

Express or Local Lanes:

On Assessing QoE over Private vs. Public Peering Links

Walter Willinger, NIKSUN Inc.

Anja Feldmann, Philipp Richter, TU Berlin

Georgios Smaragdakis, MIT/TU Berlin

Fabian Bustamante, Northwestern University

wwillinger@niksun.com

The QoE Puzzle

- ▶ **“Horizontal” view of the QoE puzzle** (spatial dimension)
 - ▶ Last-mile (e.g., home network)
 - ▶ Interconnections (e.g., crossing CP/eyeball boundary)
 - ▶ Server-side (e.g., CDN/CP infrastructure)

- ▶ **“Vertical” view of the QoE puzzle** (across TCP/IP stack)
 - ▶ Services involved (e.g., DNS)
 - ▶ Application-specific features (e.g., VOIP, video, gaming)
 - ▶ Cross-layer aspects (e.g., adaptive video streaming over HTTP over TCP)



The QoE Puzzle ...

- ▶ Pieces of the QoE puzzle (horizontal view)
 - ▶ Last-mile (e.g., home network)
 - ▶ **Interconnections (e.g., crossing CP/eyeball boundary)**
 - ▶ Server-side (e.g., CDN/CP infrastructure)

- ▶ Pieces of the QoE puzzle (vertical view)
 - ▶ Services involved (e.g., DNS)
 - ▶ Application-specific features (e.g., VOIP, video, gaming)
 - ▶ Cross-layer aspects (e.g., adaptive video streaming over HTTP over TCP)



A General Perception ...

- ▶ The distance (AS-hop count, geo-distance) between where content resides and where that content is consumed is shrinking ...
- ▶ Supporting evidence
 - ▶ “The flattening of the Internet” (PAM’08, Conext’10, Sigcomm’10)
 - ▶ “... a large IXP as an Internet vantage point” (IMC’13)
 - ▶ “Are we one hop away from a better Internet?” (IMC’15)



Interconnections – Traditional View

▶ **Private peering**

- ▶ Cross-connects from colo/data center provider (e.g., Equinix)
- ▶ Dedicated link, predicable performance, expensive(?), ...

▶ **Public peering**

- ▶ At IXPs (e.g., DE-CIX Frankfurt, ANY2 LA)
- ▶ Shared infrastructure (switch, port), unpredictable(?) performance, cheap(?), lack of control (e.g., ingress traffic), ...

- ▶ *“If you think that public peering is a good idea, you are just not large enough”* (attributed by W. Norton to an operator of a US Tier-1 ISP)
-



Interconnections – “Real World”

- ▶ **Private peering (PI) offerings at many IXPs**
 - ▶ Uses separate infrastructure from public switching fabric
 - ▶ Very few data points
 - ▶ LINX (London): a few 100 PIs vs a few 10K public peerings (~2 Tbps)
 - ▶ DE-CIX (Frankfurt): some 900+ PI's vs some 100K public peerings
- ▶ **Public peering supported at commercial colo facilities**
 - ▶ Different data centers house an IXP's colocations
 - ▶ “Deployed” networks vs “available” networks
- ▶ **A data center with (access to) an IXP is more valuable than one without**

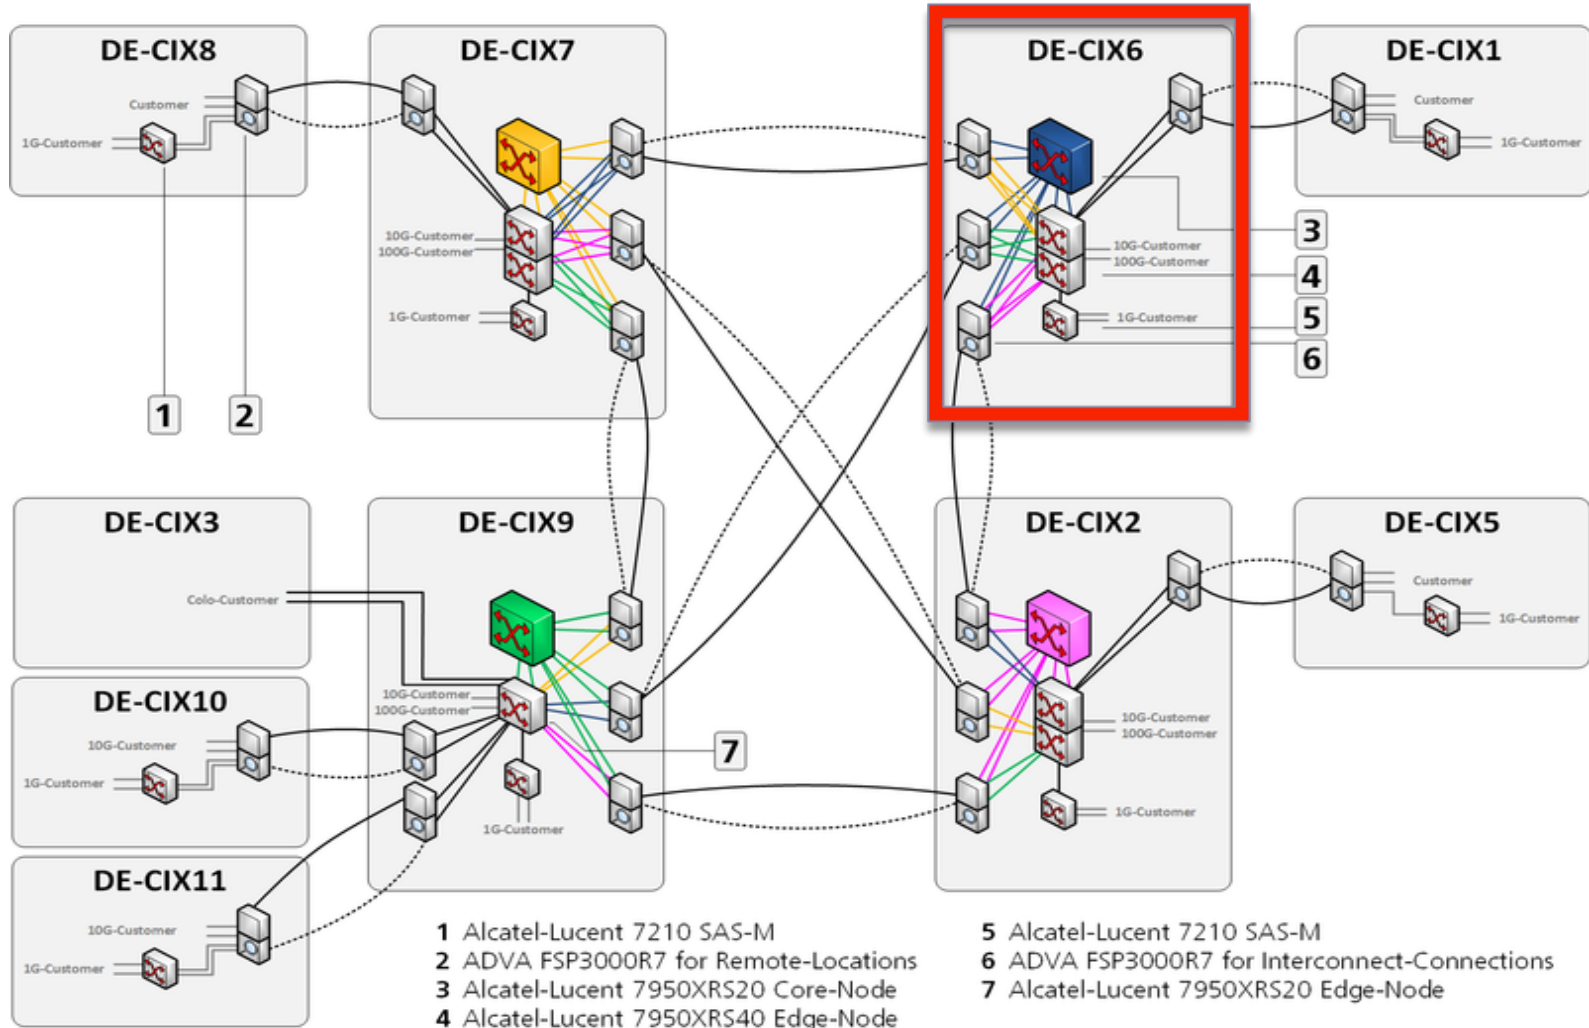


Change Changes Things

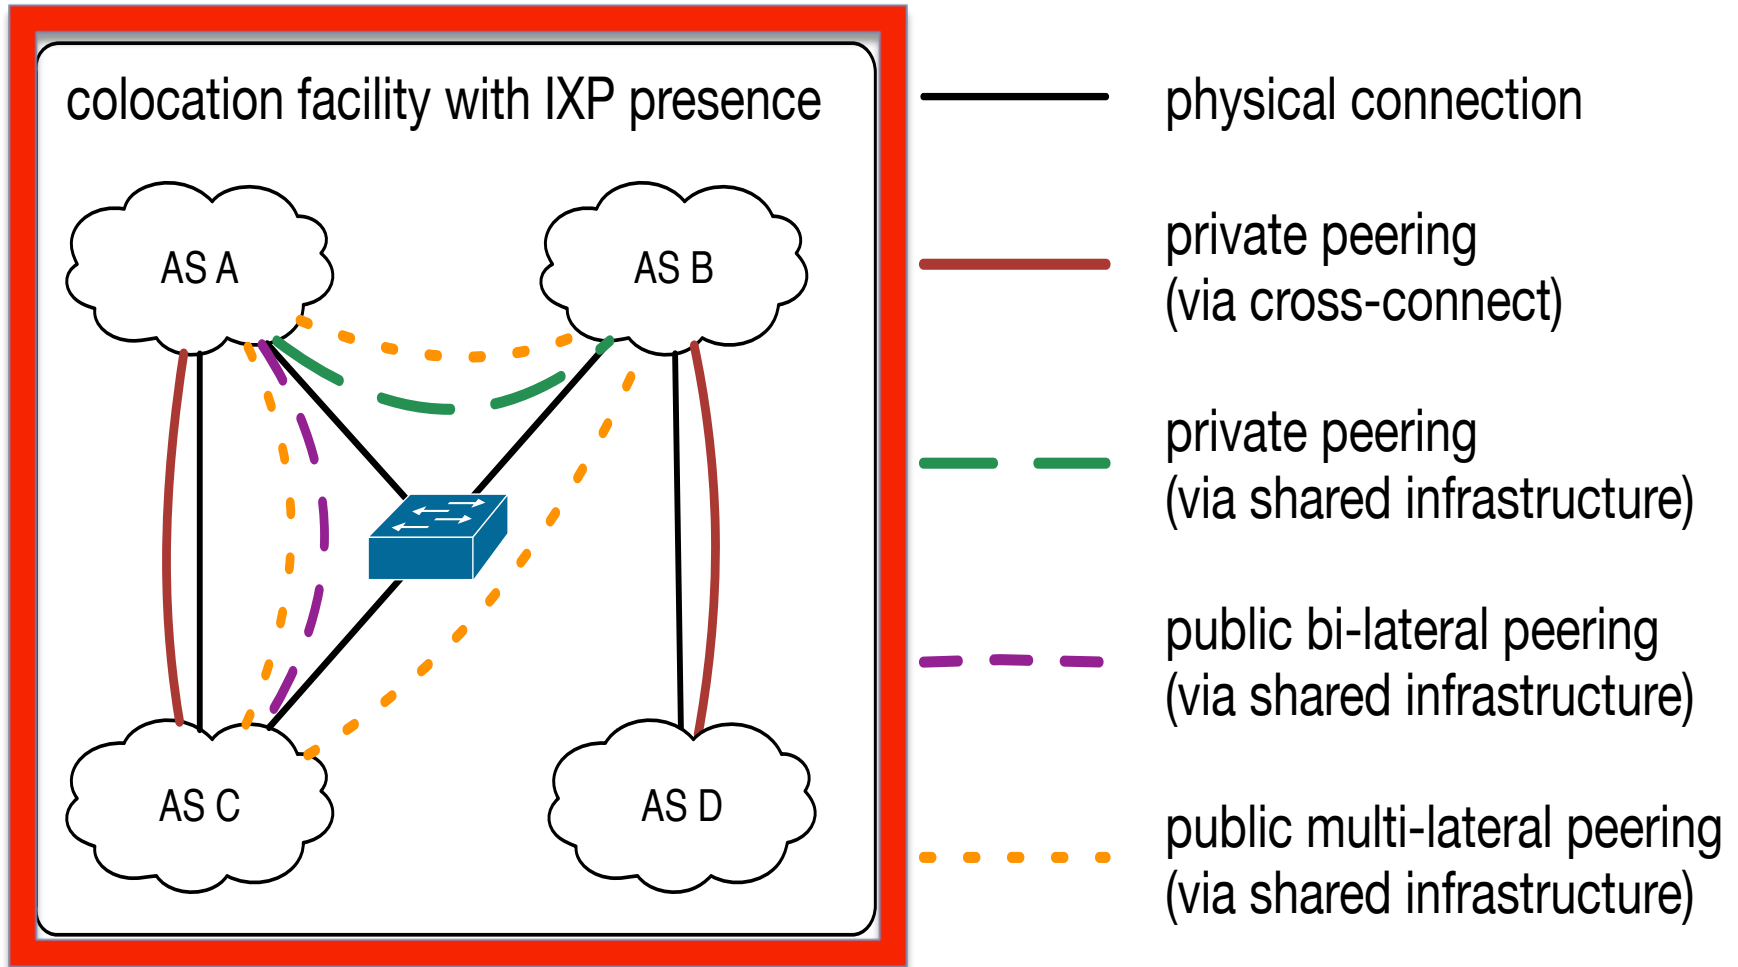
- ▶ **Free use of IXP route server: Bifurcation of public peering**
 - ▶ Multi-lateral peering (route server)
 - ▶ Bi-lateral peering (individual BGP session)
- ▶ **New IXP switching fabric: Bifurcation of private peering**
 - ▶ Virtual private peering (new IXP service offering, uses public switching fabric of IXP)
 - ▶ Physical private peering (IXP getting out of cross-connect business – leave it to data center)
- ▶ **More interconnection alternatives and options throughout the Internet peering ecosystem (e.g., compare NYC market in 2013 vs 2015; Open-IX efforts)**



Interconnections at a Modern IXP



Interconnections – “New World”



Note: The different ASes may have additional upstream connectivity

Interconnection-related QoS

- ▶ More opportunities to interconnect
- ▶ More choices for how to interconnect

- ▶ More active QoS monitoring (as part of IXP-offered services)
 - ▶ (Virtual) PIs
 - ▶ Packet loss of 0.05% on a daily average
 - ▶ One-Way Delay of 0.5ms for up to 97.5% of packets
 - ▶ Jitter of ± 0.1 ms for 97.5% of packets
 - ▶ SLAs
 - ▶ (Physical) PIs at colo/data center facilities
 - ▶ No equivalent service monitoring for cross-connects is provided



On Interconnection-related QoE

- ▶ **Step 1: Connectivity**

- ▶ Infer established peerings between relevant pairs of networks

- ▶ **Step 2: Usage**

- ▶ Infer how these networks use the different peerings

- ▶ **Step 3: Performance**

- ▶ Infer actual performance over the different inferred peerings

- ▶ **Step 4: QoE**

- ▶ Is performance (Step 3) “good enough” for usage (Step 2)?

- ▶ **Determine if (and how) interconnection-related QoE aspects matter for determining end-to-end QoE**

