

Implications of Large IXP Failures


Christoph Dietzel, Georgios Smaragdakis, Vasileios Giotsas, Emile Aben

Introduction

- » Robustness of IXP interconnection system?
 - » Dependencies between IXPs?
 - » Effect/impact on the Internet?
 - » Lessons learned?
- » Investigated large IXP outage
- » Preliminary research




Incident AMS-IX Amsterdam




- » 13th May 2015 at 12:22 pm
- » Loop with 4 x 100GE in a core switch, most traffic dropped
- » About 500 of 600 BGP sessions at the route servers dropped

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



Time Flow

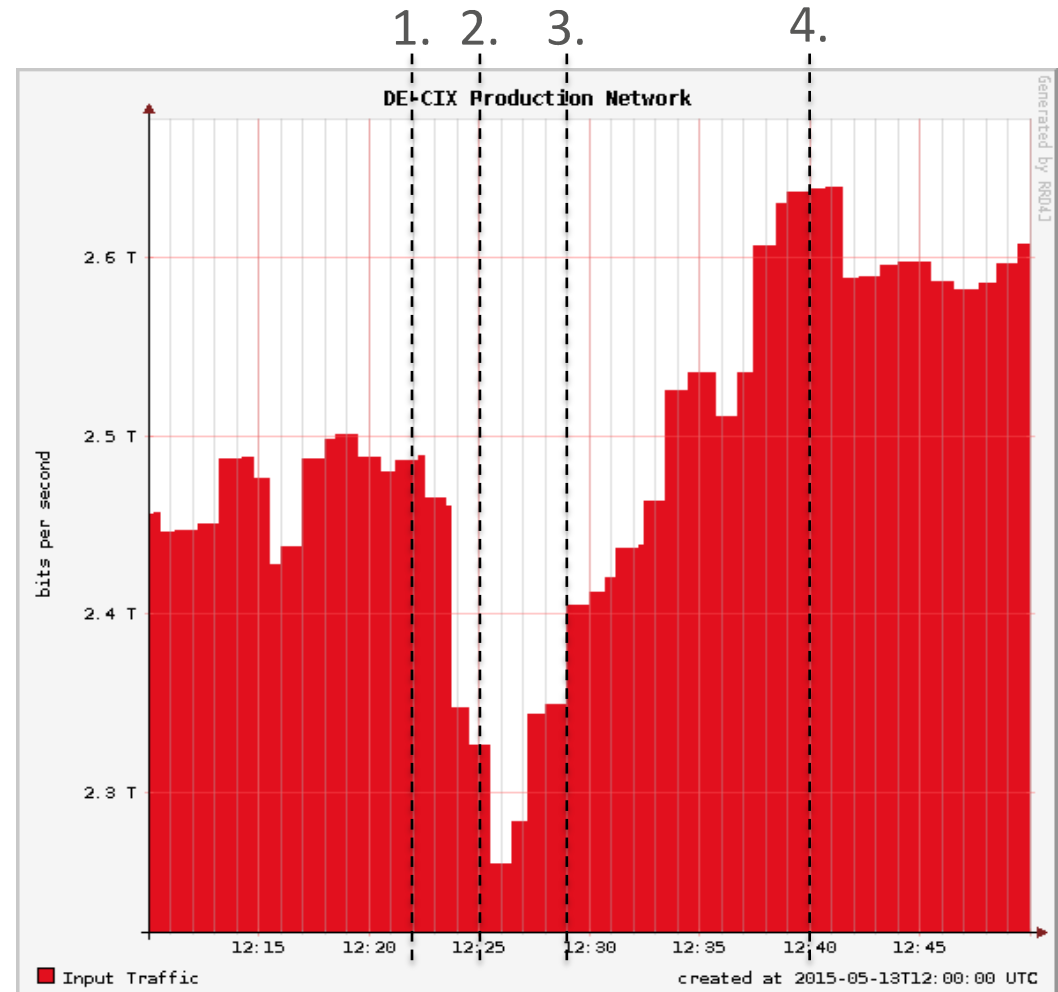
AMS-IX

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



Impact Details



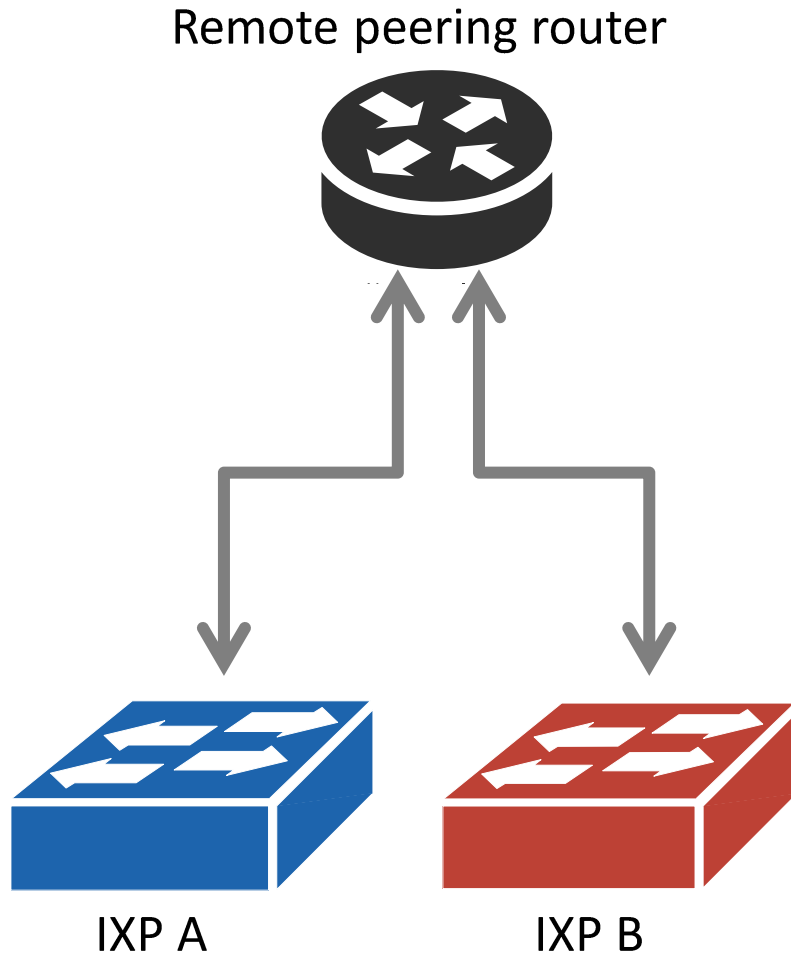
ASes with drops of more than 200 Mbit/s

What could be the reason for this behavior?

What could be the traffic dependency?

We found two 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 AMS-IX Amsterdam: *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”



Active Measurement



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