Detecting and Analyzing Peering Infrastructure Outages

Vasileios Giotsas <vgiotsas@caida.org>

Joint work with:
Christopher Dietzel, Georgios Smaragdakis,
Anja Feldmann, Arthur Berger, Emile Aben

Mission Statement

- Develop a novel methodology to detect, localize and analyze outages at peering infrastructures in near-real time:
  - Combine **BGP communities** with high-resolution co-location maps to achieve accurate **passive detection**
- Implement the proposed methodology to analyze 5 years of historical RouteViews and RIPE RIS BGP data:
  - At least **2x** more outages than what is publicly reported
  - Challenge the perception that “*local outages have local impact*”
  - Illuminate the routing behaviour during PoP-level outages

The Internet undergoes fundamental shift from a hierarchical to flattened topology

- Direct traffic exchange between edge ASes to support the increasing consumption of content, video, and cloud services [1].

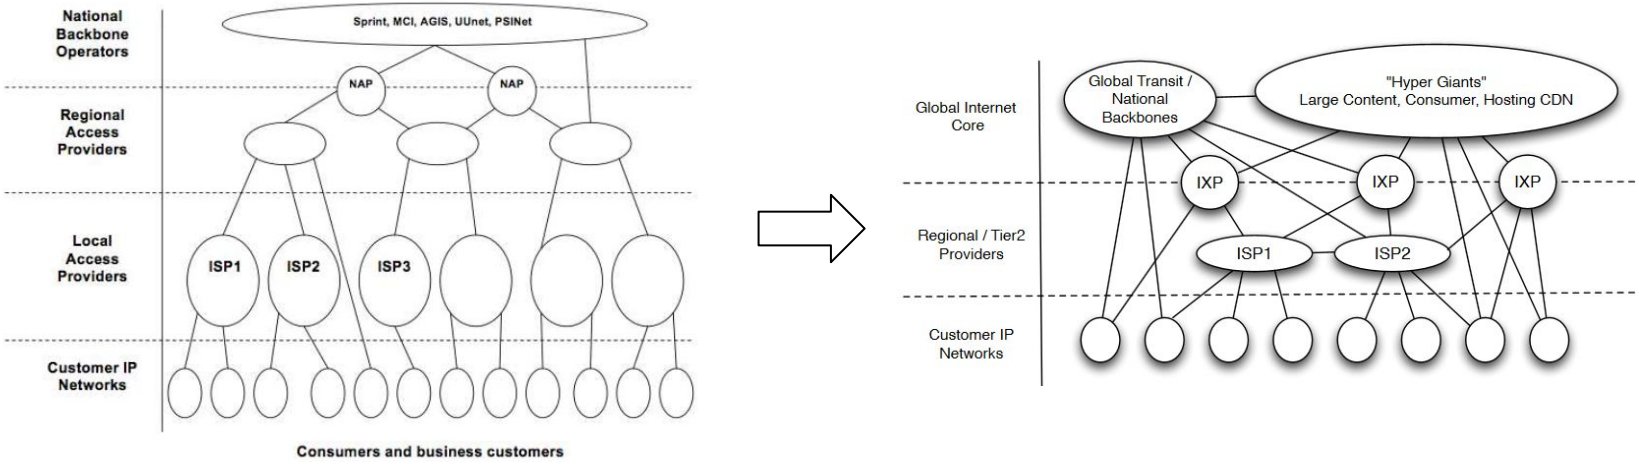

Source:Labovitz, Craig, et al. "Internet inter-domain traffic." ACM SIGCOMM 2010.

Key enablers of dense peering

- Carrier-neutral **co-location facilities** (CFs) provide infrastructure for physical co-location and cross-connect interconnections.
- **Internet Exchange Points** (IXPs) provide a shared switching fabric for layer-2 bilateral and multilateral peering.
- CFs and IXPs are largely **symbiotic**:
  - IXPs co-locate in multiple CFs to facilitate connectivity to their fabric
  - CFs subsidise the presence of IXPs in their space to increase their colocation incentives.

Example: Infrastructures interdependencies of FranceIX

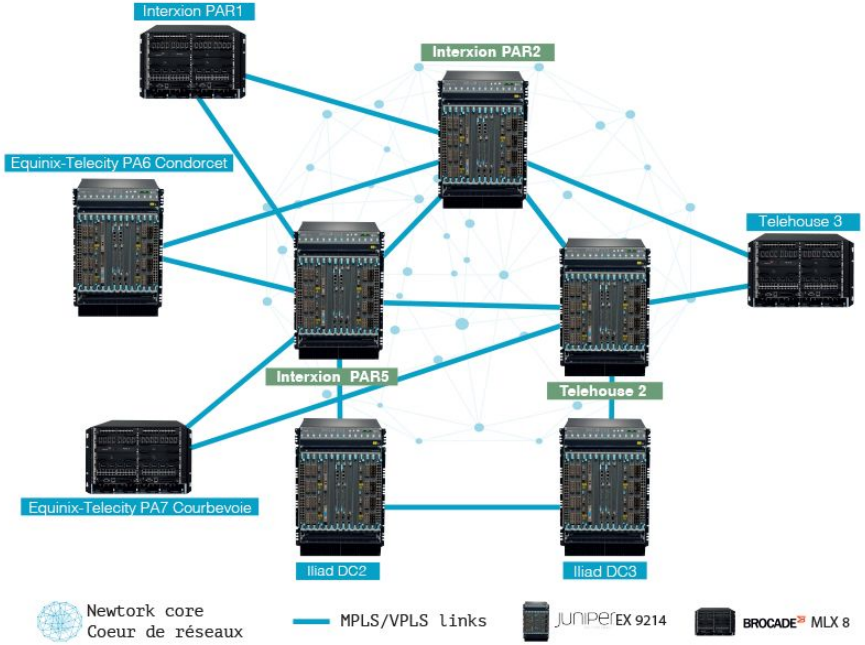

Source: https://www.franceix.net/en/technical/infrastructure/

Geographic agglomeration leads to the formation of massive interconnection hubs

- IXPs and CFs deployed in large metro areas to be close to users and creating tightly-interconnected infrastructure clusters.
- Networks effects of co-location:
  - Over 1,100 CFs and >300 IXPs globally.
  - Large IXPs (AMS-IX, DE-CIX, LINX) have over 700 members, support over **100K** peerings and carry up to **5 Tb/s** peak traffic [2].
  - The largest facilities operator (Equinix) hosts >1000 networks across its CFs that establish over **170K** cross-connects [3].

Remote peering extends the reach of IXPs and CFs beyond their local market

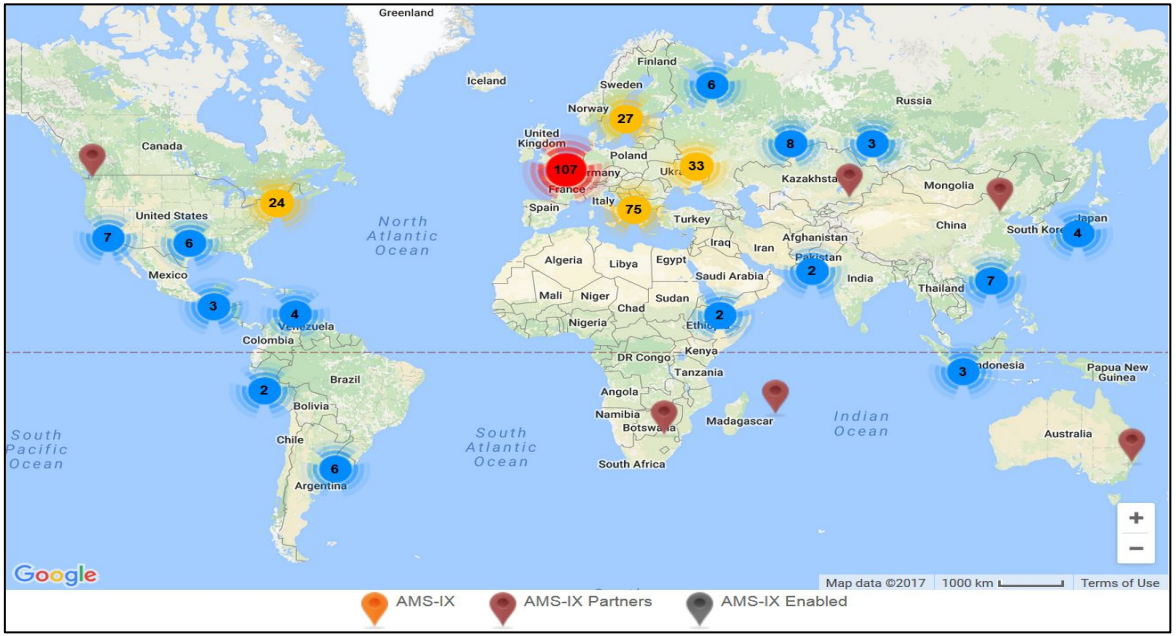

Global footprint of AMS-IX

Source: https://ams-ix.net/connect-to-ams-ix/peering-around-the-globe

Peering infrastructures are critical part of national and international ICT sector

- Governmental agencies consider IXPs and CFs critical for national cyber-security and communications:
  - DHS Regional Resiliency Assessment Program (RRAP) [4]
  - EU's Critical Information infrastructure Protection (CIIP) program [5]
- Although well-provisioned and managed prone to failures, e.g., power outages, human errors, attacks, misconfiguration and bugs
- The average cost of a single outage is estimated over **$740K**, a **38% increase** between 2010 – 2015 [6]

Outages in peering infrastructures can severely disrupt critical services and applications

DOWNTIME WOES —

BT, other ISPs hit by second major Internet outage—power failure blamed

After Telecity power outage, it seems Telehouse has had problems of its own.

KELLY FIVEASH - 21/7/2016, 03:05

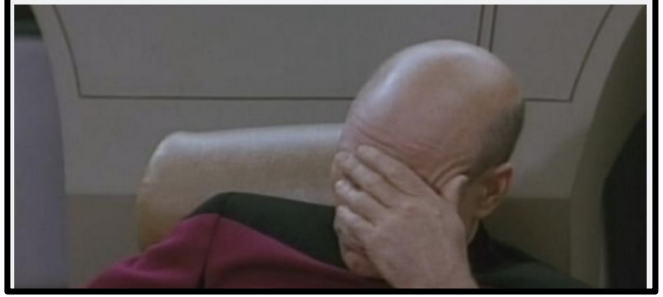

BT broadband users hit by second UK-wide outage in two days

Caroline Donnelly
Datacentre Editor
21 Jul 2016 9:50

Power supply issues at Docklands datac
behind loss of internet access for more t
broadband users

Equinix cooling outage leads to flight delays in Australia

13 November 2012 | By Penny Jones

Share Tweet Like 0

A "short interruption to utility power supply" at an Equinix data center in Sydney caused up to three hour delays for thousands of passengers flying with three major airlines from Australian airports over the weekend.

DOWNTIME

Equinix Outage Means Downtime for Zoho

BY RICH MILLER ON JANUARY 20, 2012 ADD YOUR COMMENTS

A power outage Friday morning in an Equinix data center in
California caused prob
ably Zoho, which ex
veral of its web-base
knowledged the inci
cause of the outag

TECHNOLOGY | TOP STORIES

OUTAGE AT AMSTERDAM INTERNET HUB AFFECTS MUCH OF NETHERLANDS

By Janene Pieters on May 13, 2015 - 13:11

With additional reporting by Zack Newmark.

A technical fault at the internet hub AMS-IX in Amsterdam caused online problems in several places in the Netherlands for about an hour Wednesday afternoon. The internet hub, one of the most used internet exchanges in the world, announced they resolved the problem shortly after 1:30 p.m.

The Register
Biting the hand that feeds IT

DATA CENTRE | SOFTWARE | SECURITY | TRANSFORMATION | DEVOPS | BUSINESS | PERSONAL TECH

Data Centre

Telecity London data centre outage borks VoIP, websites, AWS...

LINX reports sudden sharp traffic drop, Amazon Direct Connect goes TITSUP

Outages in peering infrastructures can severely disrupt critical services and applications

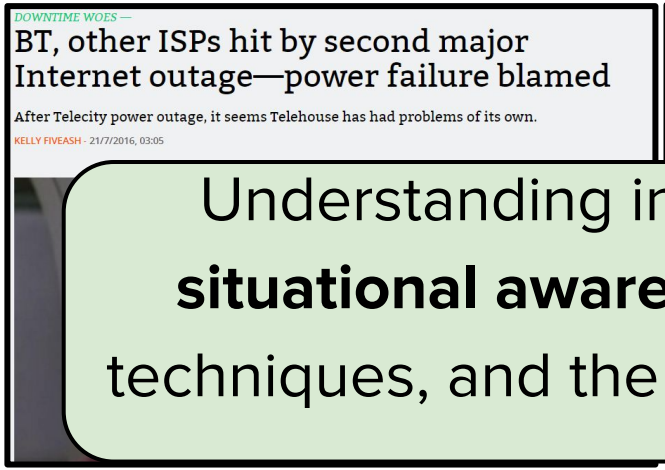
DOWNTIME WOES —

BT, other ISPs hit by second major Internet outage—power failure blamed

After Telecity power outage, it seems Telehouse has had problems of its own.

KELLY FIVEASH - 21/7/2016, 03:05


BT broadband users hit by second UK-wide outage in two days

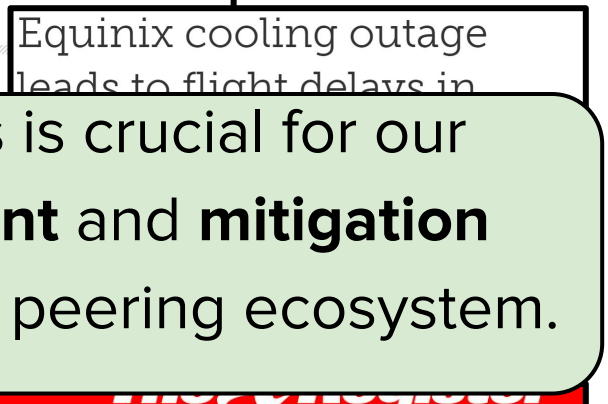
Equinix cooling outage leads to flight delays in

Understanding infrastructure outages is crucial for our **situational awareness**, **risk assessment** and **mitigation** techniques, and the **transparency** of the peering ecosystem.


TECHNOLOGY | TOP STORIES

OUTAGE AT AMSTERDAM INTERNET HUB AFFECTS MUCH OF NETHERLANDS

By Janene Pieters on May 13, 2015 - 13:11

With additional reporting by Zack Newmark.

A technical fault at the internet hub AMS-IX in Amsterdam caused online problems in several places in the Netherlands for about an hour Wednesday afternoon. The internet hub, one of the most used internet exchanges in the world, announced they resolved the problem shortly after 1:30 p.m.


Biting the hand that feeds IT

DATA CENTRE | SOFTWARE | SECURITY | TRANSFORMATION | DEVOPS | BUSINESS | PERSONAL TECH

Data Centre

Telecity London data centre outage borks VoIP, websites, AWS...

LINX reports sudden sharp traffic drop, Amazon Direct Connect goes TITSUP

Our Research Goals

- Detection of peering infrastructure outages:
  - *Timely* detection at the *finest granularity* possible
  - *Low measurement overhead* to enable real-world deployment
- Pinpoint the location of outages:
  - Distinguish *knock-on effects* from outage *source*
- Track changes and assess impact outages:
  - Determine the start and duration of an outage
  - Track the shifts in routing paths, the geographic spread, and effect on performance metrics, e.g., end-to-end delays

Current practice: “Is anyone else having issues?”

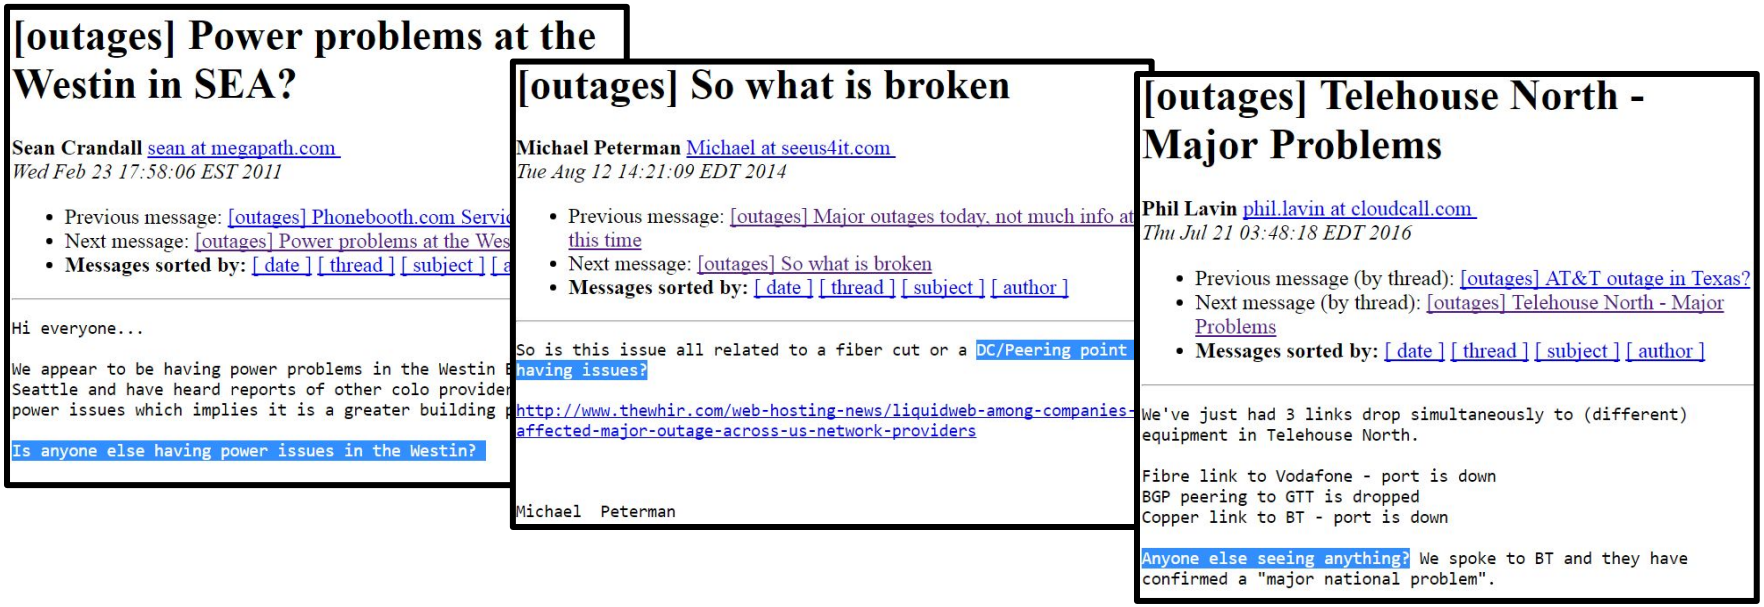

[outages] Power problems at the Westin in SEA?

Sean Crandall sean at megapath.com
Wed Feb 23 17:58:06 EST 2011

- Previous message: [outages] Phonebooth.com Servi
- Next message: [outages] Power problems at the Wes
- **Messages sorted by:** [date] [thread] [subject] [a

Hi everyone...

We appear to be having power problems in the Westin B
Seattle and have heard reports of other colo provider
power issues which implies it is a greater building p

Is anyone else having power issues in the Westin?

[outages] So what is broken

Michael Peterman Michael at seeus4it.com
Tue Aug 12 14:21:09 EDT 2014

- Previous message: [outages] Major outages today, not much info at this time
- Next message: [outages] So what is broken
- **Messages sorted by:** [date] [thread] [subject] [author]

So is this issue all related to a fiber cut or a DC/Peering point having issues?

http://www.thewhir.com/web-hosting-news/liquidweb-among-companies-affected-major-outage-across-us-network-providers

Michael Peterman

[outages] Telehouse North - Major Problems

Phil Lavin phil.lavin at cloudcall.com
Thu Jul 21 03:48:18 EDT 2016

- Previous message (by thread): [outages] AT&T outage in Texas?
- Next message (by thread): [outages] Telehouse North - Major Problems
- **Messages sorted by:** [date] [thread] [subject] [author]

We've just had 3 links drop simultaneously to (different) equipment in Telehouse North.

Fibre link to Vodafone - port is down
BGP peering to GTT is dropped
Copper link to BT - port is down

Anyone else seeing anything? We spoke to BT and they have confirmed a "major national problem".

- ASes try to crowdsource the detection and localization of outages
- Inadequate transparency/responsiveness from infrastructure operators

Detection of infrastructure-level outages is hard!

- Difficult to observe the effect of outages in BGP paths:
  - Country-level, AS-level and prefix-level outages cause global changes in network reachability/visibility.
  - Infrastructure outages affect PoPs (Points of Presence), not entire networks ⇒ Hard to distinguish from normal routing dynamics
- Mapping facility-level interconnections through traceroute campaigns incurs very high measurement cost [7,8]
- Data mining from mailing lists/social nets has varying levels of accuracy, detail and timeliness [9]

Challenges in detecting infrastructure outages

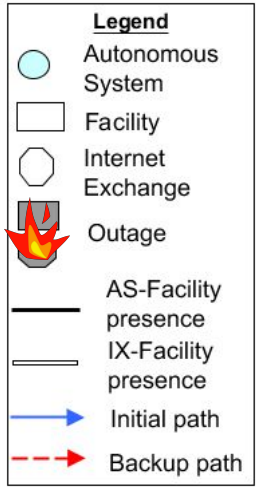

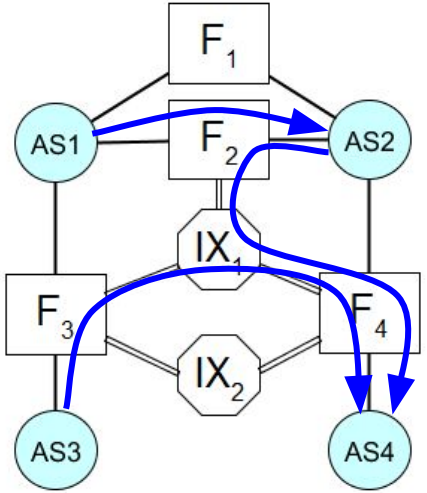

Initial paths:

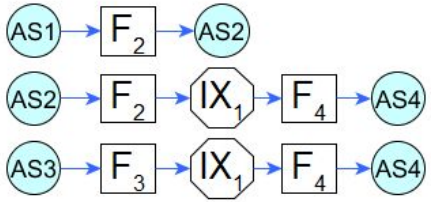

Challenges in detecting infrastructure outages

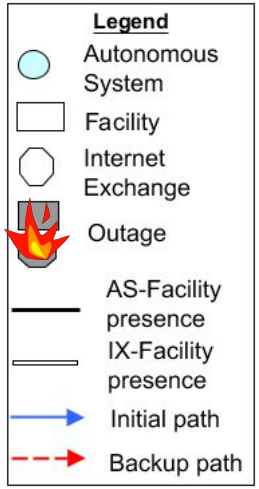

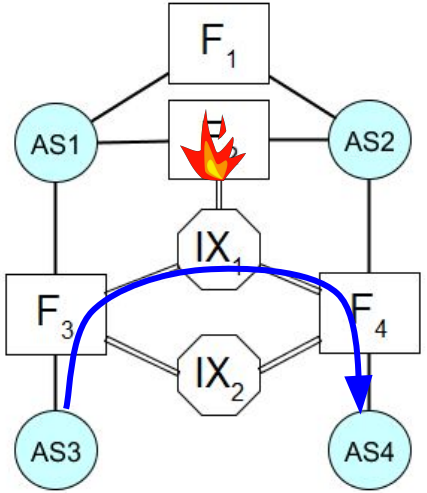

Initial paths:

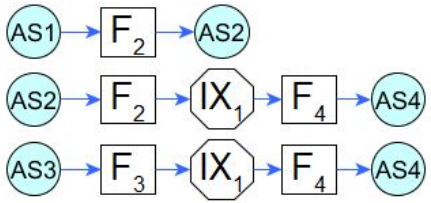

Challenges in detecting infrastructure outages

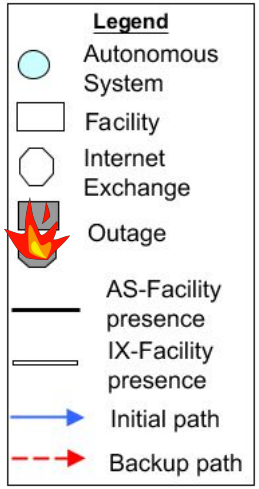

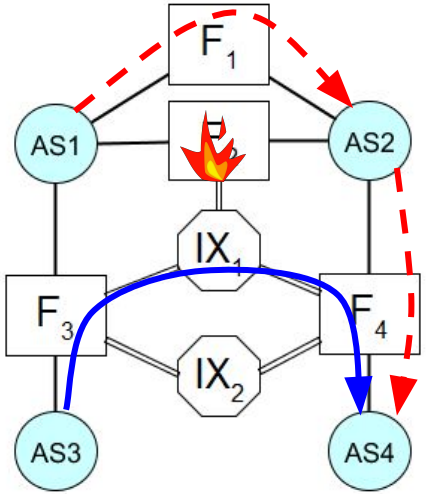

Initial paths:

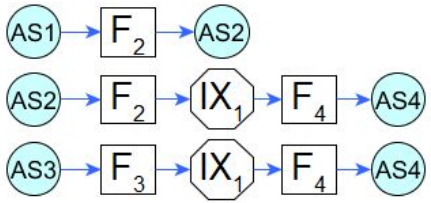

Paths after outage:

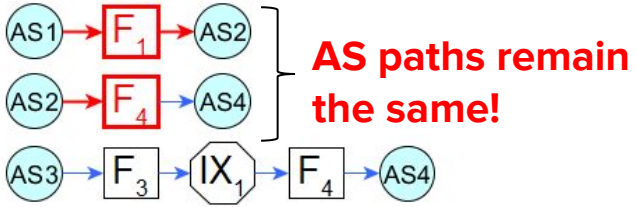

Requirement 1: Capture the PoP-level hops between ASes

Challenge: Information-hiding nature of BGP

Challenges in detecting infrastructure outages

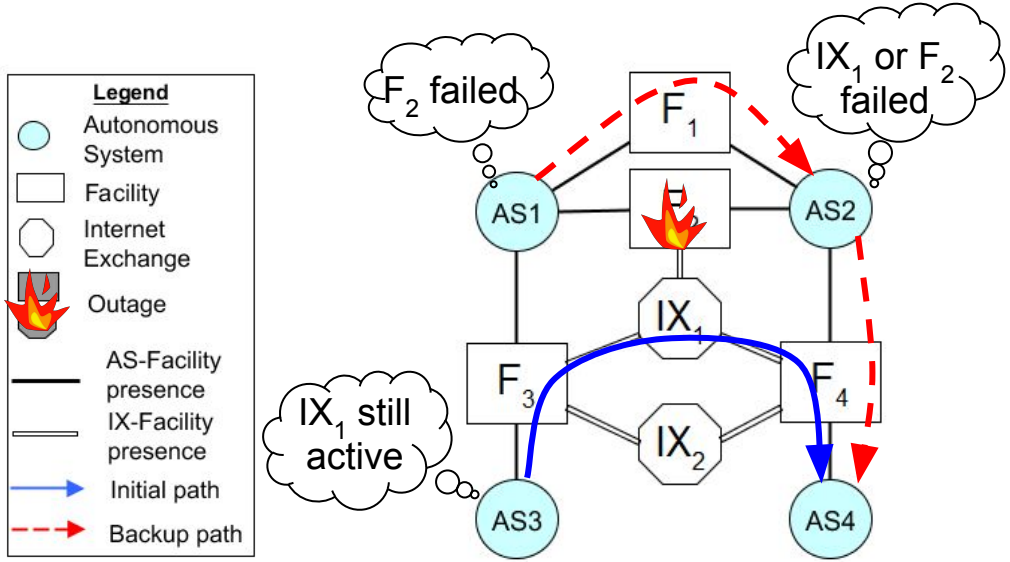

Initial paths:

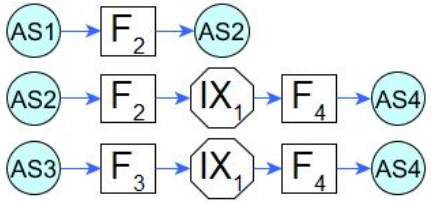

Paths after outage:

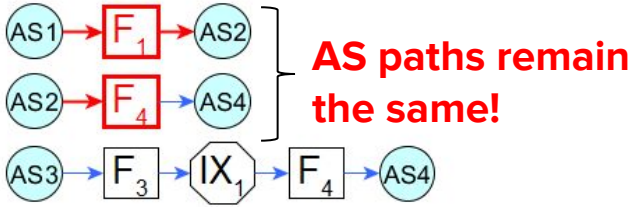

Requirement 2: Correlate the views of multiple vantage points

Challenge: Extensive mapping of PoP-level co-location

Challenges in detecting infrastructure outages

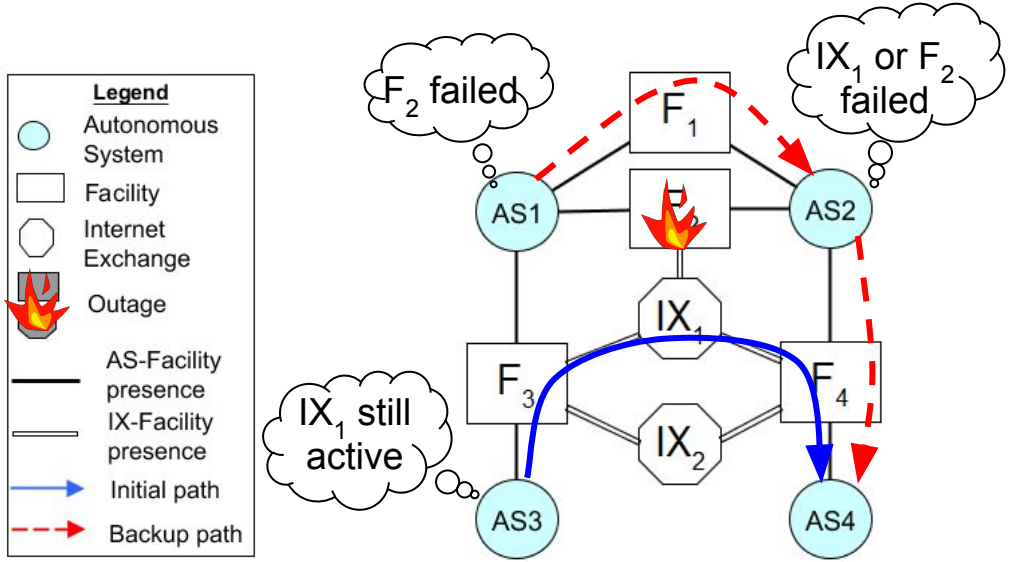

Initial paths:

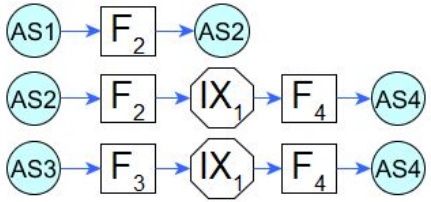

Paths after outage:

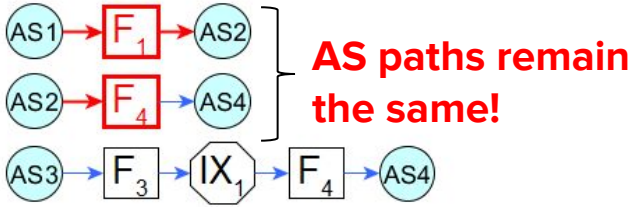

AS paths remain the same!

Requirement 3: Know both “healthy” and backup routes

Challenge: Minimize measurement overhead

Our approach: Decipher PoP-level topology using the BGP Communities attribute

- Optional attribute used to attach arbitrary metadata on BGP routes
  - Often encode the ***ingress location*** of prefixes
- 32-bit values with **X:Y** format:
  - **X** (top 16 bits) encodes the ASN that defines the Community value
  - **Y** (bottom 16 bits) encodes the actual metadata

Our approach: Decipher PoP-level topology using the BGP Communities attribute

- Optional attribute used to attach arbitrary metadata on BGP routes
  - Often encode the ***ingress location*** of prefixes
- 32-bit values with **X:Y** format:
  - **X** (top 16 bits) encodes the ASN that defines the Community value
  - **Y** (bottom 16 bits) encodes the actual metadata

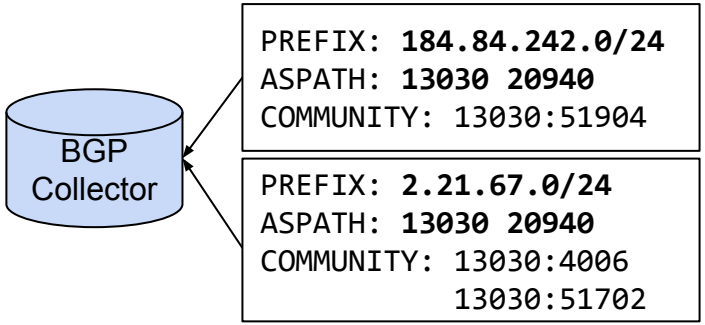

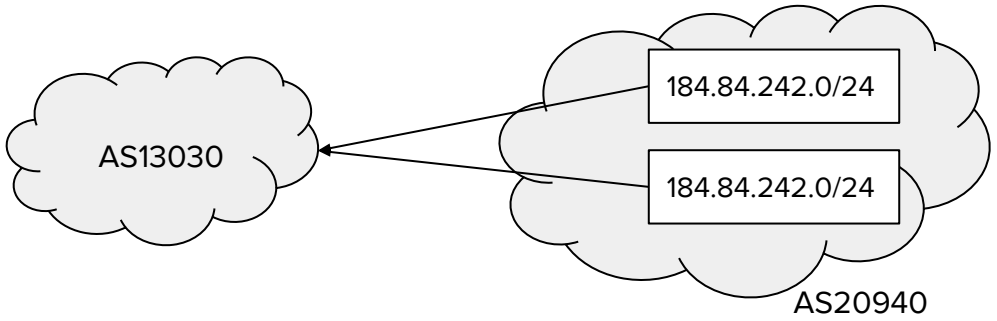

Our approach: Decipher PoP-level topology using the BGP Communities attribute

- Optional attribute used to attach arbitrary metadata on BGP routes
  - Often encode the ***ingress location*** of prefixes
- 32-bit values with **X:Y** format:
  - **X** (top 16 bits) encodes the ASN that defines the Community value
  - **Y** (bottom 16 bits) encodes the actual metadata

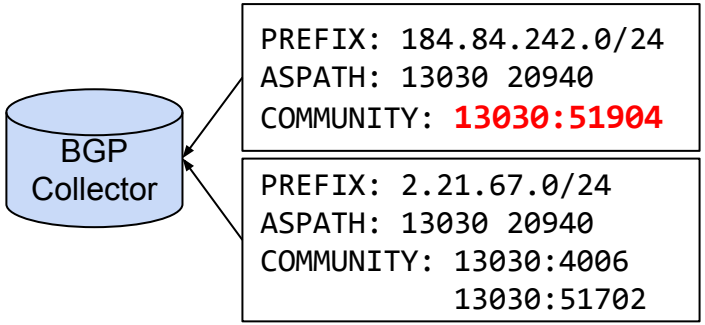

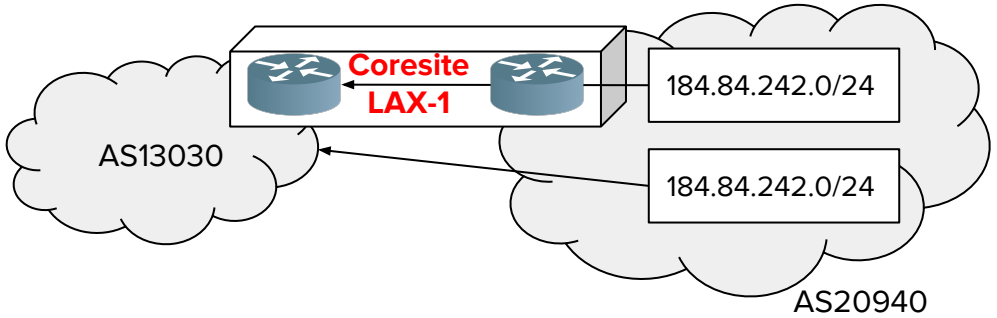

Our approach: Decipher PoP-level topology using the BGP Communities attribute

- Optional attribute used to attach arbitrary metadata on BGP routes
  - Often encode the ***ingress location*** of prefixes
- 32-bit values with **X:Y** format:
  - **X** (top 16 bits) encodes the ASN that defines the Community value
  - **Y** (bottom 16 bits) encodes the actual metadata

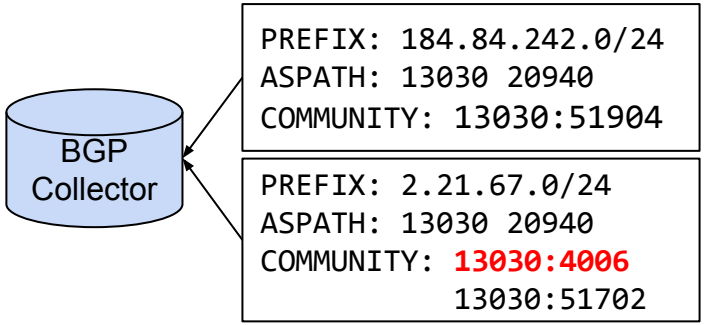

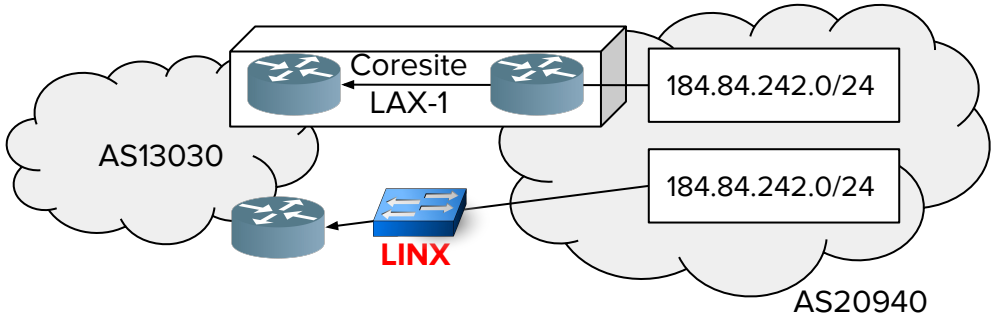

Our approach: Decipher PoP-level topology using the BGP Communities attribute

- Optional attribute used to attach arbitrary metadata on BGP routes
  - Often encode the ***ingress location*** of prefixes
- 32-bit values with **X:Y** format:
  - **X** (top 16 bits) encodes the ASN that defines the Community value
  - **Y** (bottom 16 bits) encodes the actual metadata

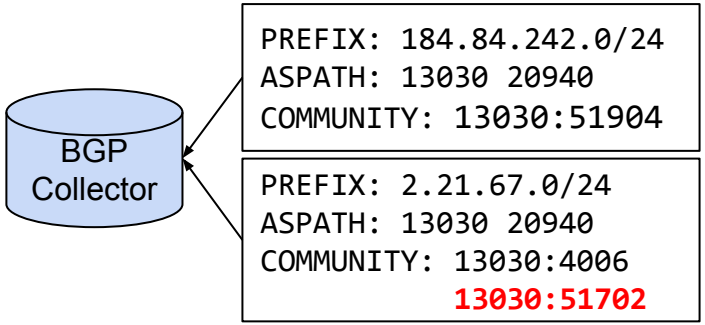

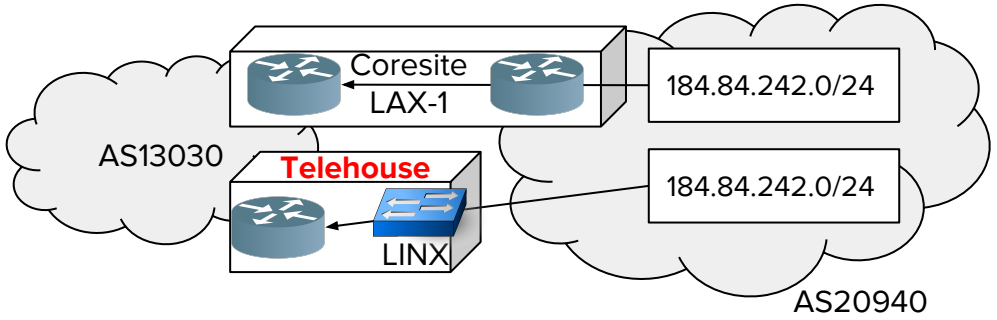

Building a BGP Communities Dictionary

- Communities **not standardized**
  - Often documented in IRR and support websites
  - Documentation sources are scattered and unstructured
- We develop an *NLP parser* to automatically discover and compile a dictionary of PoP-tagging Communities

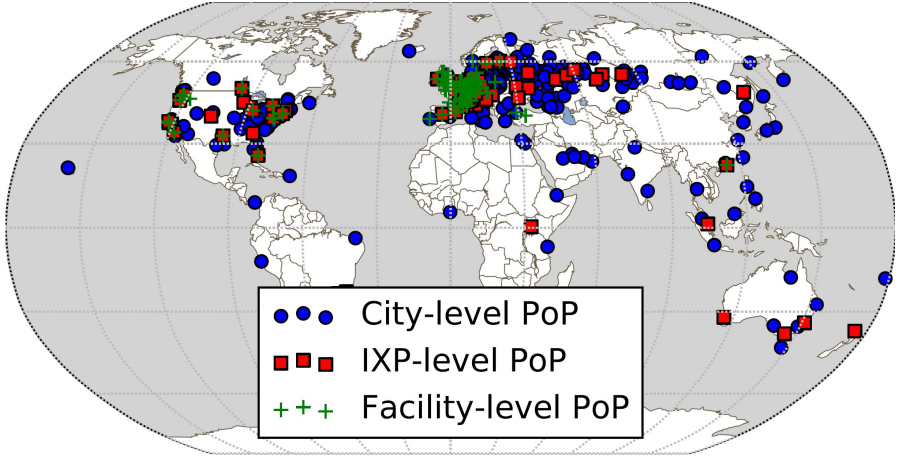

- 3,049 communities by 468 ASes
- Collected communities tag: 288 cities, 172 IXPs, 103 facilities in 72 countries

Topological coverage of interpreted BGP Communities

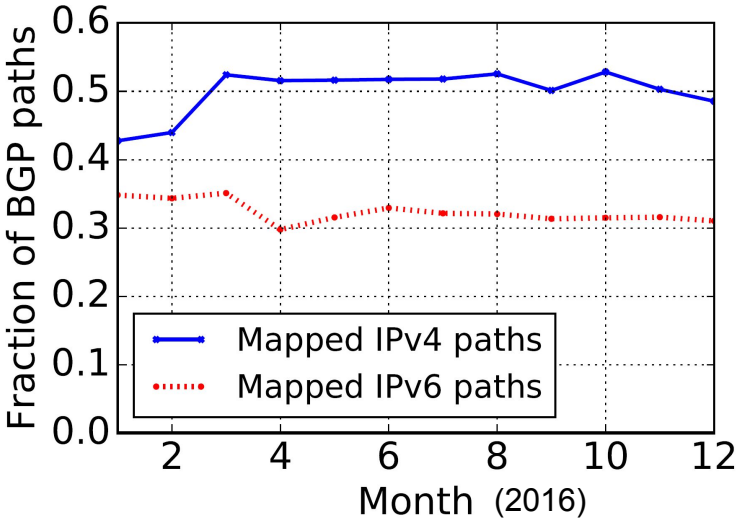

- Our dictionary mainly composed by communities of large transit providers
- Includes communities from all Tier-1 ASes except 2
- ~50% of IPv4 and ~30% of IPv6 paths annotated with at least one Community in our dictionary

Developing a Detailed Co-location Map

To improve the granularity that BGP communities offer:

- For each city in our Communities dictionary map:
  - ASes to facilities
  - ASes to IXPs
  - IXPs to facilities
- Sources: PeeringDB, CloudScene, DatacenterMap, AS websites
- Use the co-location map to de-correlate the behaviour of ASes during an incident based on their presence or absence at facilities and IXPs.

Passive outage signal detection

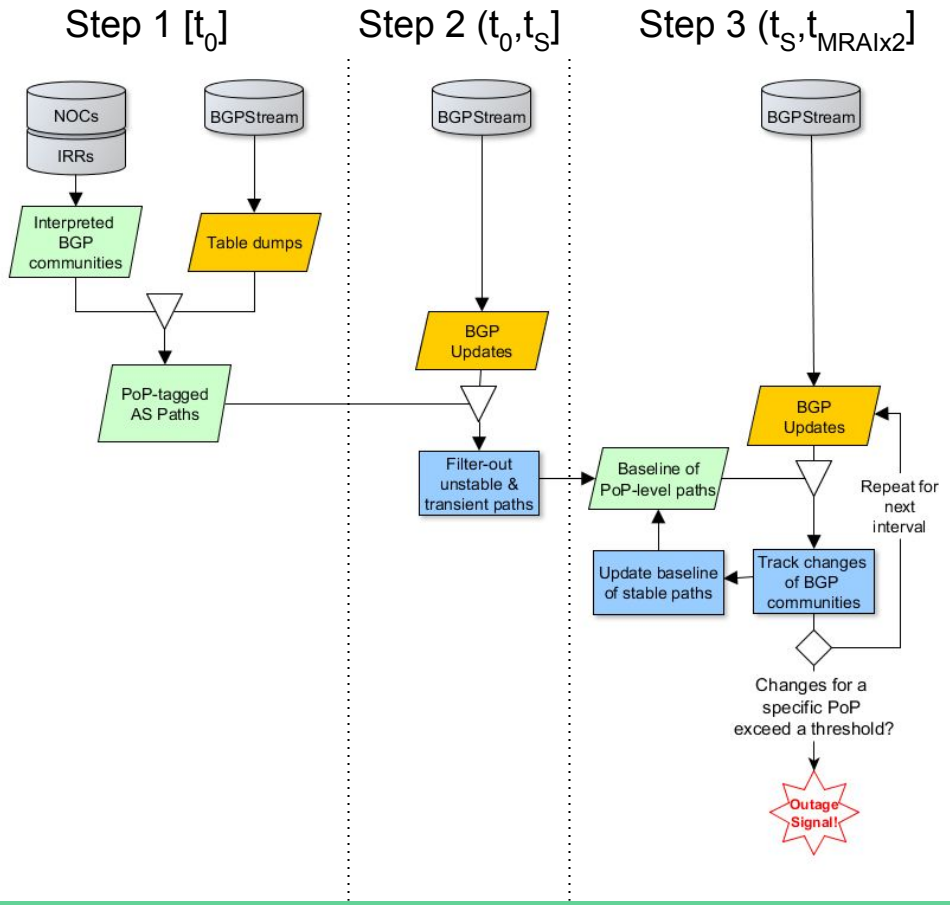

- **Step 1**: Collect BGP routes with PoP-tagging communities.
- **Step 2**: Track collected paths for a period (t_0,t_S] to remote unstable paths.
- **Step 3**: Track changes in the Communities attribute and bin them in 60 sec. intervals (twice MRAI timer).
  - **Outage signal** If changes for a specific PoP exceed a threshold

Passive outage signal detection

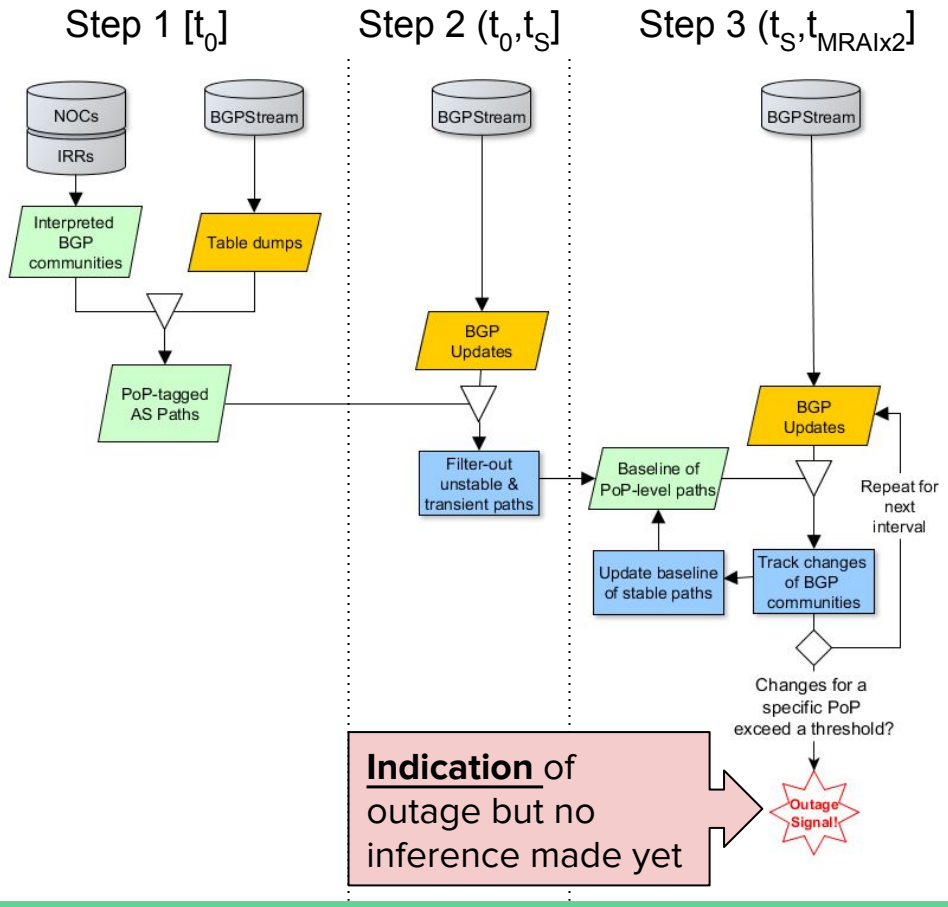

- **Step 1**: Collect BGP routes with PoP-tagging communities.
- **Step 2**: Track collected paths for a period (t_0,t_S] to remote unstable paths.
- **Step 3**: Track changes in the Communities attribute and bin them in 60 sec. intervals (twice MRAI timer).
  - **Outage signal** If changes for a specific PoP exceed a threshold

Infer granularity of outage signal

- Check which ASes were involved in the links tagged with the BGP Communities that triggered the signal:
  1. **Link-level incident**: All changes involve a single AS link, e.g. depeering between two large ASes.
  2. **AS-level incident**: All links intersect a common AS either as near-end or far-end, e.g. misconfiguration in border router.
  3. **Organizational incident**: All links intersect sibling ASes that belong to a common organization, e.g. policy change.
  4. **PoP-level incident:** Multiple affected links with disjoint near-end and far-end ASes (at least different 3 near-end and far-end).

Infer granularity of outage signal

- Check which ASes were involved in the links tagged with the BGP Communities that triggered the signal:
  1. **Link-level incident**: All changes involve a single AS link, e.g. depeering between two large ASes.
  2. **AS-level incident**: All links intersect a common AS either as near-end or far-end, e.g. misconfiguration in border router.
  3. **Organizational incident**: All links intersect sibling ASes that belong to a common organization, e.g. policy change.
  4. **PoP-level incident:** Multiple affected links with disjoint near-end and far-end ASes (at least different 3 near-end and far-end).

Naive localization of PoP-level incident

- Use location encoded in BGP Communities
- Naive approach can be **misleading**:
  - Communities tag only the near-end facility/IXP.
  - Multiple facilities may be involved in a interconnection depending on the peering paradigm.
  - Outage in intermediate PoPs will trigger Communities changes in the near-end PoP.

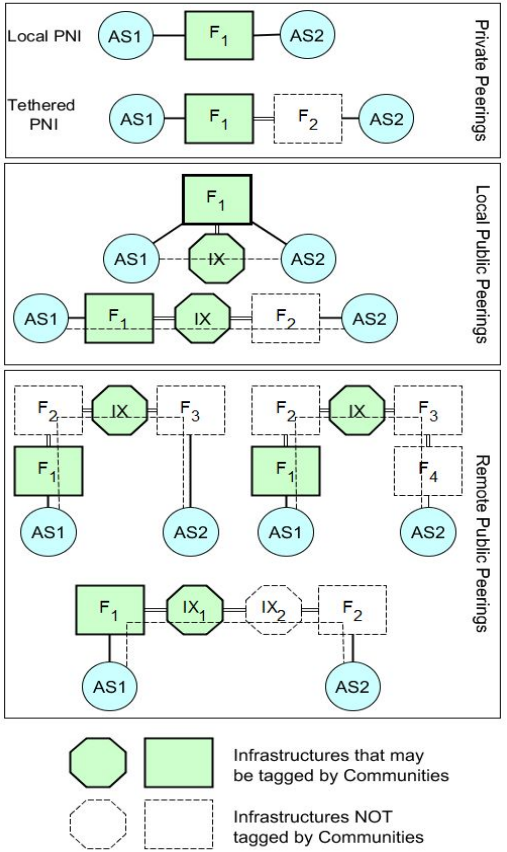

Our localization method: disambiguate outage signals based on co-location map

- **Near-end facility outage**: all links tagged with the same near-end facility are affected.
- **Far-end facility outage**: common facility among far-end ASes.
- **IXP outage**: common IXP among near-end and far-end ASes.
- **Broader outage**: multiple facilities and IXPs among affected ASes but empty intersection.

Example:

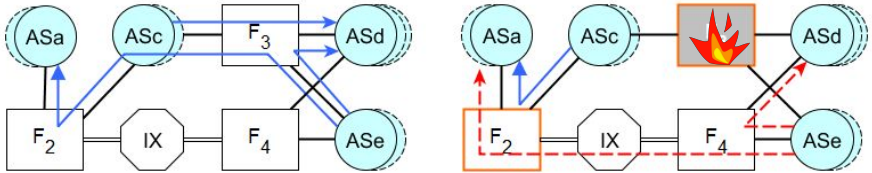

Detected signals from near-end ASes:

- ASa: Paths shifted from F_1 to F_2
- ASb: Paths shifted from F_3 to F_4

F_1,F_3 candidate outage sources

Co-location data for far-end ASes:

- ASc at [F_2,F_3], not at F_4
- ASe at [F_3,F_4], not at F_2

Correlation of signals with co-location:

- ASa: co-located peers at F_2 not affected
- ASb: co-located peers at F_3 affected

Inferred outage source

Detecting Peering Infrastructure Outages in the Wild

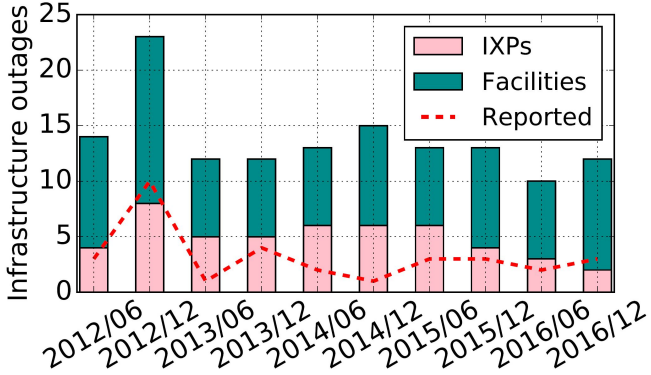

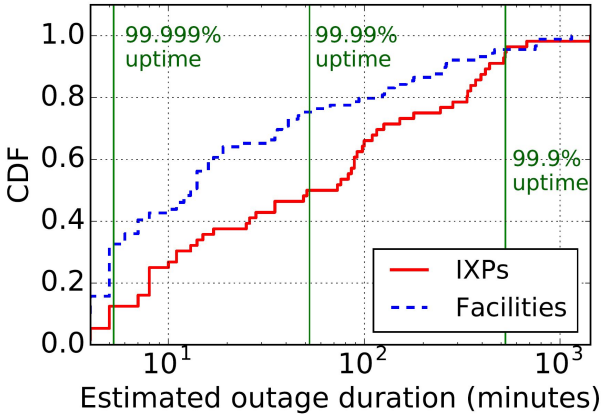

- Applied detection methodology in archived BGP data (2012-2016)
- 159 in total, 103 outages among 87 facilities, 56 outages in 41 IXPs
- Only 24% of the detected outages reported in outages.org and NANOG mailing lists and 2 specialized data center websites.
- 5% of the facilities below 99.99% uptime, 18% below 99.999% uptime.

Case study 1: Large-scale IXP outage (AMS-IX)

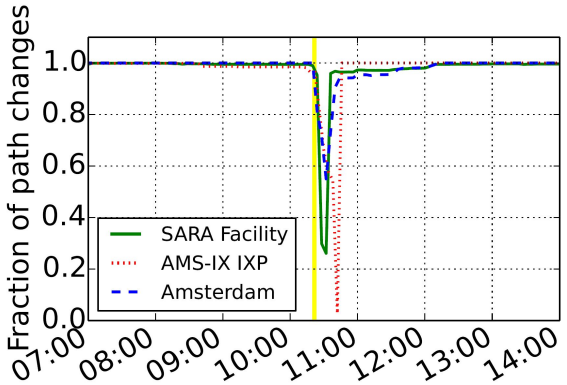

Aggregated path changes
For PoPs of 3 different granularities

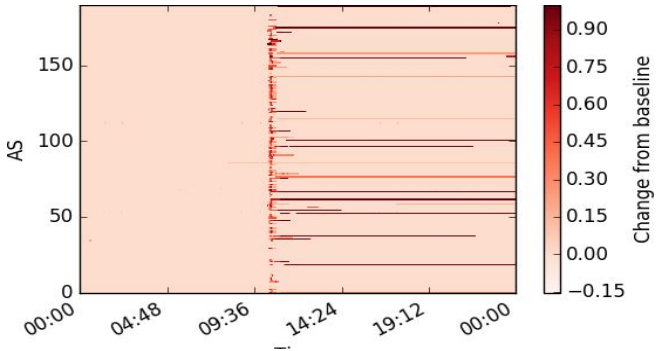

Path changes per AS/facility
involved in AMS-IX peerings

- Largest IXP in the world,
  - **>700** members, **3TBs** peak traffic
- Layer-2 loop caused **90% traffic loss** for ~15 minutes (2015/05/13) [10]
- Outage signals triggered by aggregated path changes at all PoP granularities:
  - Outage confirmed by per-AS investigation of path changes.
- **“Simple” case**:
  - Outage of high magnitude
  - Co-located BGP collector (RRC01)

Case study 2: Back-to-back partial outages at London facilities (Telecity HEX8/9, TeleHouse North)

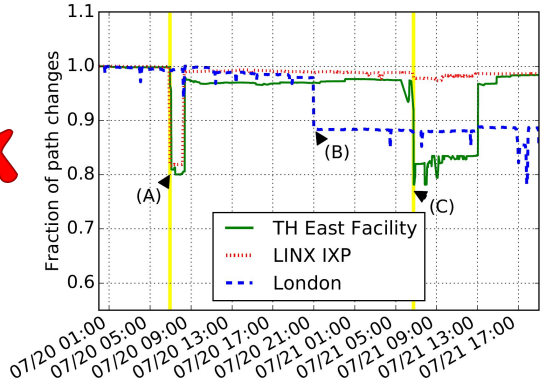

Aggregated path changes

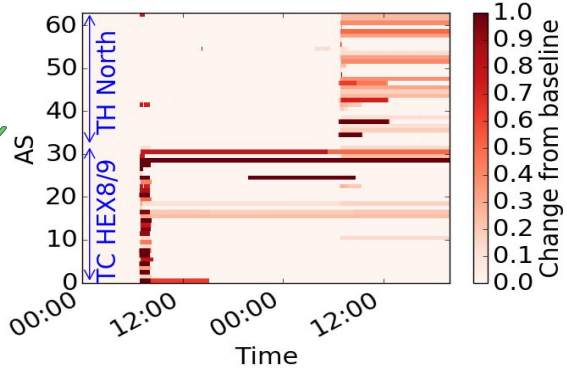

Path changes per AS/facility

- Partial power failures in two different London facilities (2016/07/20 - 21) [11]
  - TC HEX8/9 outage triggered signal in the LINX IXP and the TeleHouse East facility **(A)**
  - TH North outage triggered signal only in TH East facility **(C)**
  - AS-specific incident caused city-level signal **(B)**
- **"Hard" case**:
  - Aggregate signals triggered in PoPs different from outage sources and for non-infrastructure incident
- Signal investigation and localization **per AS and per facility** reveals the correct outage sources.

Case study 3: Medium-scale IXP outage (France-IX)

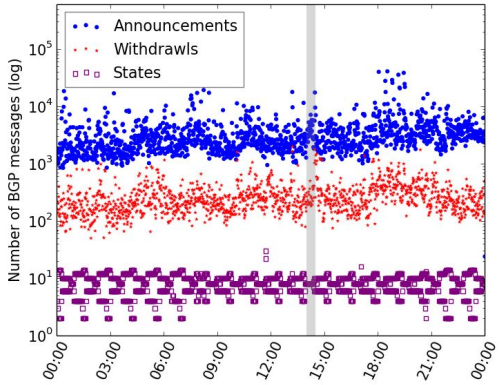

BGP activity for all paths in RRC01

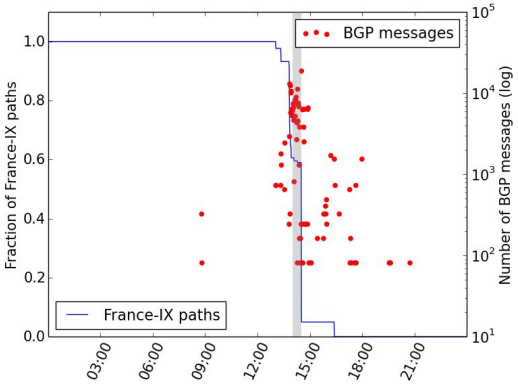

BGP activity only for paths tagged with France-IX communities in RRC01

- Outage on 2014/08/18 due to port flapping [12].
- Monitoring the aggregated activity of BGP messages (updates, withdrawals, state changes) provides no outage indication:
  - No co-located collector at the IXP
  - BGP activity due to outage masked by “noise” from normal routing dynamics.
- Monitoring the activity of BGP messages only for the paths tagged with France-IX communities provides strong outage signal.

Tracking the progress of outages

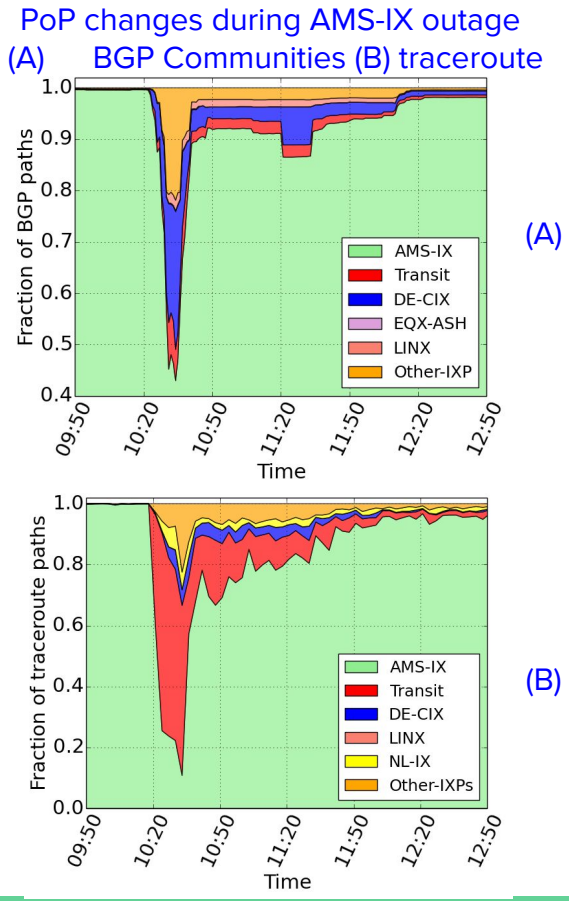

- Passive tracking:
  - Monitor how the PoP-tagging BGP Communities change.
  - Infer end of outage when >80% of the affected paths return to original PoP.
- Active tracking:
  - Collect stable publicly archived traceroute paths crossing the PoP before the outage (**Ark**, **Atlas**).
  - Execute queries between the same (source,destination) pairs.
  - Geo-locate the new border IP addresses

Measuring the impact of outages

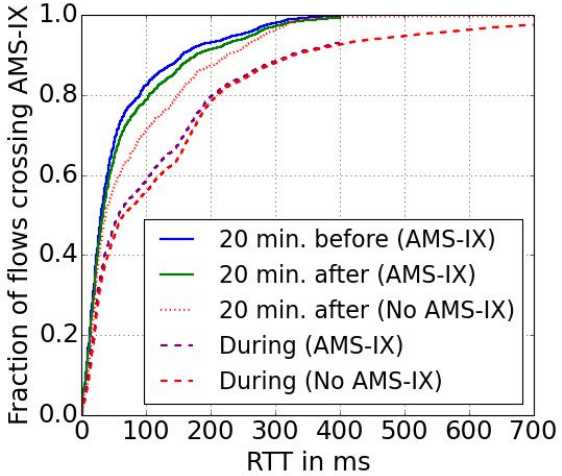

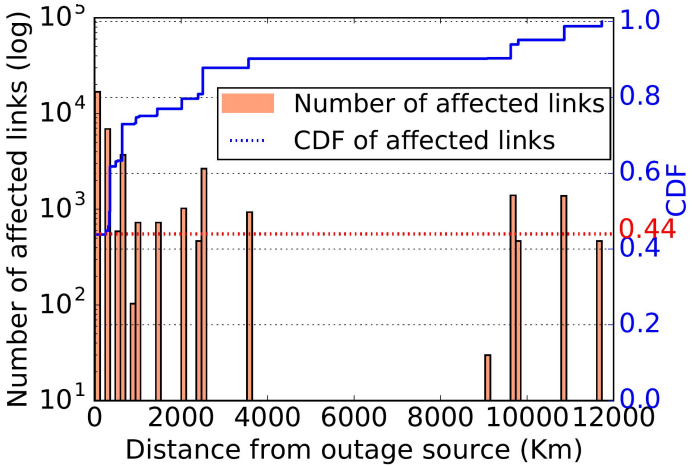

- We use traceroute measurements to asses the impact and geographic spread of outages.
- Median RTT rises by more than 100 msec for rerouted paths during AMS-IX outage

- Geolocate the IPs of far-end interfaces of the affected ASes during TC HEX8/9 outage.
- Over 45% of the interfaces are in a different country, more than 20% outside of Europe.

Conclusion

- Control-plane messages provide excellent, yet unexplored source of infrastructure-level topological information.
  - Decipher PoP-level metadata through BGP Communities
- Automate crowd-sourcing of outage detection through routing data:
  - “Hard evidence” provide accountability, transparency, accuracy
- Over a 5-year period, we detected 159 CF and IXP outages
  - 4x than what we could find in major mailing lists/news websites
  - Remote peering and complex interdependencies “globalize” the impact of local failures

References

[1] The Zettabyte Era — Trends and Analysis – Cisco
http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.html
[2] https://www.pch.net/resources/Papers/peering-survey/PCH-Peering-Survey-2016/PCH-Peering-Survey-2016.pdf
[3] “Equinix Investor Relations: Annual Report FY2015.” http://investor.equinix.com/phoenix.zhtml
[4] Department of Homeland Security, “Communications Sector-Specific Plan: An Annex to the National Infrastructure Protection Plan,” 2015.
[5] European Union Agency for Network and Information Security, “Critical Infrastructures and Services, Internet Infrastructure: Internet Interconnections.”, June 2012
[6] Ponemon Institute, “Cost of Data Center Outages”, Tech. report, January 2016,
http://planetaklimata.com.ua/instr/Liebert_Hiross/Cost_of_Data_Center_Outages_2016_Eng.pdf
[7] V. Giotsas et al., “Mapping Peering Interconnections at the Facility Level,” in CoNEXT, 2015
[8] R. Motamedi et al.. “On the Geography of X-Connects”. Technical Report CIS-TR-2014-02, University of Oregon, 2014
[9] Banerjee et al. "Internet outages, the eyewitness accounts: Analysis of the outages mailing list." PAM 2015.
[10] AMS-IX. “Follow-up on previous incident at AMS-IX platform”, 13 May 2015. https://ams-ix.net/newsitems/195
[11] Ars Technica. “BT, other ISPs hit by second major Internet outage, power failure blamed”,July 2016.
http://arstechnica.co.uk/business/2016/07/bt-isps-telehouse-north-major-outage/
[12] France-IX. “Outage Notification”, 18 August 2014.
https://www.franceix.net/en/events-and-news/news/franceix-outage-notification/