Swisscom Network Observability

Transforms and Innovates Network Operations

08.05.2025, Thomas Graf – thomas.graf@swisscom.com

Picture: Apollo 8, December 24th 1968

swisscom

Nationwide Network Outages everywhere

Increasing in impact and duration - hinting Network Visibility deficiencies

CANADA

Rogers says network upgrades after outage will cost $261M, but no timeline given

By **Staff** · **The Canadian Press**
Posted August 25, 2022 11:09 am

Rogers CEO Tony Staffieri explained to a standing committee in the House of Commons on Monday that the technology company is investing significant amounts of capital to ensure it can avoid a repeat of its Canada-wide outage of July 8 – Jul 25, 2022

BUSINESS

KDDI to spend ¥7.3 billion to compensate users for major network outage

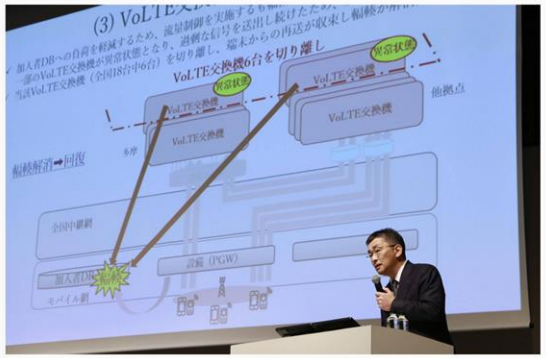

KDDI chief Makoto Takahashi speaks to reporters in Tokyo on Friday. | KYODO

BY KAZUAKI NAGATA
STAFF WRITER

SHARE Jul 29, 2022

ORANGE FRANCE UNDER FIRE FOR MISHANDLING NETWORK OUTAGE

Posted by Harry Baldock | Jul 22, 2021 | Subsea, INFRASTRUCTURE, Satellite, Towers, COMPANY NEWS, Governance, Data Centres, Networks, Wholesale, Virtualisation, Europe, Middle East & Africa, News

Optus: Telecom boss Kelly Bayer Rosmarin quits after Australian outage

6 days ago

The firm has come under fire following a nationwide network outage this month

July 14, 2021
7:57 AM GMT+2
Last Updated a year ago

Media & Telecom

Swisscom boss apologises for massive network outage - newspaper

Reuters

2 minute read

05 FEB 2023 | 08:23 AM UTC

Italy: TIM internet services interruption reported nationwide Feb. 5

TIM internet services interruption reported in Italy Feb. 5. Likely communication disruptions.

Informational | Communications/technology | Transportation | ITA

Facebook outage: what went wrong and why did it take so long to fix after social platform went down?

Billions of users were unable to access Facebook, Instagram and WhatsApp for hours while the social media giant scrambled to restore services

Facebook, Instagram and WhatsApp all went down, and reappeared online after a six-hour global outage. Photograph: Anadolu Agency/Getty Images

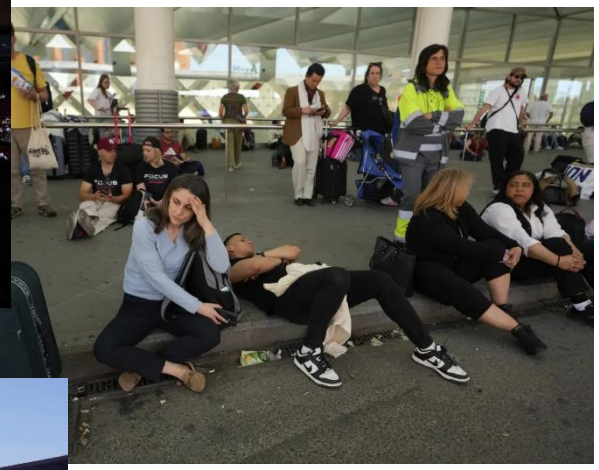

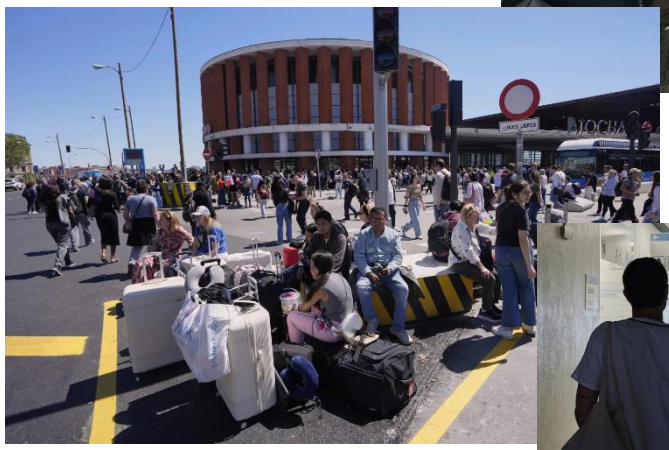

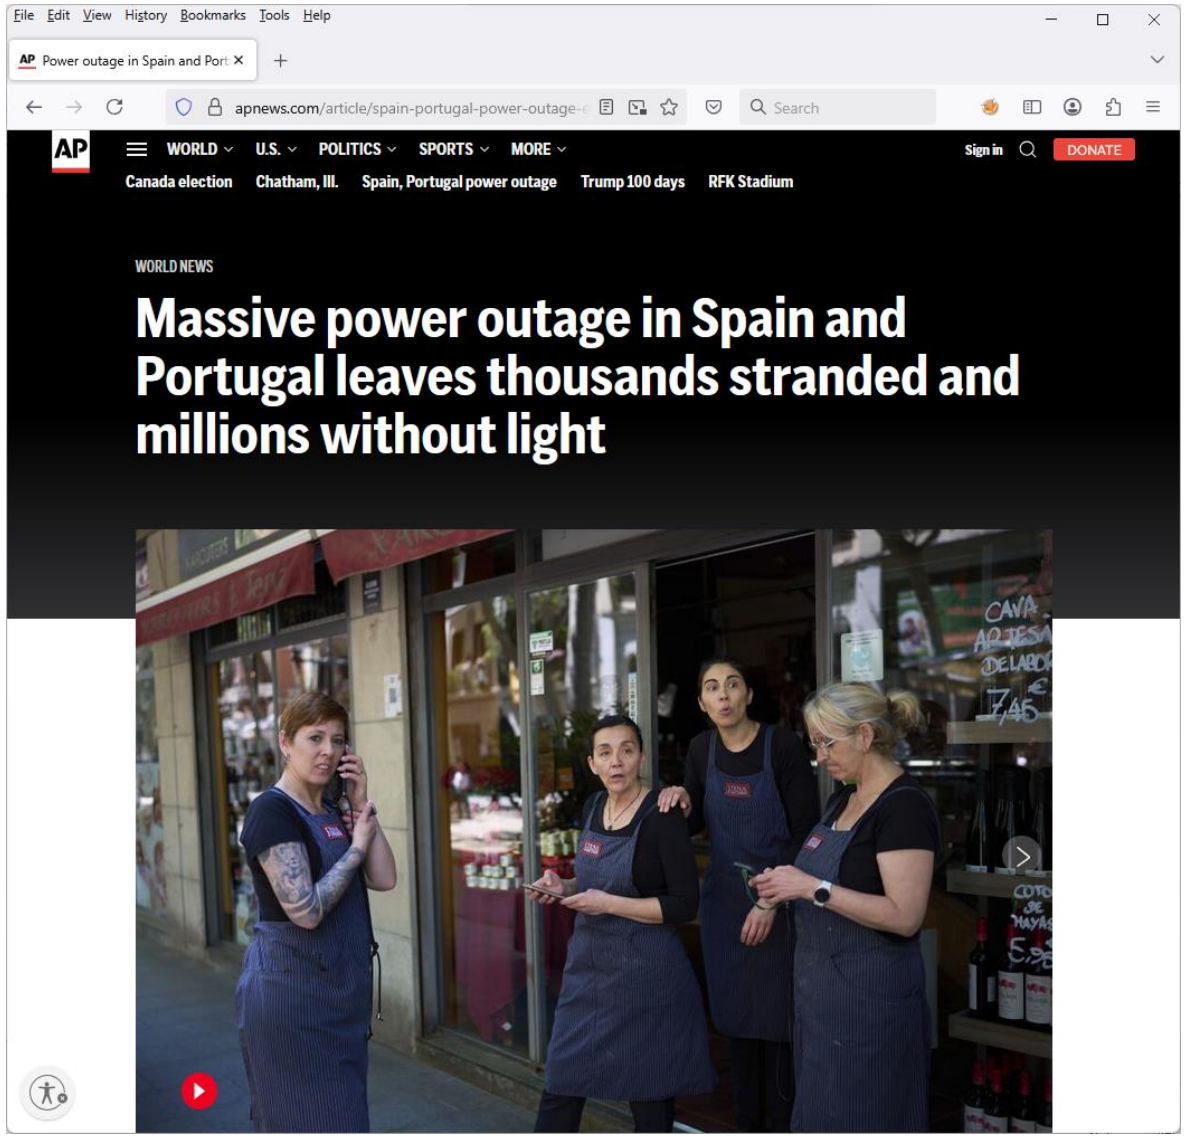
File Edit View History Bookmarks Tools Help
Power outage in Spain and Port
apnews.com/article/spain-portugal-power-outage-
Search
AP
WORLD
U.S.
POLITICS
SPORTS
MORE
Sign in
DONATE
Canada election
Chatham, Ill.
Spain, Portugal power outage
Trump 100 days
RFK Stadium
WORLD NEWS
Massive power outage in Spain and Portugal leaves thousands stranded and millions without light

The Swisscom View

Traffic and BGP Impact to **Portugal**

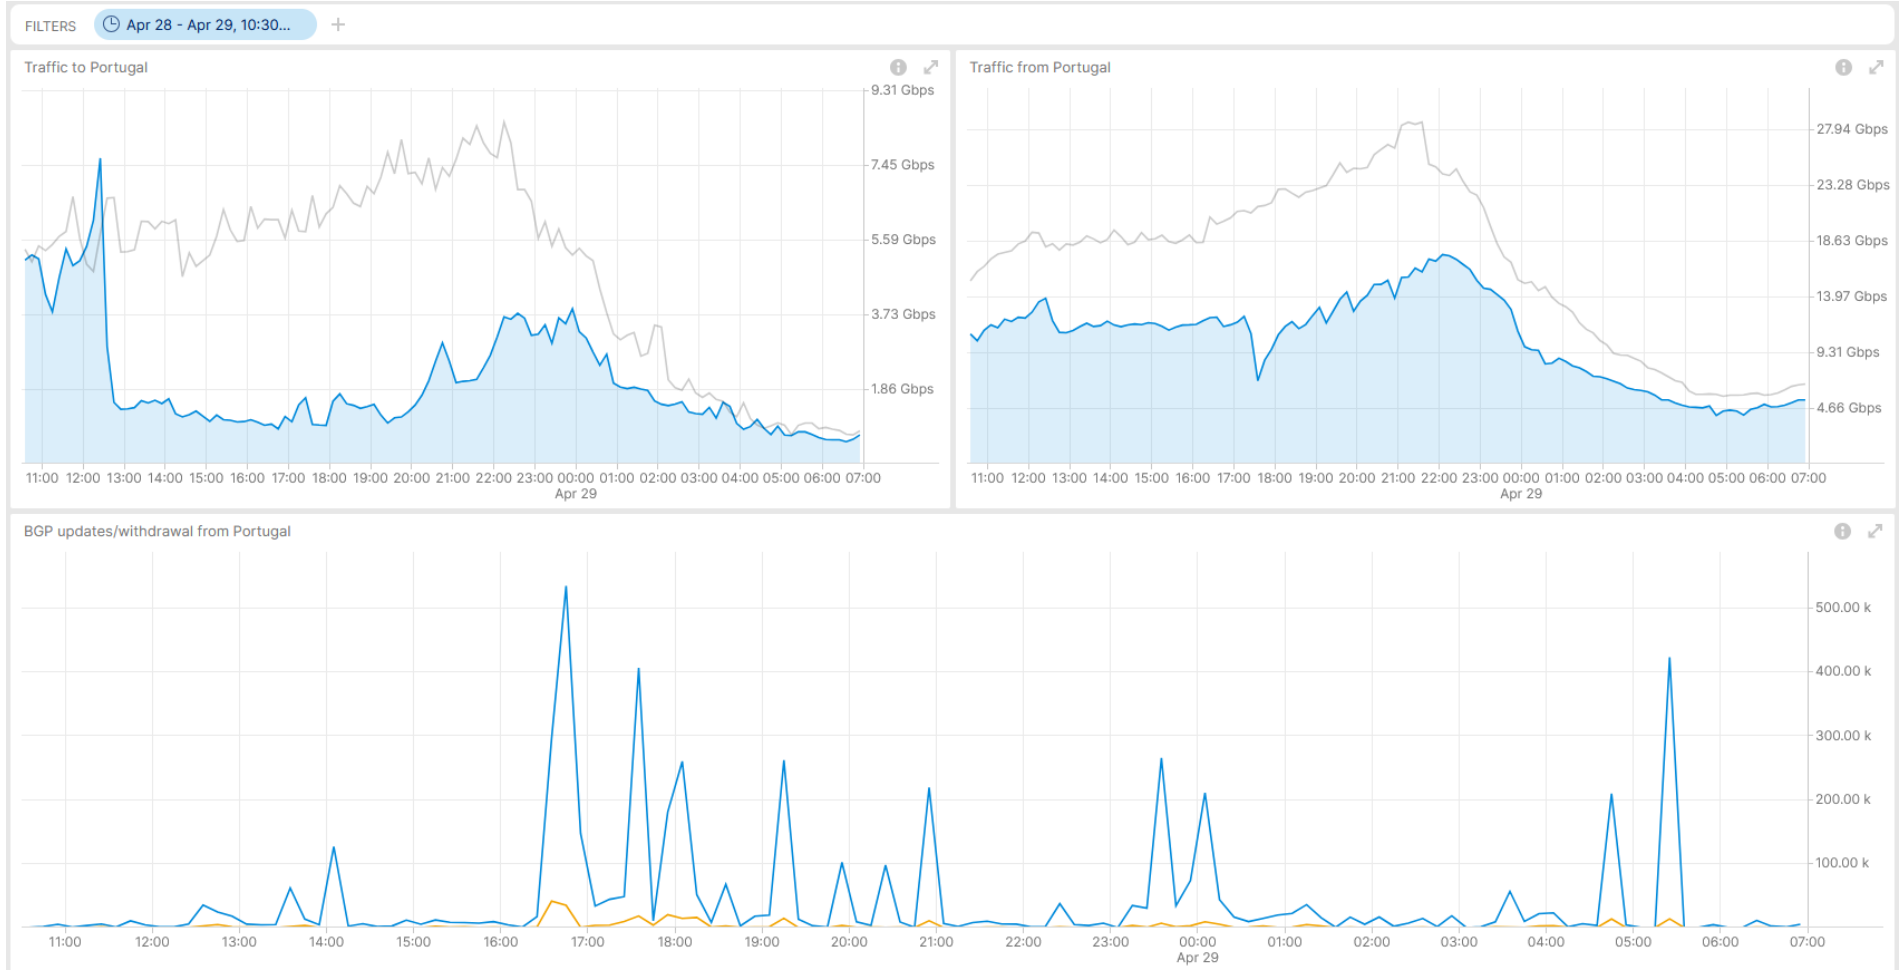

Shows >80% missing traffic between 12:30 and 03:30 for traffic to Spain.

Increased BGP topology changes between 17:00 and 00:30.

Pyramid of Technology

From envisioned to naturalized

“ Every human being has to cope with technological change, yet few of us are aware of how new technologies are introduced, accepted and discarded in our society. The Pyramid of Technology visualizes how technology becomes nature in seven steps and what we can learn from that. It helps us to dream, build and live in our next nature — the nature caused by humans “

Analytical Use Cases are already traversing these technology stages at Swisscom as we defined the vision.

Pyramid of Technology
https://nextnature.net/projects/pyramid-of-technology

How technology becomes nature
https://www.youtube.com/watch?v=EXJB4Ync82c

The Idea Factory: Bell Labs and the Great Age of American Innovation
https://en.wikipedia.org/wiki/The_Idea_Factory

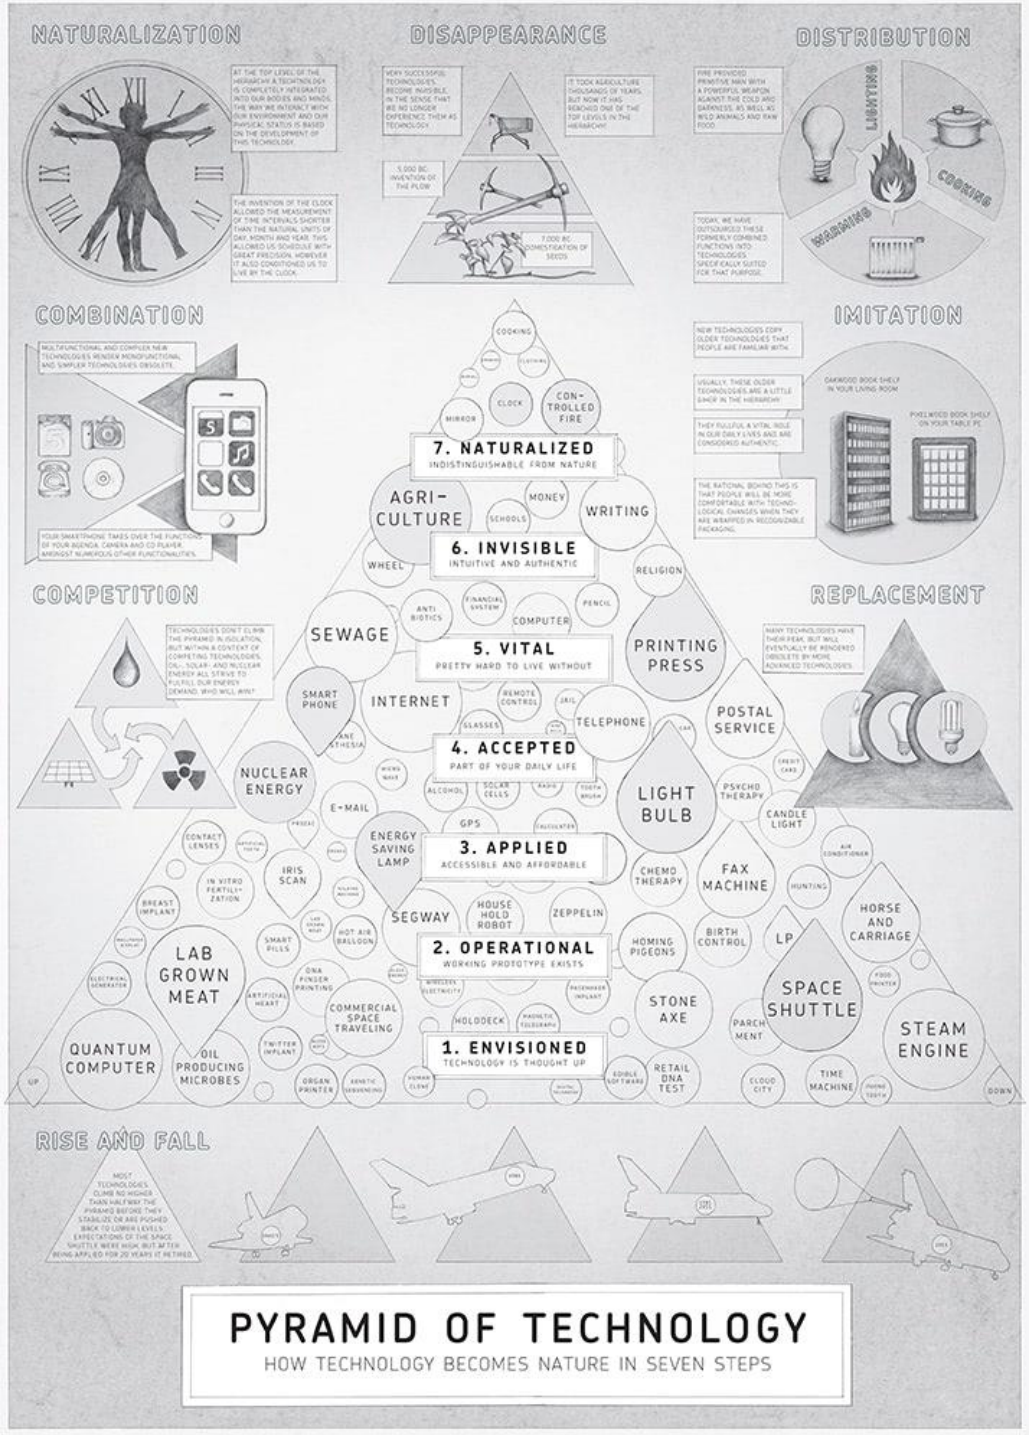

Data Mesh organizes Data in Organizations

Enables Network Analytics use cases

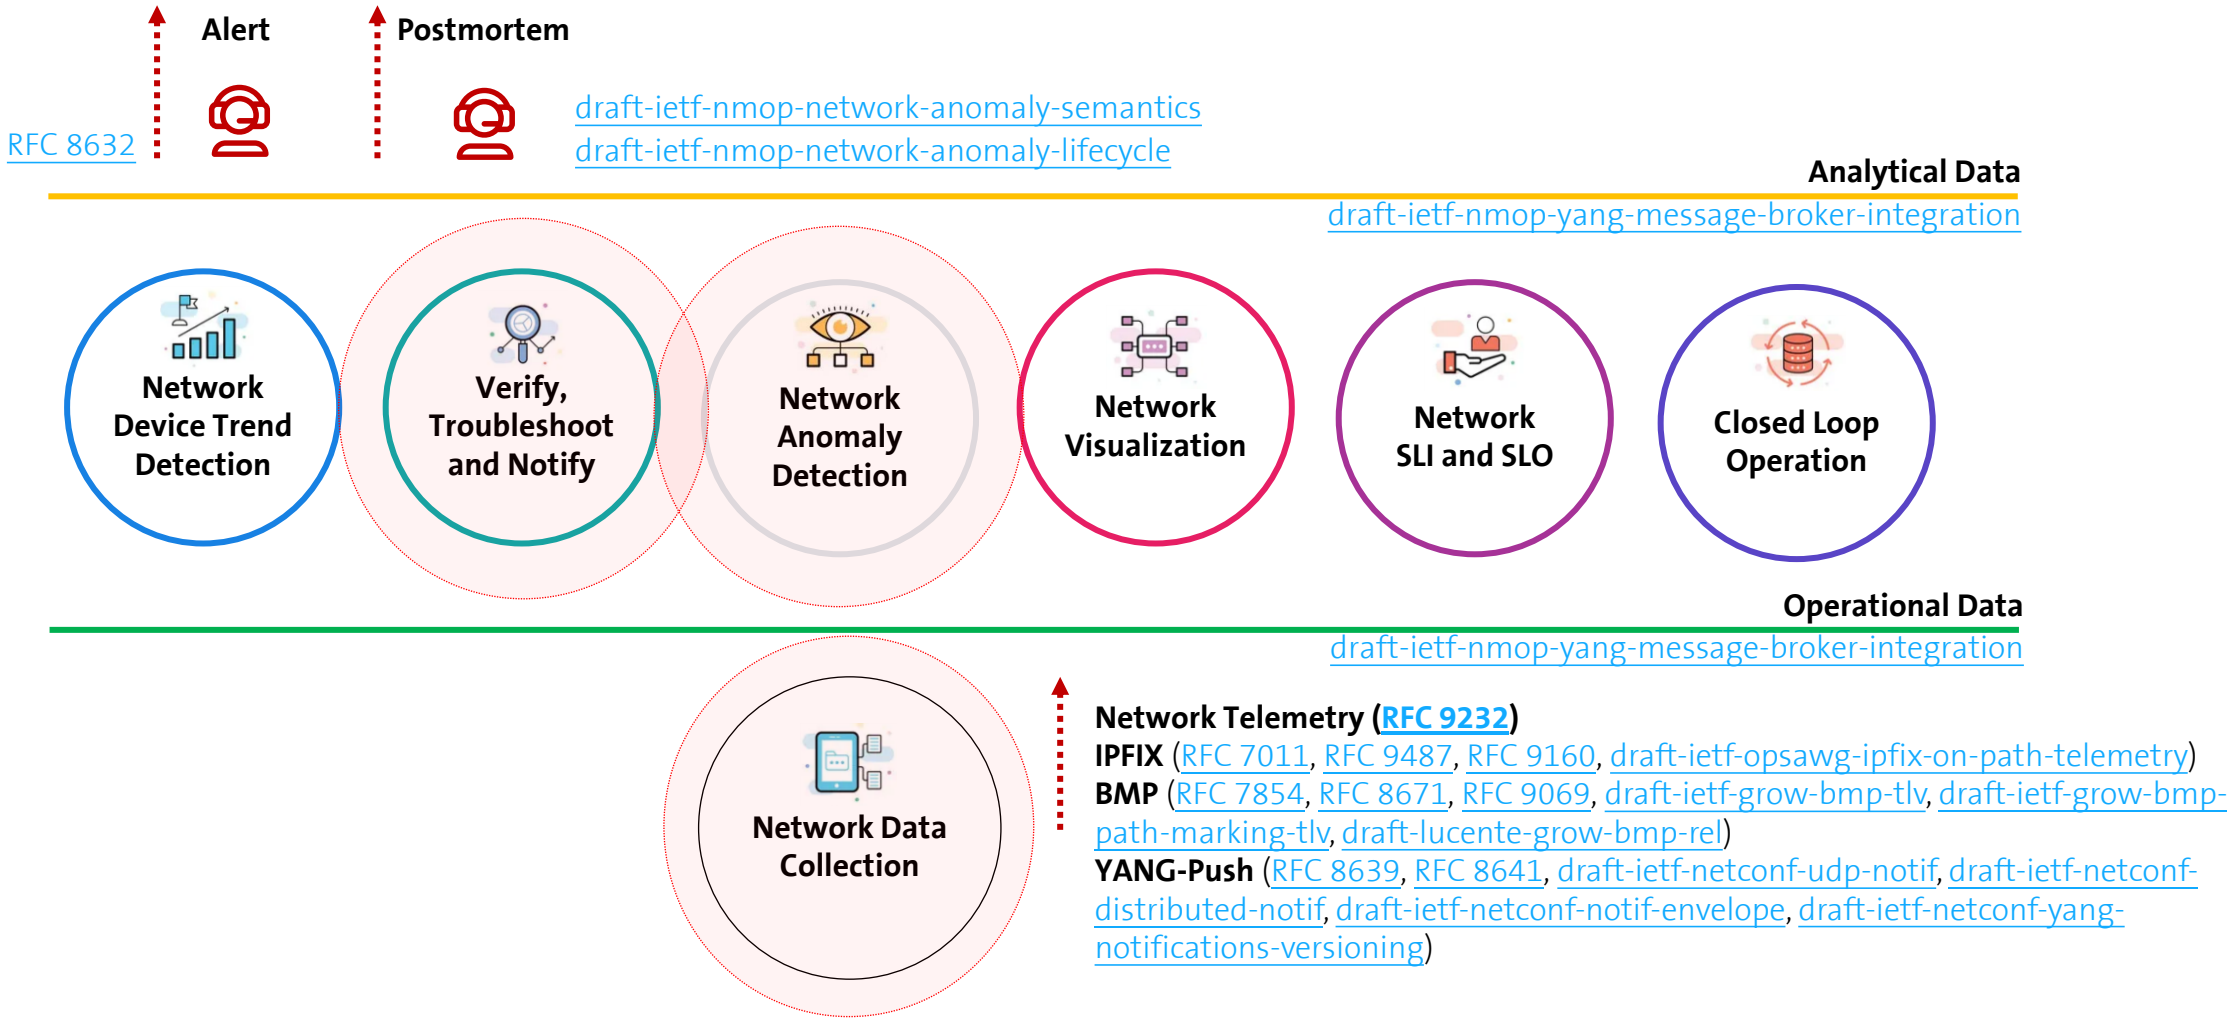

Monitoring L3 VPN's with IPFIX, BMP and YANG Push

From Connectivity Service to Realtime Network Observability

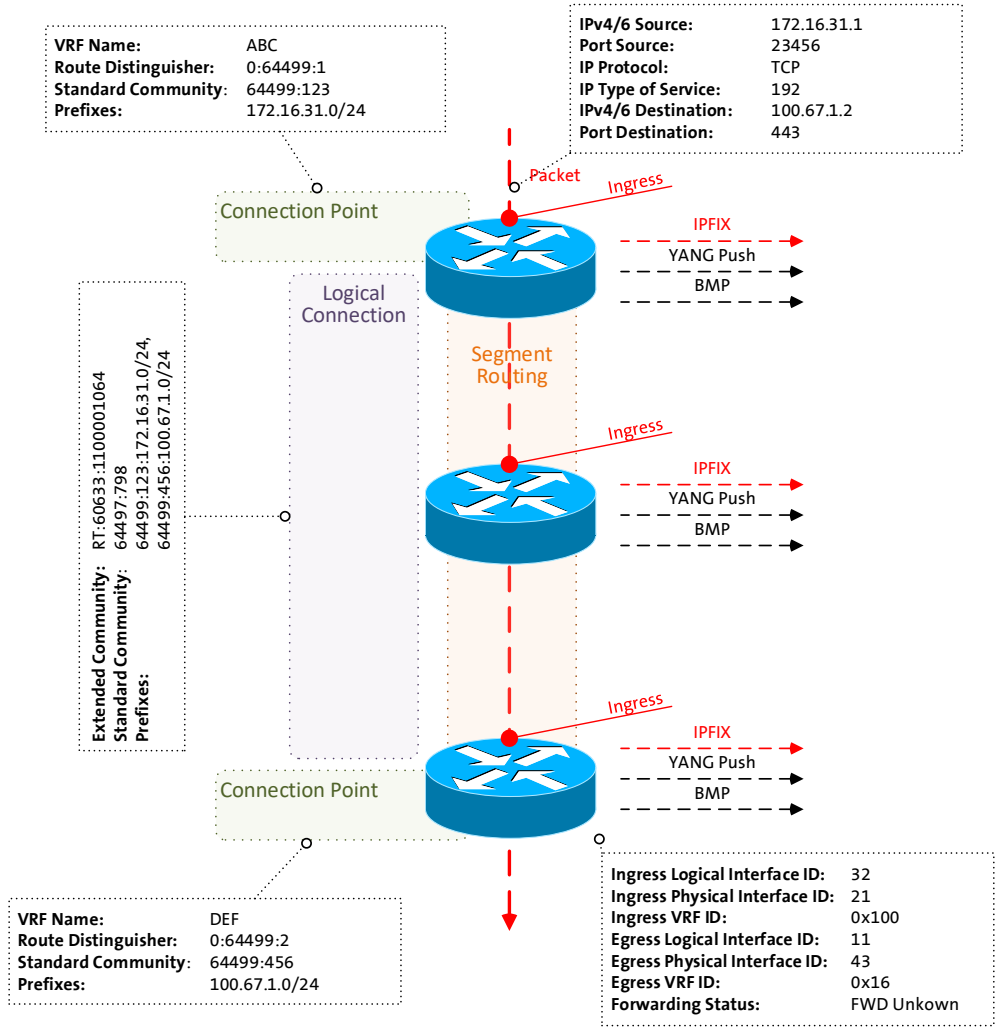

- **Connectivity Service perspective,** Connection Points are connected through Logical Connections.
- **From a BGP control-plane perspective,** IPv4/6 unicast prefixes in VRF's are tagged with BGP standard communities.
  - One BGP standard community to identify the Logical Connection. One BGP standard community to identify each Connection Point.
  - When IPv4/6 prefixes are exported from VRF's, a BGP route-distinguisher, BGP extended community route-targets and a SRv6 VPN SID for the IPv6 next-hop are allocated.
- **From a forwarding plane perspective,** when IPv4/6 unicast traffic is received from the edge at the SRv6 PE, a lookup is performed, the SRv6 VPN SID is obtained and IPv6 next-hop is added when forwarded to the core.
- **Swisscom collects** MPLS and SRv6 provider data plane, IPv4/6 unicast customer data-plane in IPFIX and at provider edge BGP VPNv4/6 unicast **in production** to perform real-time data correlation.

Problem Statement and Motivation

How it is being addressed in which document

Network Anomaly Detection

When operational or configurational changes in connectivity services are happening, **the objective is to detect interruption at network operation faster than the users using those connectivity services**.

In order to achieve this objective, **automation in network monitoring is required.** This automation needs to **monitor network changes holistically** by monitoring all 3 network planes simultaneously and detect whether that change is service disruptive.

Through network incidents postmortems we network operators **learn and improve so does network anomaly detection** and supervised and semi-supervised machine learning. **With more and more incidents the postmortem process demands automation** and with the standardization of labeled network incident collaboration among network operators, vendors and academia is facilitated.

- draft-ietf-nmop-network-anomaly-architecture describes the motivation and architecture and the relationship to other two documents.
- draft-ietf-nmop-network-anomaly-semantics defines Symptom semantics to enable standardized data exchange to validate results with network engineers and improve supervised and semi-supervised machine learning systems.
- draft-ietf-nmop-network-anomaly-lifecycle describes on managing the lifecycle process, in order to facilitate network engineers to interact with the network anomaly detection system to refine the detection abilities over time.

April 17-22th, OSPF/BGP Routing Instability

Cisco TAC Case Summary

CASE SUMMARY

STATUS	SEVERITY	CREATED
Cisco Pending	Moderate Impact (S3)	04/22/2025
REQUEST TYPE	**UPDATED**	
Diagnose and Fix my Problem	04/23/2025	

PROBLEM DESCRIPTION - Matt Snow 04/22/2025 at 21:31:27

Customer is experiencing a significant routing instability in their production network, specifically involving OSPF and BGP protocols on Catalyst 4500E switches. The issue manifests as rapid BGP updates and withdrawals (up to 600,000 state changes per minute), leading to forwarding plane drops and potential service disruption. The problem was temporarily mitigated by shutting down interface vlan 236 on device mbue51cos, which stopped OSPF flapping. However, this is only a workaround, as vlan 236 is required for normal operations and must be reactivated. The incident appears to be triggered by device mbue52cos. The business impact is high, as the instability affects the core routing infrastructure, potentially disrupting network services for users and dependent business activities. The issue is not currently a catastrophic outage due to the workaround, but normal operations cannot resume until a permanent fix is found.

Technical Indicators:
- Log message: %OSPF-3-CHKPT_STBY_LSDB_INVALID: STANDBY:Standby link-state database validation failed, expected/found count: 1733/1726 chksum: 0x369D183/0x363D367
- Log message: %OSPF-3-CHKPT_STBY_SYNC_LOST: STANDBY:Standby synchronization lost for OSPF-40 (was: synchronized)
- Log message: %OSPF-3-CHKPT_STBY_LSDB_INVALID: STANDBY:Standby link-state database validation failed, expected/found count: 1800/1785 chksum: 0x38C66E6/0x385B4CE
- Log message: %OSPF-3-CHKPT_STBY_SYNC_LOST: STANDBY:Standby synchronization lost for OSPF-40 (was: synchronized)
- Forwarding plane drops observed
- OSPF flapping and BGP instability
- Affected devices: mlss52cos, mbue52cos-m1, mbue51cos
- Temporary workaround: shutdown of interface vlan 236 on mbue51cos

CURRENT STATUS - Matt Snow 04/22/2025 at 21:31:30

ACTION PLAN - Matt Snow 04/22/2025 at 21:31:34

- Customer to provide answers to Case Owner's questions regarding the observed symptoms, BGP/OSPF state changes, timing, and any recent changes made prior to the issue.
- Customer to provide show tech outputs from the affected devices (mlss52cos, mbue52cos-m1, mbue51cos).
- Case Owner to analyze the provided data and logs to identify the root cause of the OSPF and BGP instability.
- Case Owner and Customer to coordinate on next troubleshooting steps once additional information is received.
- No specific timeframes or deadlines have been agreed upon yet; pending Customer's response with requested information.

April 17-22th, OSPF/BGP Routing Instability

Cisco IOS XR Network Telemetry Coverage

IPFIX configured on P and PE MPLS-SR nodes on MPLS and IPv4/6 VRF unicast enabled interfaces. Capturing L3 IPv4/6 overlay customer data plane and underlay MPLS-SR provider data plane metrics on MPLS enabled interfaces, and IPv4/6 overlay customer data plane metrics on IPv4/6 VRF unicast enabled interfaces.

-> Shape, means that we are engaged in IETF standardization, vendor implementations and running code. IPv4/6 unicast customer data plane visibility is in vital, MPLS data plane visibility is in applied.

BMP Adj-RIB In post-policy on BGP VPNv4 /6 and IPv4/6 VRF unicast peers and Local-RIB on all RIB's configured on SRv6 PE's. BMP Adj-RIB In post-policy on BGP VPNv4 /6 peers on Route Reflectors configured.

-> Shape, means that we are engaged in IETF standardization, vendor implementations and running code. BMP Local RIB data plane visibility is in applied, BMP Path Marking is in operational stage.

YANG Push Legacy on most nodes enabled but not relevant for this use case.

-> Take, means that current YANG-Push legacy implementation is used without any vendor code change and is in accepted stage. However, IETF YANG-Push is shape and is in operational state.

April 17-22th, OSPF/BGP Routing Instability

L3 VPN – **Real-Time** Incident Analysis

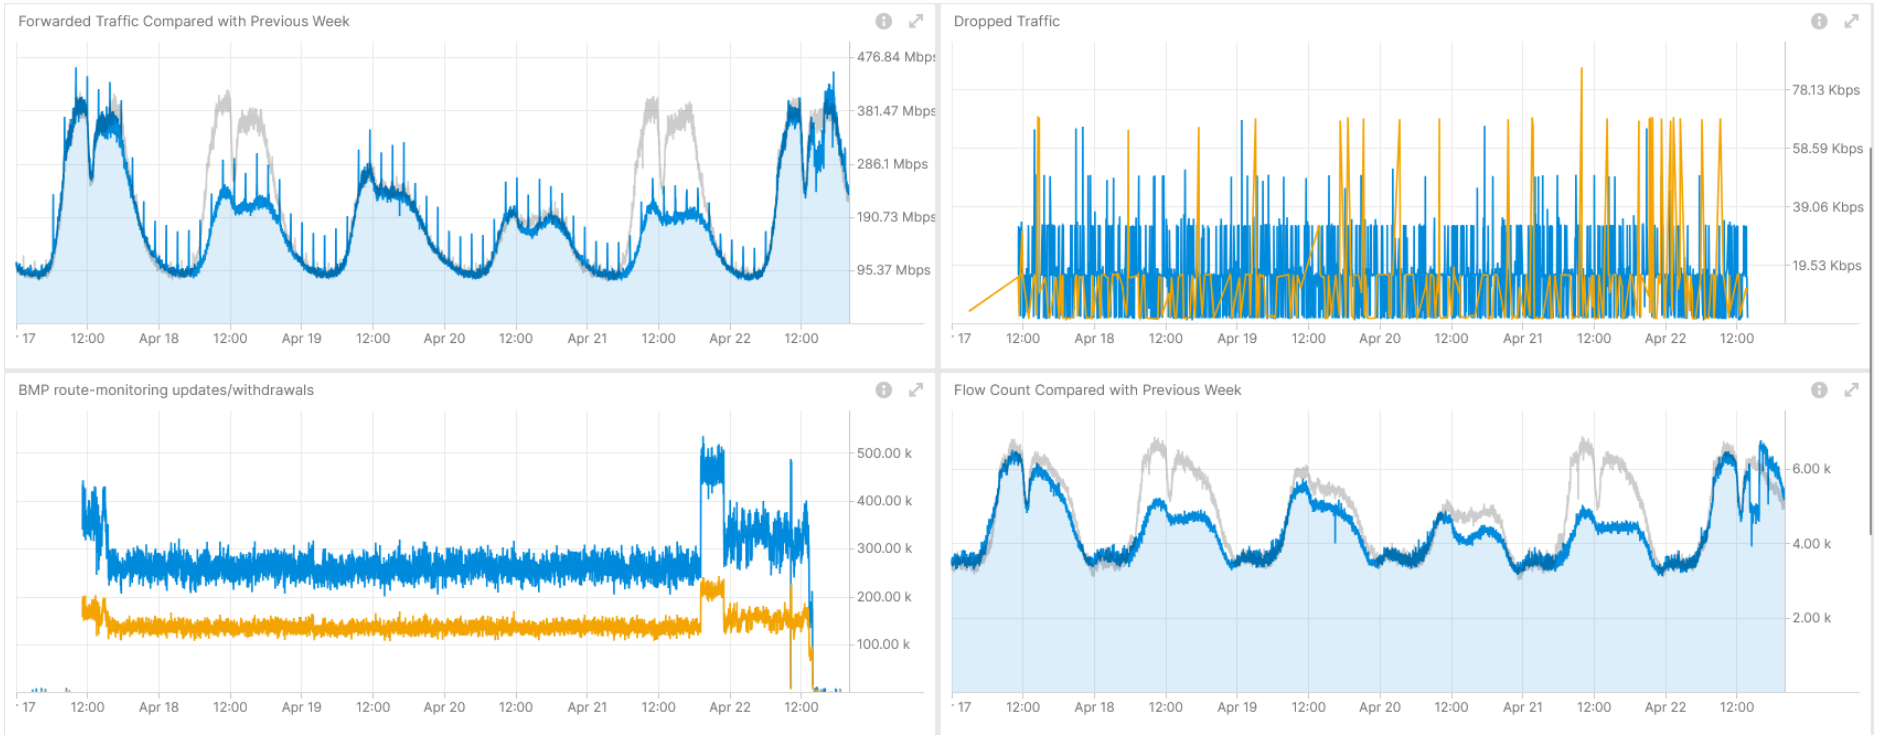

Shows **traffic bad TTL, adjacency drops** and **traffic volume changes due to public holidays, Measured with IPFIX and Correlated with BGP VPNv4/6.**

Shows **constant BGP topology changes** and **flow count changes due to public holidays. Measured with IPFIX and Correlated with BGP VPNv4/6, BMP Adj-RIB In and Local RIB.**

Operational Network Telemetry forwarding plane, IPFIX, BMP measured control plane metrics.

April 17-22th, OSPF/BGP Routing Instability

Network Anomaly Detection – **Live**

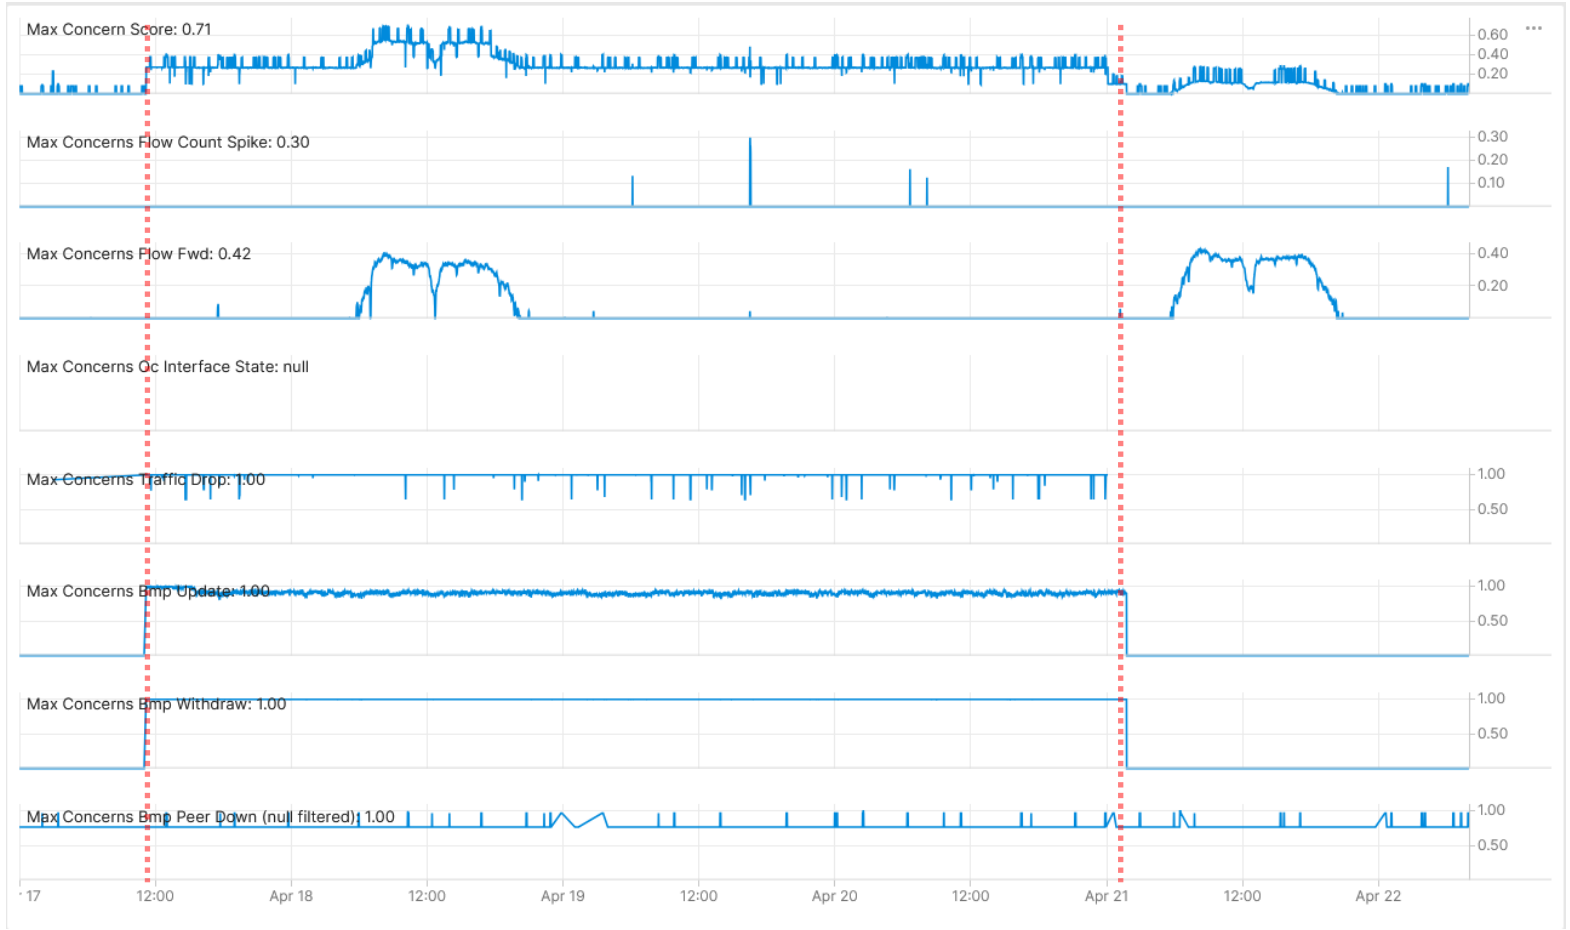

Cosmos Bright Lights monitoring 64497:471 L3 VPN in real-time during maintenance window.

Concern Score: 0.71
Flow Count Spike: **0.30**
Missing Traffic: **0.41**
Traffic Drop: **1.00**
BMP Peer/Interface Down: **0.96/0.00**
BMP Update/Withdrawal: **1.00/1.00**

BMP route-monitoring Update/Withdraw check recognized excessive topology changes.

BMP peer Down/Up check recognized issue with **unstable peer on other network platform..**

Interface Down/Up check did not apply.

Traffic Drop spike recognized drops due to instable routing topology.

Missing Traffic recognized traffic volume changes **due to public holidays.**

Increased or decreased Flow Count triggered sporadically **due to public holidays** flow count changes.

Overall: 2 out of 6 checks have detected the excessive routing topology changes with drops. Customer profiling related false positives see in conclusion.

April 17-22th, OSPF/BGP Routing Instability

Provider Impact Analysis – **BGP Churn Origination**

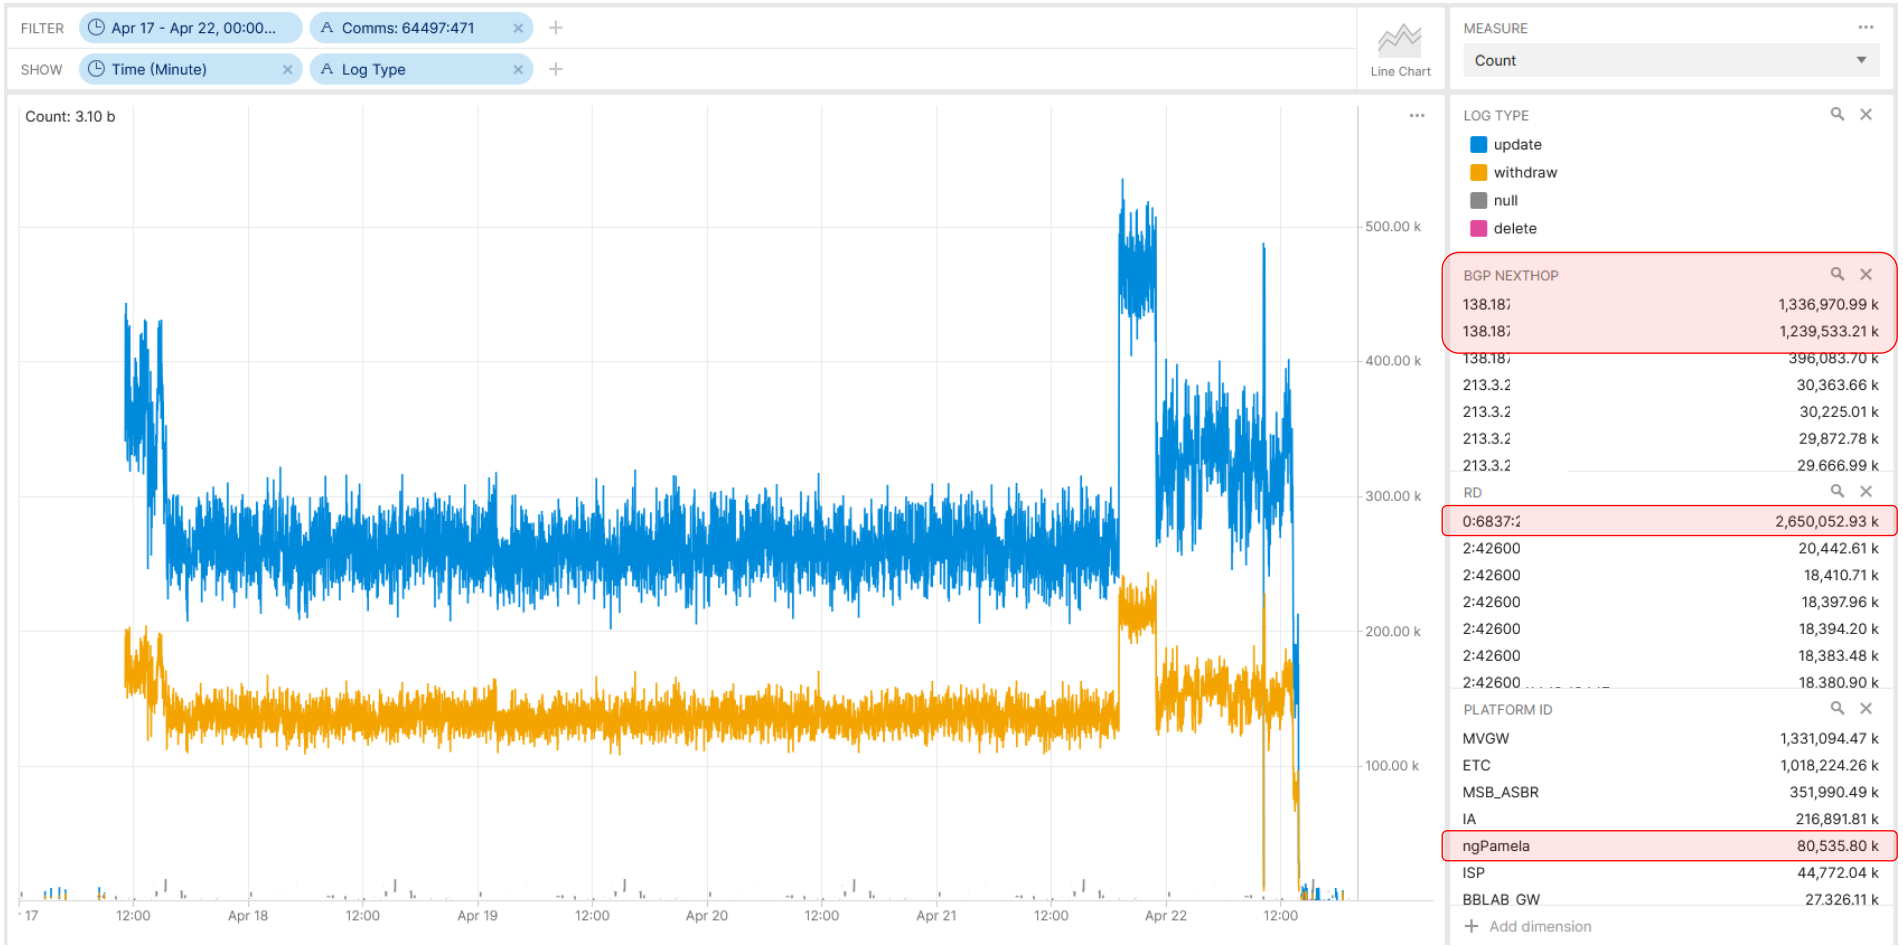

Shows BGP next-hops in updates and withdrawals and on which network platform observed.

Next-hop count shows trigger and the beginning of the causality chain.

Network platform observation count shows who is impacted.

Comparison between Next-hop and Network platform observation count shows that the **platform who originates is not the most impacted** platform

Operational Network Telemetry BMP collected metrics.

April 17-22th, OSPF/BGP Routing Instability

Provider Impact Analysis – **BGP Churn** – 64 Prefixes

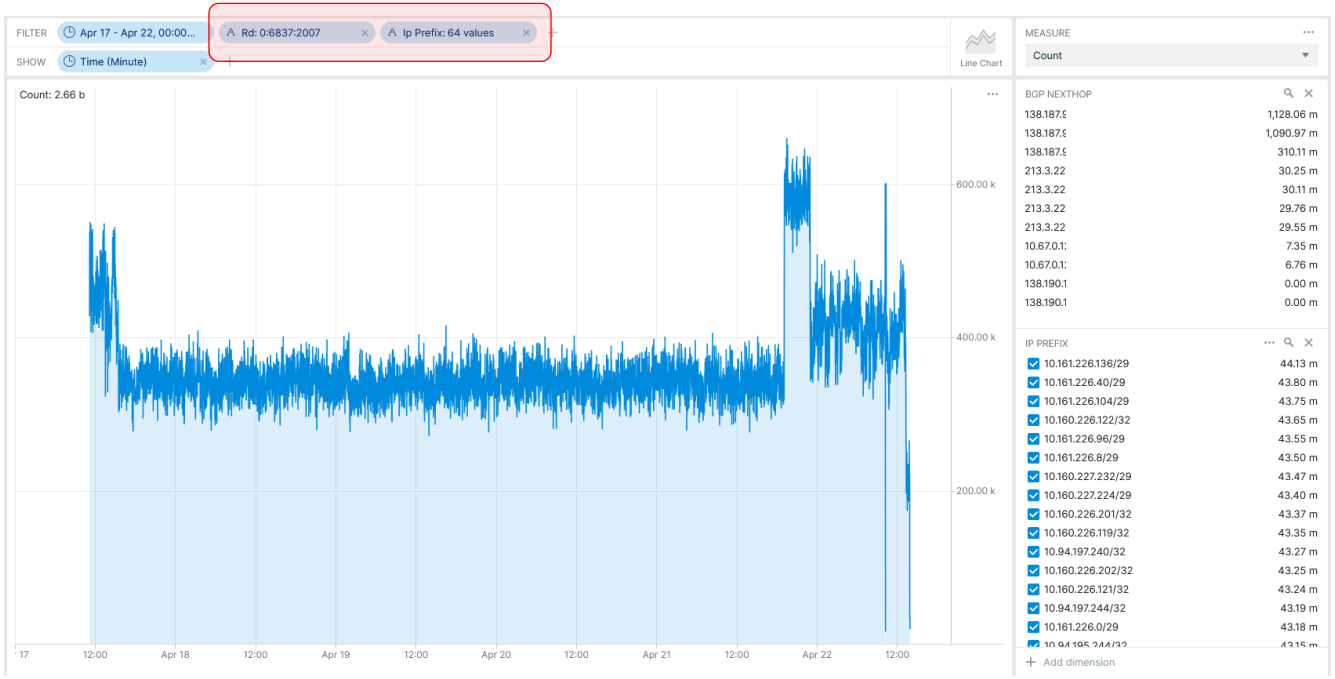

Operational Network Telemetry BMP collected metrics

```
RP/0/RSP0/CPU0:ipc-bei640-r-en-01#sh route vrf MOBILE-SIP-VRF ospf | i 00:00:0
O E2 10.94.197.112/29 [110/1] via 192.168.72.6, 00:00:02, GigabitEthernet0/0/1/4.236
O E2 10.161.226.0/29 [110/1] via 192.168.72.6, 00:00:00, GigabitEthernet0/0/1/4.236
O E2 10.161.226.8/29 [110/1] via 192.168.72.6, 00:00:01, GigabitEthernet0/0/1/4.236
O E2 10.161.226.24/29 [110/1] via 192.168.72.6, 00:00:04, GigabitEthernet0/0/1/4.236
O E2 10.161.226.56/29 [110/1] via 192.168.72.6, 00:00:01, GigabitEthernet0/0/1/4.236
O E2 10.161.226.176/29 [110/1] via 192.168.72.6, 00:00:04, GigabitEthernet0/0/1/4.236
O E2 10.161.227.64/29 [110/1] via 192.168.72.6, 00:00:00, GigabitEthernet0/0/1/4.236
O E2 10.161.227.72/29 [110/1] via 192.168.72.6, 00:00:01, GigabitEthernet0/0/1/4.236
RP/0/RSP0/CPU0:ipc-bei640-r-en-01#

RP/0/RSP0/CPU0:ipc-lss690-r-en-01#sh route vrf MOBILE-SIP-VRF ospf | i 00:00:0
O E2 10.94.195.48/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.94.195.112/29 [110/1] via 192.168.72.70, 00:00:00, GigabitEthernet0/0/1/6.236
O E2 10.94.195.240/32 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.94.195.243/32 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.94.197.96/28 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.94.197.240/32 [110/1] via 192.168.72.70, 00:00:02, GigabitEthernet0/0/1/6.236
O E2 10.160.226.121/32 [110/1] via 192.168.72.70, 00:00:02, GigabitEthernet0/0/1/6.236
O E2 10.160.226.122/32 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.160.226.194/32 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.160.226.195/32 [110/1] via 192.168.72.70, 00:00:00, GigabitEthernet0/0/1/6.236
O E2 10.160.226.198/32 [110/1] via 192.168.72.70, 00:00:01, GigabitEthernet0/0/1/6.236
O E2 10.160.226.199/32 [110/1] via 192.168.72.70, 00:00:01, GigabitEthernet0/0/1/6.236
O E2 10.160.226.200/32 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.160.226.201/32 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.160.227.224/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.161.226.48/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.161.226.104/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.161.226.120/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.161.226.152/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
O E2 10.161.226.176/29 [110/1] via 192.168.72.70, 00:00:01, GigabitEthernet0/0/1/6.236
O E2 10.161.226.184/29 [110/1] via 192.168.72.70, 00:00:03, GigabitEthernet0/0/1/6.236
```

Shows that only 64 prefixes were involved in the 600'000 BGP topology changes per minute across the Swisscom network.

April 17-22th, OSPF/BGP Routing Instability

Customer Impact Analysis – Traffic Drops

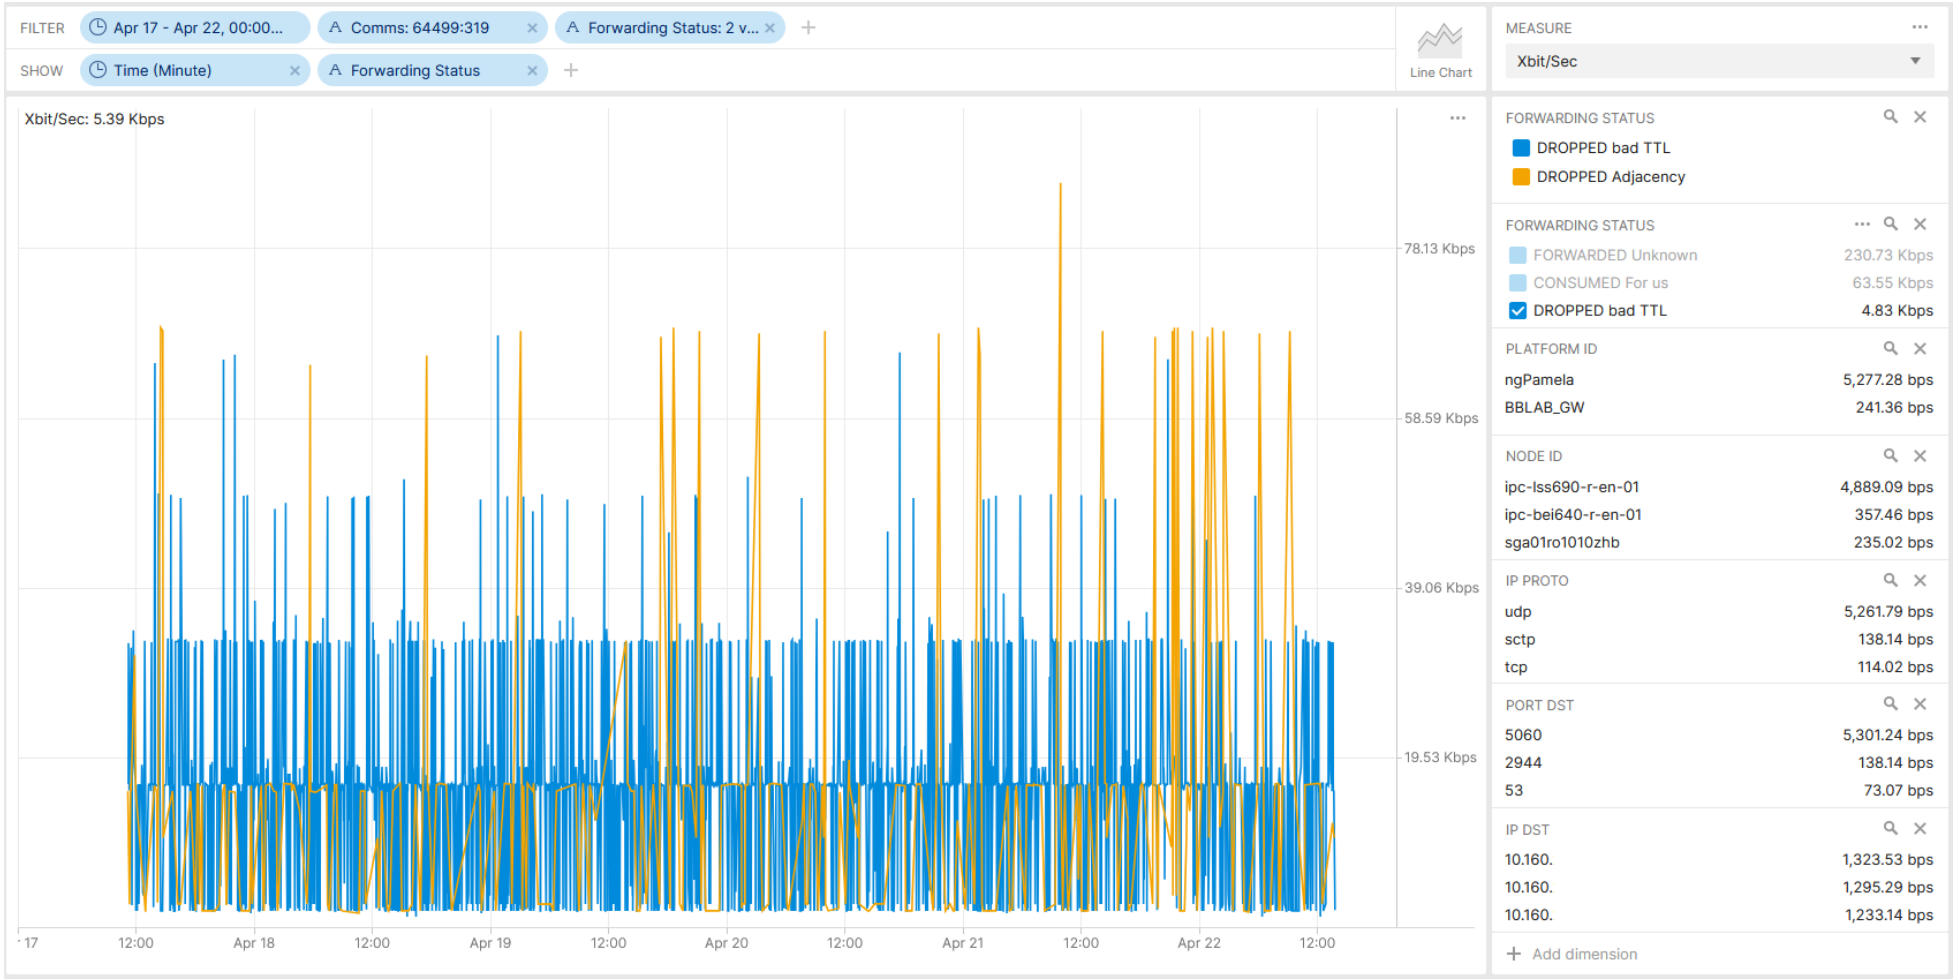

Operational Network Telemetry IPFIX and BMP collected metrics.

Shows which application transport sessions (SIP, Diameter) were affected on the unstable routing topology. Drops occurred on network nodes where OSPF routes are redistributed into BGP RIB.

IETF NMOP - Semantic Metadata Annotation for Network Anomaly Detection

draft-ietf-nmop-network-anomaly-semantics – **National Holidays – The Easter Egg**

Operational Network Telemetry forwarding plane, IPFIX, BMP measured control plane metrics.

```
+--ro symptom!
        |  +--ro id                              yang:uuid
        |  +--ro concern-score                   score
        |  +--ro smcblsymptom:action?            string
        |  +--ro smcblsymptom:reason?            string
        |  +--ro smcblsymptom:trigger?           string
        |  +--ro smcblsymptom:network-plane?     enumeration
        |  +--ro smcblsymptom:strategy?          string
        |  +--ro smcblsymptom:template?          string
        |  +--ro smcblsymptom:season?            Enumeration
```

National holiday information should be considered to improve accuracy of Contextual outliers for seasonal traffic volume and flow count change categorized profiles in the missing traffic and flow count spike strategies and declared in symptom semantics.

IETF NMOP - Semantic Metadata Annotation for Network Anomaly Detection

draft-ietf-nmop-network-anomaly-semantics – Schema Tree

```
notifications:
   +---n relevant-state-notification
      +--ro publisher
      |  +--ro id?         yang:uuid
      |  +--ro name        string
      |  +--ro version?    string
      +--ro id                                  yang:uuid
      +--ro uri?                                inet:uri
      +--ro description?                        string
      +--ro start-time                          yang:date-and-time
      +--ro end-time?                           yang:date-and-time
      +--ro smcblsymptom:strategy?              string
      +--ro confidence-score?                   score
      +--ro concern-score                       score
      +--ro (service)?
      |  +--:(smtopology:l2vpn)
      |  |  +--ro smtopology:vpn-service* [vpn-id]
      |  |     +--ro smtopology:vpn-id                string
      |  |     +--ro smtopology:uri?                  inet:uri
      |  |     +--ro smtopology:vpn-name?             string
      |  |     +--ro smtopology:site-ids*             string
      |  |     +--ro smtopology:change-id?            yang:uuid
      |  |     +--ro smtopology:change-start-time?
      |  |     |       yang:date-and-time
      |  |     +--ro smtopology:change-end-time?
      |  |             yang:date-and-time
      |  +--:(smtopology:l3vpn)
      |     +--ro smtopology:vpn-service* [vpn-id]
      |        +--ro smtopology:vpn-id                string
      |        +--ro smtopology:uri?                  inet:uri
      |        +--ro smtopology:vpn-name?             string
      |        +--ro smtopology:site-ids*             string
      |        +--ro smtopology:change-id?            yang:uuid
      |        +--ro smtopology:change-start-time?
      |        |       yang:date-and-time
      |        +--ro smtopology:change-end-time?
      |                yang:date-and-time
```

```
notifications:
   +---n relevant-state-notification
   +--ro anomaly* [id revision]
         +--ro id                                   yang:uuid
         +--ro revision                             yang:counter32
         +--ro uri?                                 inet:uri
         +--ro state                                identityref
         +--ro description?                         string
         +--ro start-time
         |       yang:date-and-time
         +--ro end-time?
         |       yang:date-and-time
         +--ro confidence-score?                    score
         +--ro pattern?                             identityref
         +--ro annotator
         |  +--ro id?                yang:uuid
         |  +--ro name               string
         |  +--ro version?           string
         |  +--ro annotator-type?    enumeration
         +--ro symptom!
         |  +--ro id                                yang:uuid
         |  +--ro concern-score                     score
         |  +--ro smcblsymptom:action?              string
         |  +--ro smcblsymptom:reason?              string
         |  +--ro smcblsymptom:trigger?             string
         |  +--ro smcblsymptom:network-plane?       enumeration
         |  +--ro smcblsymptom:template?            string
         |  +--ro smcblsymptom:season?              Enumeration
         +--ro smtopology:vpn-node-terminations*
                  [hostname route-distinguisher]
            +--ro smtopology:hostname                inet:host
            +--ro smtopology:route-distinguisher     string
            +--ro smtopology:peer-ip*                inet:ip-address
            +--ro smtopology:next-hop*               inet:ip-address
            +--ro smtopology:interface-id*           uint32
```

Shows
the observed symptoms, the network dimensions triggering and connectivity service impacted.

IETF NMOP 122 – Network Observability Development

Network Anomaly Detection and YANG-Push/Message Broker Integration

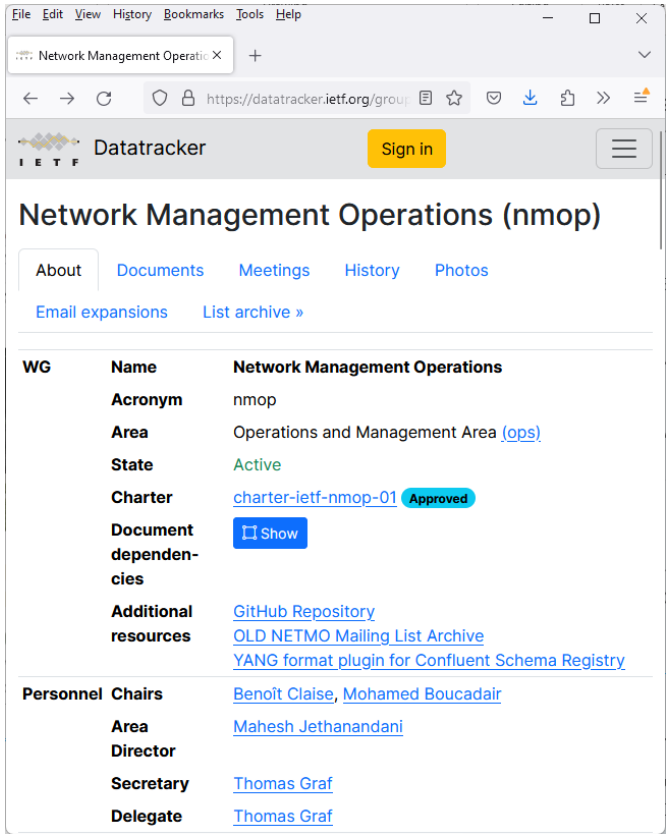

https://datatracker.ietf.org/group/nmop/about/

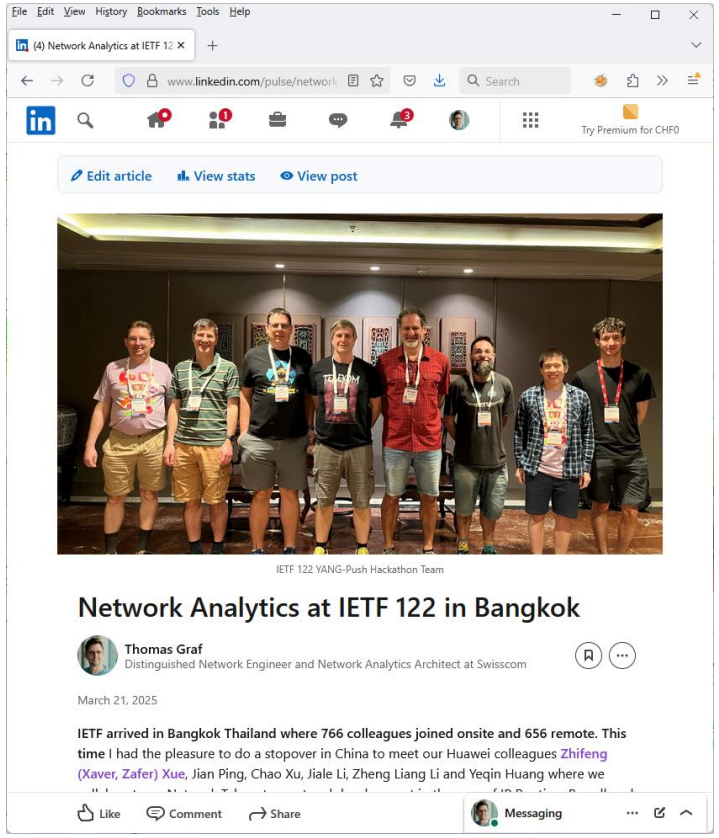

https://www.linkedin.com/pulse/network-analytics-ietf-122-bangkok-thomas-graf-nhmge/

Network Anomaly Detection

Relevant Papers for more Details

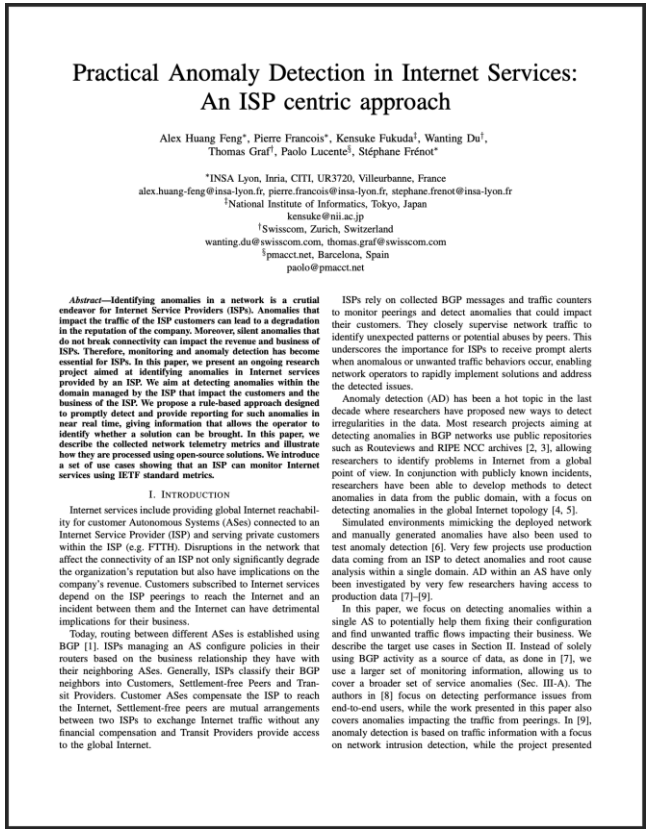

Practical Anomaly Detection in Internet Services: An ISP centric approach

Alex Huang Feng*, Pierre Francois*, Kensuke Fukuda‡, Wanting Du†, Thomas Graf†, Paolo Lucente§, Stéphane Frénot*

*INSA Lyon, Inria, CITI, UR3720, Villeurbanne, France
alex.huang-feng@insa-lyon.fr, pierre.francois@insa-lyon.fr, stephane.frenot@insa-lyon.fr
‡National Institute of Informatics, Tokyo, Japan
kensuke@nii.ac.jp
†Swisscom, Zurich, Switzerland
wanting.du@swisscom.com, thomas.graf@swisscom.com
§pmacct.net, Barcelona, Spain
paolo@pmacct.net

***Abstract*—Identifying anomalies in a network is a crutial endeavor for Internet Service Providers (ISPs). Anomalies that impact the traffic of the ISP customers can lead to a degradation in the reputation of the company. Moreover, silent anomalies that do not break connectivity can impact the revenue and business of ISPs. Therefore, monitoring and anomaly detection has become essential for ISPs. In this paper, we present an ongoing research project aimed at identifying anomalies in Internet services provided by an ISP. We aim at detecting anomalies within the domain managed by the ISP that impact the customers and the business of the ISP. We propose a rule-based approach designed to promptly detect and provide reporting for such anomalies in near real time, giving information that allows the operator to identify whether a solution can be brought. In this paper, we describe the collected network telemetry metrics and illustrate how they are processed using open-source solutions. We introduce a set of use cases showing that an ISP can monitor Internet services using IETF standard metrics.**

I. Introduction

Internet services include providing global Internet reachability for customer Autonomous Systems (ASes) connected to an Internet Service Provider (ISP) and serving private customers within the ISP (e.g. FTTH). Disruptions in the network that affect the connectivity of an ISP not only significantly degrade the organization's reputation but also have implications on the company's revenue. Customers subscribed to Internet services depend on the ISP peerings to reach the Internet and an incident between them and the Internet can have detrimental implications for their business.

Today, routing between different ASes is established using BGP [1]. ISPs managing an AS configure policies in their routers based on the business relationship they have with their neighboring ASes. Generally, ISPs classify their BGP neighbors into Customers, Settlement-free Peers and Transit Providers. Customer ASes compensate the ISP to reach the Internet, Settlement-free peers are mutual arrangements between two ISPs to exchange Internet traffic without any financial compensation and Transit Providers provide access to the global Internet.

ISPs rely on collected BGP messages and traffic counters to monitor peerings and detect anomalies that could impact their customers. They closely supervise network traffic to identify unexpected patterns or potential abuses by peers. This underscores the importance for ISPs to receive prompt alerts when anomalous or unwanted traffic behaviors occur, enabling network operators to rapidly implement solutions and address the detected issues.

Anomaly detection (AD) has been a hot topic in the last decade where researchers have proposed new ways to detect irregularities in the data. Most research projects aiming at detecting anomalies in BGP networks use public repositories such as Routeviews and RIPE NCC archives [2, 3], allowing researchers to identify problems in Internet from a global point of view. In conjunction with publicly known incidents, researchers have been able to develop methods to detect anomalies in data from the public domain, with a focus on detecting anomalies in the global Internet topology [4, 5].

Simulated environments mimicking the deployed network and manually generated anomalies have also been used to test anomaly detection [6]. Very few projects use production data coming from an ISP to detect anomalies and root cause analysis within a single domain. AD within an AS have only been investigated by very few researchers having access to production data [7]–[9].

In this paper, we focus on detecting anomalies within a single AS to potentially help them fixing their configuration and find unwanted traffic flows impacting their business. We describe the target use cases in Section II. Instead of solely using BGP activity as a source of data, as done in [7], we use a larger set of monitoring information, allowing us to cover a broader set of service anomalies (Sec. III-A). The authors in [8] focus on detecting performance issues from end-to-end users, while the work presented in this paper also covers anomalies impacting the traffic from peerings. In [9], anomaly detection is based on traffic information with a focus on network intrusion detection, while the project presented

Paper "**Practical Anomaly Detection in Internet Services: An ISP centric approach**"
Published at AnNet Workshop (In conjunction with IEEE NOMS)
Seoul, South Korea (6–10 May 2024)
Open access: https://hal.science/hal-04655324

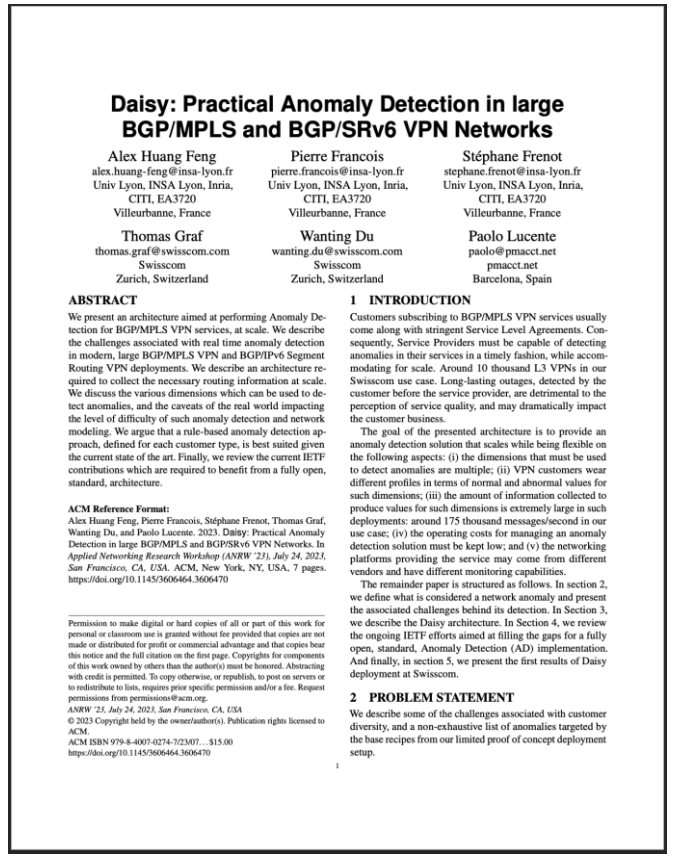

Daisy: Practical Anomaly Detection in large BGP/MPLS and BGP/SRv6 VPN Networks

Alex Huang Feng
alex.huang-feng@insa-lyon.fr
Univ Lyon, INSA Lyon, Inria, CITI, EA3720
Villeurbanne, France

Pierre Francois
pierre.francois@insa-lyon.fr
Univ Lyon, INSA Lyon, Inria, CITI, EA3720
Villeurbanne, France

Stéphane Frenot
stephane.frenot@insa-lyon.fr
Univ Lyon, INSA Lyon, Inria, CITI, EA3720
Villeurbanne, France

Thomas Graf
thomas.graf@swisscom.com
Swisscom
Zurich, Switzerland

Wanting Du
wanting.du@swisscom.com
Swisscom
Zurich, Switzerland

Paolo Lucente
paolo@pmacct.net
pmacct.net
Barcelona, Spain

ABSTRACT

We present an architecture aimed at performing Anomaly Detection for BGP/MPLS VPN services, at scale. We describe the challenges associated with real time anomaly detection in modern, large BGP/MPLS VPN and BGP/IPv6 Segment Routing VPN deployments. We describe an architecture required to collect the necessary routing information at scale. We discuss the various dimensions which can be used to detect anomalies, and the caveats of the real world impacting the level of difficulty of such anomaly detection and network modeling. We argue that a rule-based anomaly detection approach, defined for each customer type, is best suited given the current state of the art. Finally, we review the current IETF contributions which are required to benefit from a fully open, standard, architecture.

ACM Reference Format:
Alex Huang Feng, Pierre Francois, Stéphane Frenot, Thomas Graf, Wanting Du, and Paolo Lucente. 2023. Daisy: Practical Anomaly Detection in large BGP/MPLS and BGP/SRv6 VPN Networks. In *Applied Networking Research Workshop (ANRW '23), July 24, 2023, San Francisco, CA, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3606464.3606470

ANRW '23, July 24, 2023, San Francisco, CA, USA

ACM ISBN 979-8-4007-0274-7/23/07…$15.00
https://doi.org/10.1145/3606464.3606470

1 INTRODUCTION

Customers subscribing to BGP/MPLS VPN services usually come along with stringent Service Level Agreements. Consequently, Service Providers must be capable of detecting anomalies in their services in a timely fashion, while accommodating for scale. Around 10 thousand L3 VPNs in our Swisscom use case. Long-lasting outages, detected by the customer before the service provider, are detrimental to the perception of service quality, and may dramatically impact the customer business.

The goal of the presented architecture is to provide an anomaly detection solution that scales while being flexible on the following aspects: (i) the dimensions that must be used to detect anomalies are multiple; (ii) VPN customers wear different profiles in terms of normal and abnormal values for such dimensions; (iii) the amount of information collected to produce values for such dimensions is extremely large in such deployments: around 175 thousand messages/second in our use case; (iv) the operating costs for managing an anomaly detection solution must be kept low; and (v) the networking platforms providing the service may come from different vendors and have different monitoring capabilities.

The remainder paper is structured as follows. In section 2, we define what is considered a network anomaly and present the associated challenges behind its detection. In Section 3, we describe the Daisy architecture. In Section 4, we review the ongoing IETF efforts aimed at filling the gaps for a fully open, standard, Anomaly Detection (AD) implementation. And finally, in section 5, we present the first results of Daisy deployment at Swisscom.

2 PROBLEM STATEMENT

We describe some of the challenges associated with customer diversity, and a non-exhaustive list of anomalies targeted by the base recipes from our limited proof of concept deployment setup.

Paper "**Daisy: Practical Anomaly Detection in large BGP/MPLS and BGP/SRv6 VPN Networks**" published at ACM/IRTF ANRW'23
San Francisco, USA (24 July 2023)
Open access: http://hal.science/hal-04307611

19