

Investigation of Dependencies between IXPs


Daniel Kopp
R&D DE-CIX

Introduction

- » How robust is the IXP interconnection system?
 - » What happens if a large IXP fails?
 - » Does it affect other IXPs and how?
- » There was an incident, which we investigated
- » This presentation is about the results
- » What can we learn from this?

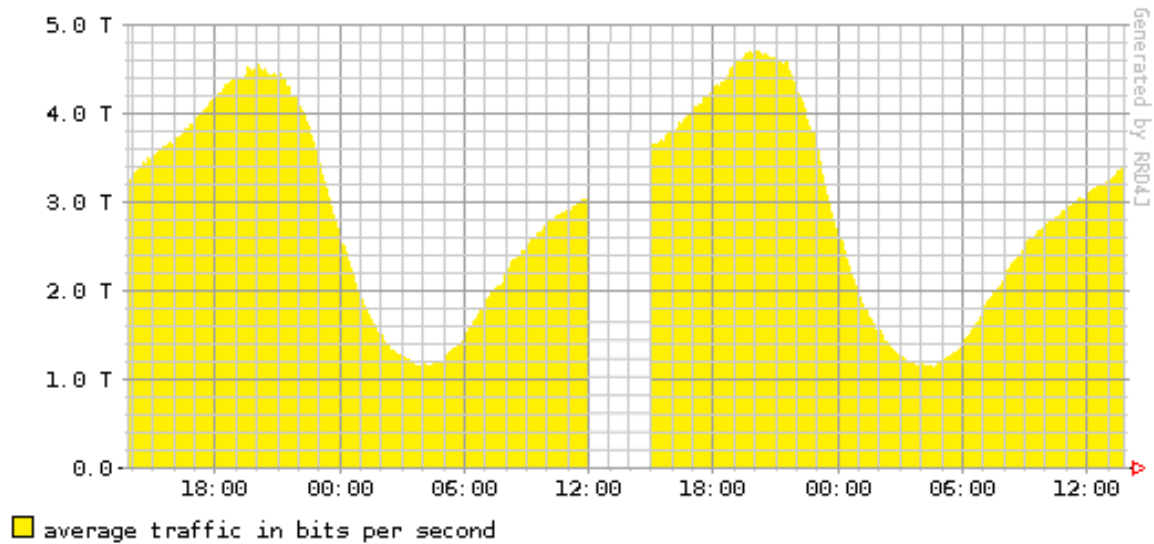


Incident AMS-IX Amsterdam




- » 13th May 2015 at 12:22 pm
- » Loop with 4 x 100GE created. Traffic was blackholed.
- » About 500 of 600 BGP sessions at the route servers dropped

What would be the impact to DE-CIX?



Impact at DE-CIX Frankfurt

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



Time Flow


AMS-IX

13th May 2015:
(information from public sources)

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

DE-CIX

- 1.
- 2.
- 3.
- 4.



What could be the reason to this behavior?


What could be a dependencies?

We found three answers... so far...

1. Remote Peering Routers Overloaded


- » A single remote peering router is connected to more than one IXP
- » The remote peering router is overloaded with broadcast traffic coming from one IXP
- » 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

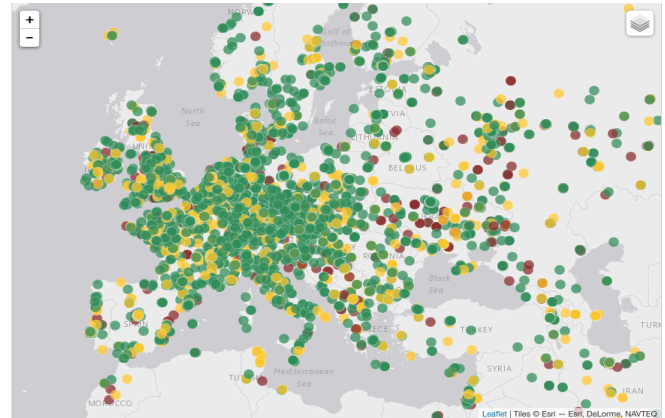
- » Are there routing paths that contain different IXPs on the up- and downstream?
- » Example:
 - » Upstream (gray) contains IXP A
 - » Downstream (red) contains IXP B



Selecting the Right Tool

» DE-CIX selected RIPE Atlas because of:

1. Extensive coverage of probes
2. Built-in traceroute measurement
3. Easy to access REST-API
4. Easy to obtain measurement results



» **Make jAtlasX available as open source:**

» <https://github.com/de-cix/jAtlasX>

» **Apache 2.0 license**



2. Asymmetric Routing Paths II

Measurement study (RIPE Atlas):


- » Number of AS-to-AS paths with a traffic drop > 200Mbit/s at DE-CIX Frankfurt: 183
- » ASes connected to DE-CIX Frankfurt and AMS-IX Amsterdam: 323
 - » ASes hosting RIPE Atlas probes: 171

→ 50 AS-to-AS routing paths which fulfill all above requirements

Measurement results:


- » 38% of all AS-to-AS paths with at least one asymmetric IXP path
- » 8% of all AS-to-AS paths traversed no IXP at all

2. Asymmetric Routing Paths

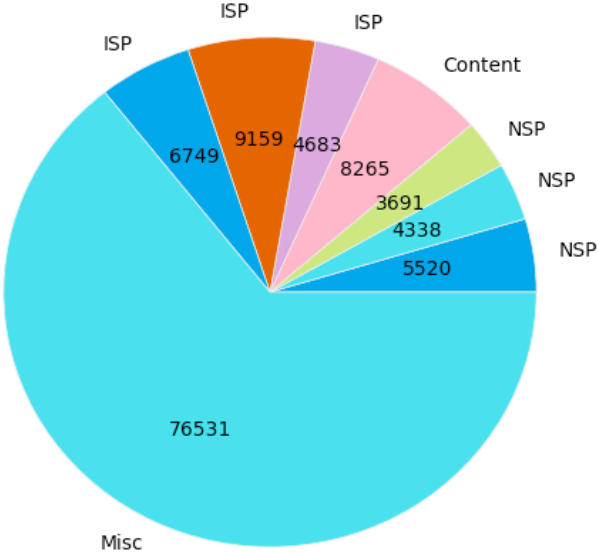


Impact Details

Source ASN with Traffic Loss > 5%



Destination ASN with Traffic Loss > 3%

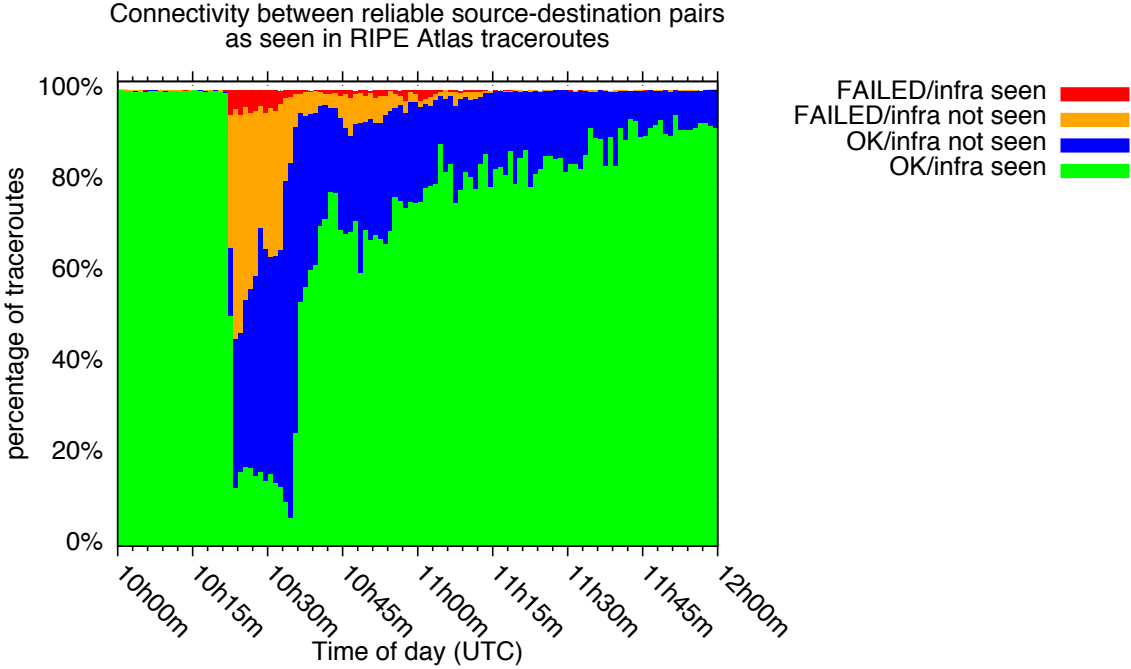


3. Layer 8: Less Users



- » 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 measure

Other Results



Source: <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: The Internet infrastructure is not hit 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



Awesome for
20
years

By joining DE-CIX, you become
part of a universe of networks.
Everywhere.

DE-CIX. Where networks meet.

A stylized world map with glowing nodes and connecting lines, symbolizing a global network.

**Where
networks
meet**

DE-CIX Management GmbH
Lindleystr. 12
60314 Frankfurt
Germany
Phone +49 69 1730 902 0

rnd@de-cix.net

www.de-cix.net