# About ThousandEyes

- Network Intelligence startup based in San Francisco, CA
- Founded by a couple of UCLA PhDs
- Leverages thousands of agents to measure performance of networks mostly through active probing

# Cloud Agent Footprint

- ~125 locations worldwide
- Each location is a cluster with many physical instances
- Add to this 300+ locations w/ enterprise agents

# Collected data

- Application
  - Http, dns, voice (sip+rtp)
  - Real browser tests (chrome)
- Network
  - End-to-end: loss, delay and jitter
  - Path trace: fwd loss, link delay
- Routing
  - BGP data from RouteViews

# Internet Outage Detection

## Traffic Outage Detection

- Detect outages in ISPs and understand their impact both globally and as it relates to a specific customer

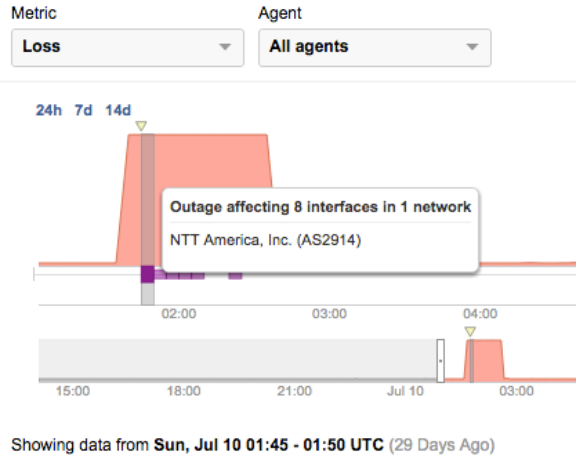


## Routing Outage Detection

- See the global and account scope, as well as likely root cause of BGP reachability outages

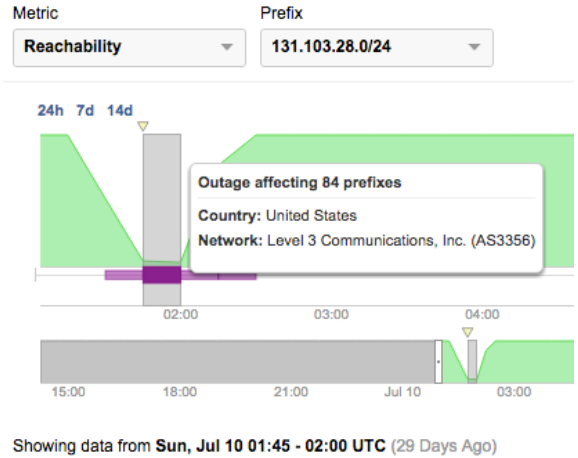

# Traffic Outage Detection

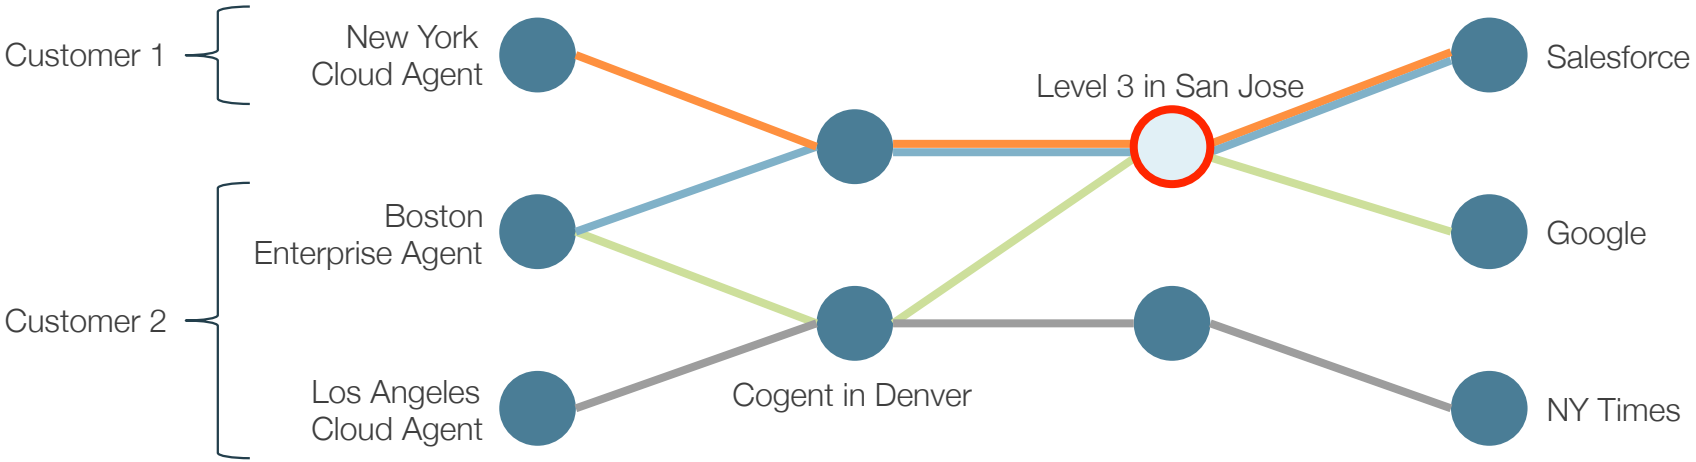


1. Anonymized (http) traffic data is aggregated from all tests across the entire user base
2. Algorithms then look for patterns in path traces terminating in the same ISP
3. Exclude: noisy interfaces and non-ISP networks

# Traffic Outage Detection

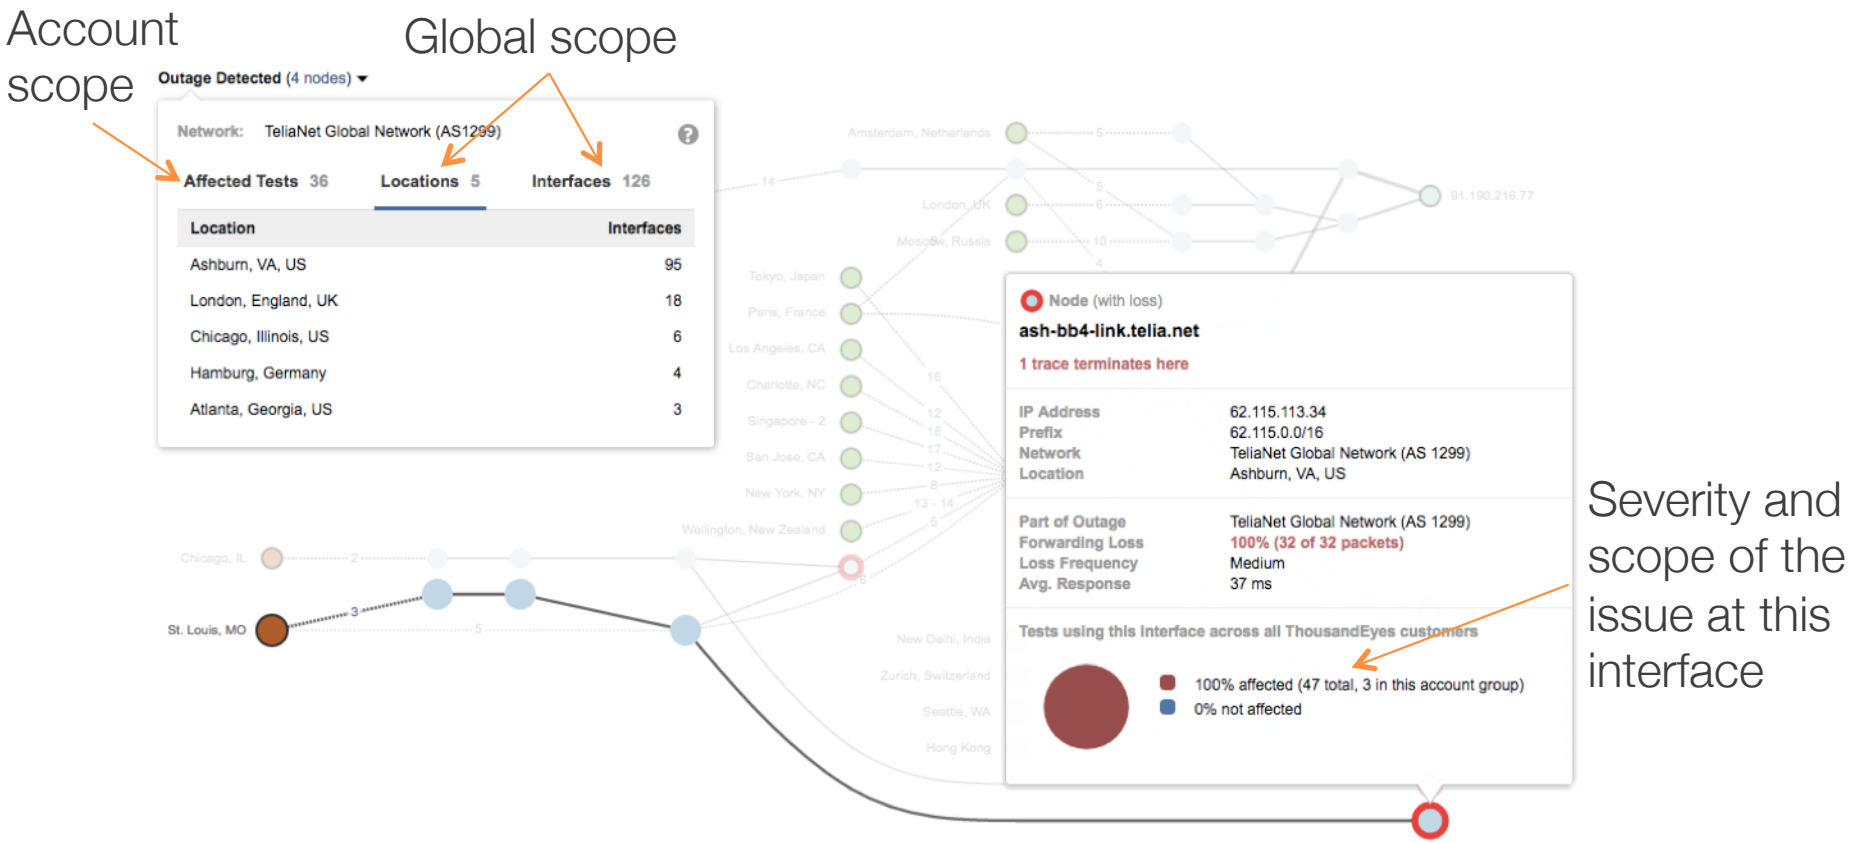

# Routing Outage Detection

Aggregates reachability issues in routing data from 350 routers

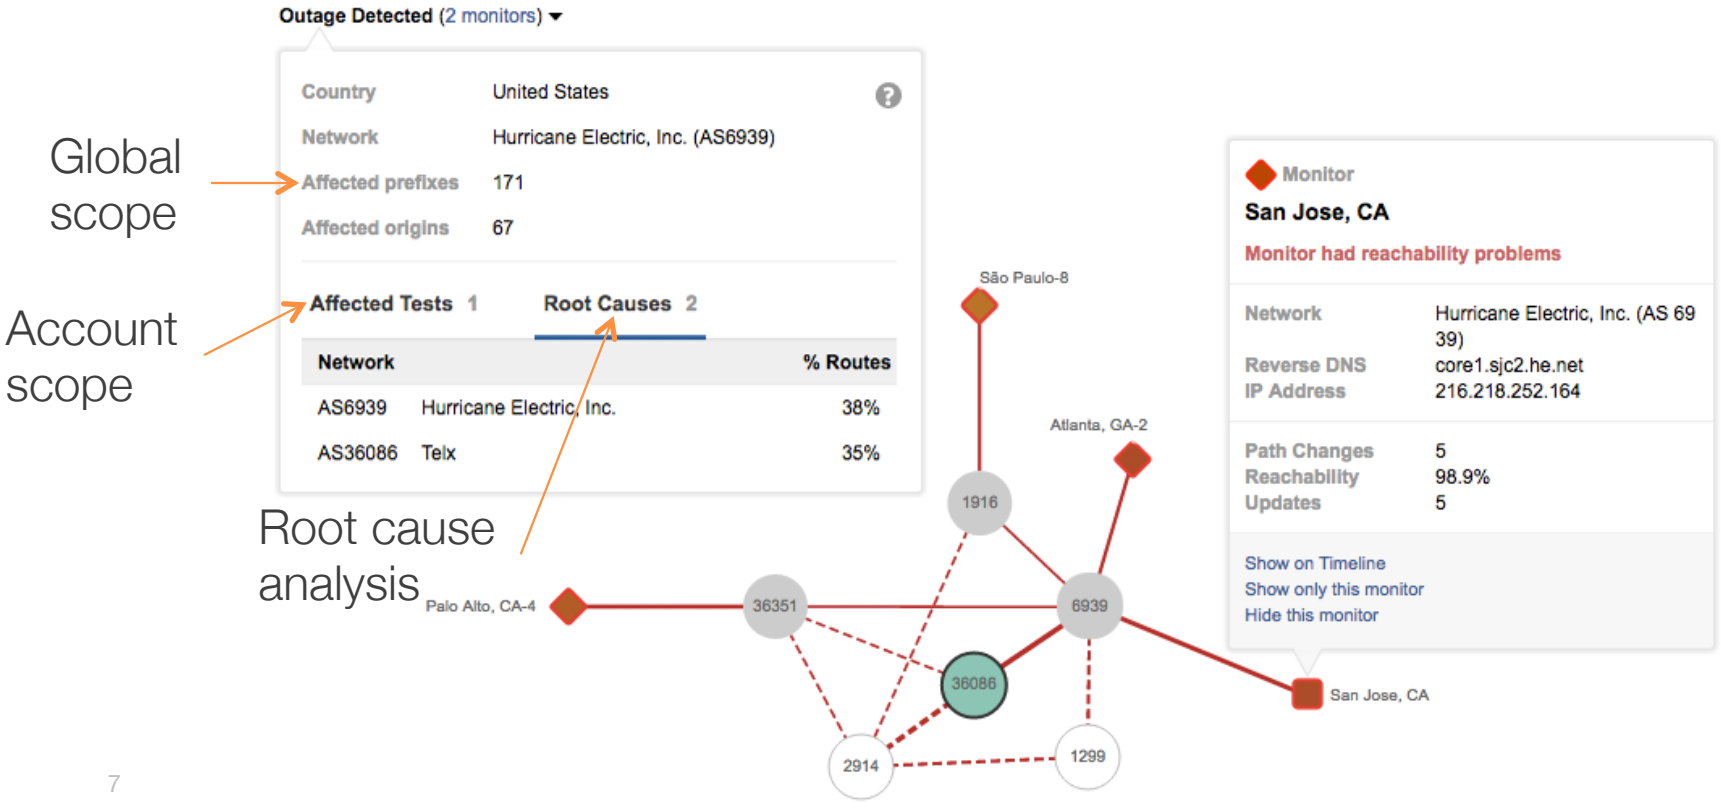

# Internet Outages Happen All the Time

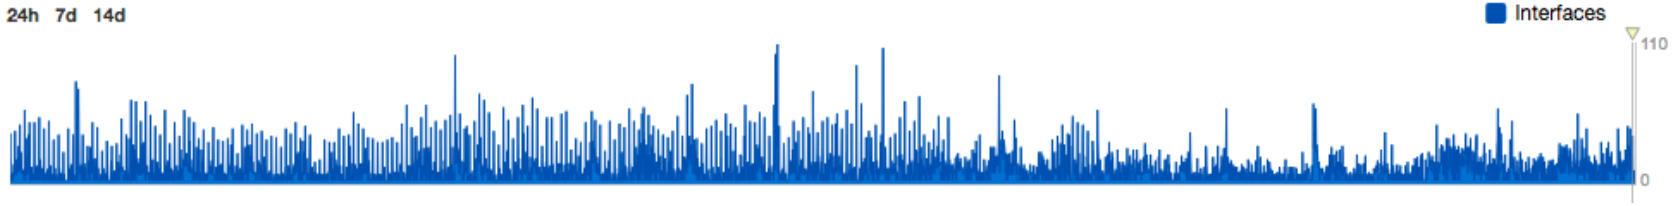


~ 170 affected interfaces / hour

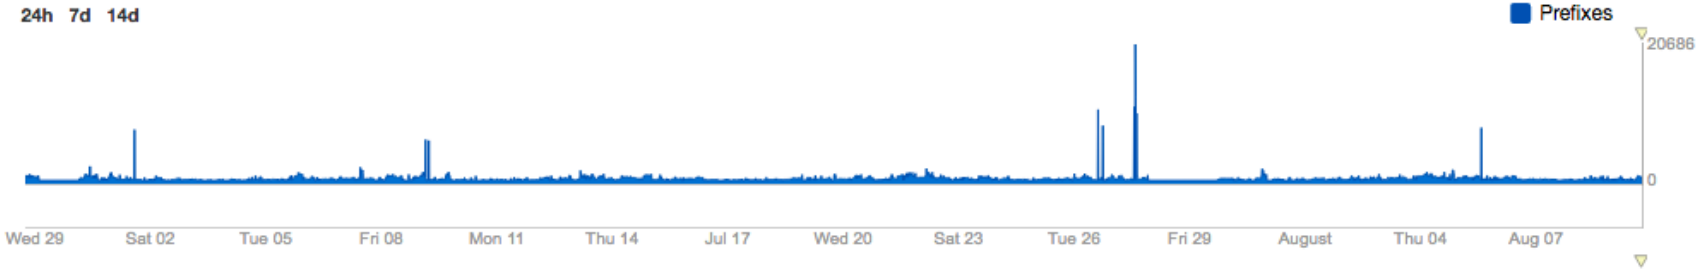


~ 1.6K prefixes / hour

# Recent Major Outages Detected

- April 23: Hurricane Electric route leak affecting AWS
  - https://jfqtbmvy.share.thousandeyes.com
- May 3: Trans-Atlantic issues in Level 3
  - https://blog.thousandeyes.com/trans-atlantic-issues-level-3-network/
- May 20: Tata and TISparkle issues with submarine cable
  - https://blog.thousandeyes.com/smw-4-cable-fault-ripple-effects-across-networks/
- June 6: Hurricane Electric removed >500 prefixes
- June 24: Tata cable cut in Singapore affecting Dropbox
  - https://gedopd.share.thousandeyes.com
- July 10: Level 3, NTT routing issues affecting JIRA
  - https://blog.thousandeyes.com/identifying-root-cause-routing-outage-detection/
- July 17: Widespread issues in Telia's network in Ashburn
  - https://blog.thousandeyes.com/analyzing-internet-issues-traffic-outage-detection/

# Examples of outages

# 1. Network Layer Issues in Telia in Ashburn

https://fvqmu.share.thousandeyes.com/

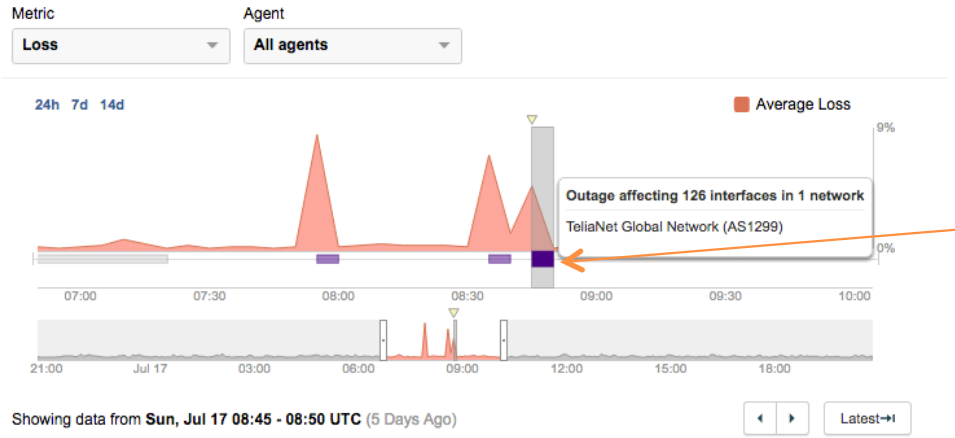


Detected outage coincides with packet loss spikes

Ashburn, VA is “ground zero” for this outage

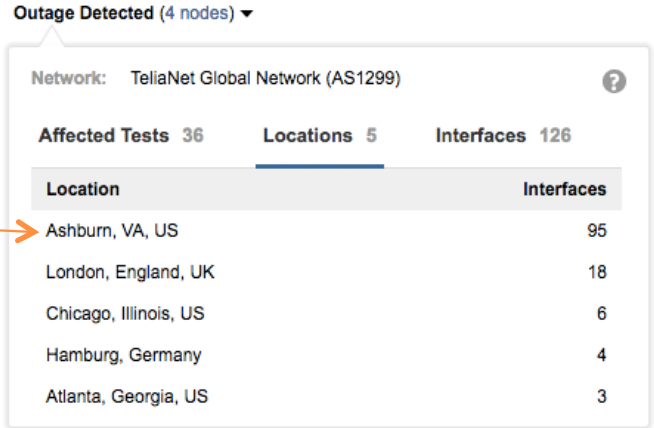

# Specific Failure Points in Telia

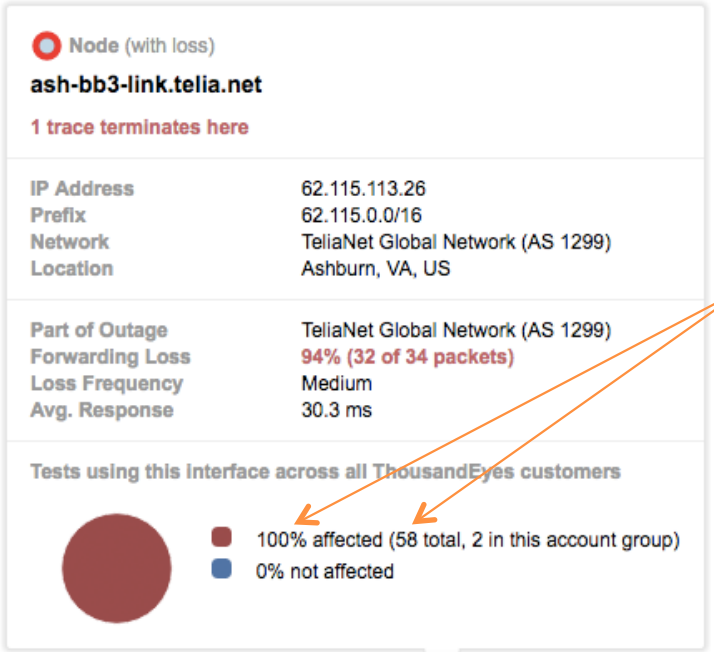


High severity and wide scope (Outages affecting at least 20 tests for a NA/EU interface are likely to be wide in scope)

Terminal nodes in Telia

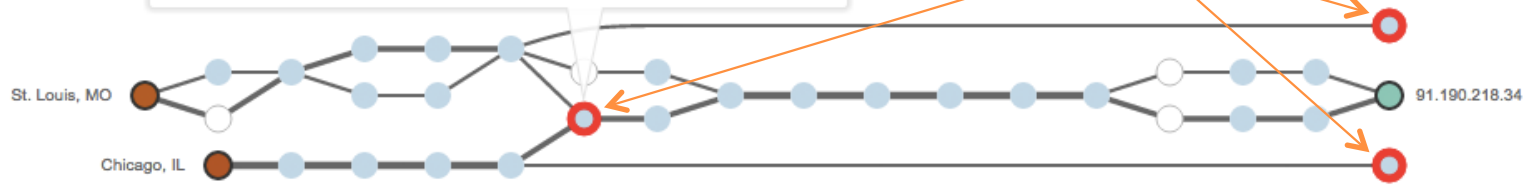

# 2. Hurricane Electric Route Flap

https://njjgkif.share.thousandeyes.com/

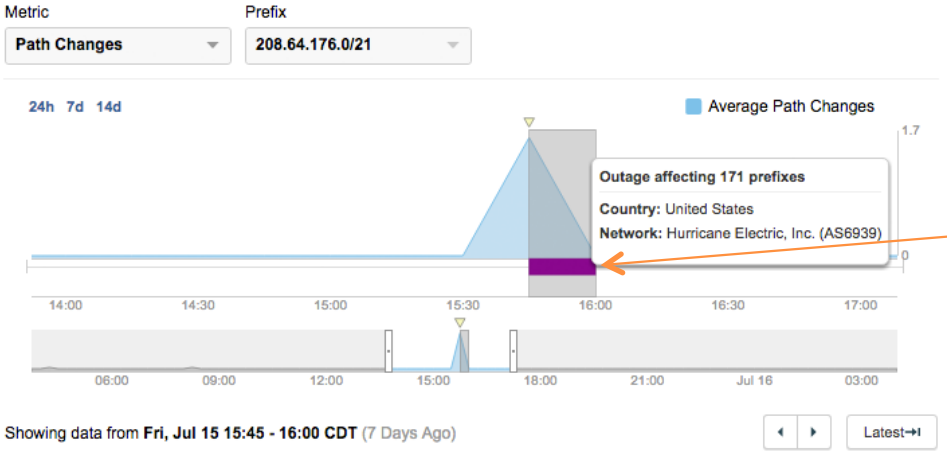


Detected outage coincides with spike in AS path changes

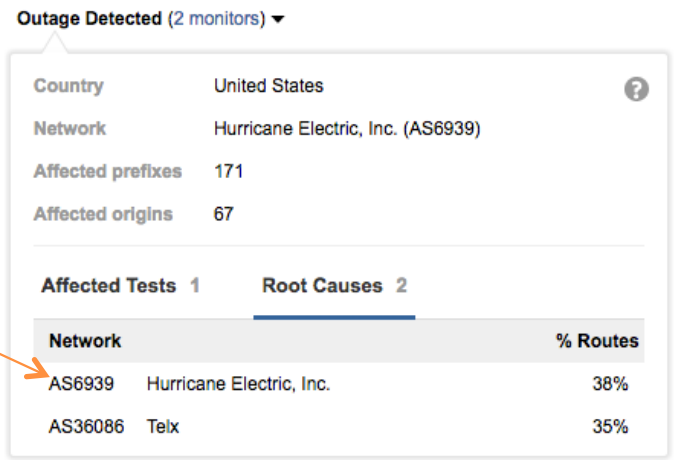


Root cause analysis points to Hurricane Electric and Telx

# Route Flap by Hurricane Electric


Hurricane Electric


Routes flap from using HE to NTT, then back to HE

# Traffic Issues in Hurricane Electric

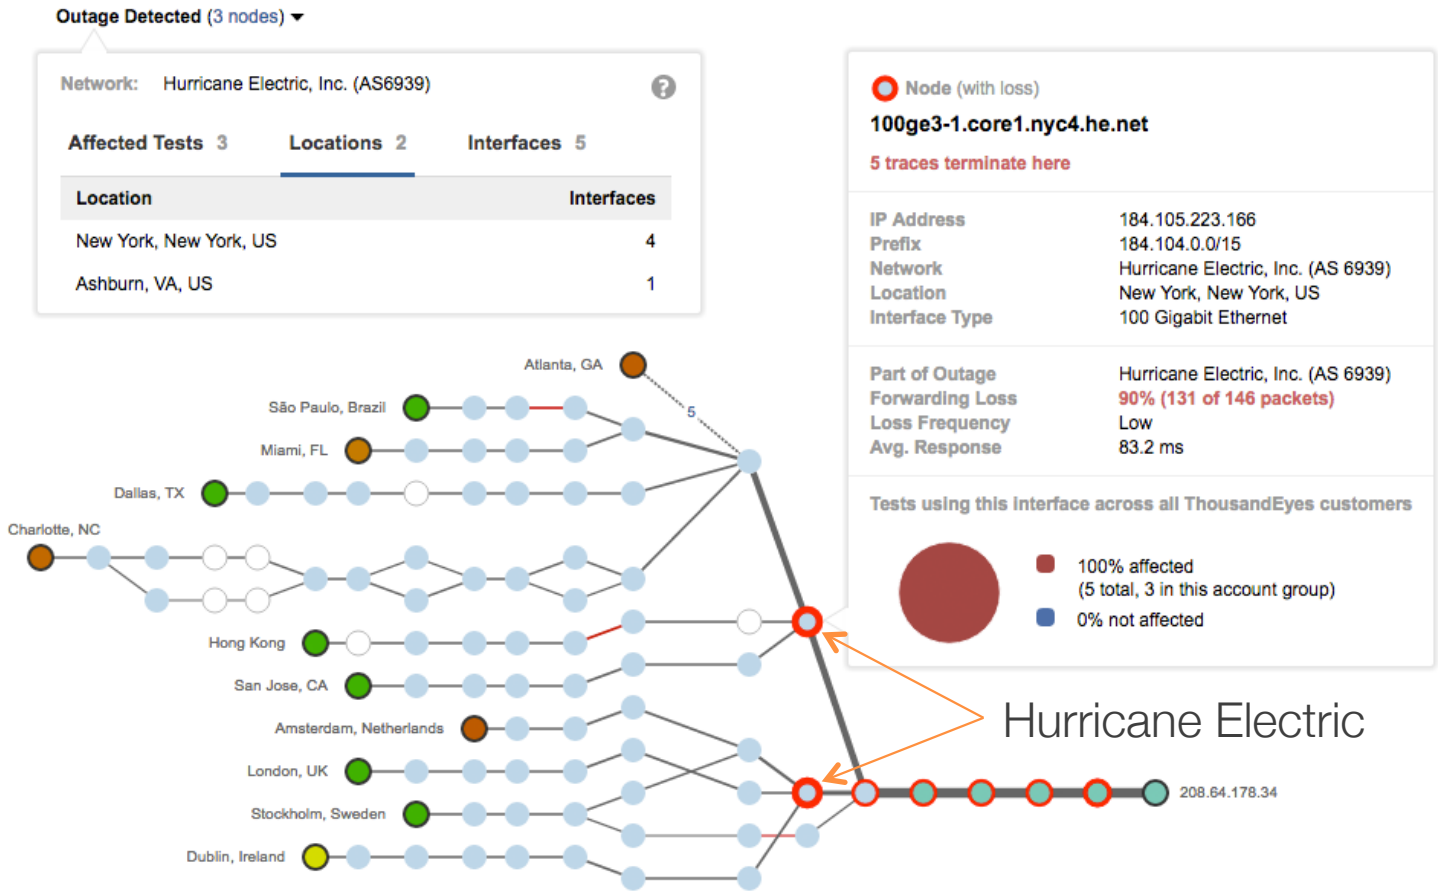

# 3. NTT and Level 3 Routing Issues Affect JIRA

https://ncigwwph.share.thousandeyes.com/

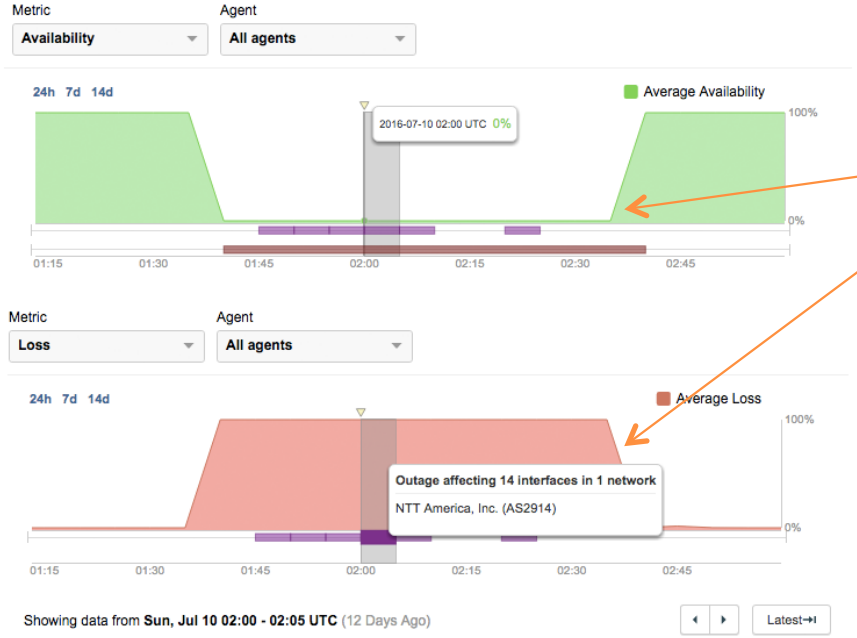


JIRA saw 0% availability and 100% packet loss

**Outage Detected** (6 nodes)

| Affected Tests 1 | Locations 7 | Interfaces 14 |
|---|---|---|

| Location | Interfaces |
|---|---|
| Ashburn, VA, US | 6 |
| San Francisco, California, US | 2 |
| San Jose, California, US | 2 |
| Los Angeles, California, US | 1 |
| Seattle, Washington, US | 1 |
| Singapore, Singapore | 1 |
| Tokyo, Tōkyō, Japan | 1 |

Most affected interfaces are in Ashburn, VA

# JIRA's /24 Prefix Becomes Unreachable

**Outage Detected** (32 monitors)

| | |
|---|---|
| Country | United States |
| Network | Level 3 Communications, Inc. (AS3356) |
| Affected prefixes | 35 |
| Affected origins | 21 |

**Affected Tests** 1 **Root Causes** 1

| Network | | % Routes |
|---|---|---|
| AS3356 | Level 3 Communications, Inc. | 34% |

As the primary upstream ISP, Level 3 is associated with the most affected routes

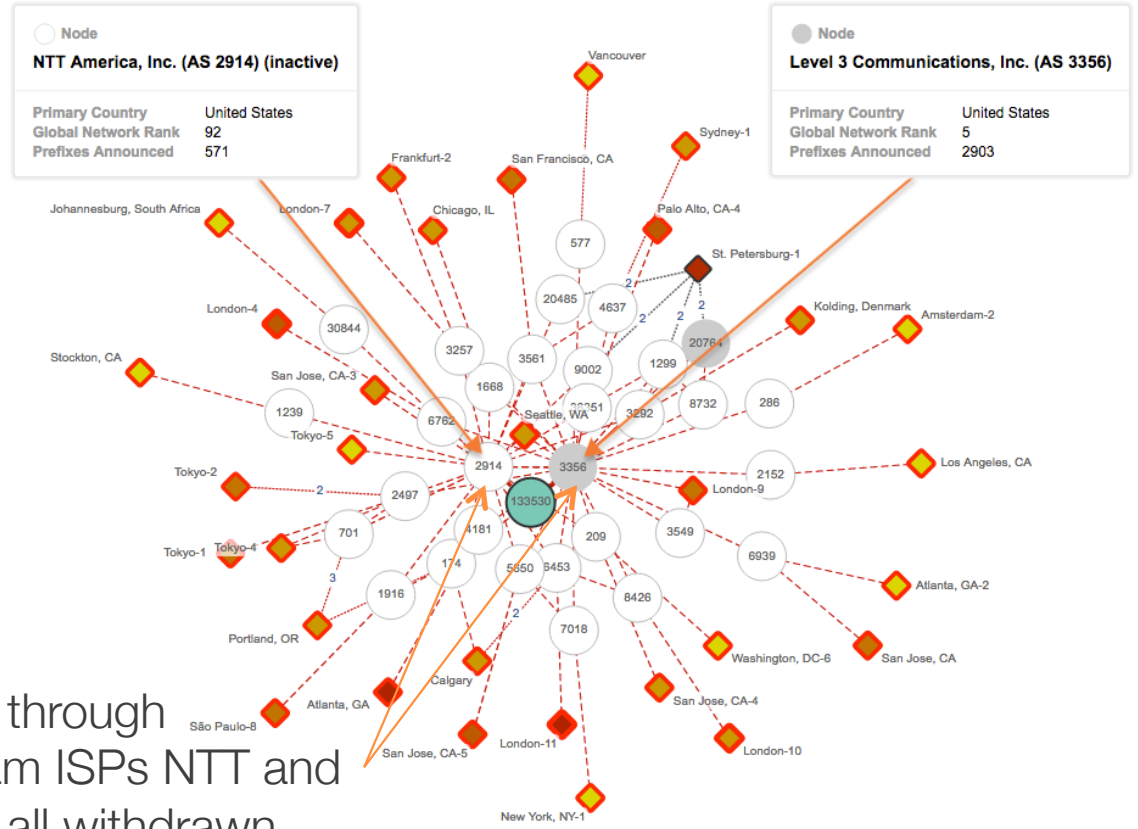


Routes through upstream ISPs NTT and Level 3 all withdrawn

# Routers Begin Using Misconfigured /16 Prefix


The /16 prefix directs to NTT, not JIRA's network. This is why the traffic path changed to traverse only NTT, terminating there when JIRA's IP couldn't be found in NTT's network.

# Traffic Terminating in NTT

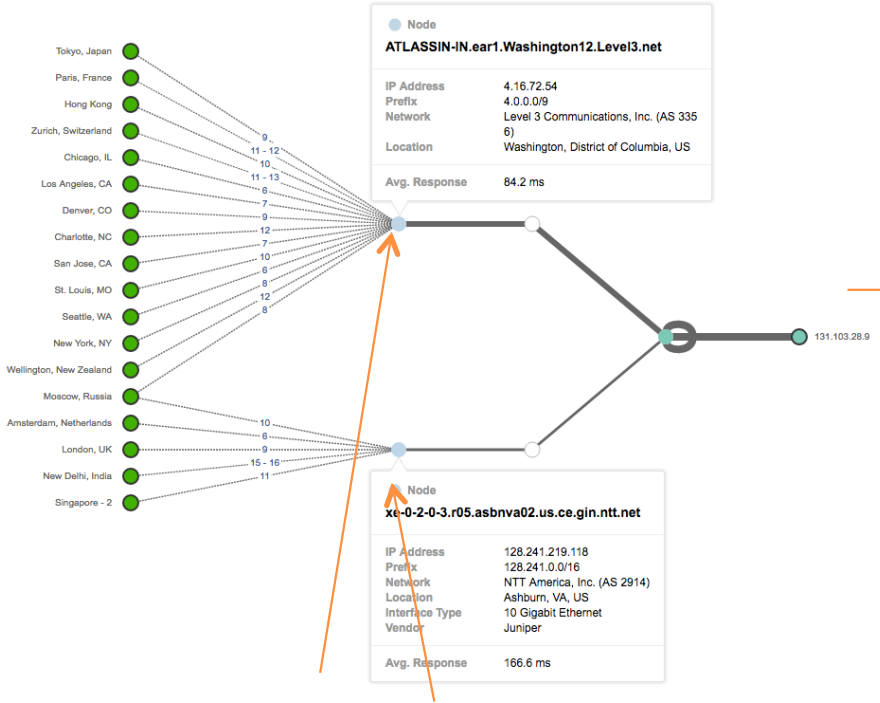


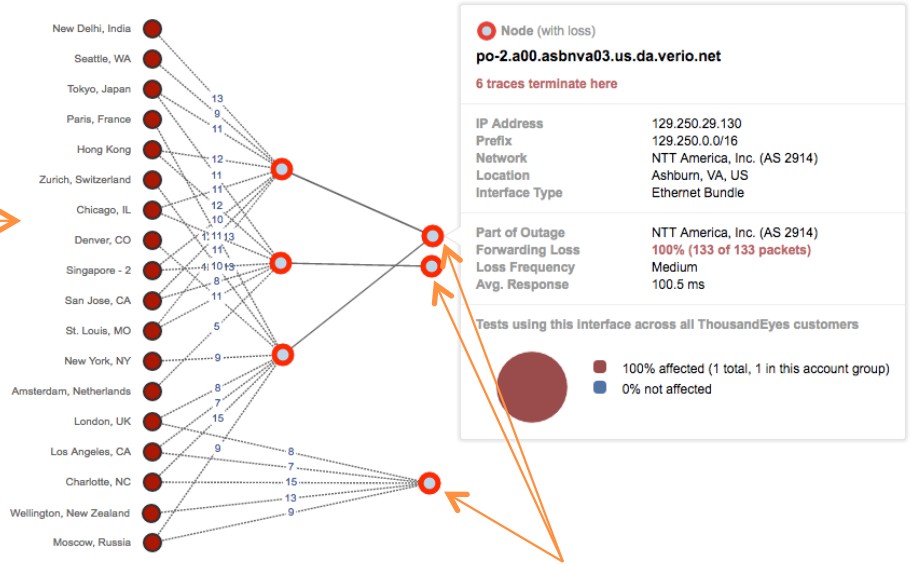


Traffic paths originally traversed Level 3 and NTT to prefix 131.103.28.0/24

Traffic paths then change to traverse only NTT, terminating there


# 4. SEA-ME-WE-4 Cable Fault

# Tata Backbone Under Normal Conditions

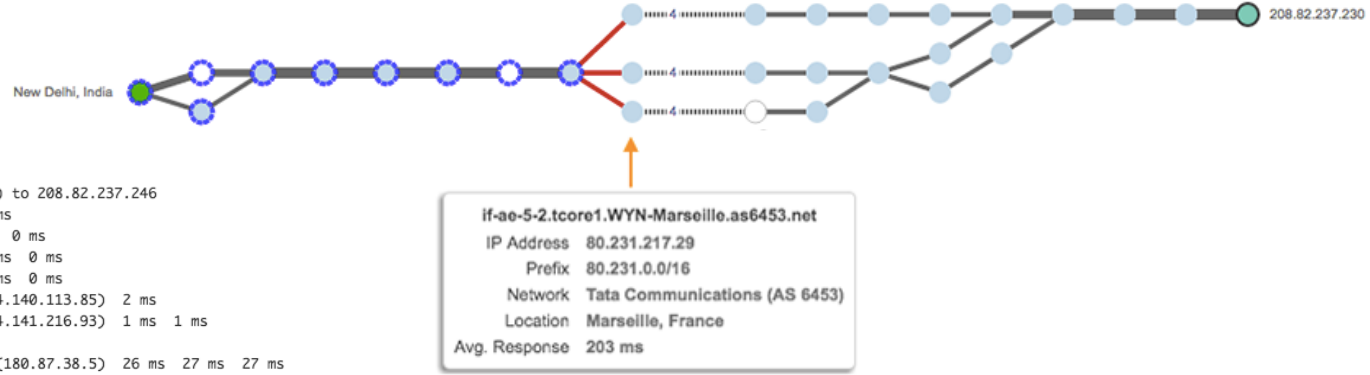


```
Path trace from New Delhi, India (10.10.10.70) to 208.82.237.246
 1  10.10.10.65 (10.10.10.65)  0 ms  0 ms  0 ms
 2  180.179.204.2 (180.179.204.2)  0 ms  0 ms  0 ms
 3  180.179.192.73 (180.179.192.73)  0 ms  0 ms  0 ms
 4  180.179.197.37 (180.179.197.37)  0 ms  0 ms  0 ms
 5  14.140.113.85.static-Delhi-vsnl.net.in (14.140.113.85)  2 ms
    14.141.216.93.static-Delhi-vsnl.net.in (14.141.216.93)  1 ms  1 ms
 6  * * *
 7  ix-ae-0-100.tcore1.MLV-Mumbai.as6453.net (180.87.38.5)  26 ms  27 ms  27 ms
 8  if-ae-9-6.tcore1.WYN-Marseille.as6453.net (80.231.217.77) <MPLS:L=1349058567,E=0,S=1,T=1> 210 ms
    if-ae-5-2.tcore1.WYN-Marseille.as6453.net (80.231.217.29) <MPLS:L=1349058567,E=0,S=1,T=1> 205 ms
    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 m
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
```

# 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


```
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

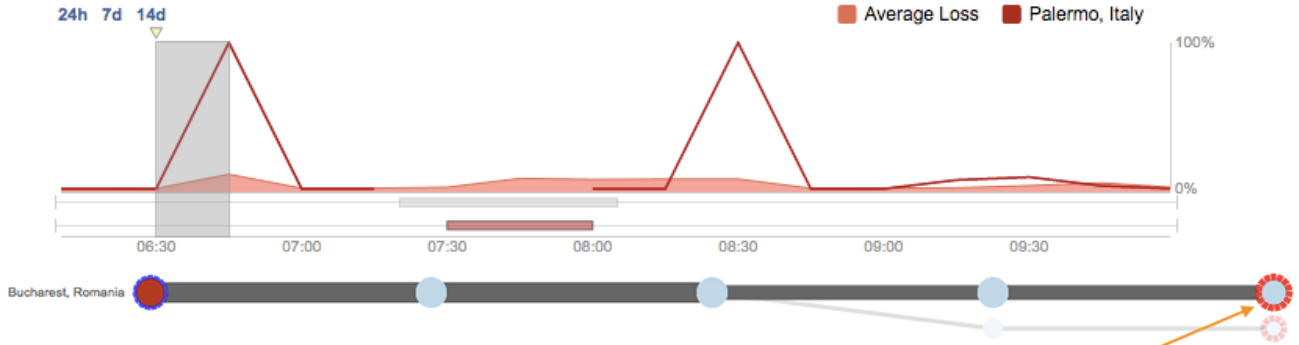


```
Path trace from Bucharest, Romania (93.115.82.43) to 45.57.3.131
 1  buc-ird-35sw.voxility.net (93.115.82.33)  10 ms  0 ms  7 ms
 2  buc-ird-01c.voxility.net (93.114.40.153)  0 ms  0 ms  0 ms
 3  buc-ird-01gw.voxility.net (109.163.235.49)  0 ms  1 ms
    buc-ird-02gw.voxility.net (109.163.235.61)  4 ms
 4  xe-9-0-0.bucarest1.buc.seabone.net (93.186.132.48)  0 ms  0 ms
    xe-10-0-0.bucarest1.buc.seabone.net (93.186.132.50)  14 ms
```

# European Detour In Effect

- Netflix (AS2906) begins to route via Level 3 (AS3356) instead of directly through TISparkle (AS6762)
- Traffic flowed via Frankfurt rather than Paris (and Marseilles)

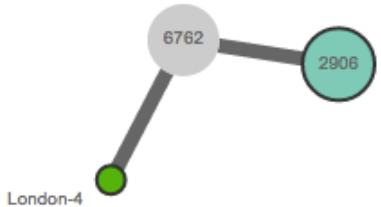


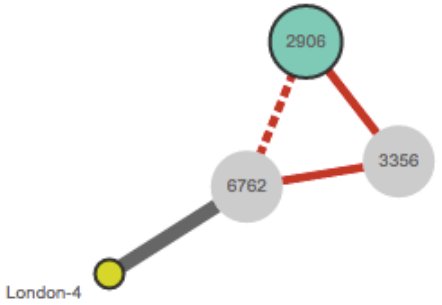

# What Commonalities between Tata and TIS?

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

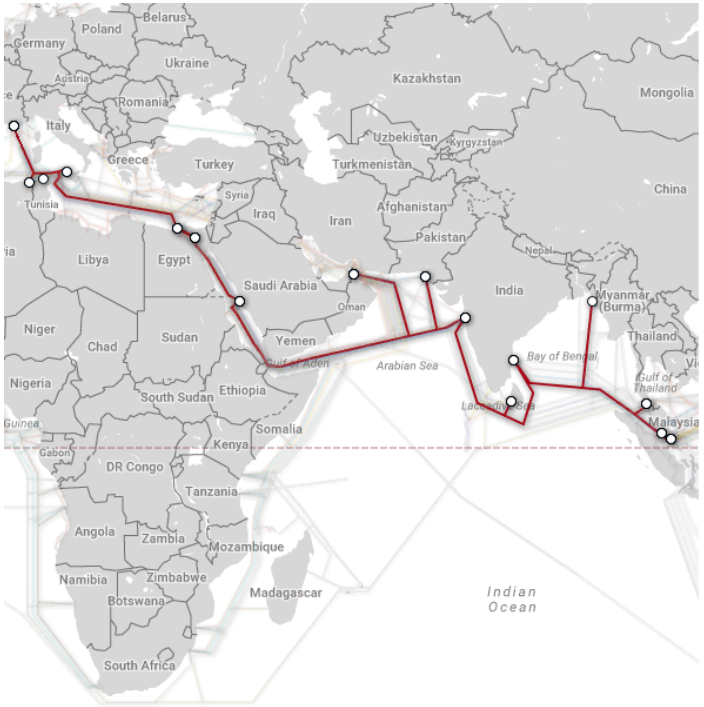


Submarine Cable List

**SeaMeWe-4**

Email link

RFS: December 2005
Cable Length: 20,000 km
Owners: Bangladesh Telegraph and Telephone Board (BTTB), 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

**Landing Points**

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


http://www.submarinecablemap.com/#/submarine-cable/seamewe-4

# SEA-ME-WE-4 Cable

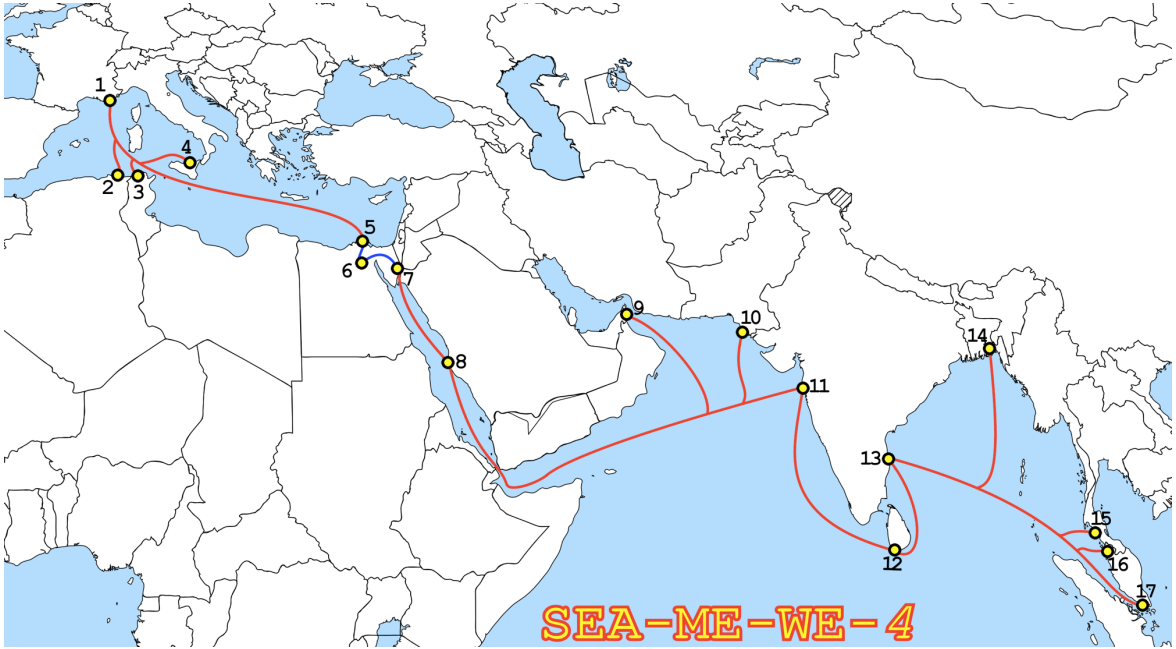


## Background

- 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

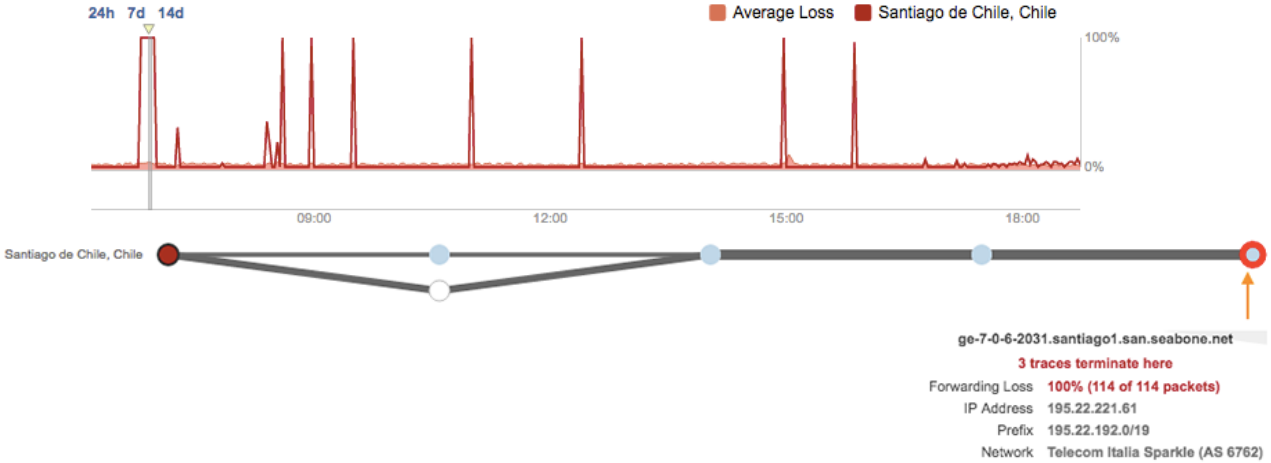

# Why Would This Impact the Americas?

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

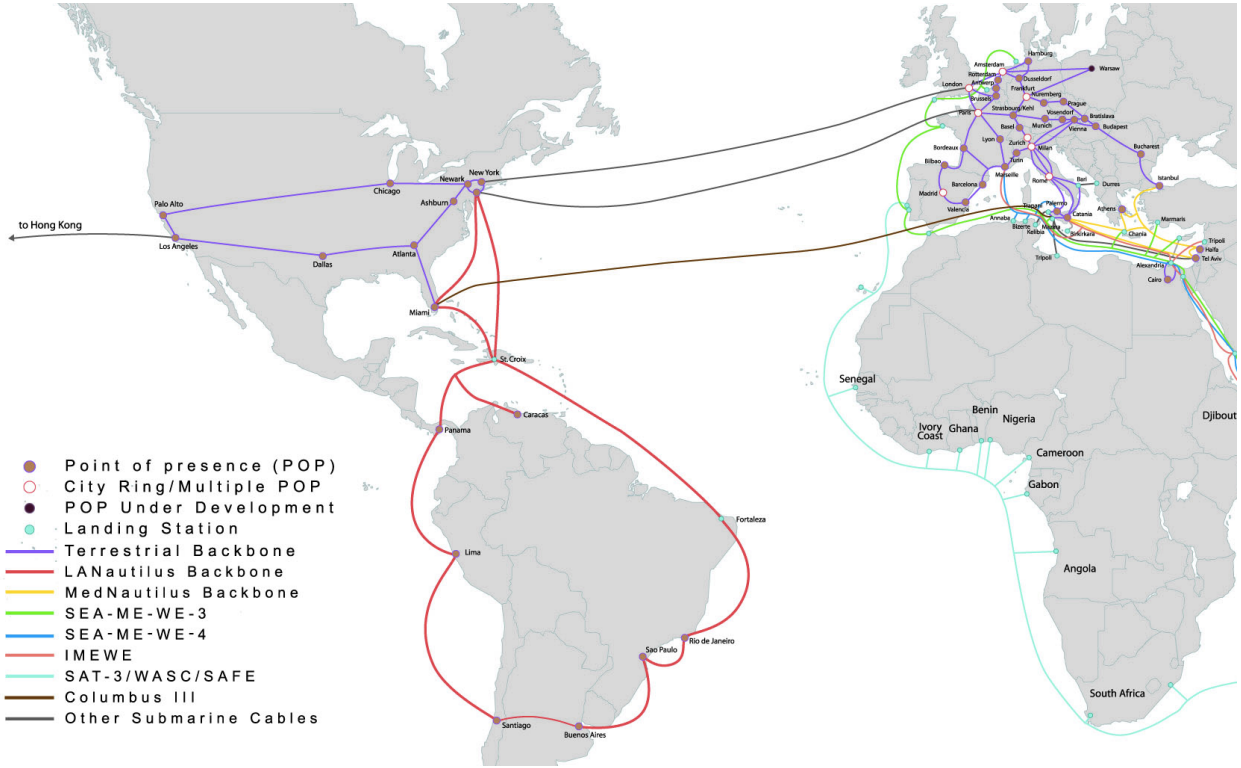

# And BGP Sessions Begin to Fail

- 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)
- Hence, faults in a Mediterranean cable can disrupt BGP sessions that cause dropped traffic worldwide

# 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 |



http://www.gossamer-threads.com/lists/nsp/outages/58551?do=post_view_threaded

ThousandEyes

Thank You
@rveloso

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