



The Dataset

- Active Probing
 - From 100+ geographic locations
 - Application Layer: HTTP/DNS availability, load times
 - Network layer end-to-end: latency, loss, jitter
 - Network layer hops: L3 hop-by-hop measurements on latency, loss and capturing connectivity and changes
- Routing
 - BGP data from RouteViews, RIPE collectors

Internet Outages Happen All the Time




~ 170 affected interfaces / hour



~ 1.6K prefixes / hour

Let's Start with 3 Events from 2016

DNS Root DDoS



June 26, 2016

Submarine Cable Fault



May 17, 2016

AWS Route Leak



April 22, 2016


DNS Root DDoS

https://en.wikipedia.org/wiki/Root_name_server#/media/File:Ams-ix.k.root-servers.net.jpg

DNS Root Server DDoS


The Attack

- June 25th 2016
2:45-5:50 PDT
(21:45-0:50 UTC)
- All DNS roots affected
- TCP SYN and ICMP flood
- 10M packets/sec,
17Gbps per root



With Real World Impact

Impacts




- 50% loss in availability
- 3.7ms → 13ms response time
- Based on 273 measurements every 5 mins


The Varying Impact on Root Servers

Map	Servers	Agents	
Server	Date (UTC)	Number of Errors ↓	Avg. Resolution Time (ms)
h.root-servers.net	2016-06-25 21:51:54	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 21	
b.root-servers.net	2016-06-25 21:51:23	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 20	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 122
g.root-servers.net	2016-06-25 21:51:49	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 18	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 136
c.root-servers.net	2016-06-25 21:51:43	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 13	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 50
m.root-servers.net	2016-06-25 21:51:51	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 12	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 132
d.root-servers.net	2016-06-25 21:52:20	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 11	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 44
e.root-servers.net	2016-06-25 21:51:45	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 10	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 35
i.root-servers.net	2016-06-25 21:51:51	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 10	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 103
k.root-servers.net	2016-06-25 21:51:54	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 8	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 174
f.root-servers.net	2016-06-25 21:51:18	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 5	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 32
a.root-servers.net	2016-06-25 21:51:38	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 3	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 404
j.root-servers.net	2016-06-25 21:51:30	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 2	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 30
l.root-servers.net	2016-06-25 21:51:32	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 2	<div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div><div></div></div> 37

Corroborated from RIPE Atlas DNSMON Data



Impact Not As Widespread as Dec 2015 DDoS



H Root Server – Two Anycast Sites

Root Servers


A B C D E F G **H** I J K L M


Operator:

U.S. Army Research Lab

Locations:

Sites: 2

 Aberdeen Proving Ground, US

 San Diego, US

IPs:

IPv4: 198.97.190.53

IPv6: 2001:500:1::53

ASN:

1508

B Root Server – One Anycast Site

Root Servers


A**B**CDEFGHIJKLM

Operator:

Information Sciences Institute

Locations:

Sites: 1

 **Los Angeles, USA**

IPs:

IPv4: 192.228.79.201

IPv6: 2001:500:84::b

J Root Server – 113 Anycast Sites

Root Servers

[Archives](#)

[A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) **[J](#)** [K](#) [L](#) [M](#)

Operator: Verisign, Inc.

[Homepage](#)[Statistics](#)[RSSAC](#)

Locations:

Sites: 116

[Amsterdam, NL](#) [Ashburn, US](#) [Athens, GR](#) [Atlanta, US](#) [Bangalore, IN](#) [Bangkok, TH](#) [Barcelona, ES](#) [Barueri, BR](#) [Beijing, CN](#)
[Belgrade, RS](#) [Boston, US](#) [Brasilia, BR](#) [Bratislava, SK](#) [Brisbane, AU](#) [Bucharest, RO](#) [Buenos Aires, AR](#) [Cairo, EG](#)
[Campina Grande, BR](#) [Cape Town, ZA](#) [Chicago, US](#) [Cochabamba, BO](#) [Cordoba City, AR](#) [Dakar, SN](#) [Dallas, US](#) [Dar es Salaam, TZ](#)
[Des Moines, US](#) [Dhaka, BD](#) [Djibouti City, DJ](#) [Dublin, IE](#) [Geneva, CH](#) [Hangzhou, CN](#) [Hong Kong, CN](#) [Honolulu, US](#) [Ipoh, MY](#)
[Itabaiana, BR](#) [Jakarta, ID](#) [Joao Pessoa, BR](#) [Johannesburg, ZA](#) [Juazeiro do Norte, BR](#) [Kalamazoo-Battle Creek, US](#) [Kaunas, LT](#)
[Kigali City, RW](#) [Klagenfurt, AT](#) [Kuala Lumpur, MY](#) [Lagarto, BR](#) [Lagos, NG](#) [Lapu-Lapu City, PH](#) [Leeds, UK](#) [Limonest, FR](#) [Lisbon, PT](#)
[Ljubljana, SI](#) [Luanda, AO](#) [Luxembourg, LU](#) [Madison, US](#) [Madrid, ES](#) [Male, MV](#) [Melbourne, AU](#) [Miami, US](#) [Milan, IT](#) [Moscow, RU](#)
[Mumbai, IN](#) [Nairobi, KE](#) [New Castle, US](#) [New Delhi, IN](#) [North Sydney, AU](#) [Oslo, NO](#) [Oulu, FI](#) [Pakkret, TH](#) [Paris, FR](#) [Perth, AU](#)
[Porto Alegre, BR](#) [Prague, CZ](#) [Reno, US](#) [Reykjavik, IS](#) [Ribeirao Preto, BR](#) [Riga, LV](#) [Rio de Janeiro, BR](#) [Roma, IT](#) [Romblon, PH](#)
[Salzburg, AT](#) [San Francisco, US](#) [San Jose, CR](#) [San Juan, PR](#) [Seattle, US](#) [Seoul, KP](#) [Singapore, SG](#) [Sofia, BG](#) [St George, US](#)
[St Petersburg, RU](#) [Stockholm, SE](#) [Sydney, AU](#) [Taipei, TW](#) [Tamuning, GU](#) [Tel Aviv, IL](#) [Tokyo, JP](#) [Turin, IT](#) [Warsaw, PL](#)
[Wellington, NZ](#) [Yerevan, AM](#) [Zagreb, HR](#) [Zurich, CH](#)

L Root Server – 154 Anycast Sites

Root Servers

[A](#)[B](#)[C](#)[D](#)[E](#)[F](#)[G](#)[H](#)[I](#)[J](#)[K](#)[L](#)[M](#)

Operator:

ICANN

[Homepage](#)[Statistics](#)[Peering Policy](#)[Contact Email](#)[RSSAC](#)














Locations:

Sites: 158

[Abidjan, Cote d'Ivoire](#)[Al Muharraq, Bahrain](#)[Amman, Jordan](#)[Anchorage, United States](#)[Ankara, Turkey](#)[Apia, Samoa](#)[Asuncion, Paraguay](#)[Atlanta, United States](#)[Baku, Azerbaijan](#)[Bangkok, Thailand](#)[Beijing, People's Republic of China](#)[Beirut, Lebanon](#)[Belem, Brazil](#)[Belgrade, Serbia](#)[Belo Horizonte, Brazil](#)[Blantyre, Malawi](#)[Bogota, Colombia](#)[Bouake, Cote d'Ivoire](#)[Brasilia, Brazil](#)[Bratislava, Slovakia](#)[Brisbane, Australia](#)[Brussels, Belgium](#)[Callao, Peru](#)[Campinas, Brazil](#)[Cape Town, South Africa](#)[Chicago, United States](#)[Christchurch, New Zealand](#)[Cochabamba, Bolivia](#)[Copenhagen, Denmark](#)[Curitiba, Brazil](#)[Dakar, Senegal](#)[Dammam, Saudi Arabia](#)[Dar es Salaam, Tanzania](#)[Denver, United States](#)[Dortmund, Germany](#)[Dubai, United Arab Emirates](#)[Dundee, United Kingdom](#)[Dusseldorf, Germany](#)[El Prat de Llobregat, Spain](#)[Ezeiza, Argentina](#)[Florence, Italy](#)[Florianopolis, Brazil](#)[Fortaleza, Brazil](#)[Geneva, Switzerland](#)[Haarlemmermeer, Netherlands](#)[Hagatna, Guam](#)[Hamburg, Germany](#)[Heraklion, Greece](#)[Honiara, Solomon Islands](#)[Honolulu, United States](#)[Incheon, South Korea](#)[Islamabad-Rawalpindi, Pakistan](#)[Istanbul, Turkey](#)[Jakarta, Indonesia](#)[Jeddah, Saudi Arabia](#)[Johannesburg, South Africa](#)[Kalamazoo-Battle Creek, United States](#)[Kharkiv, Ukraine](#)[Kiev, Ukraine](#)[Kolkata, India](#)[Kolonja, Federated States of Micronesia](#)[Kuwait City, Kuwait](#)[Lahore, Pakistan](#)[Lawrence, United States](#)[Leeds-Bradford, England](#)[London, United Kingdom](#)[Londrina, Brazil](#)[Los Angeles, United States](#)[Lyon, France](#)[Mahe, Seychelles](#)[Maiquetia, Venezuela](#)[Malmo, Sweden](#)[Mandalay, Myanmar](#)[Mangere, New Zealand](#)[Marseille, France](#)[Mascot, Australia](#)[Melbourne, Australia](#)[Metro Manila, Philippines](#)[Miami, United States](#)[Minsk, Belarus](#)[Mississauga, Canada](#)[Monterrey, Mexico](#)[Montevideo, Uruguay](#)[Moscow, Russia](#)[Mumbai, India](#)[Muscat, Oman](#)[Nadi, Fiji](#)[Nairobi, Kenya](#)[Natal, Brazil](#)[Noumea, New Caledonia](#)[Odessa, Ukraine](#)[Ottawa, Canada](#)[Papeete, French Polynesia](#)[Paris, France](#)[Paris-Orly, France](#)[Perth, Australia](#)[Phoenix, United States](#)[Plaisance, Mauritius](#)[Port Moresby, Papua New Guinea](#)[Portland, United States](#)[Porto, Portugal](#)[Porto Alegre, Brazil](#)[Prague, Czech Republic](#)[Punta Caucedo, Dominican Republic](#)[Quito, Ecuador](#)[Rabat, Morocco](#)[Reno, United States](#)[Reston, United States](#)[Rio de Janeiro, Brazil](#)[Riyadh, Saudi Arabia](#)[Rochester, England](#)[Rostov-on-Don, Russia](#)[Salvador, Brazil](#)[San Jose, United States](#)[San Jose, Costa Rica](#)[San Juan, Puerto Rico](#)[San Salvador, El Salvador](#)[Sana'a, Yemen](#)[Santiago, Chile](#)[Sao Jose dos Campos, Brazil](#)[Sao Paulo, Brazil](#)[SeaTac, United States](#)[Semey, Kazakhstan](#)[Sofia, Bulgaria](#)[St Denis, Reunion](#)[Stockholm, Sweden](#)[Suva-Nausori, Fiji](#)[Tokyo, Japan](#)[Toronto, Canada](#)[Tunis-Carthage, Tunisia](#)[Uberlandia, Brazil](#)[Ullensaker, Norway](#)[Vancouver, Canada](#)[Wilmington, United States](#)[Winnipeg, Canada](#)[Yangon, Myanmar](#)[Yekaterinburg, Russia](#)[Yerevan, Armenia](#)[Yogyakarta, Indonesia](#)

13

The Varying Impact on Root Servers

Map	Servers	Agents
Server	Date (UTC)	Number of Errors ↓
h.root-servers.net	2016-06-25 21:51:54	 21
b.root-servers.net	2016-06-25 21:51:23	 20
g.root-servers.net	2016-06-25 21:51:49	 18
c.root-servers.net	2016-06-25 21:51:43	 13
m.root-servers.net	2016-06-25 21:51:51	 12
d.root-servers.net	2016-06-25 21:52:20	 11
e.root-servers.net	2016-06-25 21:51:45	 10
i.root-servers.net	2016-06-25 21:51:51	 10
k.root-servers.net	2016-06-25 21:51:54	 8
f.root-servers.net	2016-06-25 21:51:18	 5
a.root-servers.net	2016-06-25 21:51:38	 3
j.root-servers.net	2016-06-25 21:51:30	 2
l.root-servers.net	2016-06-25 21:51:32	 2

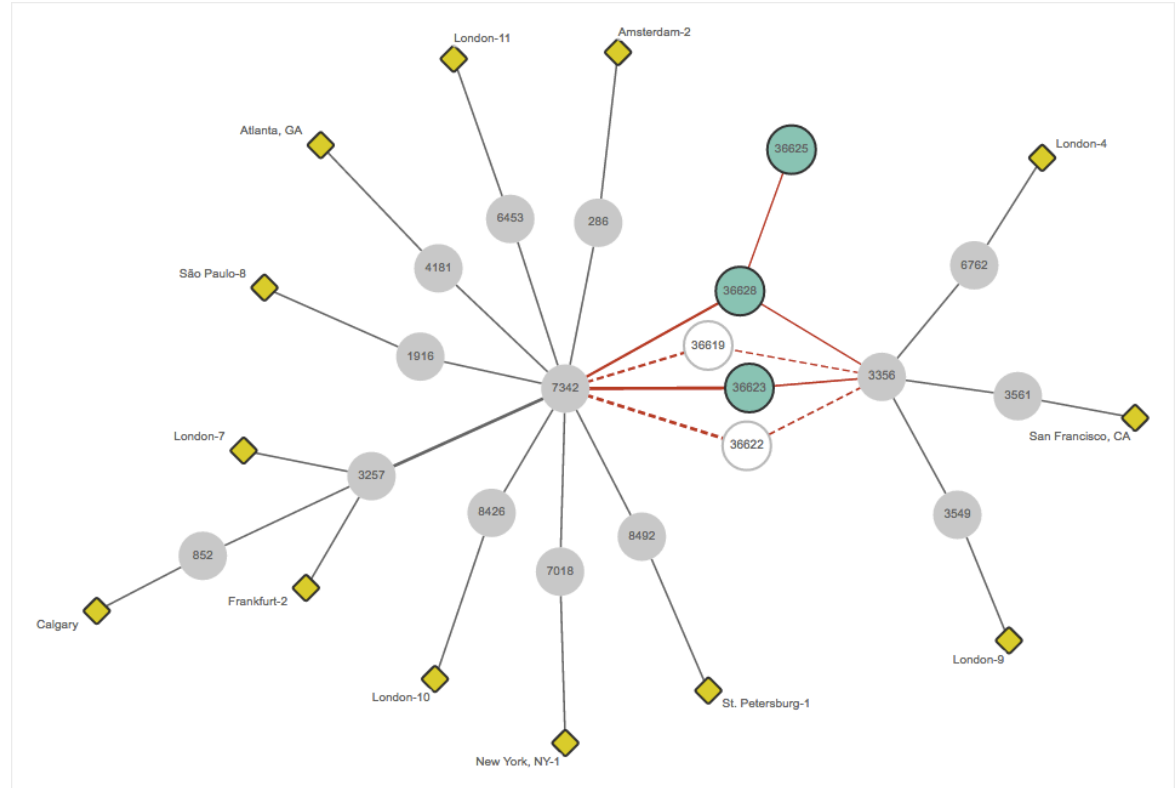
Anycast made a difference but capacity at each site was also an important factor

With Real World Impact

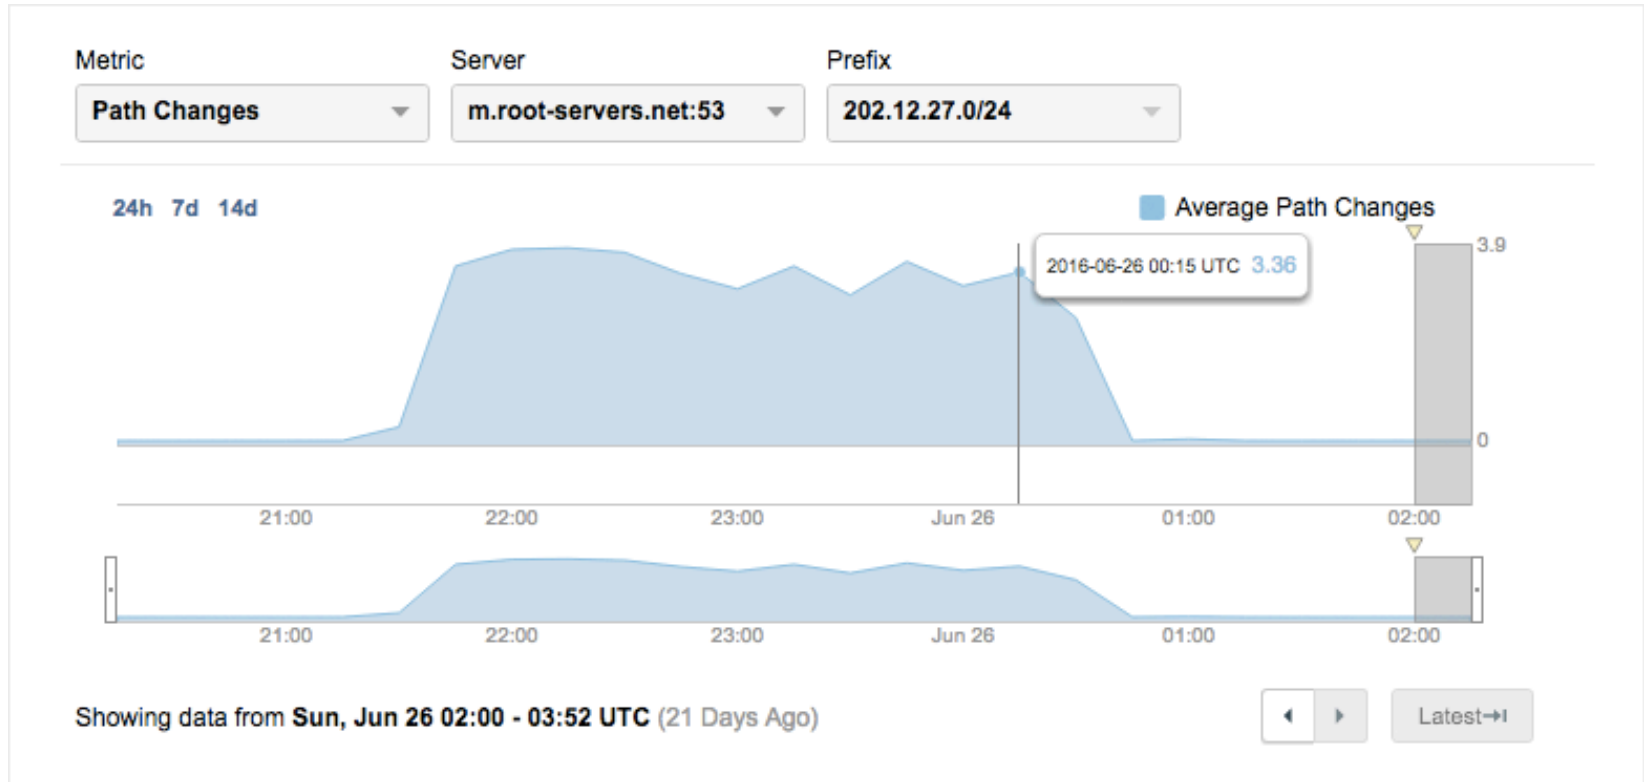
Servers	DNS Availability (%)	DNS Resolution Time (ms)	Network - Agent to Server Packet Loss (%)
	Mean	Mean	Mean
g.root-servers.net:53	13 ▼87	170.84 ▲39.38	89.9 ▲89.87
h.root-servers.net:53	14.16 ▼85.84	201.6 ▲83.76	92.18 ▲91.83
b.root-servers.net:53	14.67 ▼85.33	175.06 ▲54.32	97.12 ▲97.11
c.root-servers.net:53	41.06 ▼58.82	65.98 ▲3.16	62.22 ▲54.46
e.root-servers.net:53	52.25 ▼47.75	59.62 ▲33.41	46.12 ▲46.1
m.root-servers.net:53	54.18 ▼45.82	138.68 ▲34.06	64.41 ▲64.41
i.root-servers.net:53	54.44 ▼45.56	210.53 ▲138.49	53.84 ▲53.81
d.root-servers.net:53	70.01 ▼29.99	81.64 ▲49.46	33.75 ▲33.73
k.root-servers.net:53	73.49 ▼26.51	124.51 ▲54.6	46.51 ▲46.5
f.root-servers.net:53	81.6 ▼18.4	36.17 ▲7.9	26.04 ▲23.08
l.root-servers.net:53	90.48 ▼9.52	53.83 ▼2.96	27.09 ▲27.05
j.root-servers.net:53	94.21 ▼5.79	42.36 ▲3.4	10.64 ▲10.61
a.root-servers.net:53	97.81 ▼2.19	149.22 ▲84.61	45.78 ▲45.73

Operators Mitigate

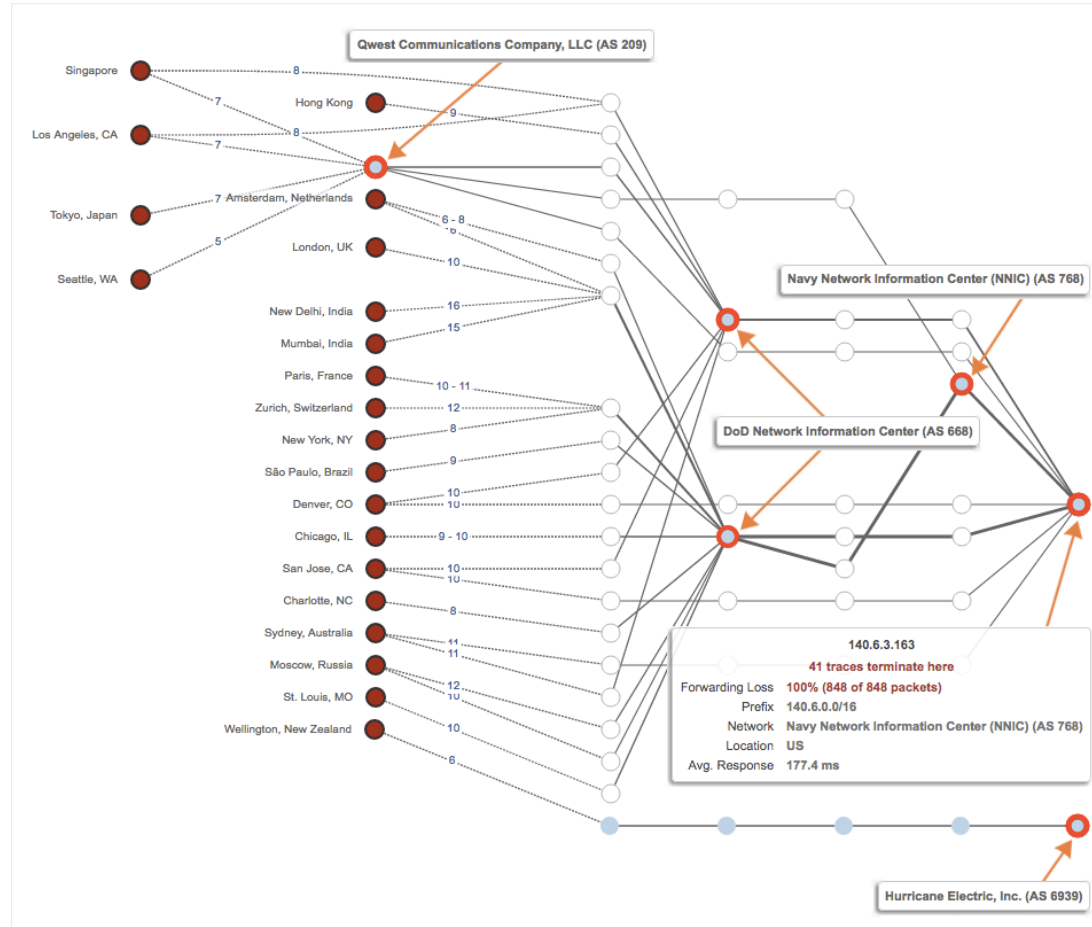
- A-Root makes some BGP changes to deal with the attack



Lots of BGP changes with other Roots as well



Packet loss observed at Upstream as well



DDoS Indicators To Look Out For


1. Availability (DNS root query) below 90%
2. Resolution time (DNS root query) $>1\sigma$
3. Multiple roots impacted
4. Multiple anycast POPs impacted
5. Multiple upstreams impacted

Many of these indicators correlate to metrics in DDoS attacks on other targets and services as well.

SEA-ME-4 Cable Fault



Tata Backbone Under Normal Conditions



Path trace from New Delhi, India (10.10.10.70) to 208.82.237.246

```

1 10.10.10.65 (10.10.10.65) 0 ms 0 ms 0 ms
2 180.179.204.2 (180.179.204.2) 0 ms 0 ms 0 ms
3 180.179.192.73 (180.179.192.73) 0 ms 0 ms 0 ms
4 180.179.197.37 (180.179.197.37) 0 ms 0 ms 0 ms
5 14.140.113.85.static-Delhi-vsnl.net.in (14.140.113.85) 2 ms
  14.141.216.93.static-Delhi.vsnl.net.in (14.141.216.93) 1 ms 1
6 * * *
7 ix-ae-0-100.tcore1.MLV-Mumbai.as6453.net (180.87.38.5) 26 ms 2
8 if-ae-9-6.tcore1.WYN-Marseille.as6453.net (80.231.217.77) <MPLS:
  if-ae-5-2.tcore1.WYN-Marseille.as6453.net (80.231.217.29) <MPLS:
    if-ae-5-6.tcore1.WYN-Marseille.as6453.net (180.87.38.126) <MPLS:L=1349058567,E=0,S=1,T=1> 212 ms
9 if-ae-8-1600.tcore1.PYE-Paris.as6453.net (80.231.217.6) <MPLS:L=1613169420,E=0,S=1,T=1> 216 ms 216 ms 215
10 if-ae-3-6.tcore1.L78-London.as6453.net (80.231.130.85) <MPLS:L=2150042894,E=0,S=1,T=1> 219 ms 215 ms 219 n
11 if-ae-17-2.tcore1.LDN-London.as6453.net (80.231.130.130) <MPLS:L=277285894,E=0,S=1,T=1> 212 ms 212 ms *
12 * * *
13 if-ae-1-3.thar2.NJY-Newark.as6453.net (216.6.57.2) <MPLS:L=1343685129,E=0,S=1,T=1> 210 ms 210 ms *
14 if-ae-18-2.tcore2.NTO-New-York.as6453.net (66.198.111.7) <MPLS:L=2684618752,E=0,S=1,T=1> 207 ms 207 ms
  if-ae-14-14.tcore2.NTO-New-York.as6453.net (66.198.111.126) <MPLS:L=2684618752,E=0,S=1,T=1> 213 ms
15 if-ae-12-2.tcore1.N75-New-York.as6453.net (66.110.96.5) 210 ms 210 ms 212 ms
16 66.110.96.146 (66.110.96.146) 210 ms
   66.110.96.138 (66.110.96.138) 217 ms
   66.110.96.142 (66.110.96.142) 209 ms
17 hu-1-3-0-8-cr02.newyork.ny.ibone.comcast.net (68.86.84.241) 209 ms
   hu-1-3-0-2-cr02.newyork.ny.ibone.comcast.net (68.86.83.97) 218 ms
   hu-1-4-0-0-cr02.newyork.ny.ibone.comcast.net (68.86.84.249) 215 ms
18 et-15-1-0-0-ar01.whitemarsh.md.bad.comcast.net (68.86.94.102) 216 ms 216 ms 219 ms
19 te-8-1-ur01.michiganave.dc.bad.comcast.net (68.85.133.70) 219 ms 220 ms 221 ms
20 50-203-200-110-static.hfc.comcastbusiness.net (50.203.200.110) 216 ms 216 ms 218 ms
21 post.craigslist.org (208.82.237.246) 214 ms 217 ms 220 ms
    
```

if-ae-5-2.tcore1.WYN-Marseille.as6453.net
 IP Address 80.231.217.29
 Prefix 80.231.0.0/16
 Network Tata Communications (AS 6453)
 Location Marseille, France
 Avg. Response 203 ms

Trouble in the Tata Backbone

- May 17th 2016 06:10-8:30 PDT (13:10-15:30 UTC)
- Performance degradation in Tata India to Europe backbone

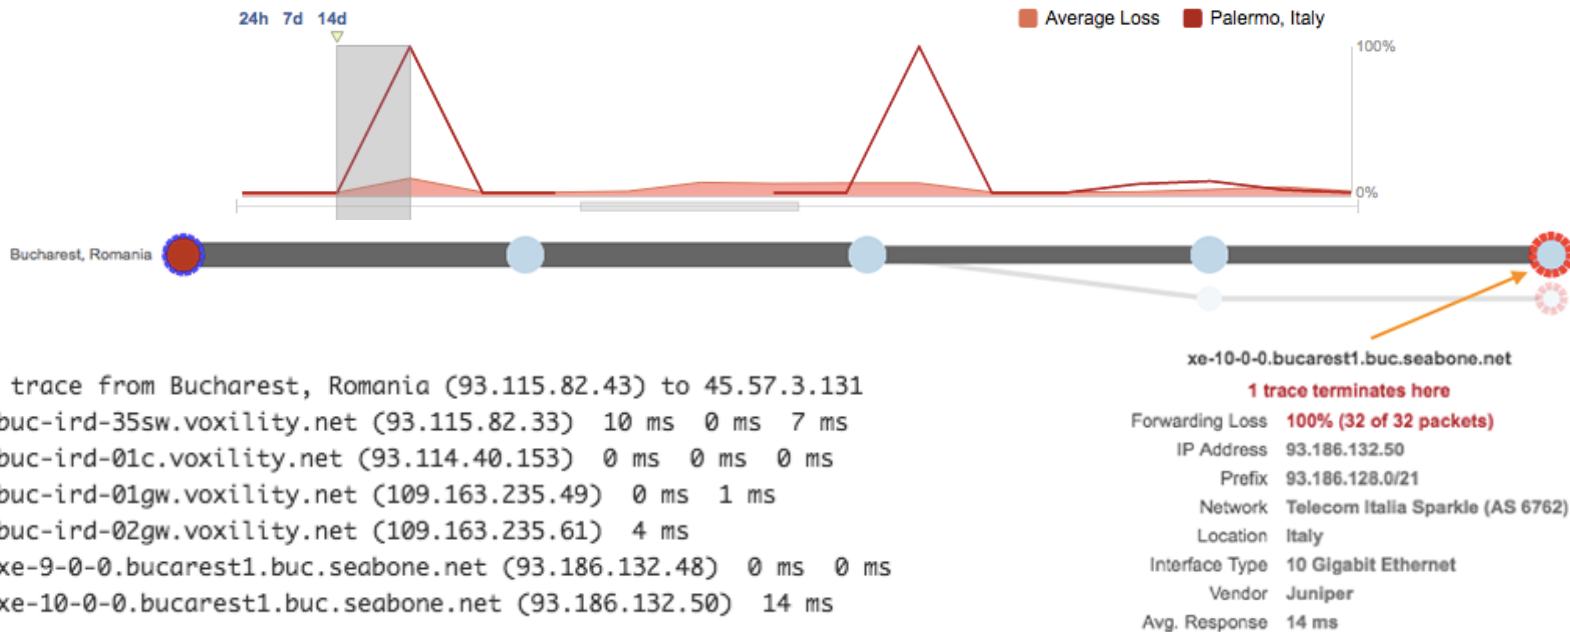


Path trace from New Delhi, India (10.10.10.70) to 208.82.237.6

```
1 10.10.10.65 (10.10.10.65) 0 ms 0 ms 0 ms
2 180.179.204.2 (180.179.204.2) 5 ms 1 ms 1 ms
3 180.179.192.73 (180.179.192.73) 0 ms 0 ms 0 ms
4 180.179.197.41 (180.179.197.41) 1 ms 0 ms 0 ms
5 219.65.44.177.static-delhi.vsnl.net.in (219.65.44.177) 1 ms 1 ms 1 ms
6 * * *
7 ix-ae-0-100.tcore1.MLV-Mumbai.as6453.net (180.87.38.5) 26 ms 26 ms 26 ms
```

And Also in Telecom Italia Sparkle

- 06:35-6:40 PDT (13:35-13:40 UTC)
- TISparkle Mediterranean backbone sees complete loss



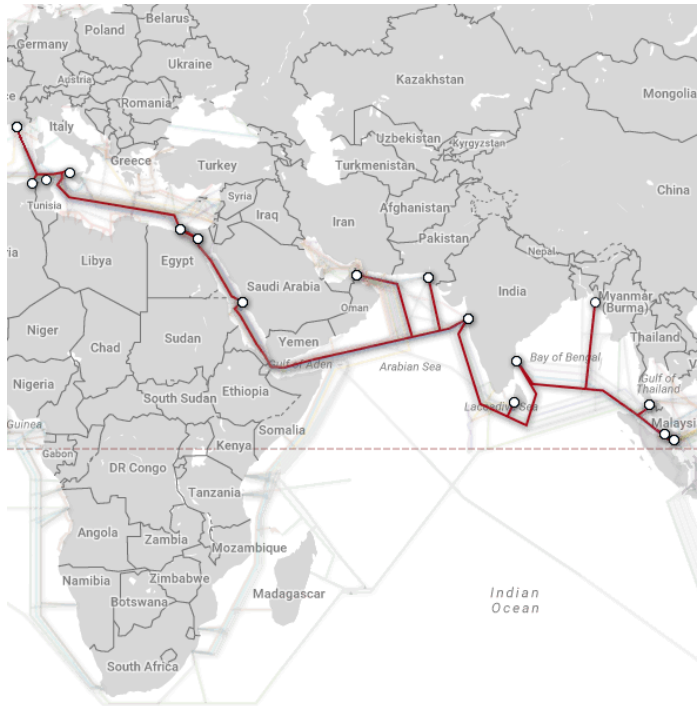
European Detour In Effect

- Netflix (AS2906) drops TISparkle (AS6762) and begins to route via Level3 (AS3356) instead
- Traffic flows via Frankfurt rather than Paris (and Marseilles)



Commonalities between Tata and TIS

- Multiple, geographically correlated backbone outages
- Both share Mediterranean transit paths on Sea-Me-We-3 and Sea-Me-We-4



SeaMeWe-4

RFS: December 2005

Owners: Bangladesh Tele

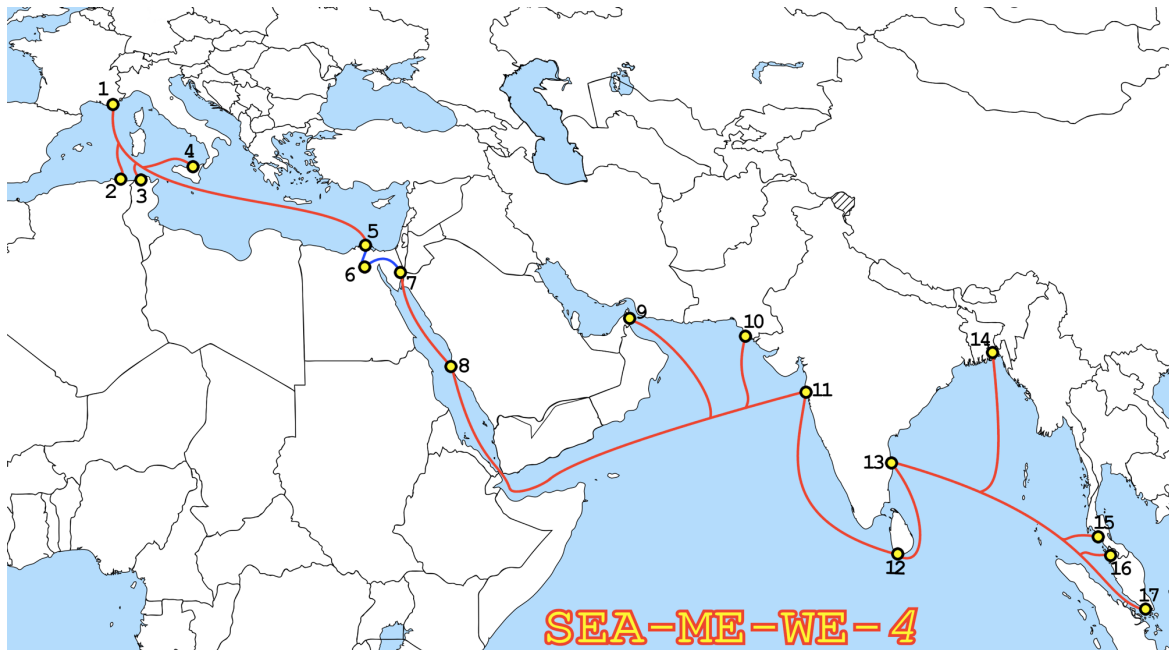
Owners: Bangladesh Telegraph and Telephone Board (BTB), Orange, SingTel, Telecom Italia Sparkle, Tata Communications, PT Indonesia Satellite Corp., Telekom Malaysia, Airtel (Bharti), Sri Lanka Telecom, Etisalat, Saudi Telecom, Communications Authority of Thailand, Tunisia Telecom, Verizon, Pakistan Telecommunications Company Ltd., Telecom Egypt, Telstra

URL: <http://www.seamewe4.net>

Alexandria, Egypt

Alexandria, Egypt
Annaba, Algeria
Bizerte, Tunisia
Chennai, India
Colombo, Sri Lanka
Cox's Bazar, Bangladesh
Fujairah, United Arab Emirates
Jeddah, Saudi Arabia
Karachi, Pakistan
Marseille, France
Melaka, Malaysia
Mumbai, India
Palermo, Italy
Satun, Thailand
Suez, Egypt
Tuas, Singapore

SEA-ME-WE-4 Cable



Background

- Connects Europe to Middle East, South and SE Asia
- 4.6 Tbps
- Has suffered more than a dozen major faults

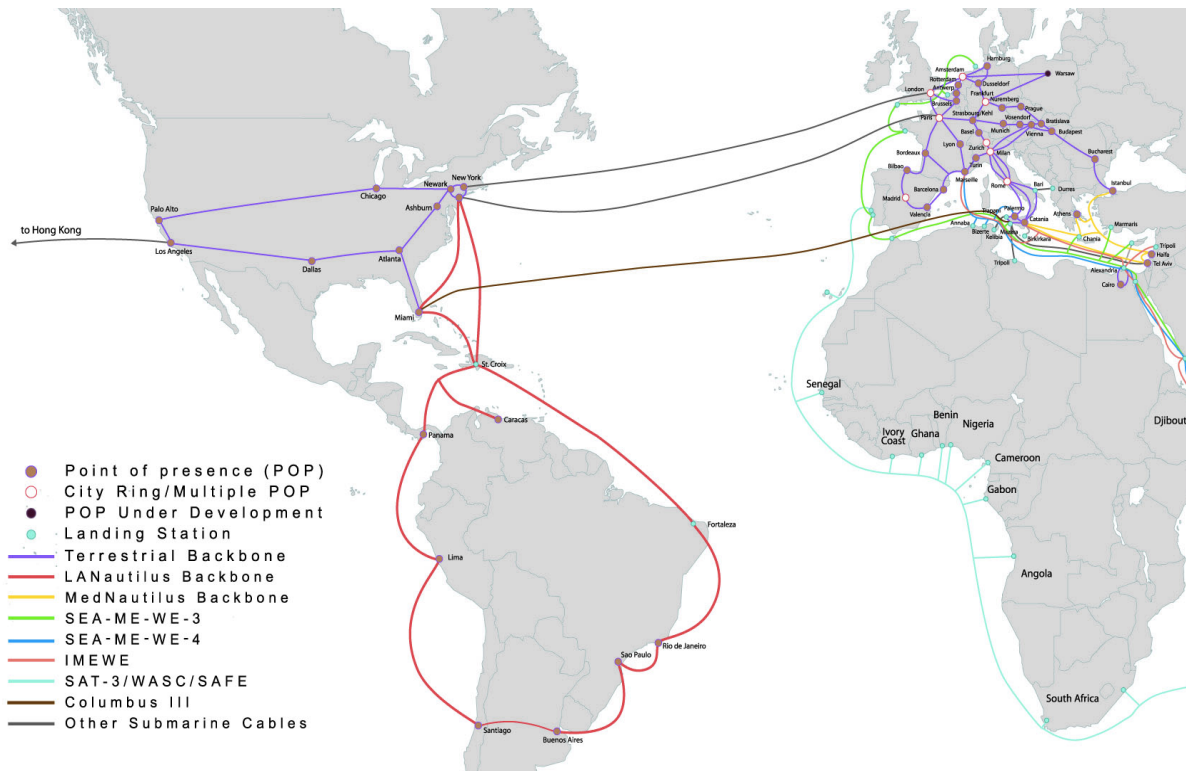
Issues Spread in the TI Sparkle Backbone

- 6:45-8:10 PDT (13:45-15:10 UTC)
- POPs affected in Europe and Americas
 - Palermo, Santiago, Milan, Catania, Baires, Frankfurt, Paris, Dallas, London, Miami, New York
 - Peering drop with Level 3 Paris



Why Would This Impact the Americas?

- TI Sparkle backbone connects Latin America to Europe via Miami and New Jersey



And BGP Sessions Begin to Fail

- 6:45-7:15 PDT (13:45-14:15 UTC)
- Reachability affected for thousands of prefixes due to TI Sparkle network
 - 85 prefixes in the Netherlands (BIT, Akamai)
 - 1479 prefixes in Argentina (Telecom Argentina)
 - 95 prefixes in Greece (FORTHnet)
- Ripple effect, Impact beyond Europe/Asia

Eventual Announcement as to Root Cause

- Segment 4 (Cairo to Marseilles) faulty repeater acknowledged 2 days later
- Likely cause between Palermo and Marseilles based on the data we've seen

SEA-ME-WE-4 outage

Schedule

Activity-1

Start Date/Time 13th May 2016, 11:00pm, Pakistan Standard Time
End Date/Time 14th May 2016, 3:00am, Pakistan Standard Time
Scope of work Power re-configuration in SMW4 Segment-4
Impact Degradation of Service on all international IP services

Activity-2

Start Date/Time 15th May 2016, 5:00am, Pakistan Standard Time
End Date/Time 22nd May 2016, 4:59am, Pakistan Standard Time
Scope of work Replacement of faulty repeater (R4113) in SMW4 Segment-4
Impact Degradation of Service on all international IP services

Activity-3

Start Date/Time 22nd May 2016, 5:00am, Pakistan Standard Time
End Date/Time 1st June 2016, 4:59am, Pakistan Standard Time
Scope of work Replacement of faulty repeater (R4103) in SMW4
Impact Degradation of Service on all international IP services

Cable Fault Indicators To Look Out For

1. Many path traces impacted in adjacent POPs on the same network
2. Jitter can be an even more convincing measure than loss
3. Multiple networks impacted suggest a cable fault, IXP failure or peering failure
 - Cable fault: Elevated loss, elevated jitter
 - IXP failure: Elevated loss on many interfaces in the same POP
 - Peering: Terminal loss, path changes
4. Dropped BGP sessions may occur when problems persist


AWS Route Leak



https://upload.wikimedia.org/wikipedia/commons/6/68/Zurich_in_night1.jpg

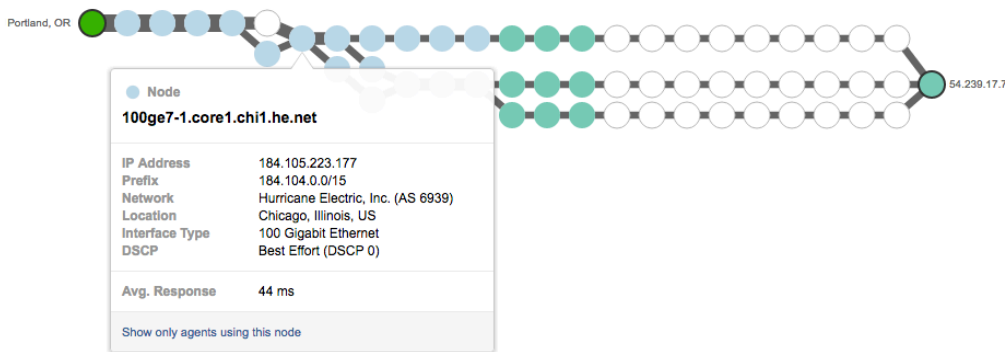
AWS Routes on a Normal Day

- 54.239.16.0/20 prefix
Amazon.com (AWS US East)
- Peering with the expected providers: NTT, TI Sparkle, Telia, CenturyLink, HE



Traffic to AWS

- From Portland OR to AWS US East, traffic normally transits HE Chicago and peers with AWS

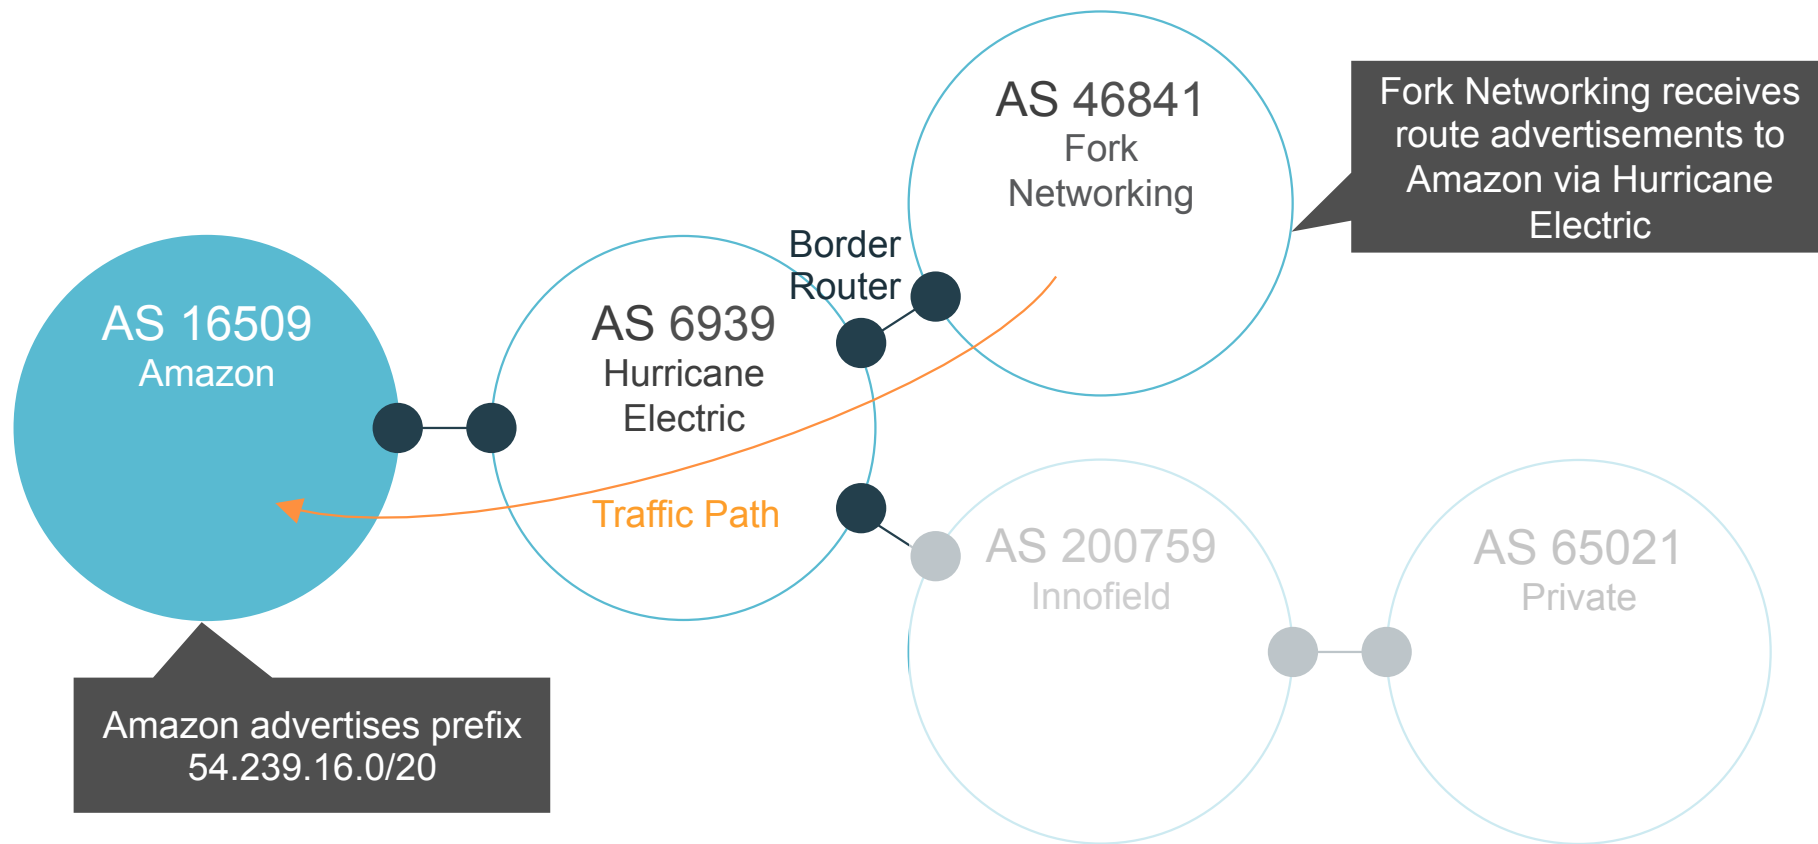


Path trace from Portland, OR (162.218.67.132) to 54.231.10.96

```

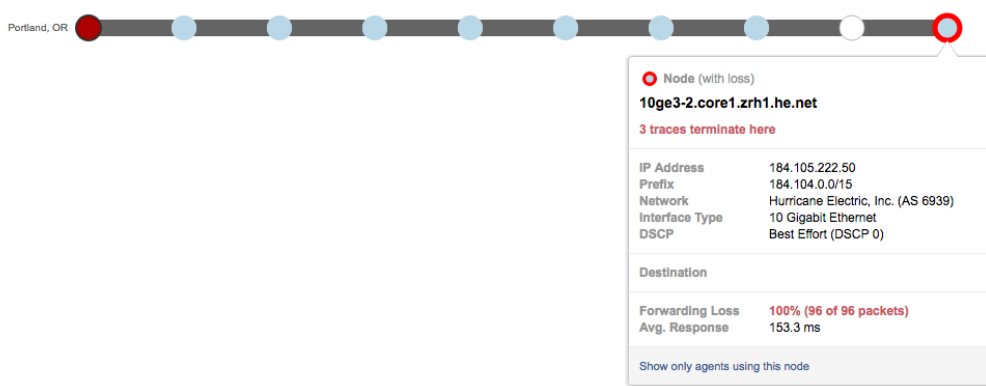
1 162.218.67.129 (162.218.67.129) 0 ms 0 ms 0 ms
2 pdx-edge-rtr02.forked.net (199.87.231.45) 0 ms 0 ms 10 ms
3 gi0-2-1-3.rcr21.b006468-1.pdx02.atlas.cogentco.com (38.104.104.141) 0 ms 0 ms 0 ms
4 te0-0-1-2.rcr12.pdx02.atlas.cogentco.com (154.24.49.213) 1 ms
  te0-0-2-2.rcr12.pdx02.atlas.cogentco.com (154.24.49.205) 0 ms
  te0-0-1-2.rcr12.pdx02.atlas.cogentco.com (154.24.49.213) 1 ms
5 be2670.ccr22.sea01.atlas.cogentco.com (154.54.42.149) 4 ms 4 ms 5 ms
6 be2084.ccr21.sea02.atlas.cogentco.com (154.54.0.254) 4 ms 5 ms 5 ms
7 ae-11.r05.sttlwa01.us.bb.gin.ntt.net (154.54.11.10) 4 ms 6 ms 4 ms
8 ae-3.r20.sttlwa01.us.bb.gin.ntt.net (129.250.2.44) <MPLS:L=6947586,E=0,S=1,T=1> 21 ms 20 ms 21 ms
9 ae-2.r20.chcgil09.us.bb.gin.ntt.net (129.250.3.43) <MPLS:L=2686717186,E=0,S=1,T=1> 48 ms 48 ms 61 ms
10 ae-1.r06.chcgil09.us.bb.gin.ntt.net (129.250.4.146) 49 ms 58 ms 59 ms
11 ae-3.amazon.chcgil09.us.bb.gin.ntt.net (129.250.201.106) 48 ms 48 ms 59 ms
12 52.95.62.74 (52.95.62.74) <MPLS:L=3236167948,E=0,S=0,T=1/L=2422475008,E=0,S=1,T=1> 66 ms
   52.95.62.26 (52.95.62.26) <MPLS:L=3223585286,E=0,S=0,T=1/L=2696154368,E=0,S=1,T=1> 65 ms
   52.95.62.138 (52.95.62.138) <MPLS:L=2431913476,E=0,S=0,T=1/L=2696154368,E=0,S=1,T=1> 86 ms
13 52.95.62.79 (52.95.62.79) <MPLS:L=2422475008,E=0,S=1,T=1> 62 ms
   52.95.62.33 (52.95.62.33) <MPLS:L=2696154368,E=0,S=1,T=1> 66 ms
   52.95.62.145 (52.95.62.145) <MPLS:L=2696154368,E=0,S=1,T=1> 74 ms
14 54.239.42.61 (54.239.42.61) <MPLS:L=1621493006,E=0,S=0,T=1/L=3771007232,E=0,S=1,T=1> 65 ms
   54.239.42.63 (54.239.42.63) <MPLS:L=1614349314,E=0,S=0,T=1/L=1351155968,E=0,S=1,T=1> 64 ms 74 ms
15 54.239.42.66 (54.239.42.66) <MPLS:L=3771007232,E=0,S=1,T=1> 65 ms
   54.239.42.69 (54.239.42.69) <MPLS:L=1351155968,E=0,S=1,T=1> 64 ms 74 ms
16 54.239.109.96 (54.239.109.96) <MPLS:L=2965637379,E=0,S=0,T=1/L=2961442048,E=0,S=1,T=1> 81 ms
   54.239.109.234 (54.239.109.234) <MPLS:L=2417233160,E=0,S=0,T=1/L=2962755334,E=0,S=1,T=1> 89 ms
   54.239.110.46 (54.239.110.46) <MPLS:L=2431911687,E=0,S=0,T=1/L=2962755334,E=0,S=1,T=1> 92 ms
17 54.239.109.111 (54.239.109.111) <MPLS:L=2961442048,E=0,S=1,T=1> 63 ms
   54.239.111.75 (54.239.111.75) <MPLS:L=2962755334,E=0,S=1,T=1> 64 ms
   54.239.111.79 (54.239.111.79) <MPLS:L=2962755334,E=0,S=1,T=1> 74 ms
18 205.251.245.226 (205.251.245.226) 63 ms
   205.251.244.191 (205.251.244.191) 64 ms
   205.251.244.219 (205.251.244.219) 72 ms
  
```

Route Propagation from AWS



Why Is Our AWS Traffic in Switzerland?

- Traffic goes all the way to Zurich in HE
- And stops there
- Along with a lot of AWS traffic! Also causing outages in Level 3 and Cogent POPs at the same time.

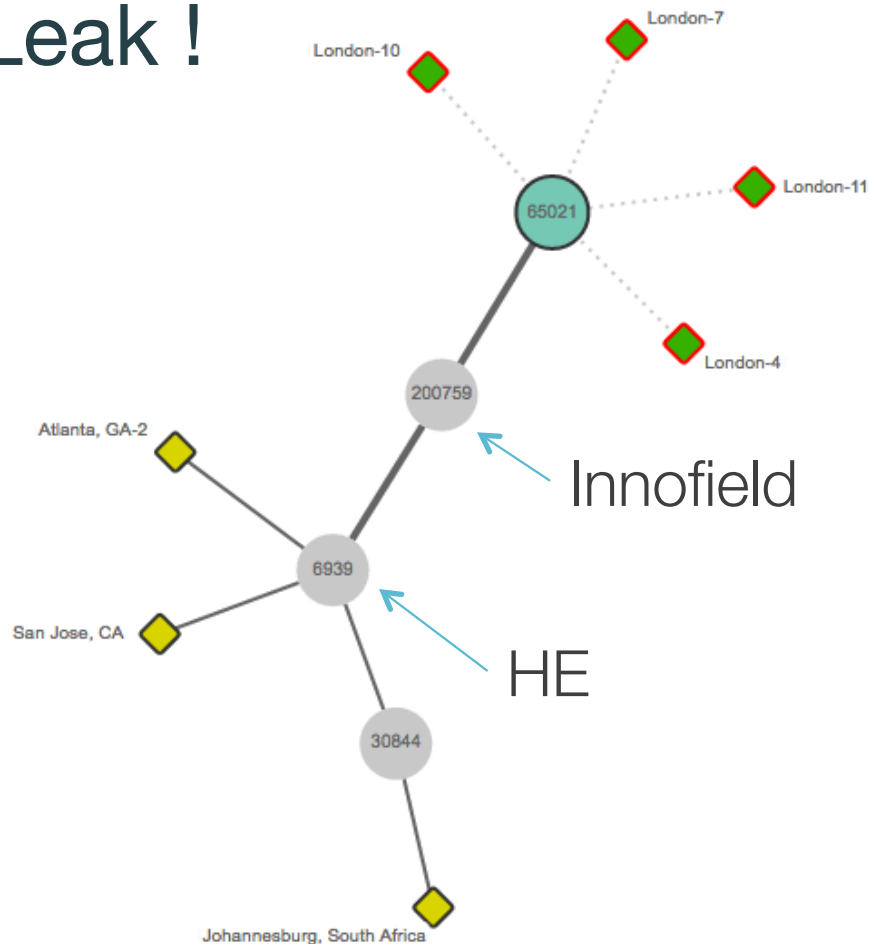


Path trace from Portland, OR (162.218.67.132) to 54.239.17.7

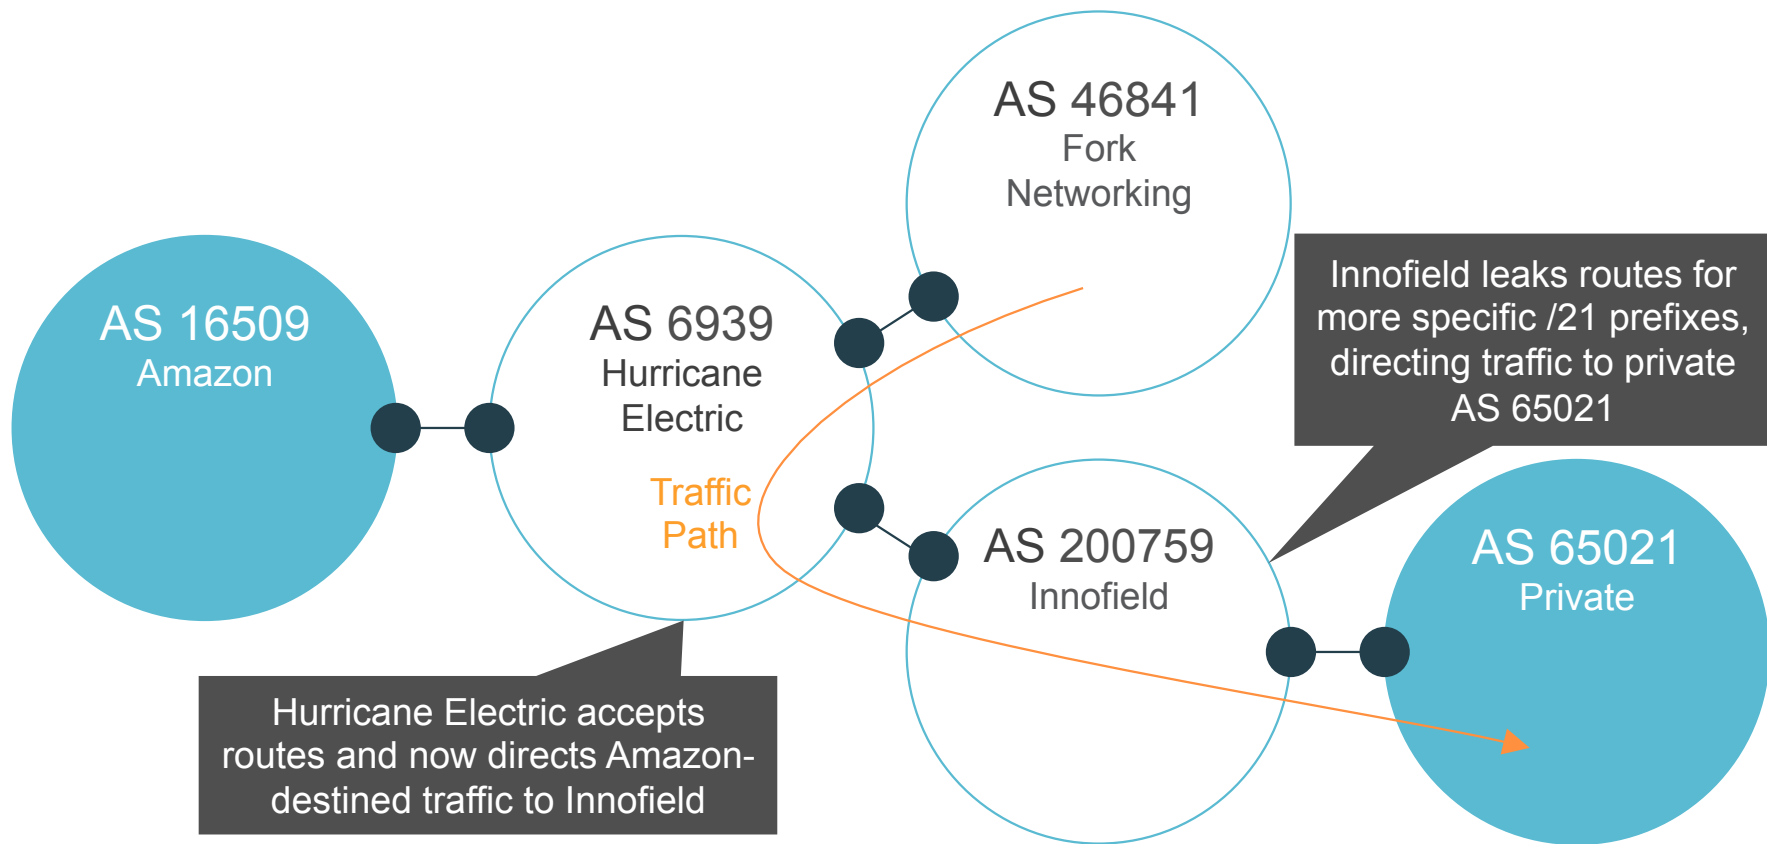
1	162.218.67.129 (162.218.67.129)	2 ms	0 ms	0 ms
2	pdx-edge-rtr01.forked.net (199.87.231.17)	0 ms	0 ms	0 ms
3	ge2-20.core1.pdx1.he.net (216.218.244.225)	3 ms	2 ms	1 ms
4	10ge1-14.core1.sea1.he.net (72.52.92.9)	4 ms	11 ms	48 ms
5	100ge10-2.core1.msp1.he.net (184.105.223.194)	36 ms	36 ms	115 ms
6	100ge7-1.core1.chi1.he.net (184.105.223.177)	51 ms	45 ms	43 ms
7	100ge5-2.core1.nyc4.he.net (184.105.223.162)	60 ms	60 ms	60 ms
8	* * *			
9	10ge3-2.core1.zrh1.he.net (184.105.222.50)	148 ms	156 ms	156 ms

Route Leak !

- 10:10-12:10 PDT (17:10-19 UTC) two new prefixes are advertised
 - 54.239.16.0/21
 - 54.239.24.0/21
- Advertised by AS200759 (Innofield AG)
- Origin AS65021 (private)
- Via AS6939 (HE)



A Route Leak in the Wild



Survey Says... Route Optimizer

- Prefixes leaked in SwissIX and onto HE
- Route optimizer is likely cause
- Similar cause as July 2015 incident with Enzu in Los Angeles

Background of this incident:

- Yesterday (2016-04-22) SwissIX (Swiss Internet Exchange) has performed a maintenance on our port and we had to de-activate and re-activate the BGP sessions.
- After re-activating the sessions at approx. 17:10 UTC, for a currently unknown reason, we have redistributed other prefixes into SwissIX.
- This definitely shouldn't happen, as a matter of course we have filters in place for SwissIX (which only allow to send our prefixes).
- Immediately after we received the first complains about the route-hijack, we have de-activated the BGP sessions again (at approx. 17:25 UTC)

Next steps at our side:

- We are in permanent contact with the router and the route-optimizer vendors. Together, we are deeply analyzing the system/logs why the filters were not considered after re-activating the BGP sessions.
- Currently (but not confirmed yet) it looks like a code-bug in the router software.
- We will definitely not re-activate the sessions until the root-cause is found!
- With this planned actions we will prevent this from happening again.

Again we are very sorry for this incident and we will do everything you can think of to avoid this situations in the future.


Respectfully,
NOC Team from innofield AG

Route Leak Indicators To Look Out For

- New Prefix or new Destination ASN
- Major BGP route changes – significant change in new path
- Involves ASNs that maybe in geos far from destination ASN
- High packet loss at one of the ASNs in the path or ASNs with a common next hop ASN

Summarizing the 3 Events from 2016

DNS Root DDoS



June 26, 2016

Submarine Cable
Fault



May 17, 2016

AWS Route Leak



April 22, 2016



Thank You
@mohitlad

<https://blog.thousandeyes.com/category/outage-reports/>