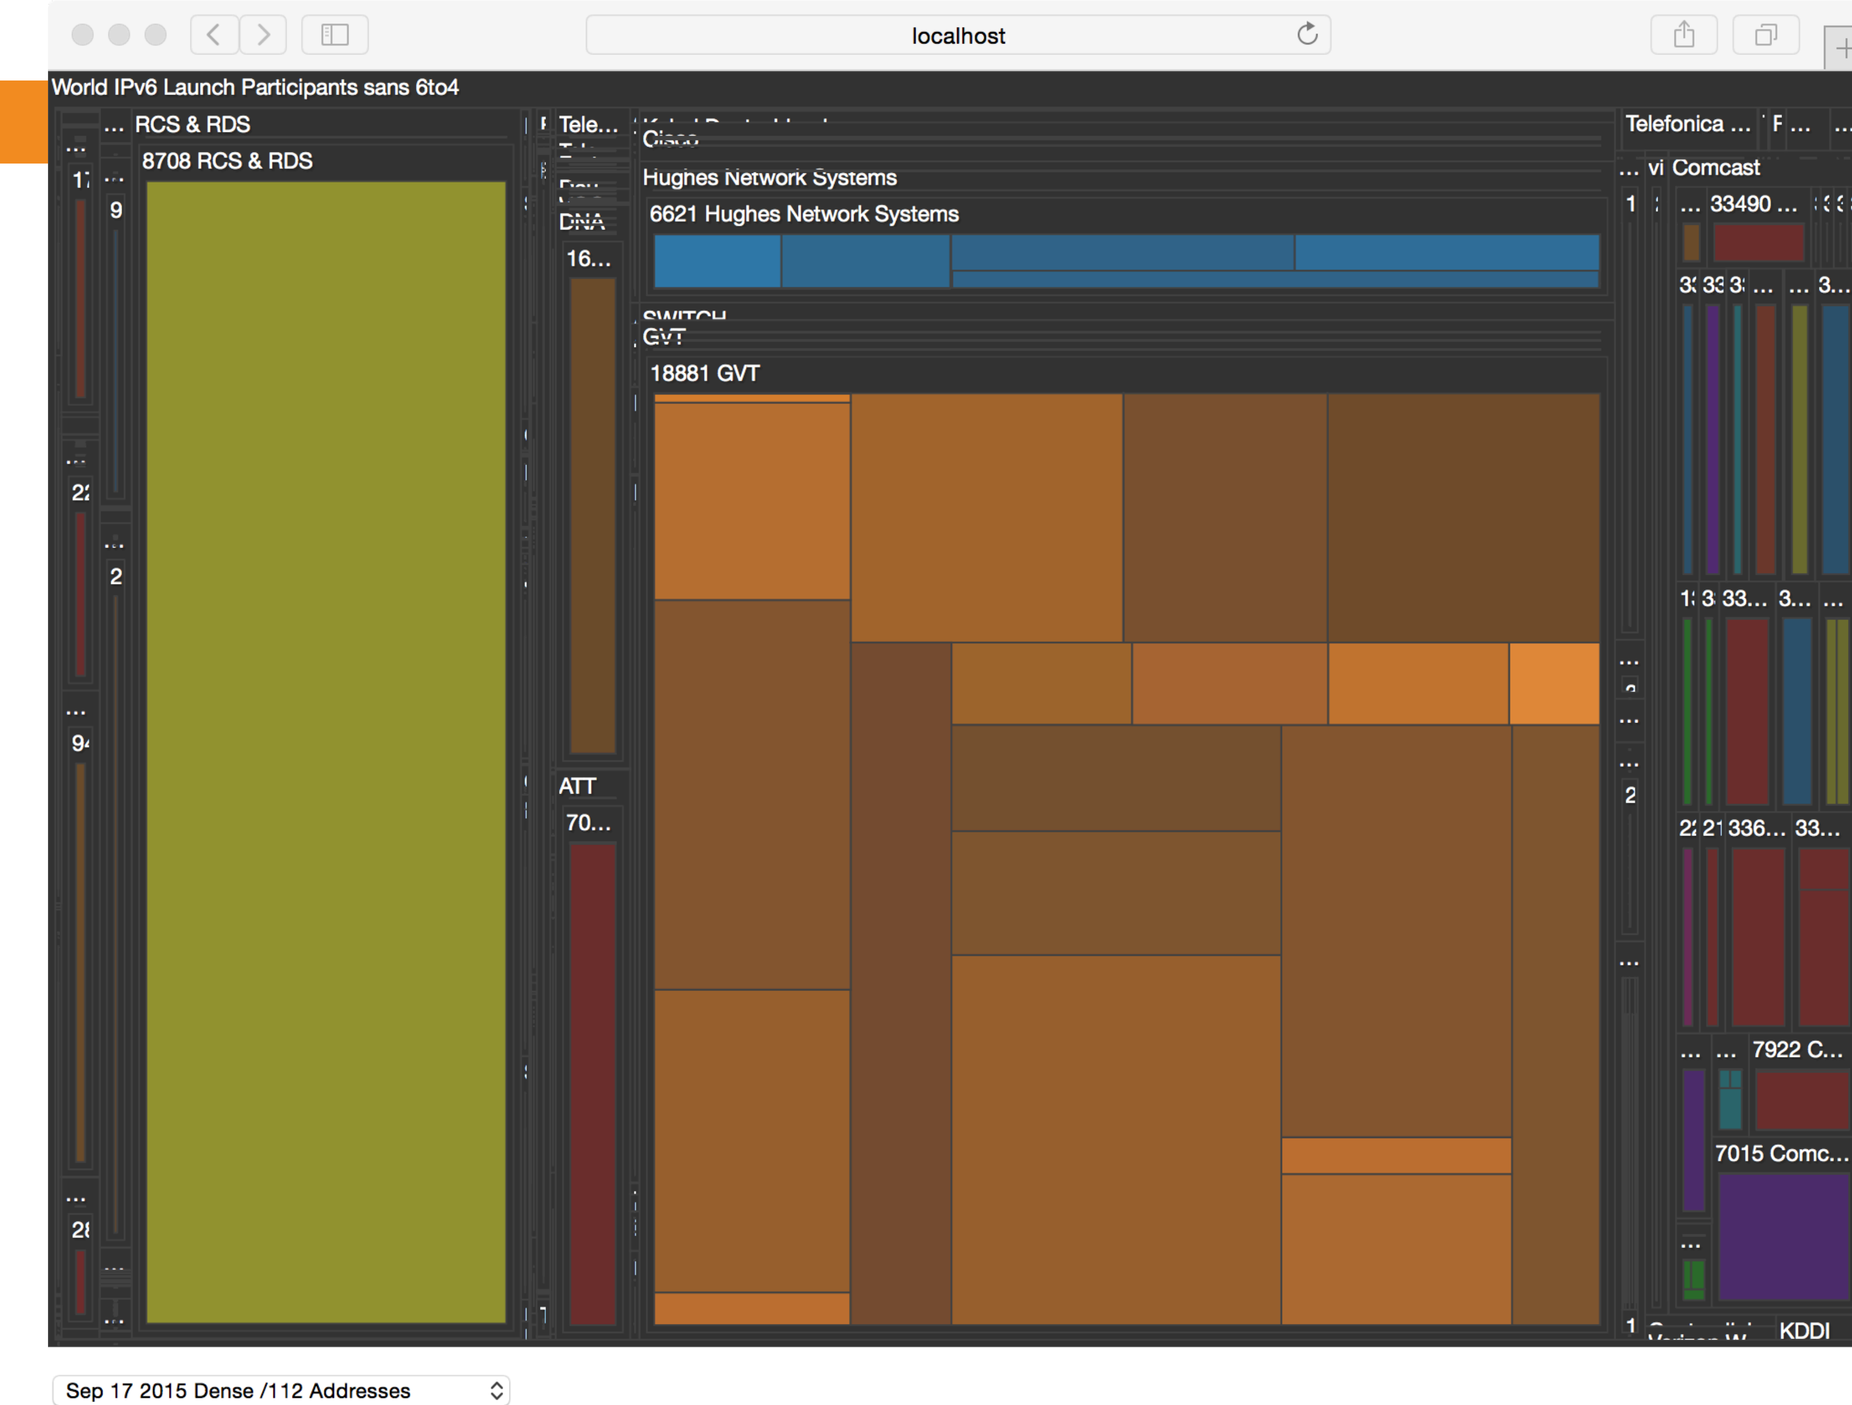
localhost
World IPv6 Launch Participants sans 6to4
RCS & RDS
8708 RCS & RDS
Tele…
DNA
16…
ATT
70…
Cisco
Hughes Network Systems
6621 Hughes Network Systems
GVT
18881 GVT
Telefonica …
Comcast
33490 …
7922 C…
7015 Comc…
KDDI
Sep 17 2015 Dense /112 Addresses

Sample IPv6 Active WWW Client Address Aggregate Counts Data
http://www.akamai.com/technical-publications

Thanks!
Questions? Comments?

David Plonka plonka@akamai.com

The subsequent slides are supporting slides…

Spatial Address Classification

- Pertinent prior works:
  - K. Sklower, ***"A Tree-Based Packet Routing Table for Berkeley Unix"*** [USENIX 1991]
  - K. Cho, R. Kaizaki, and A. Kato. ***"Aguri: An Aggregation-Based Traffic Profiler"*** [QoFIS 2001]
  - E. Kohler, J. Li, V. Paxson, and S. Shenker. ***"Observed Structure of Addresses in IP Traffic"*** [IMC 2002]

Alternate Densification Method: “*n@/p*-dense”

IPv6 addresses in presentation format:

```
2001:db8:0:1cdf:21e:c2ff:fec0:11db
2001:db8:10:1::103
2001:db8:167:1109::10:901
2001:db8:4137:9e76:3031:f3fd:bbdd:2c2a
```

In “hex_ip” format:

```
20010db800001cdf021ec2fffec011db
20010db8001000010000000000000103
20010db801671109000000000100901
20010db841379e763031f3fdbbdd2c2a
```

Alternate Densification Method: *n@/p*-dense

IPv6 addresses in presentation format:

```
2001:db8:0:1cdf:21e:c2ff:fec0:11db
2001:db8:10:1::103
2001:db8:167:1109::10:901
2001:db8:4137:9e76:3031:f3fd:bbdd:2c2a
```

In “hex_ip” format:

```
20010db800001cdf021ec2fffec011db
20010db8001000010000000000000103
20010db8016711090000000000100901
20010db841379e763031f3fdbbdd2c2a
```

`$ sort [-m] |cut -c1-$((`*`p`*`/4)) |uniq -c`

Dense Prefix Discovery from 3.2M router addresses

Density Class	Dense Prefixes	Router Addresses	Possible Addresses	Router Address Density
2 @ /124	43.1K	116K	689K	0.1678459119
3 @ /120	8.28K	81.0K	**2.12M**	0.0382372758
2 @ /120	64.2K	193K	16.4M	0.0117351137
2 @ /116	207K	568K	852M	0.0006670818
64 @ /112	187	41.2K	12.3M	0.0033593815
32 @ /112	509	54.8K	33.4M	0.0016417438
16 @ /112	3.06K	105K	201M	0.0005259994
8 @ /112	21.5K	290K	1.41B	0.0002057970
4 @ /112	101K	681K	6.63B	0.0001026403
2 @ /112	367K	1.29M	24.1B	0.0000534072
2 @ /108	289K	1.72M	303B	0.0000056895
2 @ /104	108K	1.84M	1.81T	0.0000010171

Address Density Applied: ip6.arpa PTR queries

```
2001:728:0:4000::3d ae-25.r02.amstnl02.nl.bb.gin.ntt.net.
2001:268:f204:535::1 sjk-TPChorinochiML13.v6.kddi.ne.jp.
2001:4c08:2003::f1 ge-6-17.car2.Berlin1.Level3.net.
2001:1690:0:e::21 pc21.hgl2art02.previder.net.
2001:5a0:40::32 if-9-0-0.core5.MLN-Miami.ipv6.as6453.net.
2001:688:0:3:7::51 xe2-1-10.marcr3.Marseille.opentransit.net.
2001:5000:0:d5::2 ae7-xcr1.bkl.cw.net.
2600:5:2000:4::22 sl-crs1-sj-be250.v6.sprintlink.net.
2001:418:0:1000::e r22.snjsca04.us.bb.gin.ntt.net.
2001:558:180:145::2 te-4-3-ur01.bluefield.wv.richmond.comcast.net.
2001:558:82:2ee::2 ge-1-19-ur02.sf19th.ca.sfba.comcast.net.
2001:1890:ff:ffff:12:122:28:53 sffca22crs.ipv6.att.net.
2a03:7c00:0:6c69:6e6b:6164:6472:1 kau-toh-l.ip6.kasenet.fi.
2001:4478:4000:10::1 po4-100.ade-pipe-bdr1.iinet.net.au.
2001:da8:c803:ffff:ffff:ffff:11:4f cernet2.net.
2607:f390:0:3000::16 ULL-G2-TE-5-2.loni.org.
...
```

Aguri Tree-based Densification Example

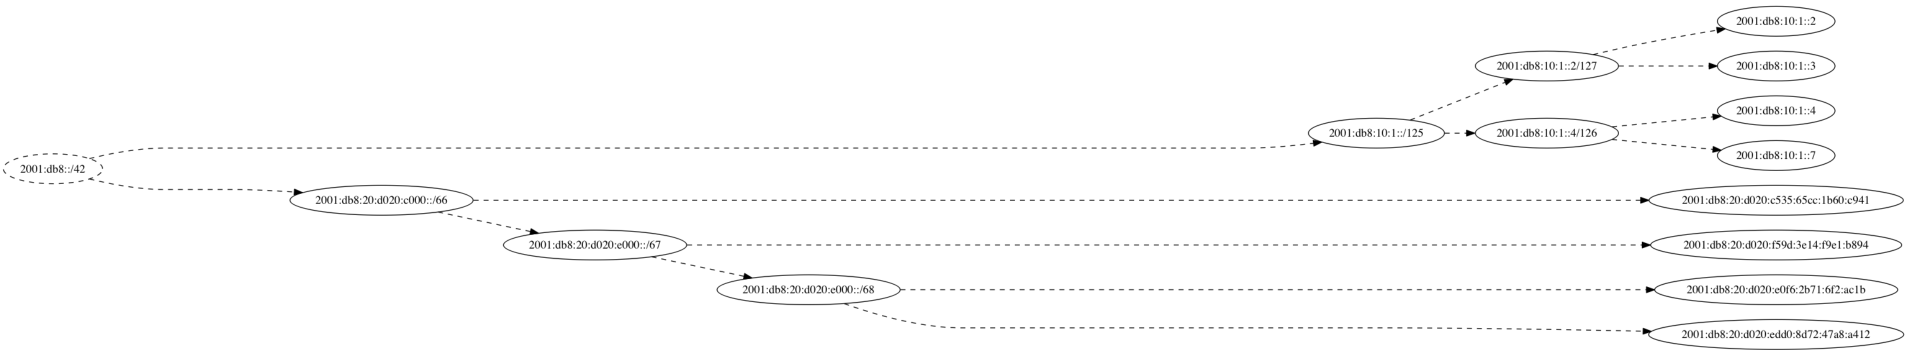

Densification Step 1: populate with addresses.

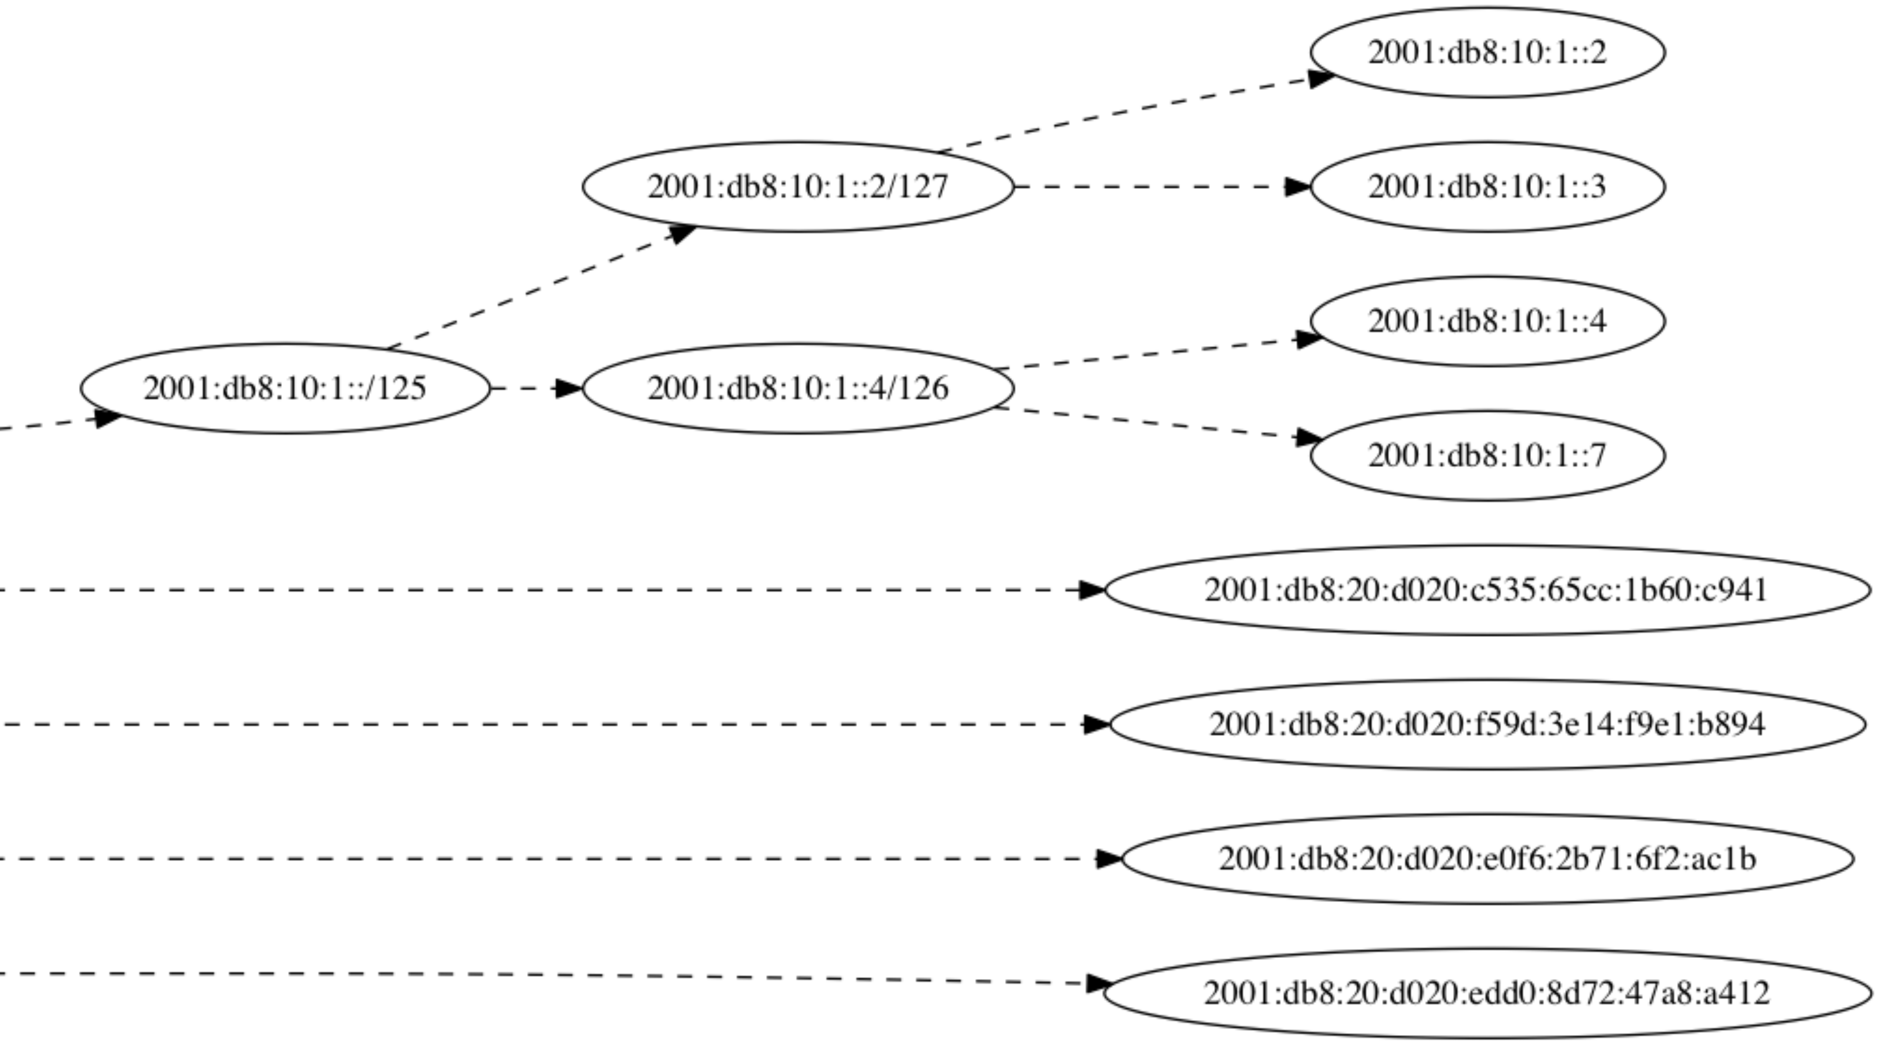

Densification Step 1: add addresses.

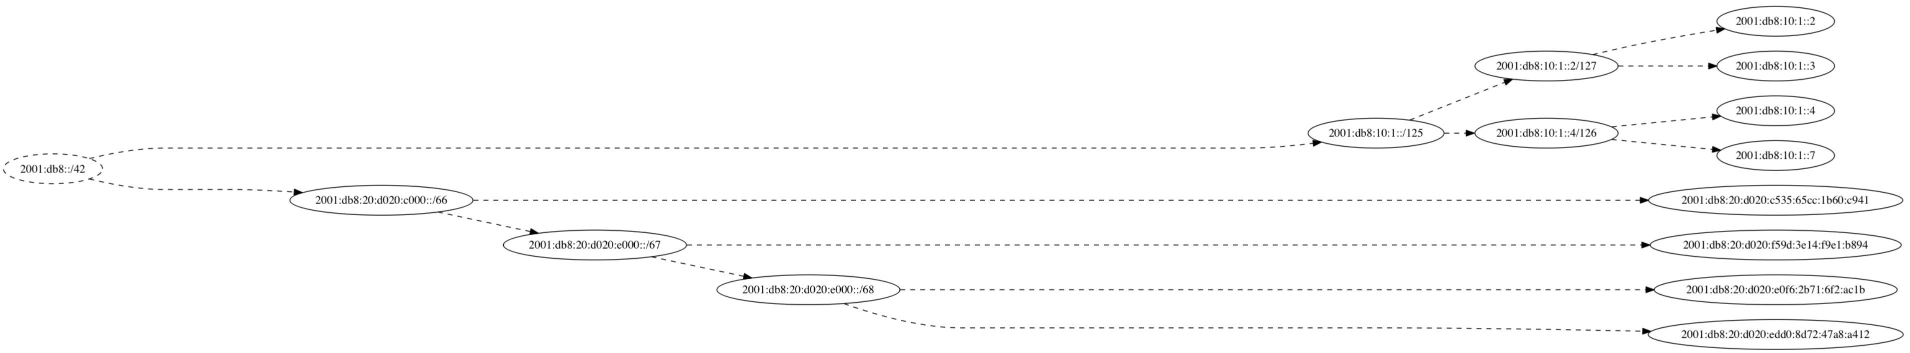

Densification Step 2: “densify” to $2/2^{24}$.

Densification Step 2: densify to 2@104.

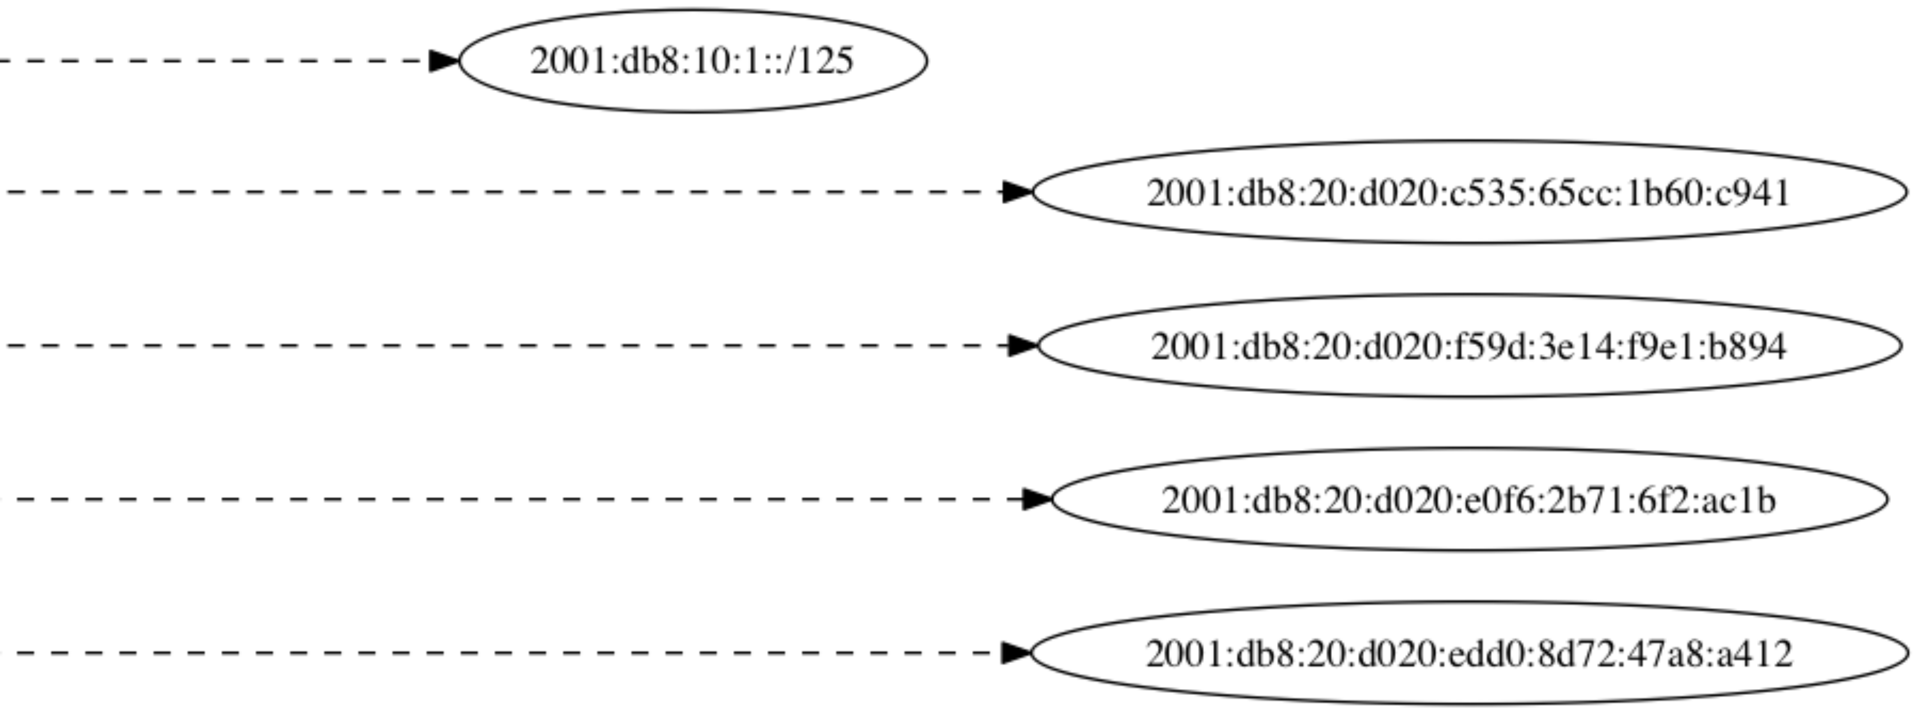

Densification Step 3: “prune” singletons.

Densification Result: 1 dense, 2 sparse prefixes.