Investigation of Traffic Dependencies between IXPs in Failure Scenarios

MENOG16

Christoph Dietzel

R&D DE-CIX / TU Berlin

Introduction

- How robust is the IXP interconnection system?
  - What happens if a large IXP fails?
  - Does it affect other IXPs and how?
  - How are other regions affected?
- Analyzed a large IXP outage
- This talk is about the results
- Lessons Learned?

Incident at IXP A – Data Plane

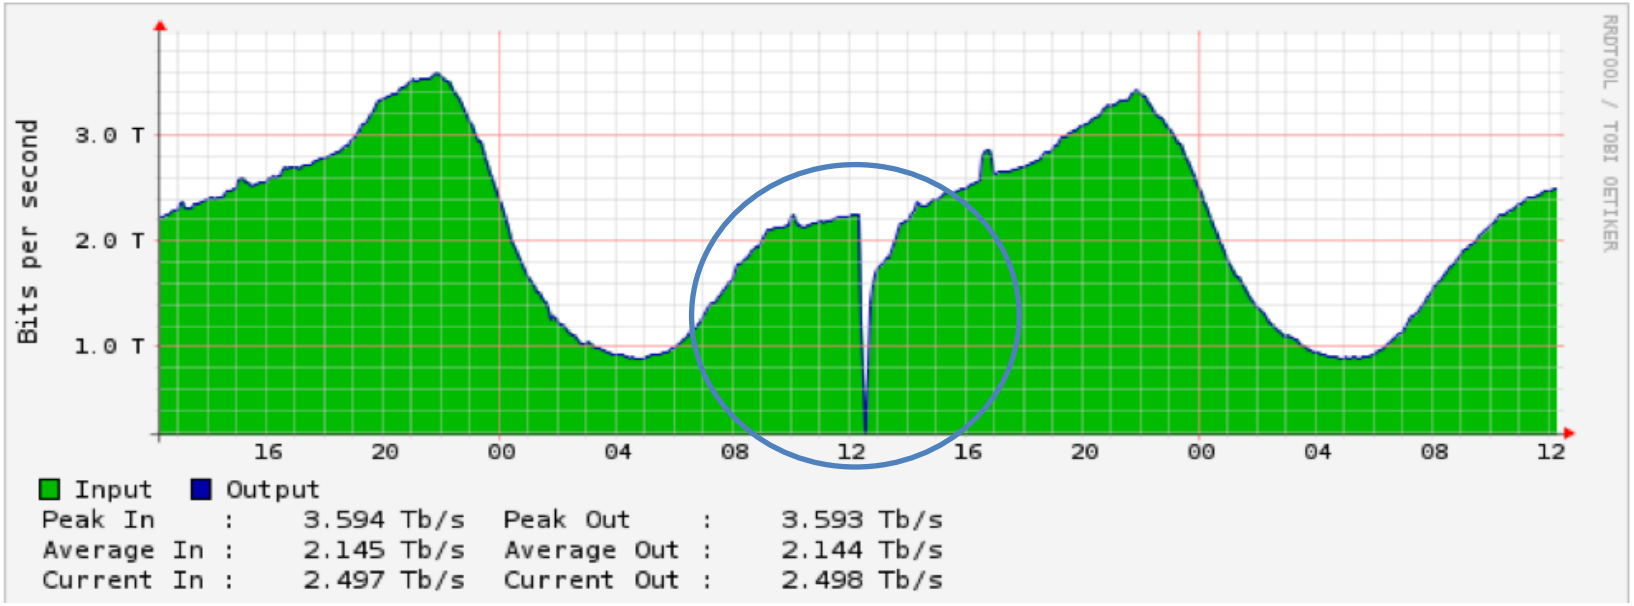

- 13th May 2015 at 12:22 pm (CET)
- Loop with 4 x 100GE, traffic was blackholed
- About 500 of 600 BGP session drops at the route servers

Incident at IXP A – Control Plane

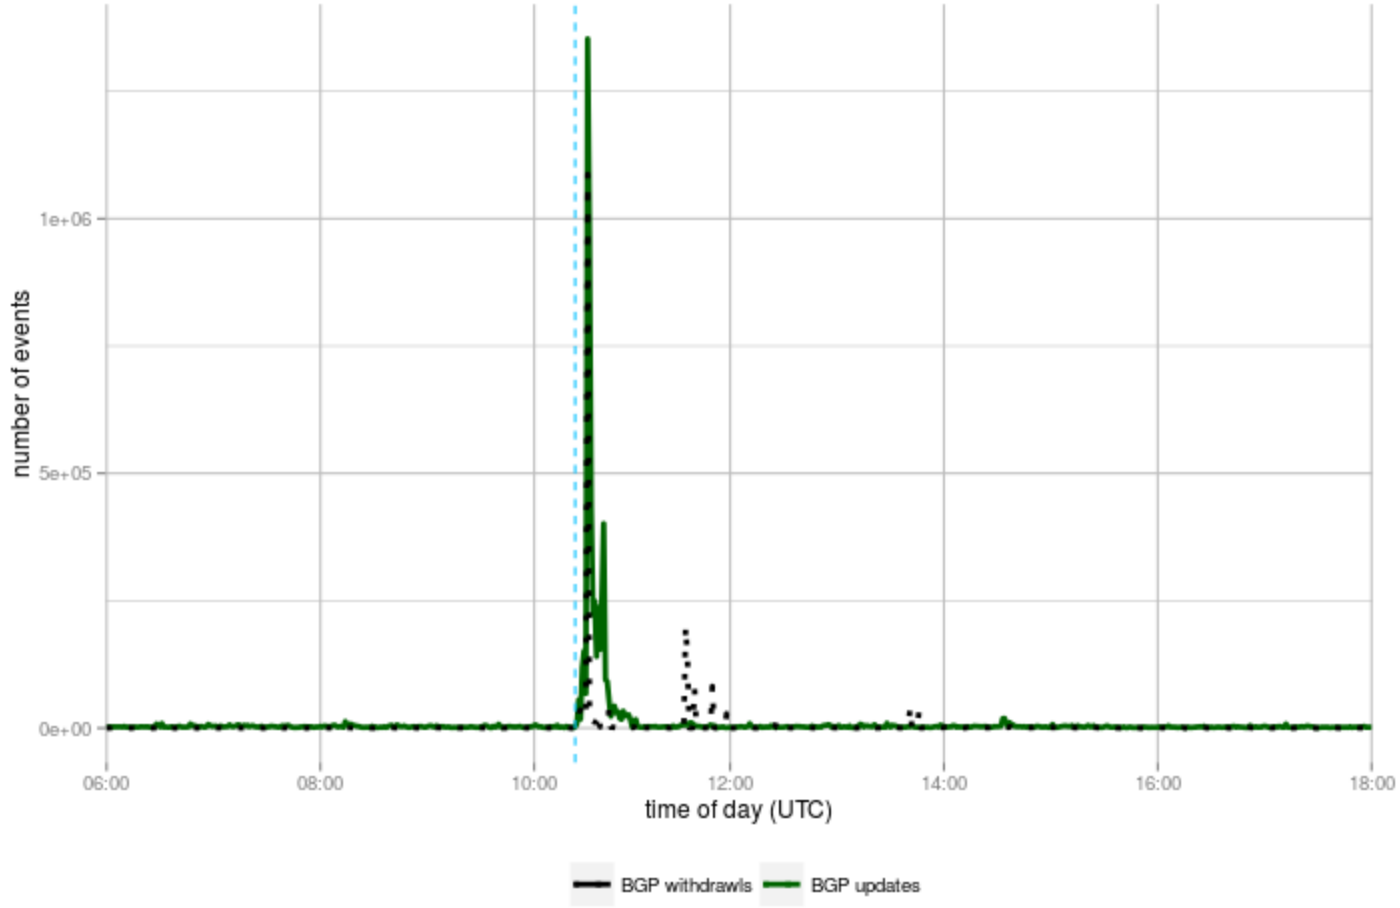

- More than 1.4 million prefix updates
- More than 1 million prefix withdrawals

Impact on other IXPs, i.e., DE-CIX?

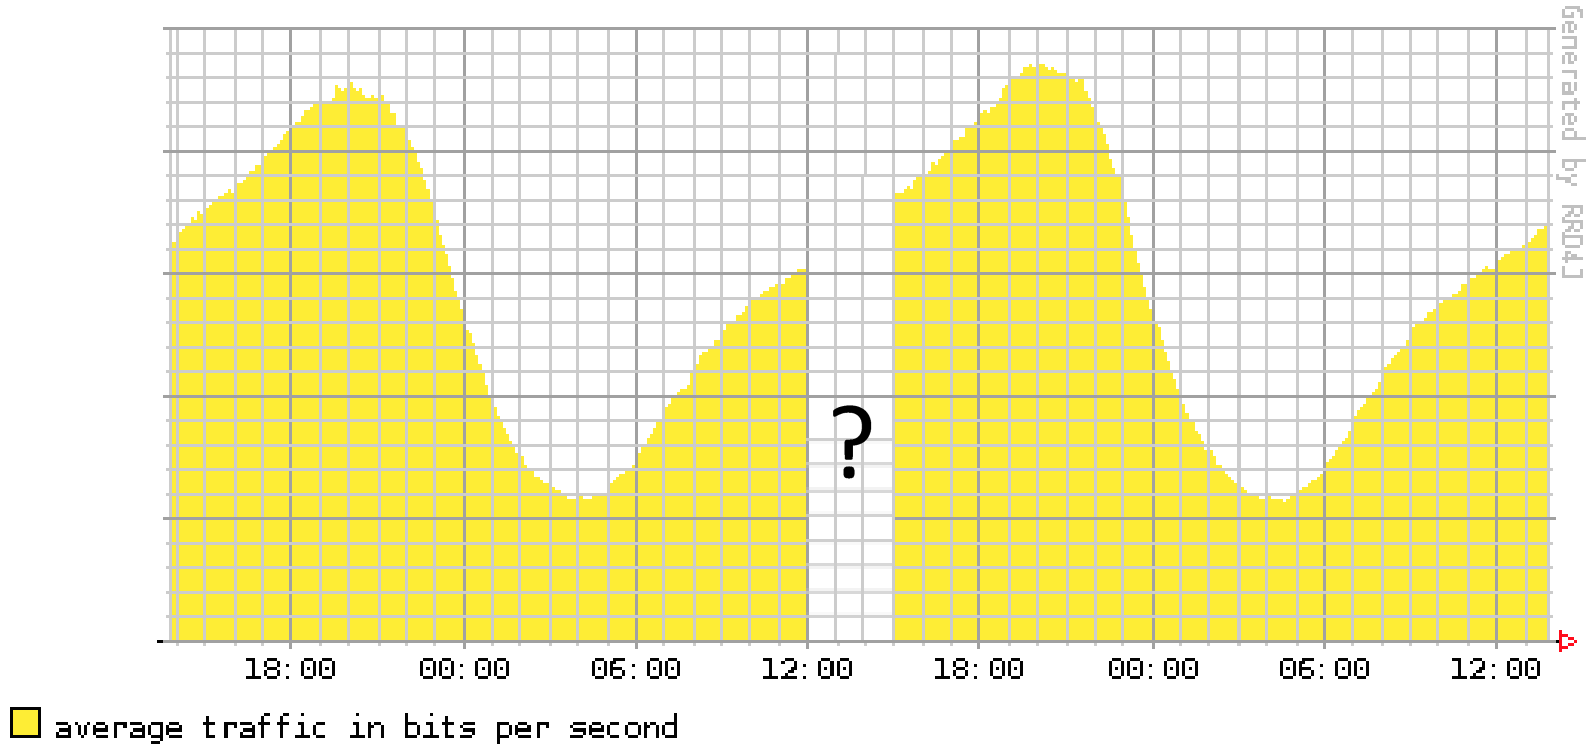

Impact on DE-CIX Frankfurt

- Decreased traffic volume
- Drop of about 240 Gbit/s within 5 minutes
- Recovering after about 10 minutes

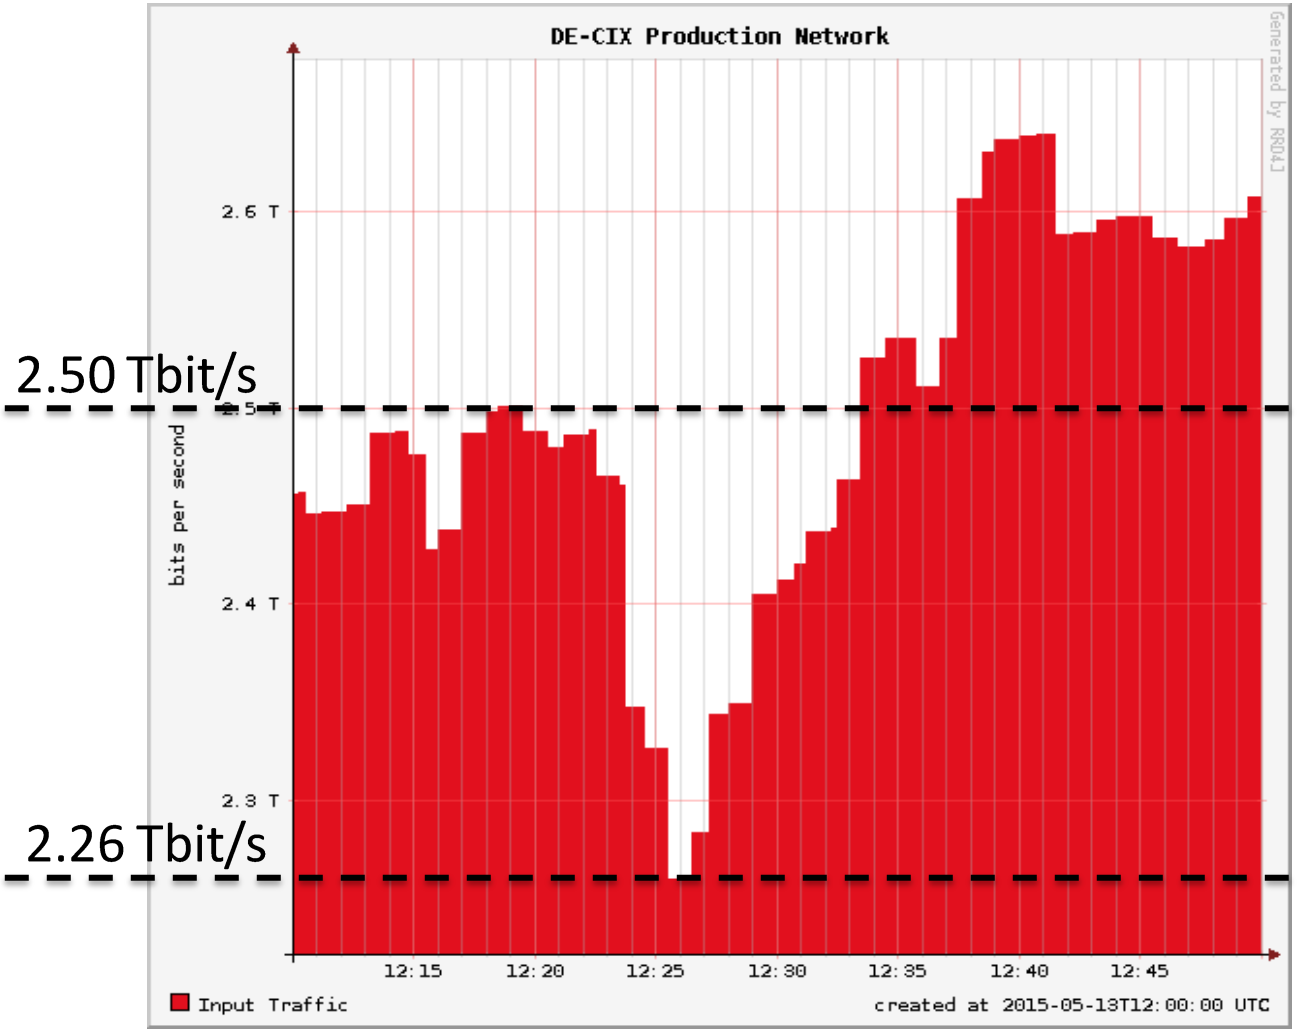

Impact on DE-CIX Frankfurt -- Time Flow

IXP A

DE-CIX

The day of the incident:
(information from public sources [CET times])

1. 12:22 pm – Loop with 4 x 100GE created. Traffic was blackholed.
2. 12:25 pm – About 500 of 600 BGP sessions at the route servers dropped
3. 12:29 pm – NOC reacted and deactivated ports responsible for loop
4. 12:40 pm – BGP sessions to route server are back online

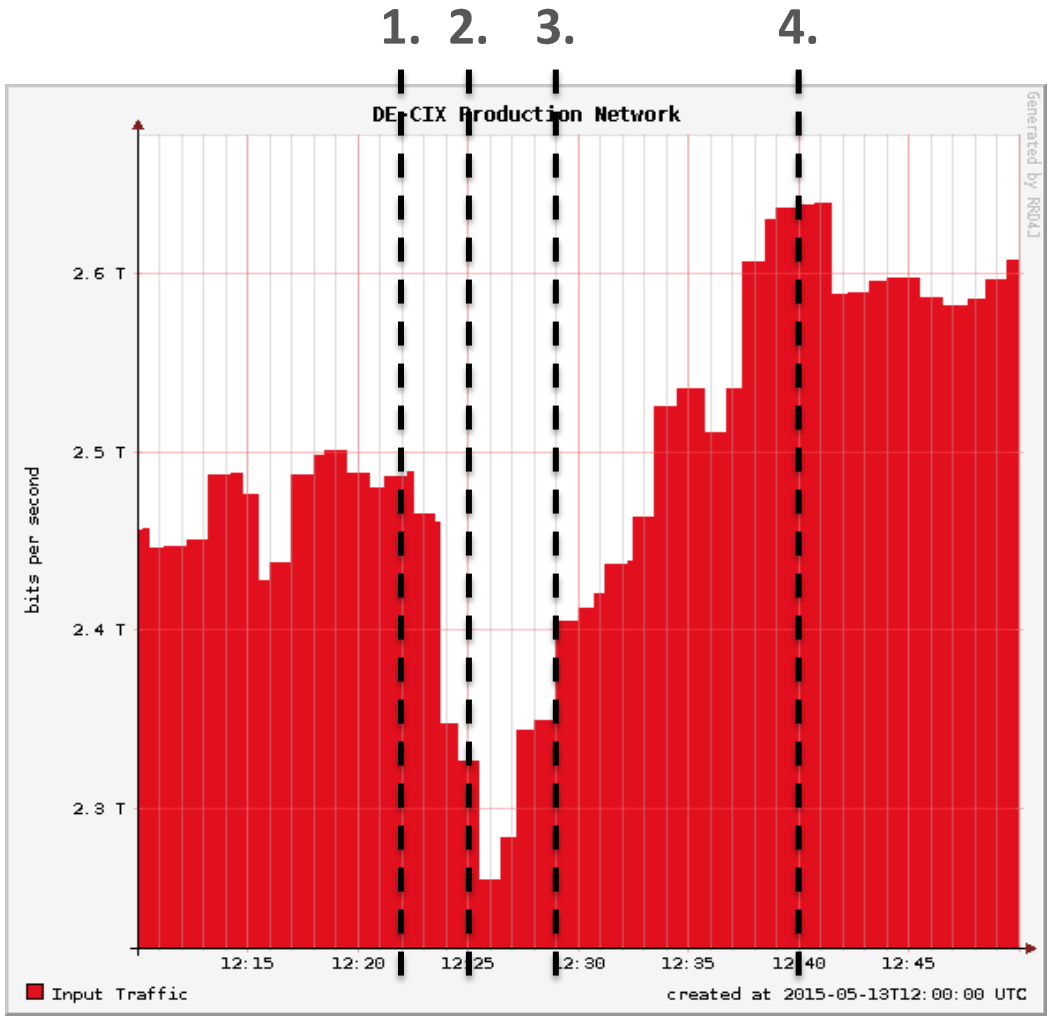

Impact on DE-CIX Frankfurt

- Decreased traffic volume
- Drop of about 240 Gbit/s within 5 minutes
- Recovering after about 10 minutes

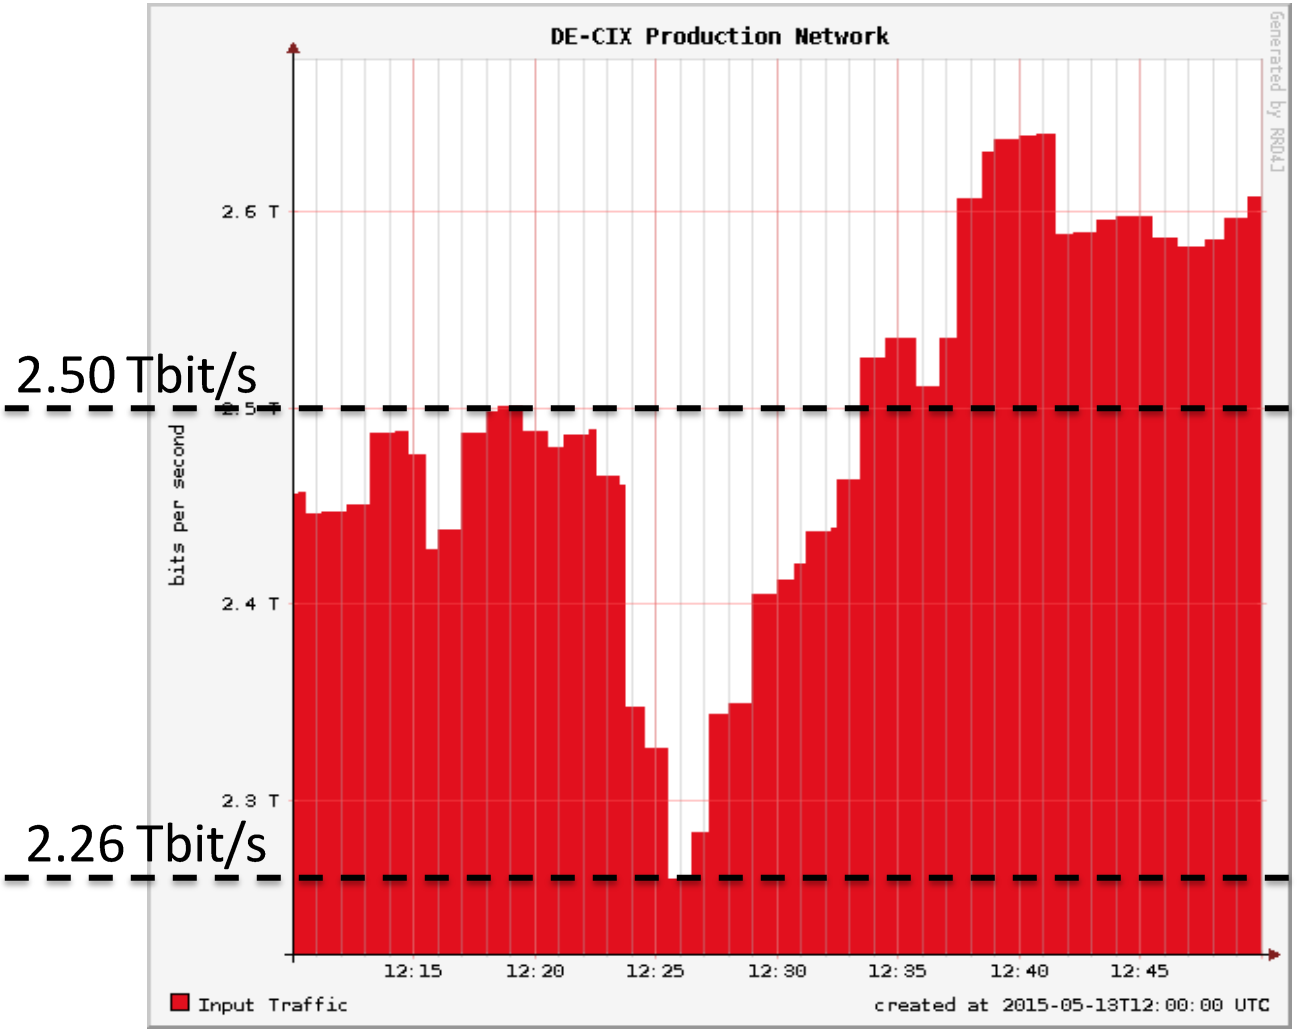

What could be the reason for this behavior?

What are dependencies?

We found three answers… so far…

1. Remote Peering Routers Overloaded

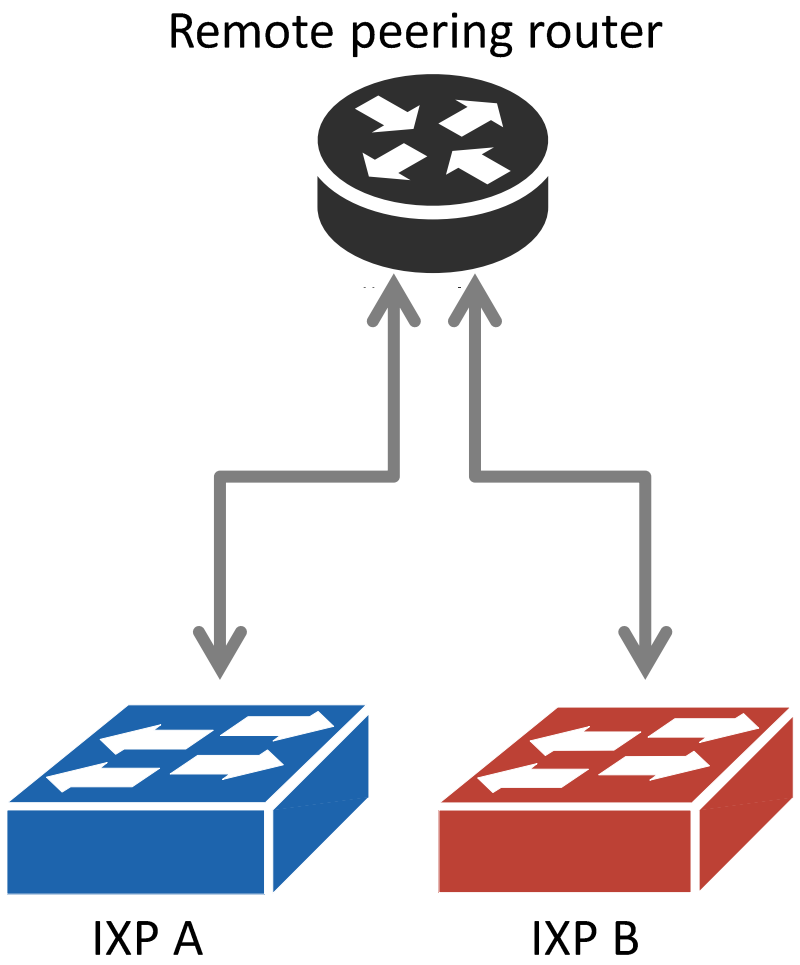

- Overloaded remote peering router drops all BGP sessions
- Four customers at DE-CIX Frankfurt affected with a traffic volume drop of 0.92 Gbit/s

2. Asymmetric Routing Paths

Is there a significant number of asymmetric paths traversing both IXPs?

Example:

- » Upstream traverses IXP A
- » Downstream traverses IXP B

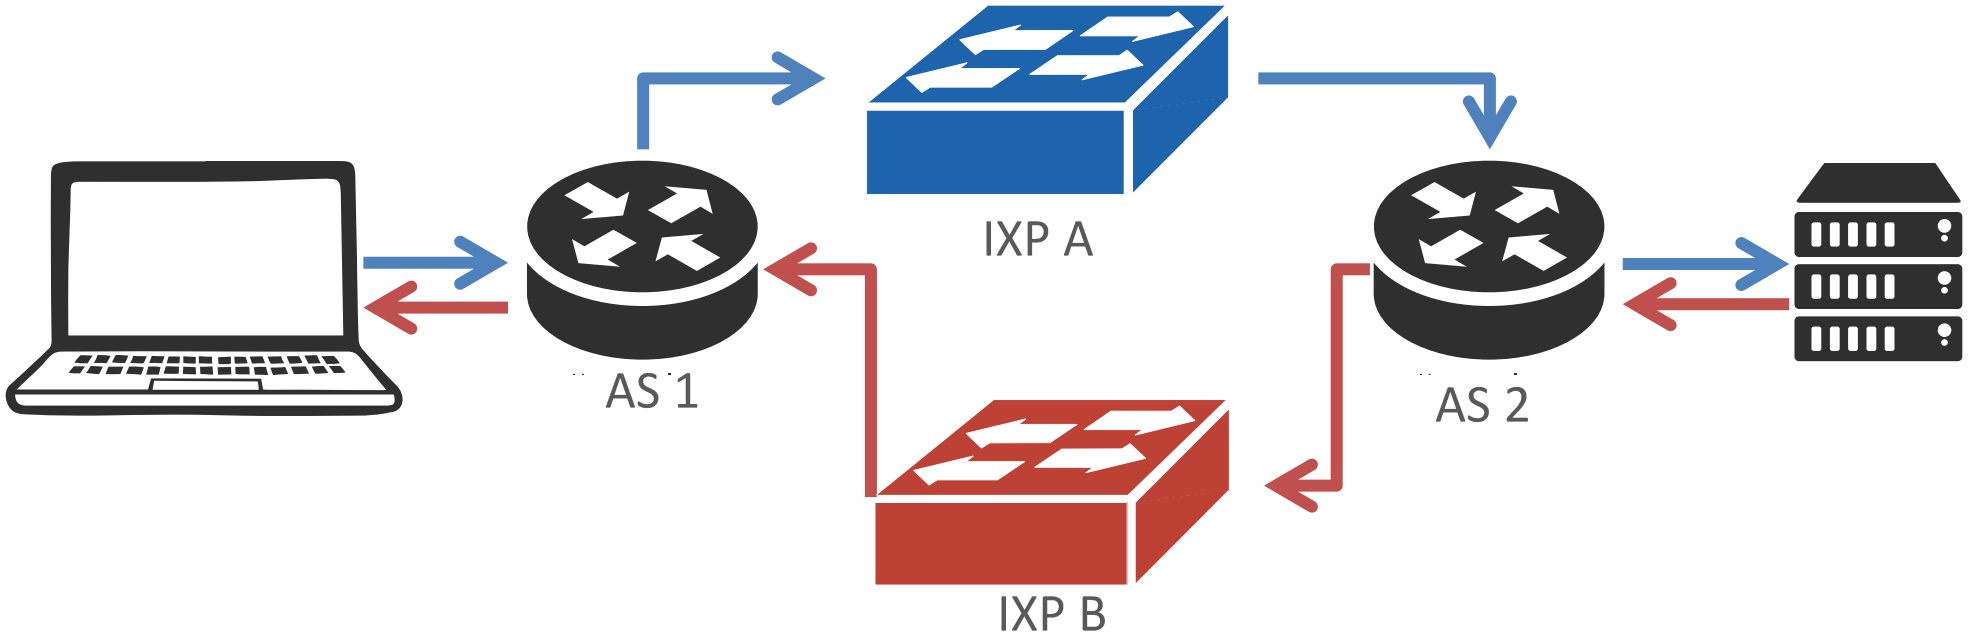

2. Asymmetric Routing Paths II

RIPE Atlas measurements:

» ASes connected to DE-CIX Frankfurt and IXP A: *323 (40-50%)*

» “Peerings” with a traffic drop > 200Mbit/s at DE-CIX Frankfurt: *183*

» ASes hosting (at least one) RIPE Atlas probe: *171*

» Intersection: 50 ASes

Measurement results of full mesh traceroutes:

» 38% of all “connections” are asymmetric

» 8% of all connections traversed no IXP

2. Asymmetric Routing Paths “Validation”

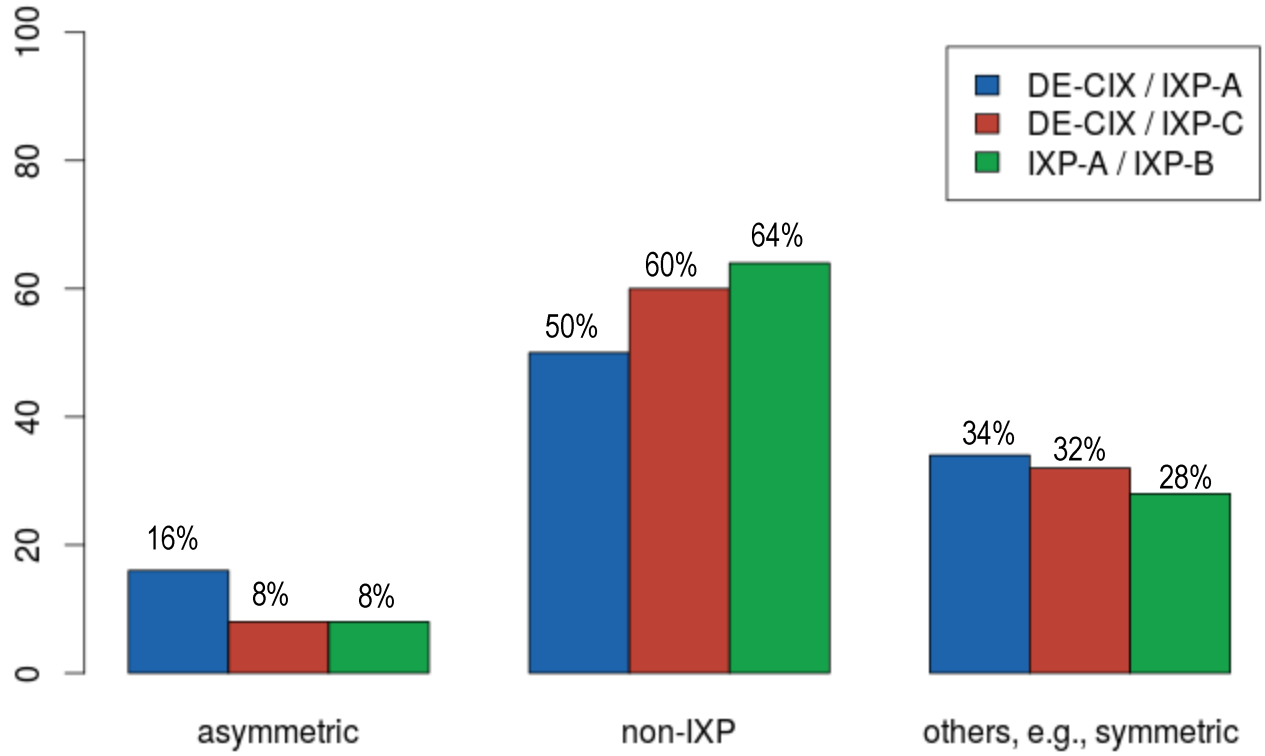

Impact Details

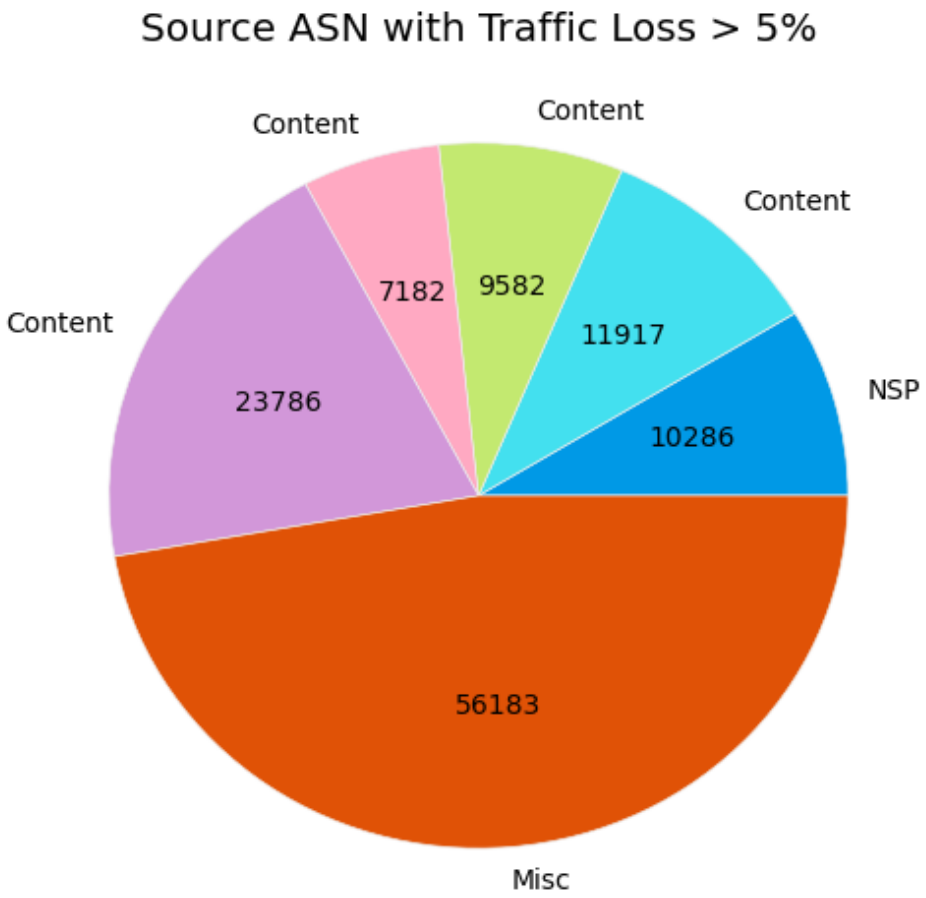

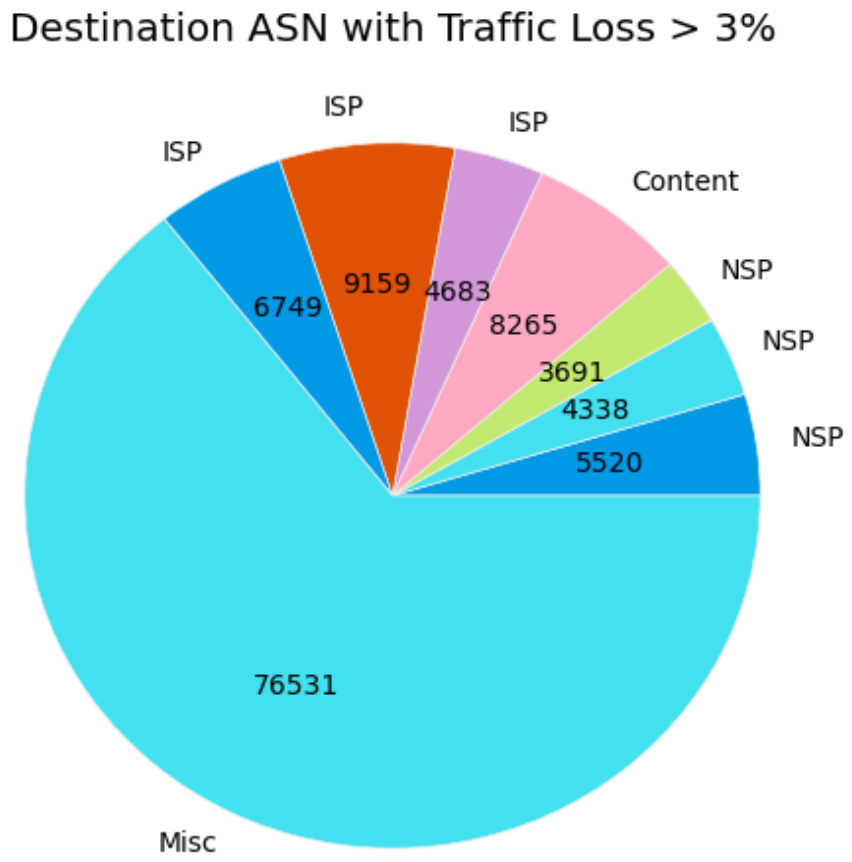

3. Layer 8: Less Users

- Users experienced connection errors
- Users were annoyed by broken “Internet” and switched activities
- Less users resulted in less traffic

- Impact on traffic volume is hard to quantify

Passive Measurement during Outage (RIS collector)

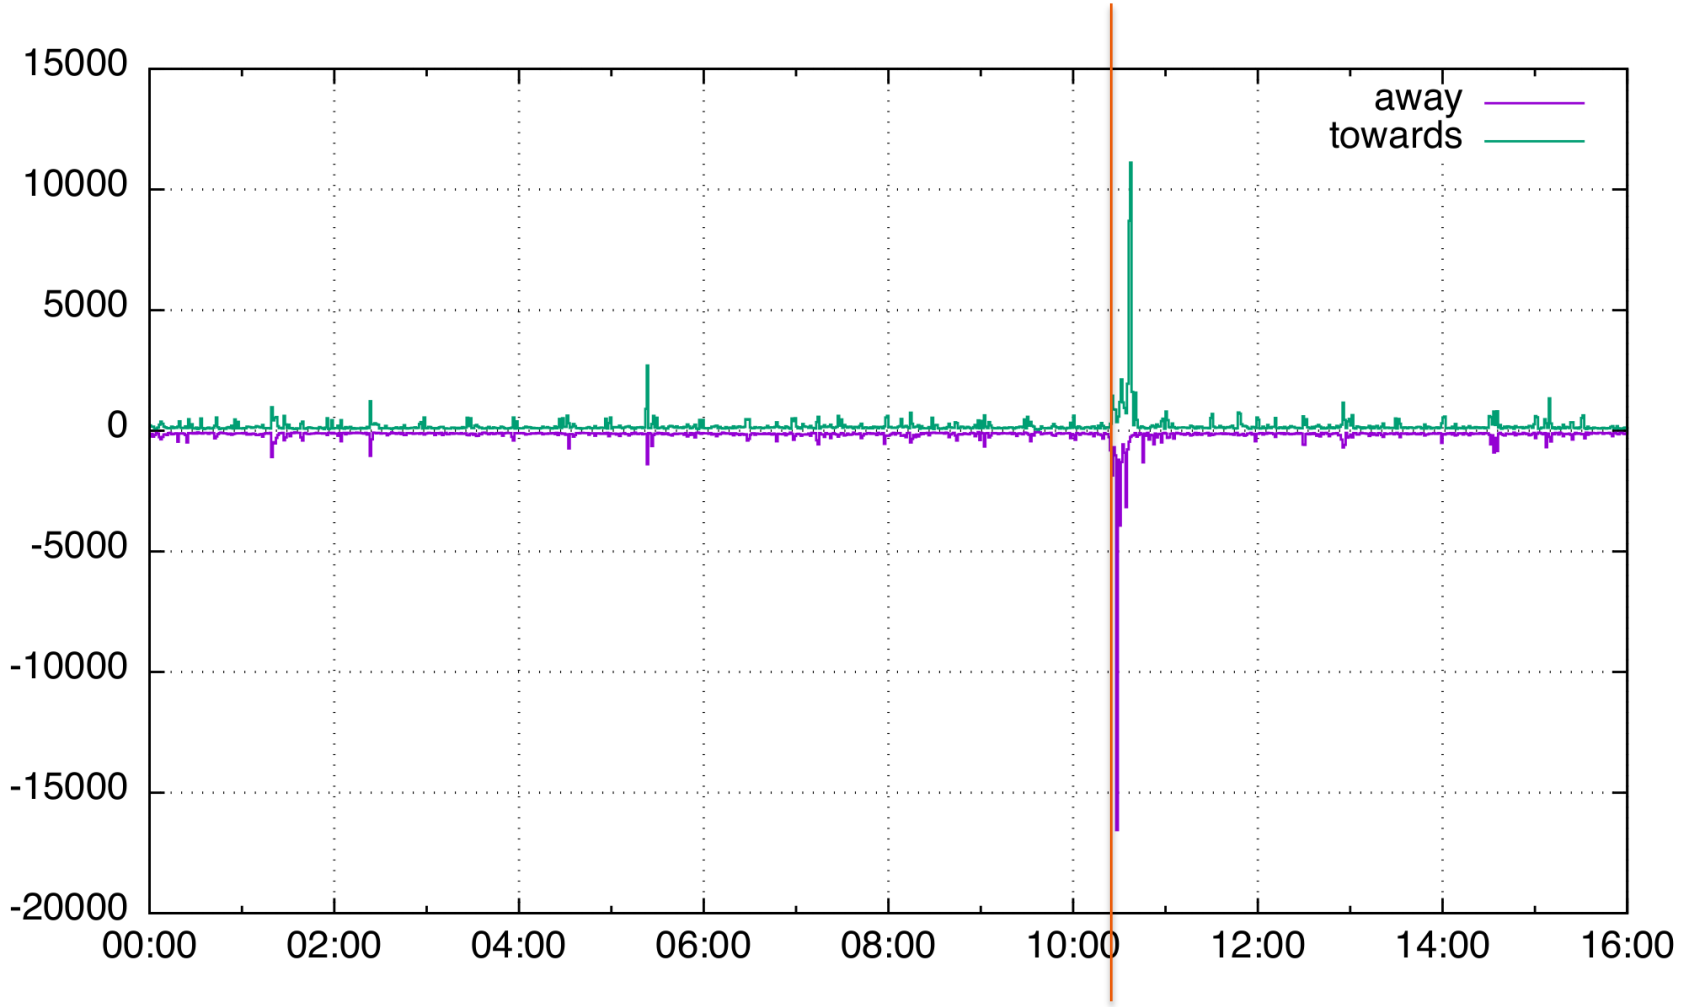

RIPE RCC01 RIS Collector at IXP-C (UTC time)

Active Measurement during Outage (traceroutes)

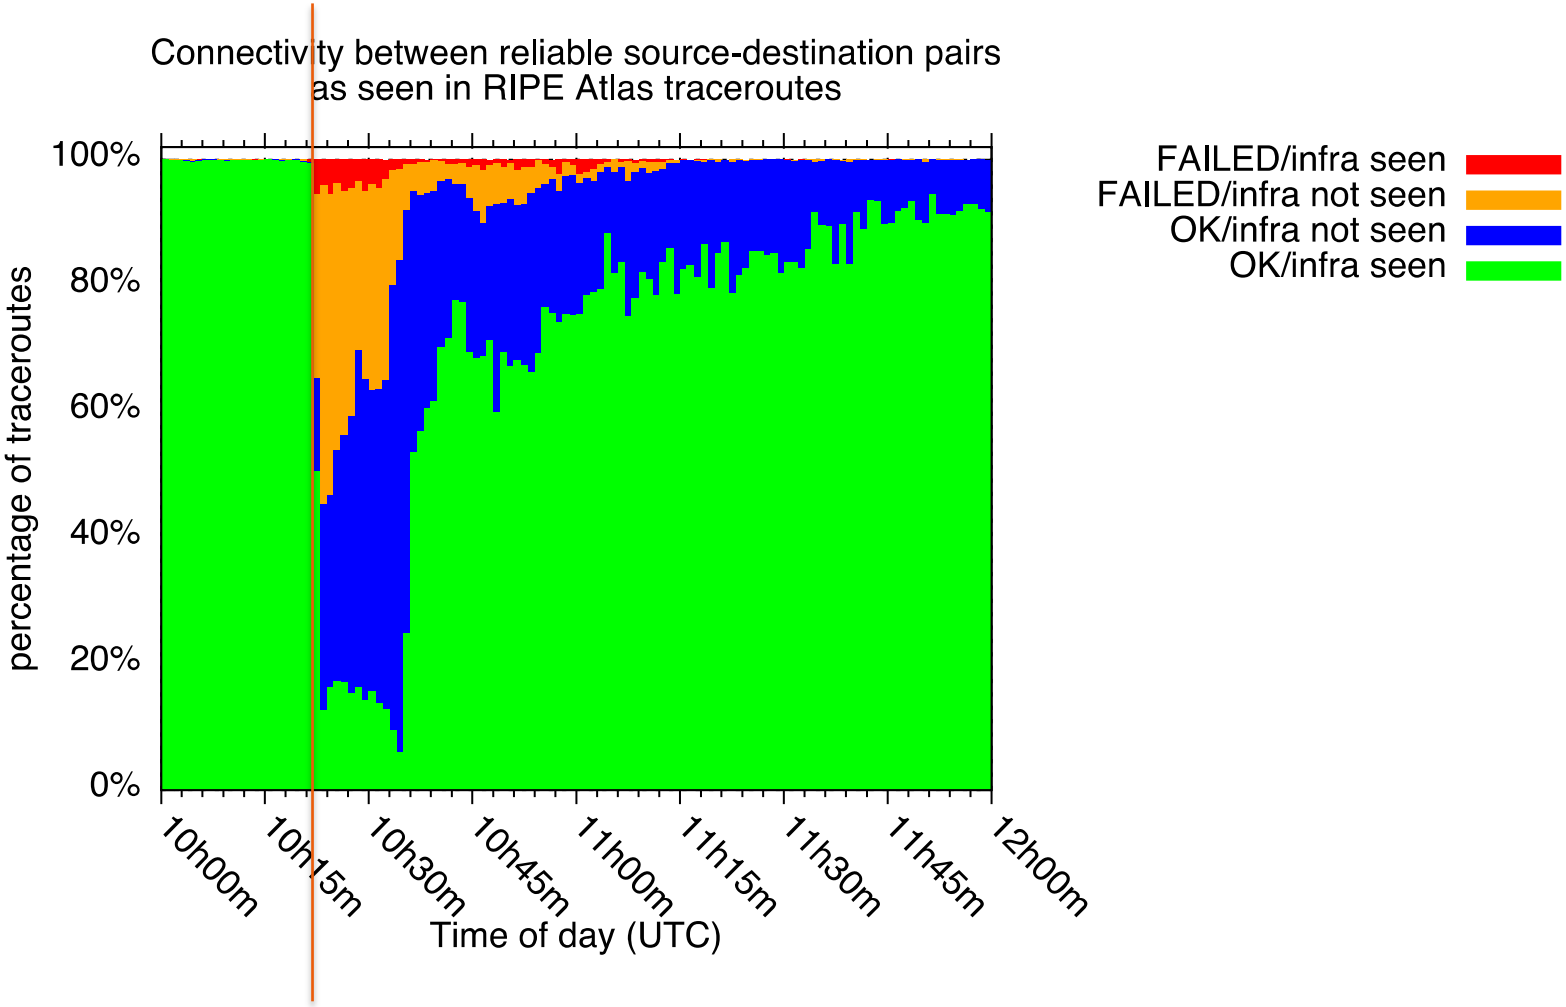

https://labs.ripe.net/Members/emileaben/does-the-internet-route-around-damage

Summary and Takeaway

Reasons for traffic volume dependencies between IXPs:

1. Remote peering routers overloaded
2. Asymmetric routing paths
3. Layer 8: Less users

Good news: Internet infrastructure is not affected largely if a large IXP fails.

Takeaway:

» Knowledge of traffic dependencies of IXPs
» Useful for designing peering and especially remote peering
» Improve recovery time e.g. route server BFD

DE-CIX

Awesome for 20 years

Comments? Questions?

rnd@de-cix.net

Christoph Dietzel

R&D DE-CIX / TU Berlin