

Traffic Volume Dependencies between IXPs

Thomas King
R&D, DE-CIX

Important

This talk is about understanding
and learning.


It is not about blaming.

Introduction

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



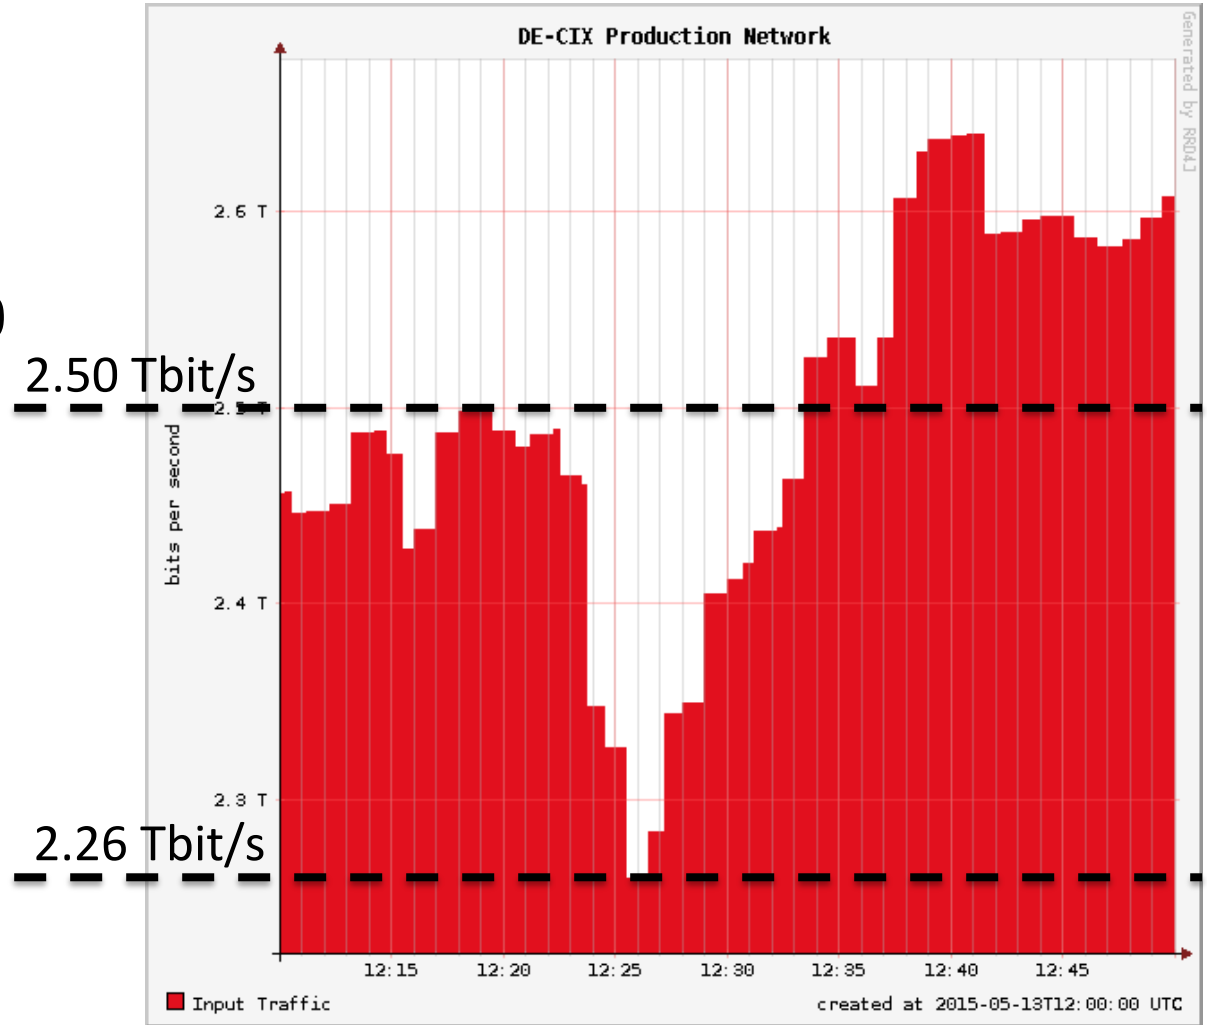
Incident AMS-IX Amsterdam



Impact at DE-CIX Frankfurt

Decreased traffic volume:

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



Time Flow

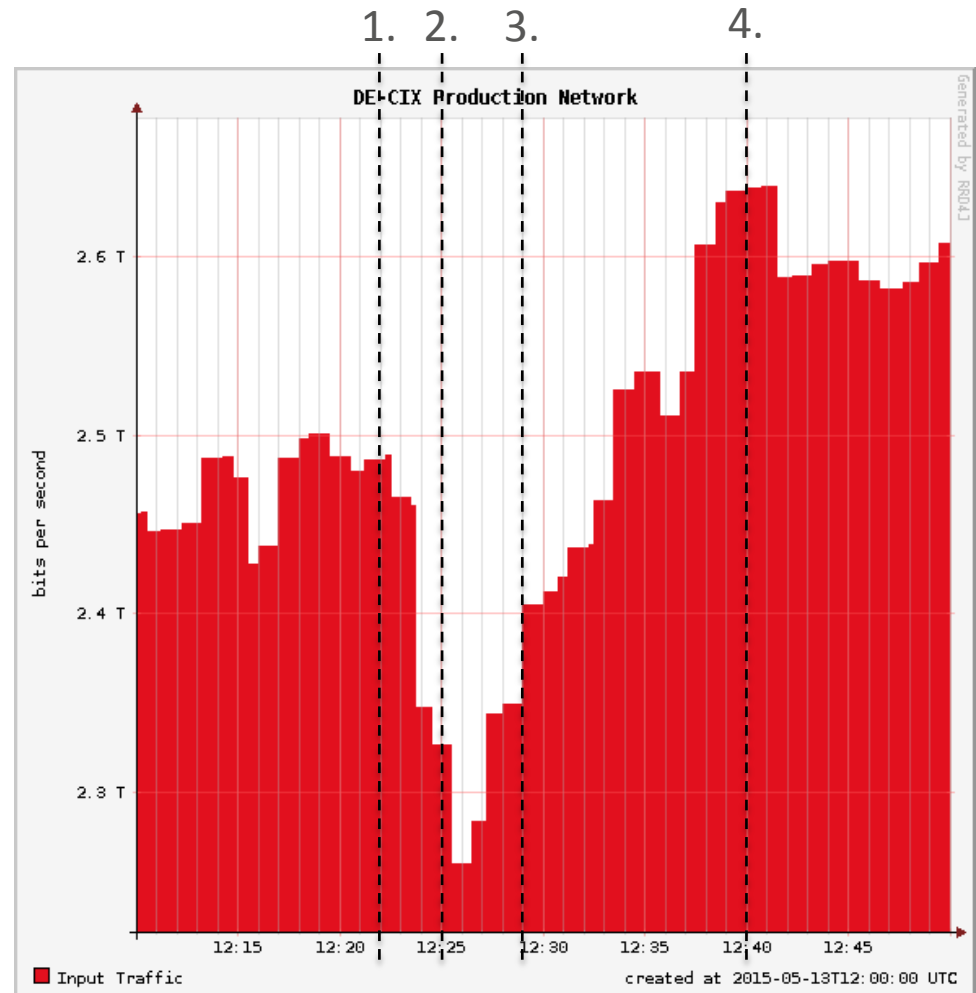
AMS-IX:

13th May 2015:

(information from public sources)

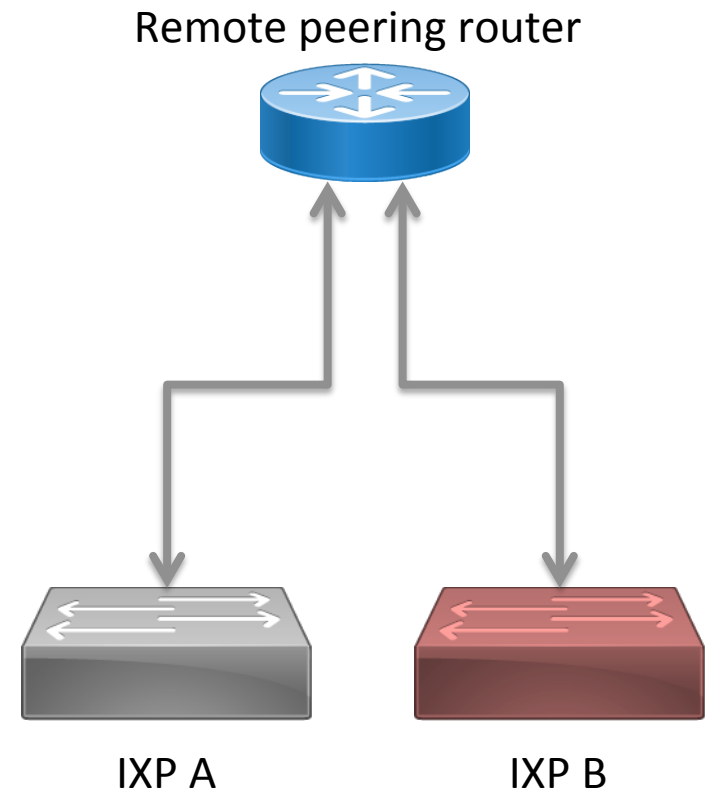
1. 12:22 pm – Loop with 4 x 100GE created. Traffic was blackholed.
2. 12:25 pm – About 500 of 600 BGP sessions at the route servers dropped
3. 12:29 pm – NOC reacted and deactivated ports responsible for loop
4. 12:40 pm – BGP sessions to route server are back online

DE-CIX:




1. 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



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 which are 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

Traceroute measurement results:

- 38% of all AS-to-AS paths were asymmetric
- 8% of all AS-to-AS paths traversed no IXP at all

3. Layer 9: 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

Summary and Outlook

Reasons for traffic volume dependencies between IXPs:

1. Remote peering routers overloaded
2. Asymmetric routing paths
3. Layer 9: Less users

Outlook:

- Research asymmetric routing paths between large IXPs
- Come up with recommendations in order to reduce the impact of traffic volume dependencies

Thank you!

Questions? Comments?

Please vote for me: RIPE PC