DE·CIX

Where networks meet

IX networking challenges

Tuesday, June 7th, 2011
ENOG1, Moscow

Arnold Nipper
CTO/COO and Founder
arnold.nipper@de-cix.net

Where networks meet DE·CIX

Growth rate 300% per year

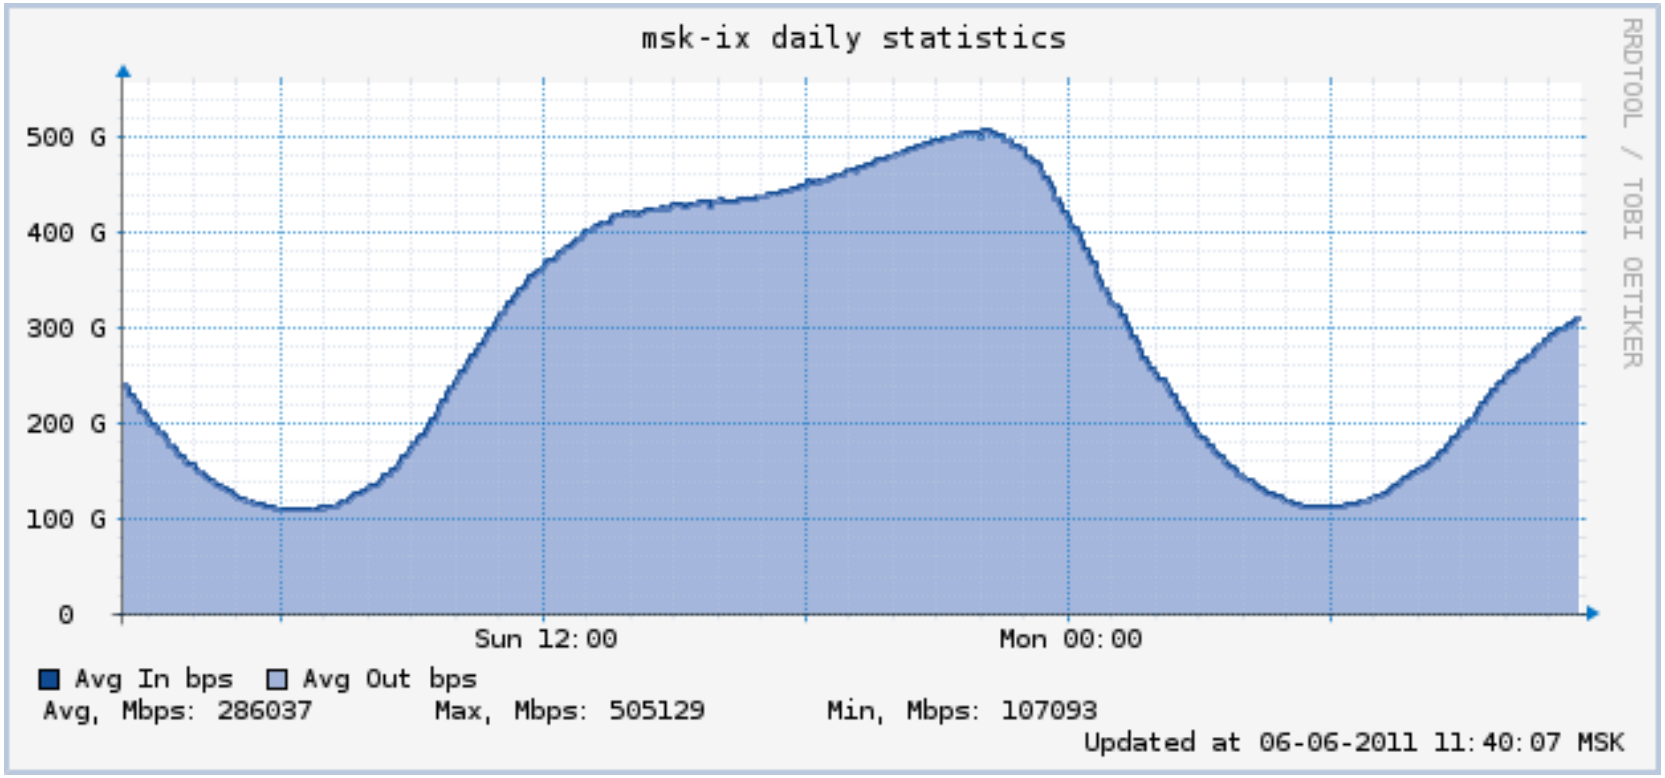

Growth rate 100% per year, new architecture

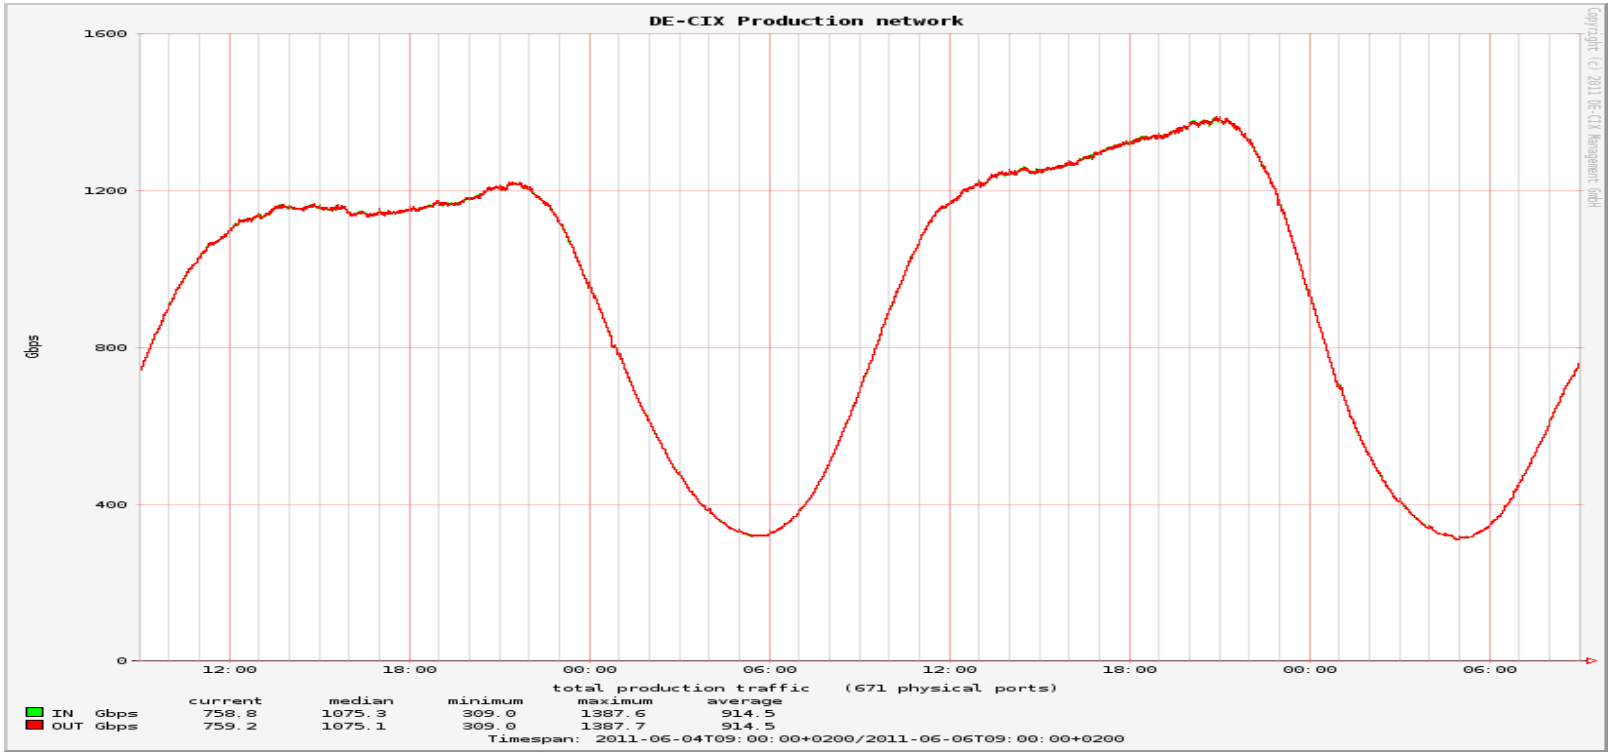

Introducing new technology

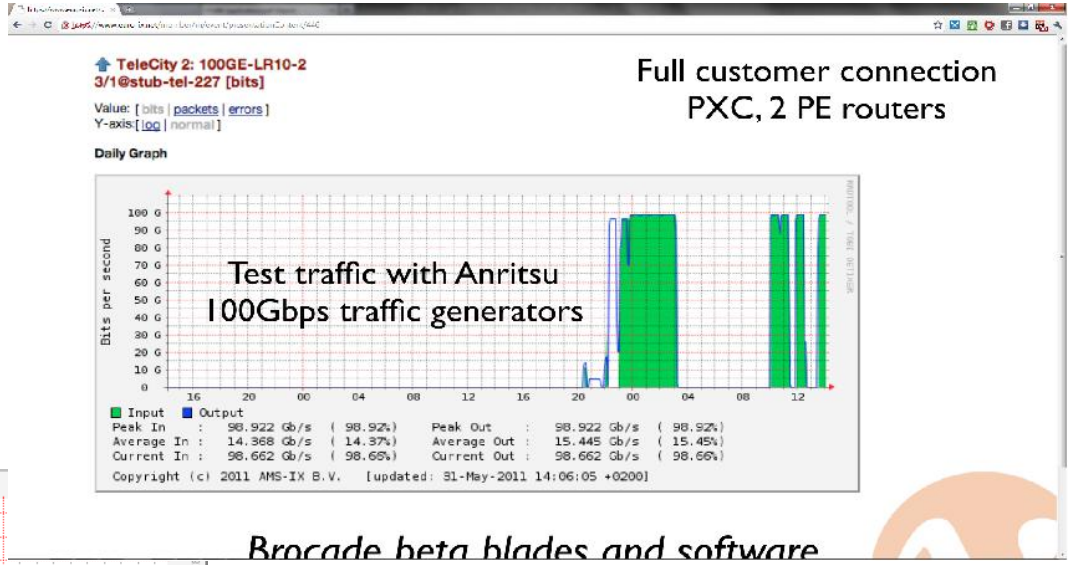

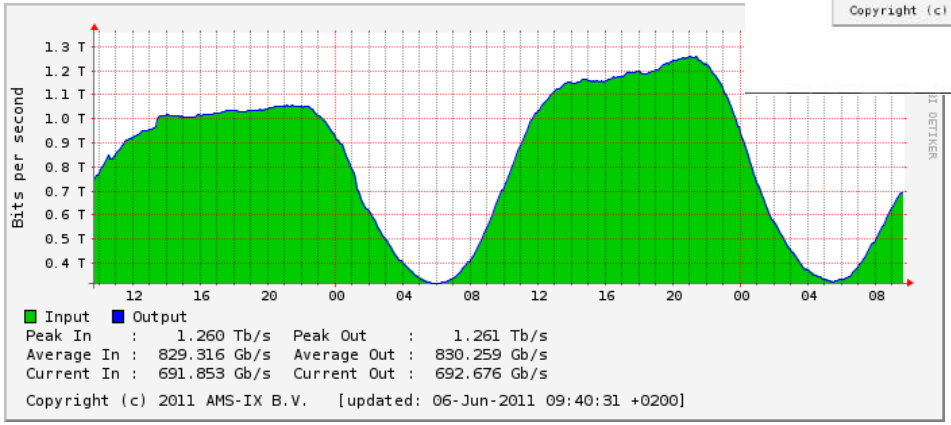

New architecture

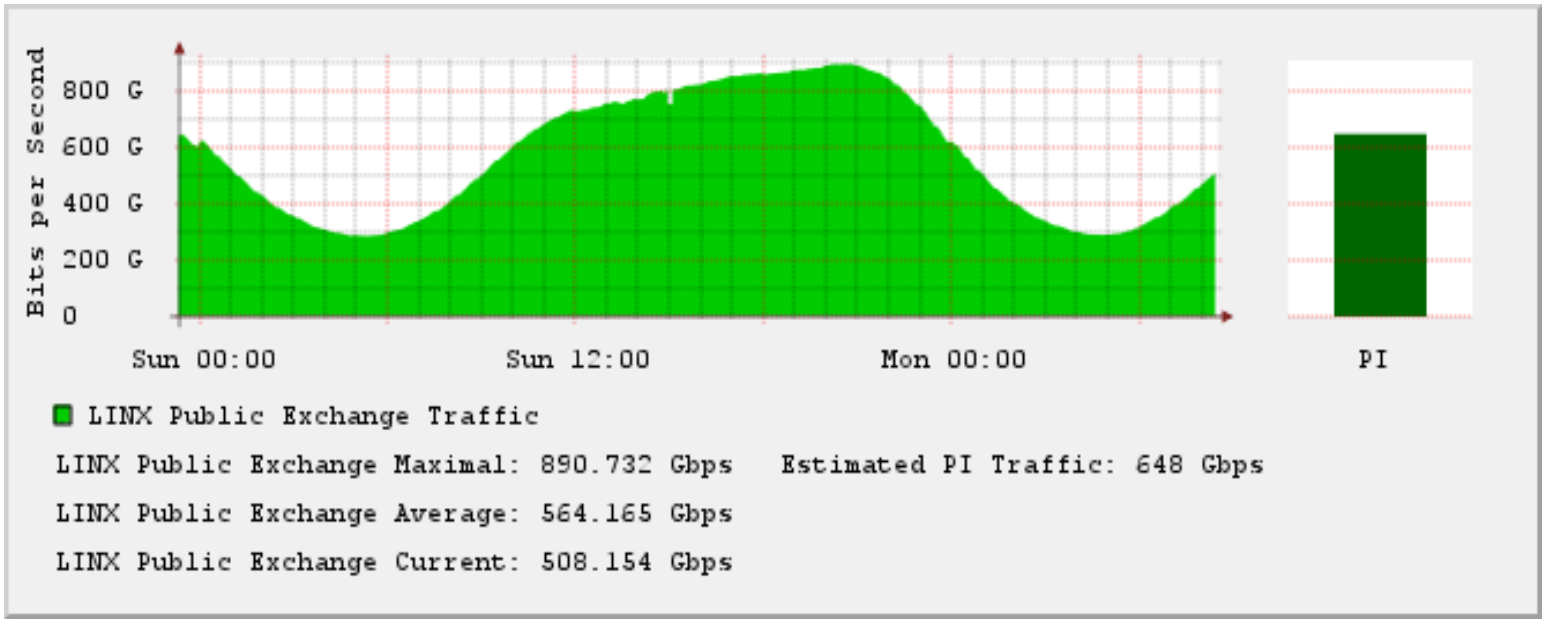

Scaling the infrastructure

- The ideal IXP only has one big switch
- Reality is that IXP's have to serve a lot of locations
- The more sites the more backbone ports
- More backbone ports ➔ less customers ports ➔ more switches ➔ more backbone ports

# switches	local traffic
1	100%
2	50%
3	33%
4	25%
5	20%
6	17%
7	14%
8	13%
9	11%
10	10%

Scaling the infrastructure

- Simple Layer2 technologies (STP like) are too unefficient
- Waste of links
- Slow convergence
- Ring topology is simple and efficient
- However limited in scalability
  - What to do if a switch is full?
  - The more switches the higher the hop count

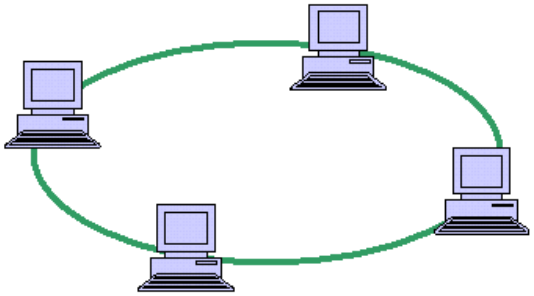

Scaling the infrastructure

- No need for loop avoidance protocol
- Scales very well
- Core is the most critical part ➔ dual core
  - Half of the core capacity unused
- Mechanism needed to switch between cores
  - Optical switches

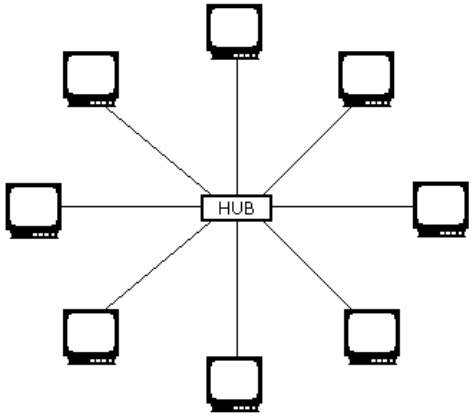

Scaling the infrastructure

- Clos architecture
- Change paradigm from active/standby to always active
- New technologies needed
  - VPLS (Virtual Private LAN Service)
    - Connecting LANs via MPLS
  - TRILL (Transparent Interconnect of Lots of Links)
    - Use sort of IS-IS for topology detection
    - TRILL devices are called Rbridges
    - Frames are encapsulated and sent to destination Rbridge
    - Allows load sharing
    - Backwards compatible

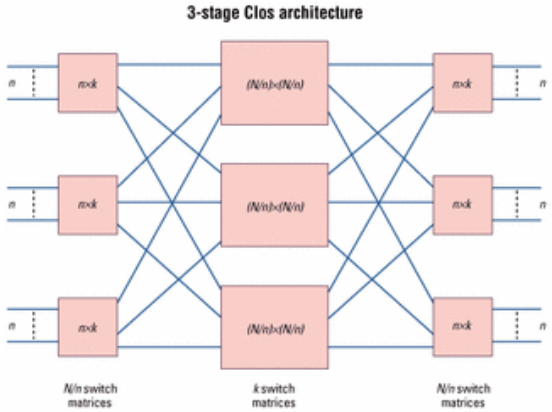

Introducing new technologies (100GbE)

- Standard (802.3ba) approved since almost one year
- products are available from some vendors
- prices are still very high
  - –7x the price of 10GbE would be expected/acceptable
- work on the transceivers is still going on
  - – 100GBASE-LR4 vs. 100GBASE-LR10
  - – long term vs. short term
  - – expensive vs. Cheap
- 100GbE at larger IXP will most probably need active optical DWDM equipment

Bringing customers up to speed aka routeservers

- per peer/customer routing table is state of the art (BIRD, Cisco, OpenBGPD, Quagga)
- giving customers full control is still work in progress
  - using communities breaks with AS32
  - inbound filters have to be on routeservers
- Optimal solution will be add path (Advertisement of Multiple Paths in BGP)
  - does need changes in client BGP stack as well

Financial challenges

- transit prices are dropping constantly
- costs for bigger exchanges grow slightly faster than revenue
  - more need for backbone ports and interconnects
- investment in next generation technology is expensive

Thank you

Join DE-CIX now!

DE-CIX Competence Center
Lindleystrasse 12
60314 Frankfurt/Germany

Phone +49 69 1730 902 - 0
info@de-cix.net

DE-CIX Competence Center @
Kontorhaus Building

Frankfurt Osthafen (Docklands)