100G Applications

Donn Lee
Network Engineering Team
Facebook Inc.
24 February 2010
Ethernet Summit

Top 10 Global Web Parent Companies, Home & Work

December 2009

RANK	PARENT	UNIQUE AUDIENCE (000)	ACTIVE REACH %	TIME PER PERSON (HH:MM:SS)
1	GOOGLE	353,851	83.91	2:38:50
2	MICROSOFT	315,490	74.81	3:01:38
3	YAHOO!	228,711	54.23	2:12:36
4	FACEBOOK	206,878	49.06	5:57:17
5	EBAY	163,844	38.85	1:41:31
6	WIKIMEDIA FOUNDATION	141,239	33.49	0:16:01
7	AMAZON	137,364	32.57	0:32:11
8	AOL LLC	129,360	30.67	2:21:03
9	NEWS CORP. ONLINE	120,316	28.53	0:59:17
10	INTERACTIVECORP	115,131	27.30	0:11:36

Source: Nielsen NetView

User growth is driving bandwidth, datacenters

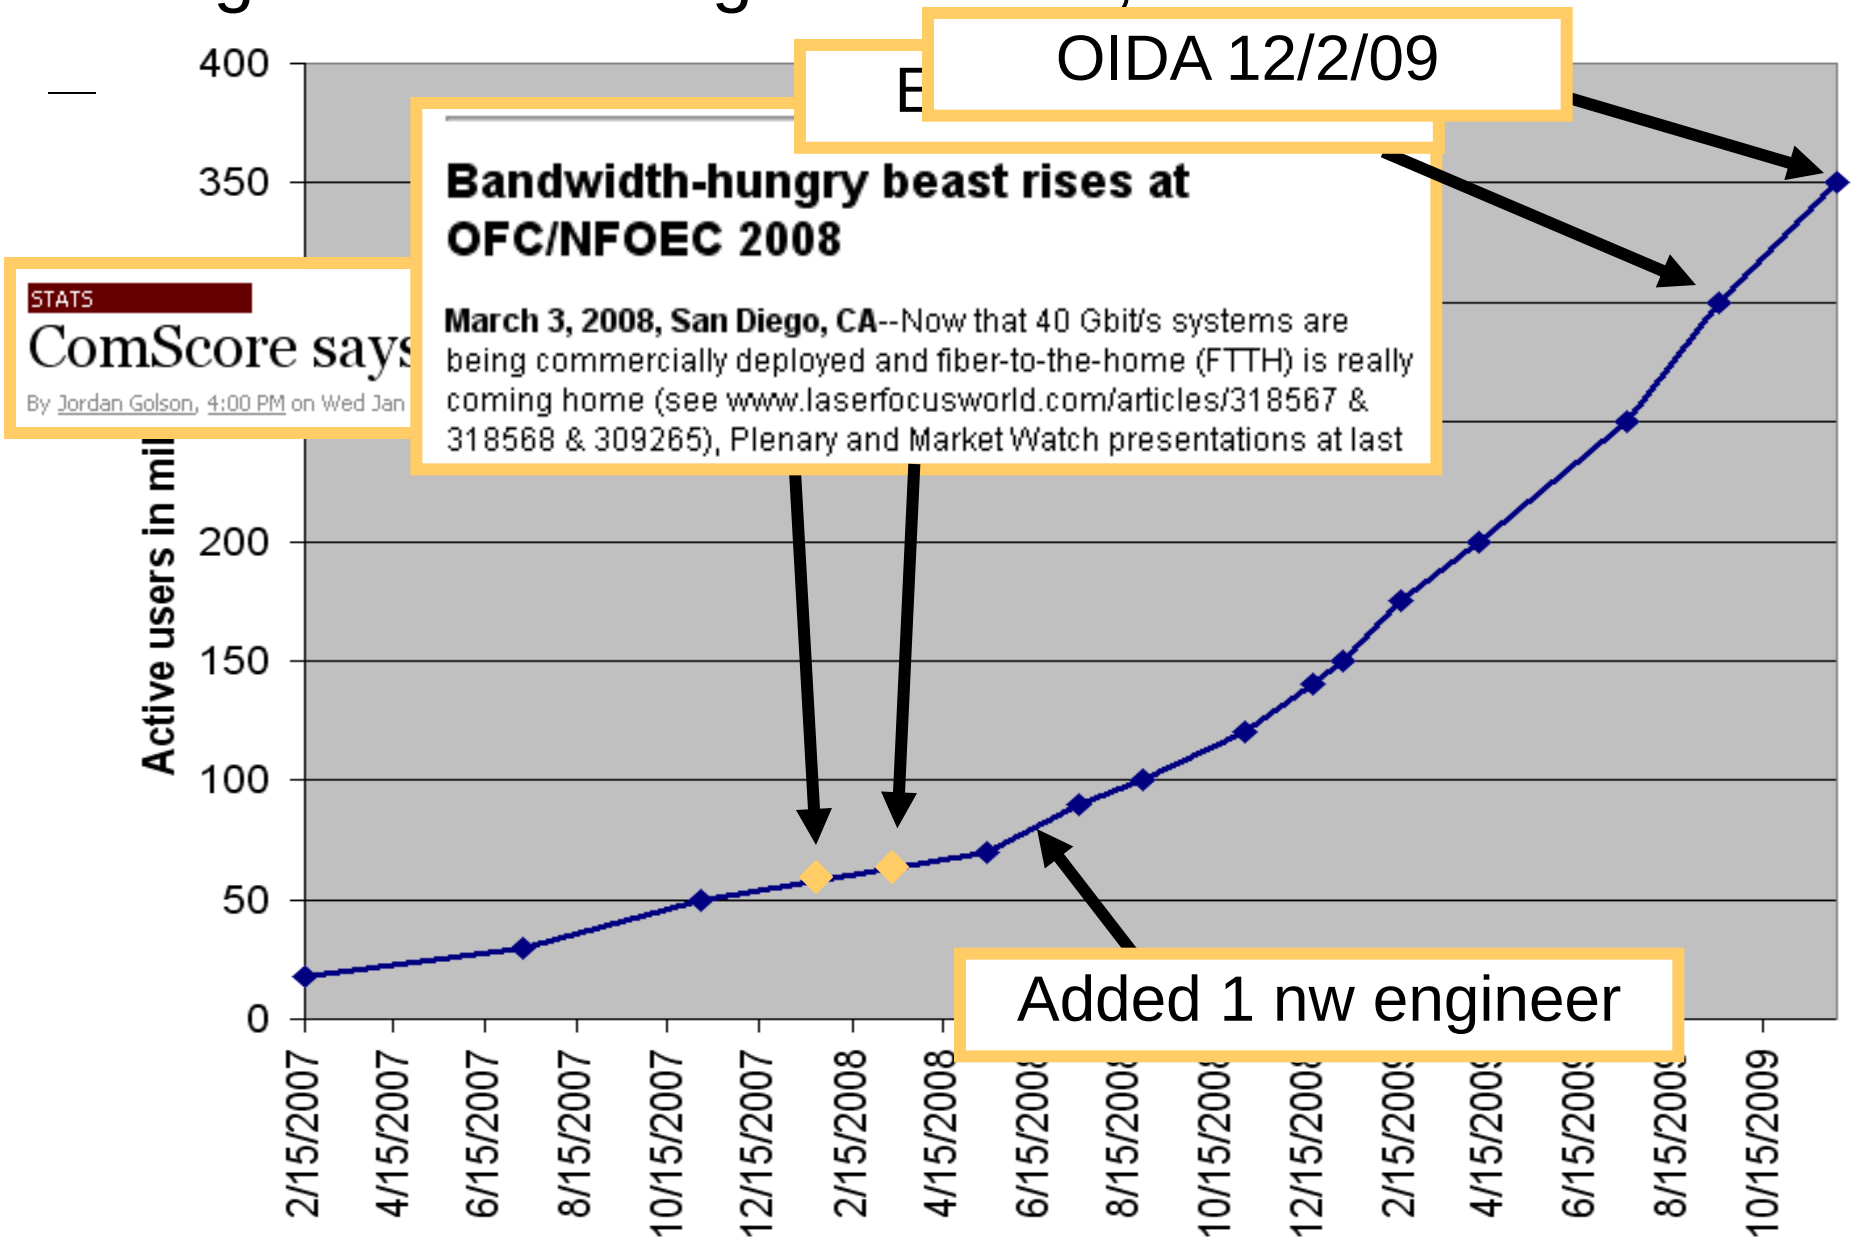

(not a Facebook data center)

2007

Next

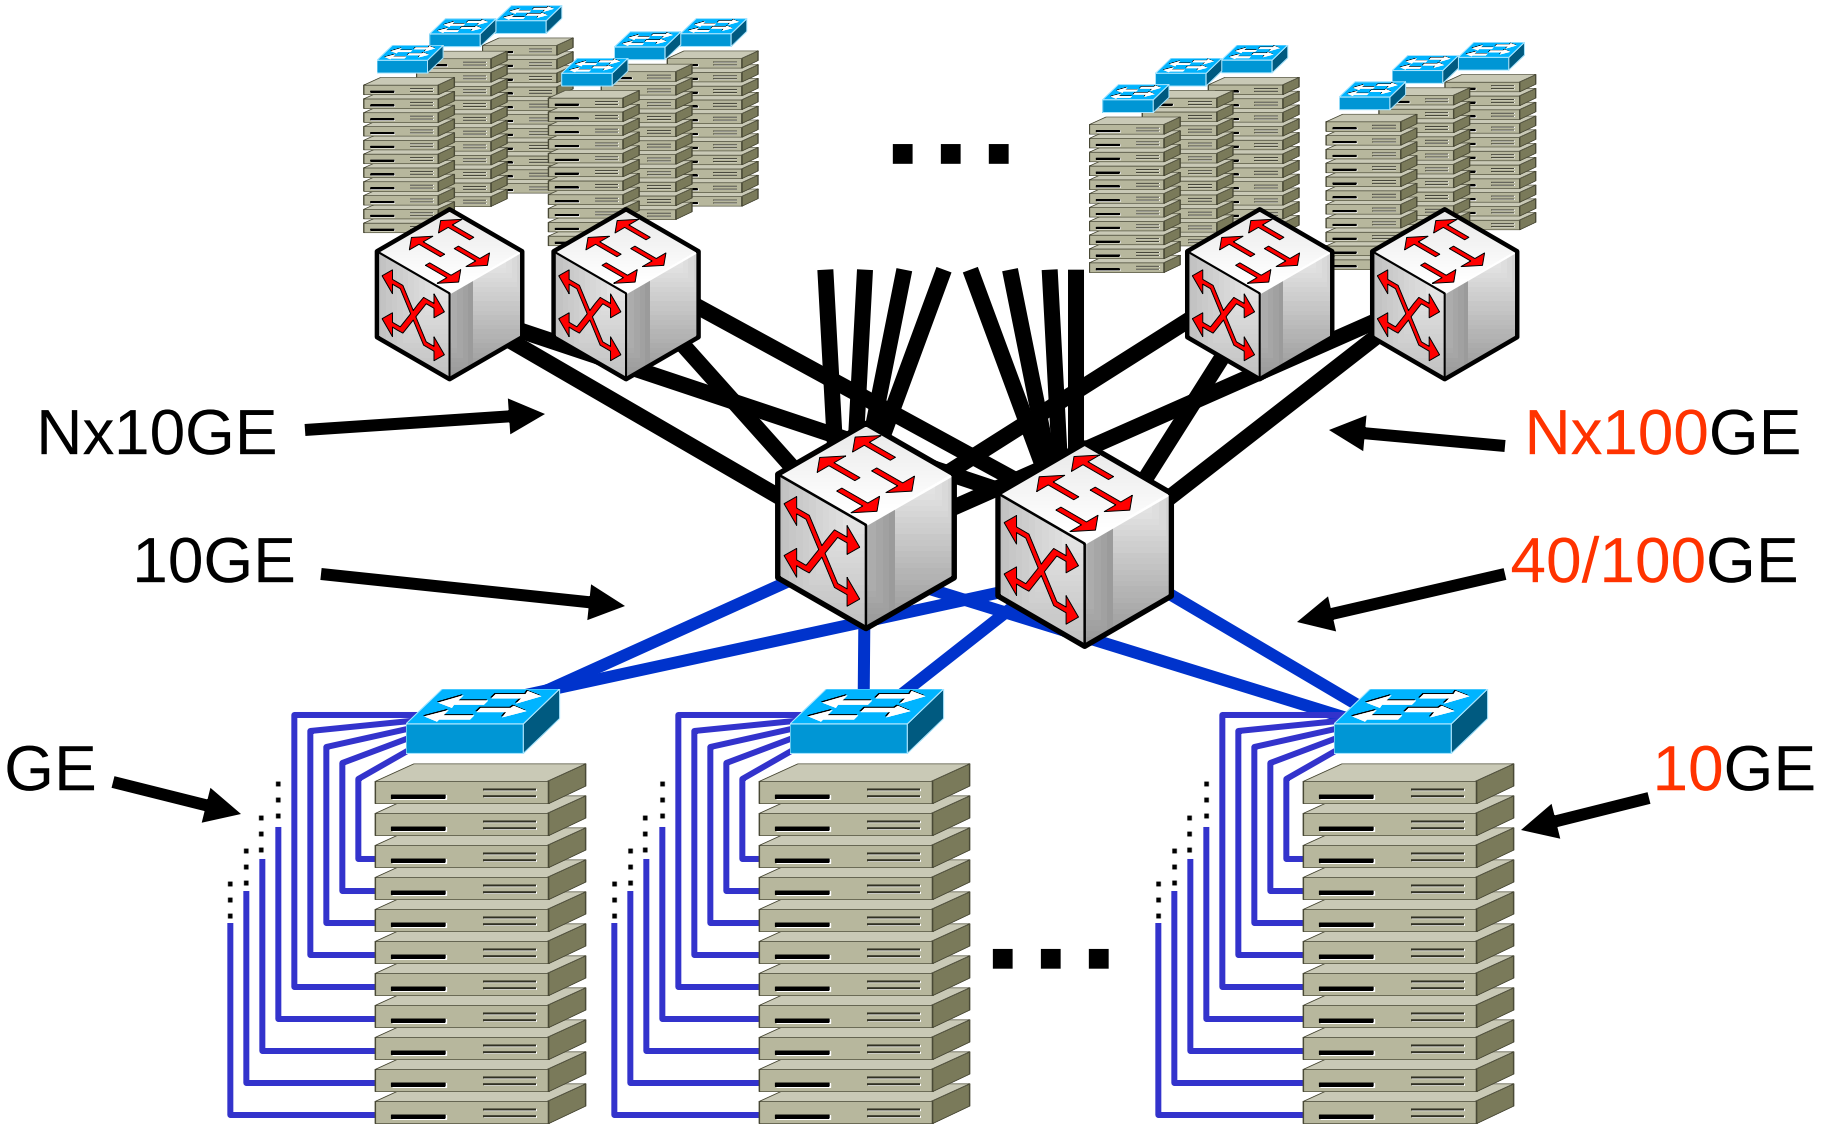

Toothbrush Graph

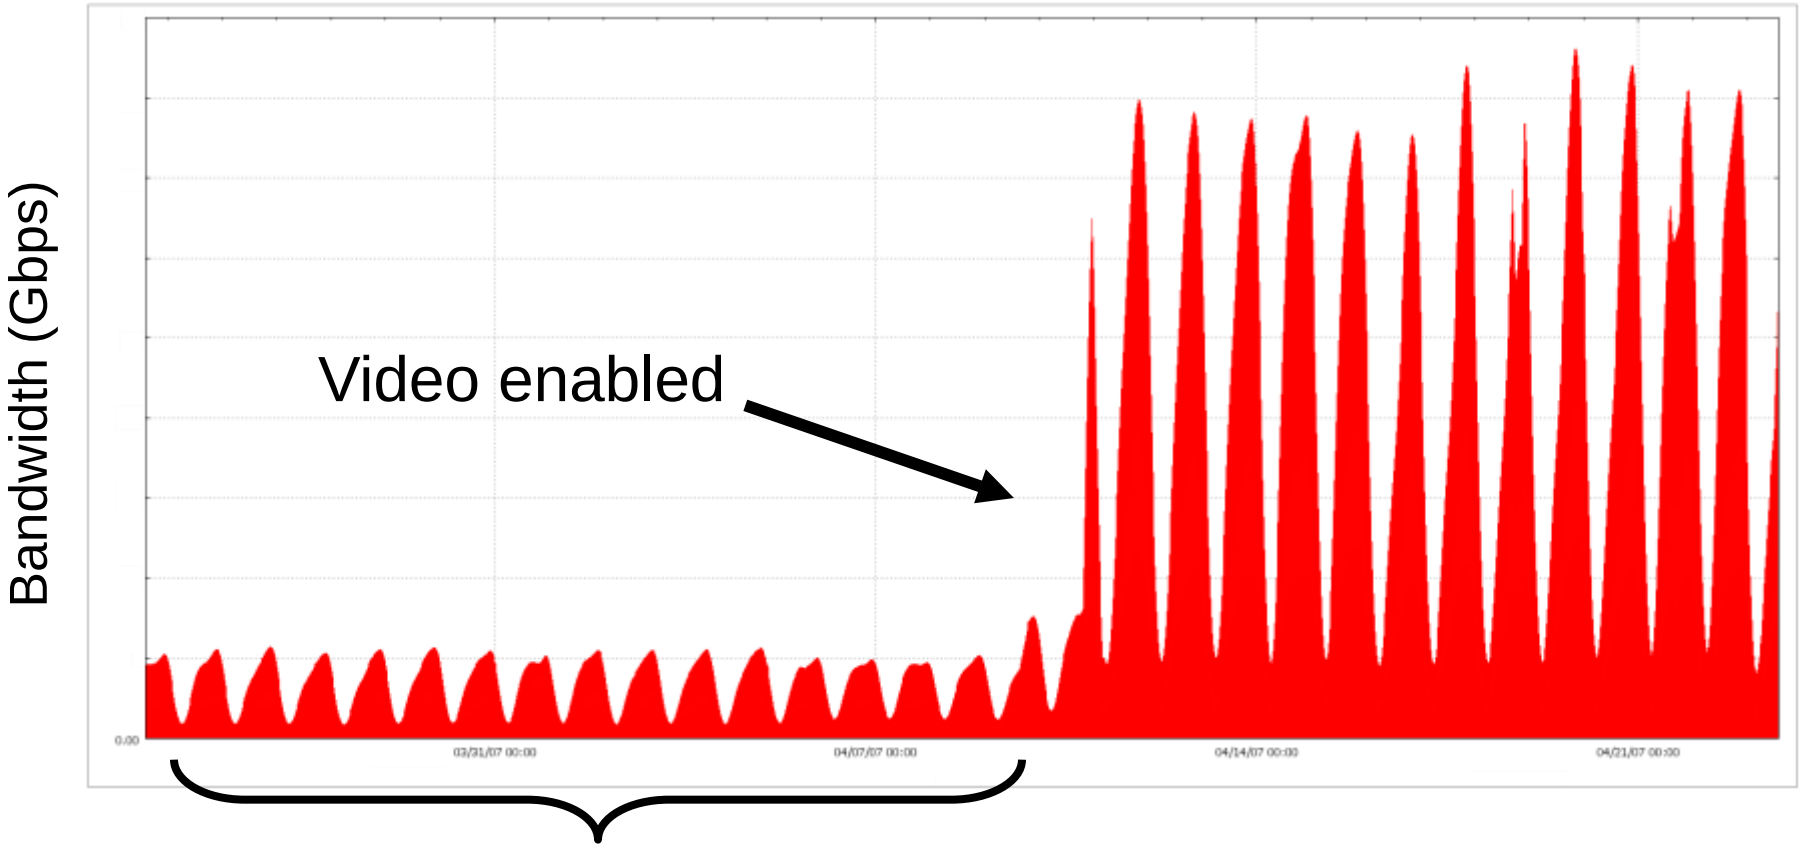

This *was* a huge amount of traffic
(multiple web services, millions of users)

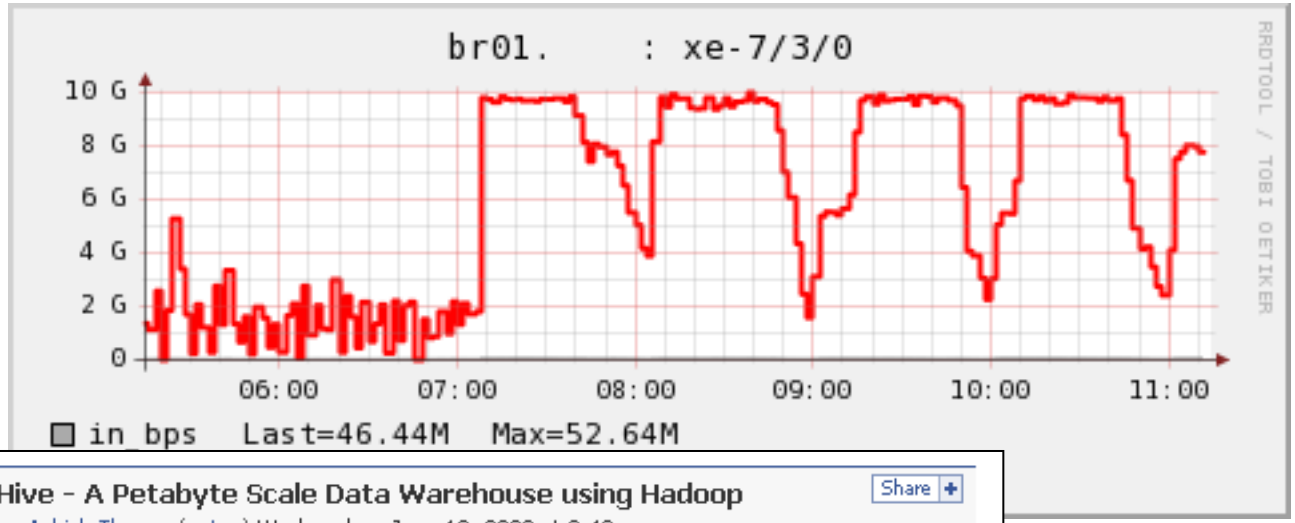

8/21/2009

Hive - A Petabyte Scale Data Warehouse using Hadoop

Share

by Ashish Thusoo (notes) Wednesday, June 10, 2009 at 8:49am

A number of engineers from Facebook are speaking at the Yahoo! Hadoop Summit today about the ways we are using Hadoop and Hive for analytics. Hive is an open source, peta-byte scale date warehousing framework based on Hadoop that was developed by the Data Infrastructure Team at Facebook. In this blogpost we'll talk more about Hive, how it has been used at Facebook and its unique architecture and capabilities.

Scalable analysis on large data

Facebook - both engineering ar

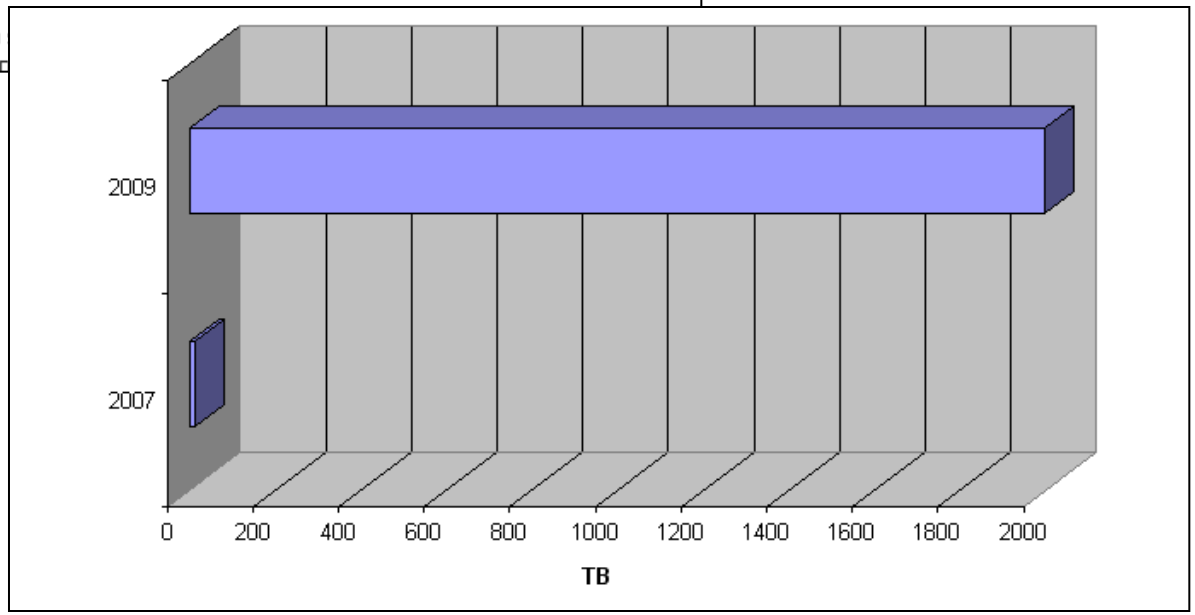

Non-blocking 10GE per Chassis vs. Need

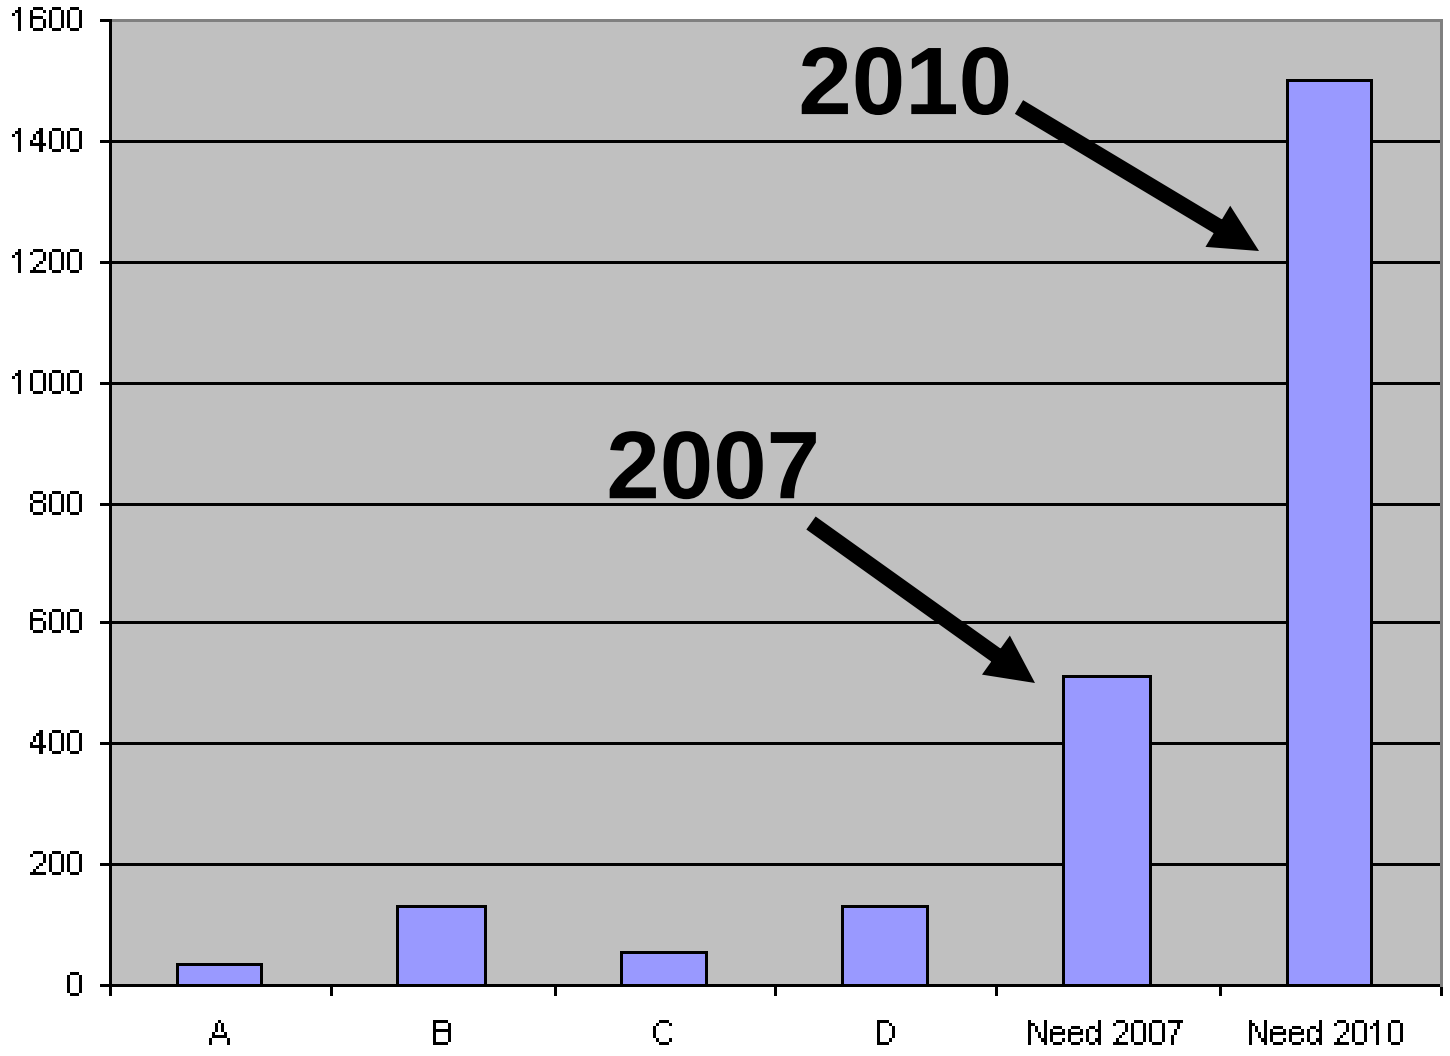

Case Study

54,000 servers

43,200 servers (assume 80% in main datacenter)

1,080 racks

40,000 Gbps bisectional bandwidth, datacenter

400 x 100GE links

(100 x 400GE links)

Case Study #2

36 MW per datacenter

Assume, 600 racks / 5 MW*

4,320 racks / datacenter

173,000 Gbps bisectional bandwidth, datacenter

1,730 x 100GE links

(432 x 400GE links)

Again, with only GE hosts

* Hypothetically, not the best, not the worst™

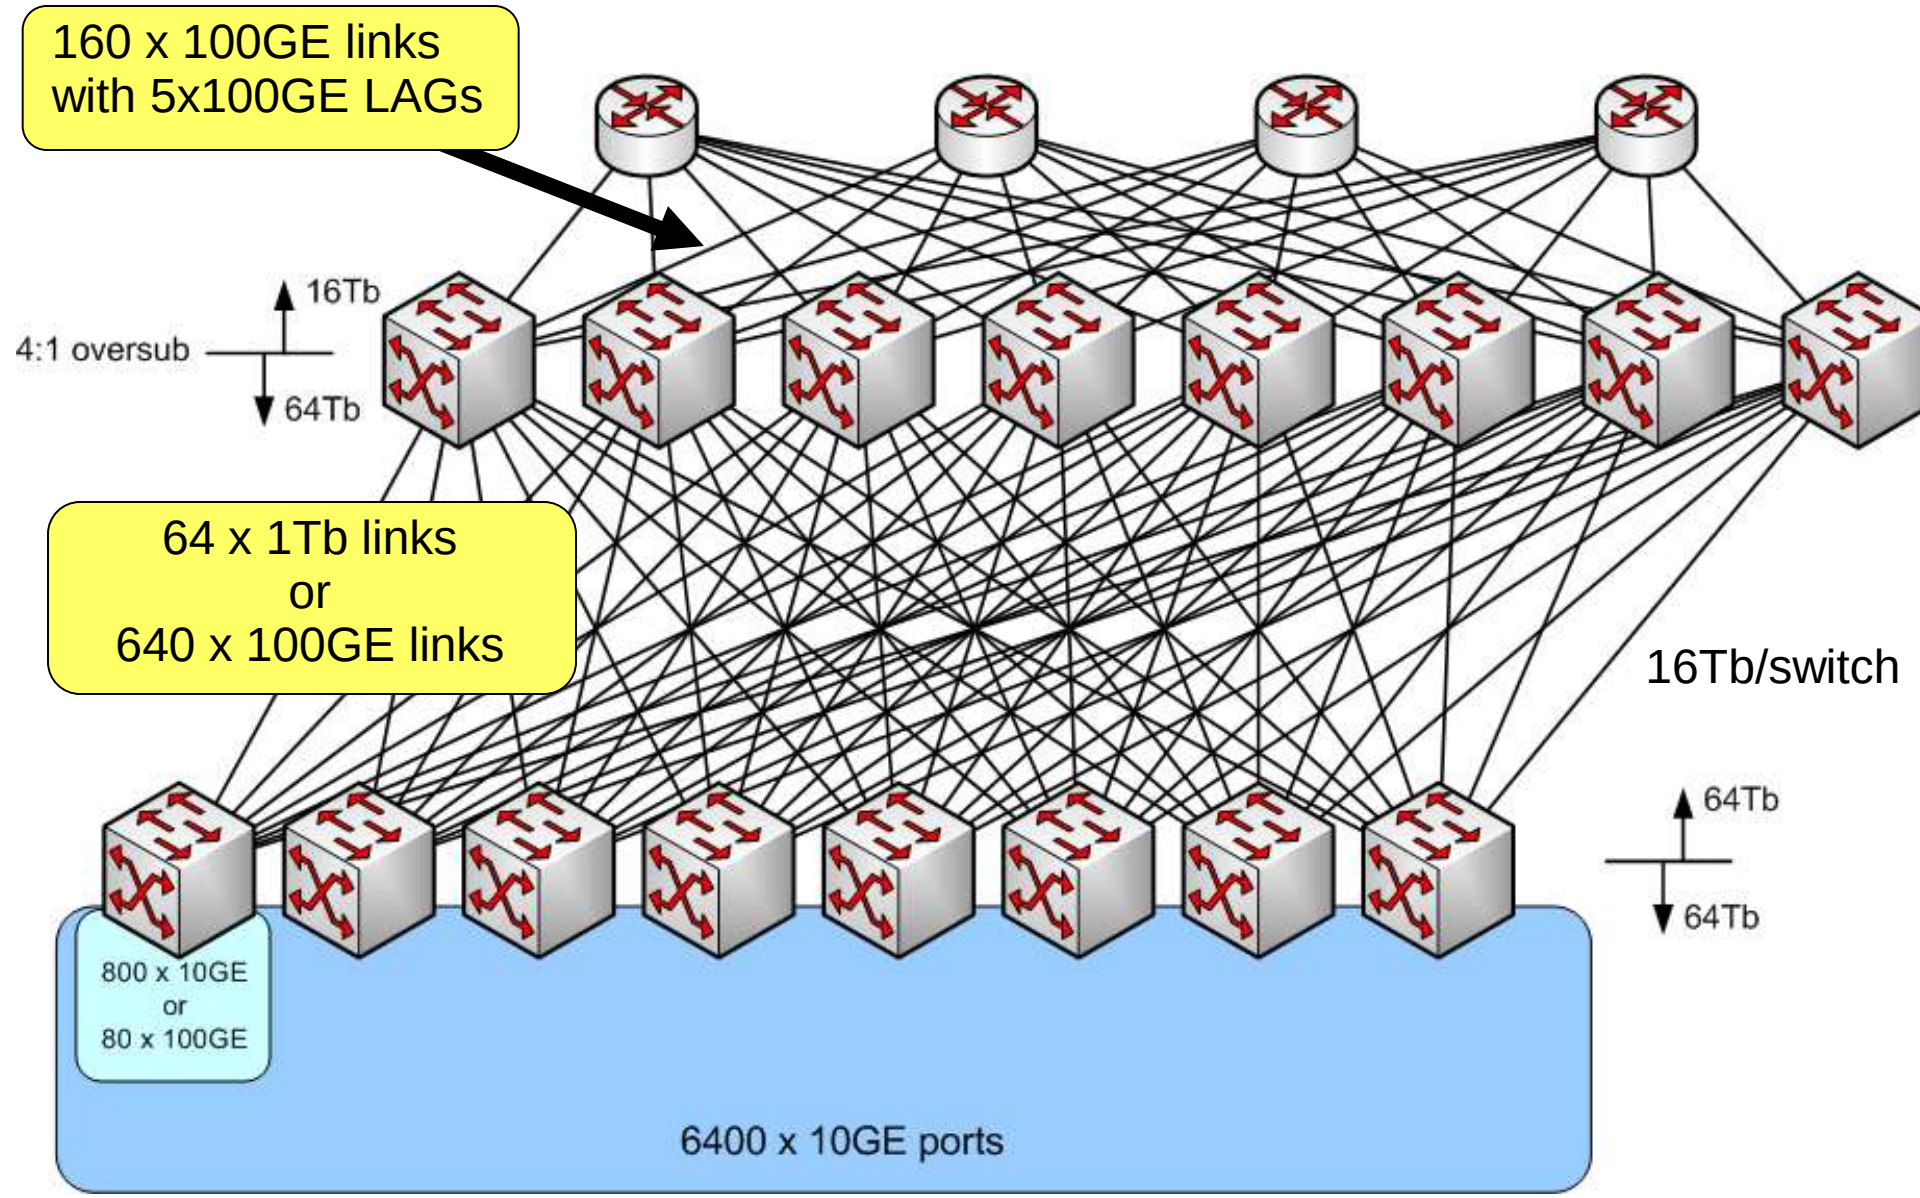
160 x 100GE links
with 5x100GE LAGs
16Tb
4:1 oversub
64Tb
64 x 1Tb links
or
640 x 100GE links
16Tb/switch
64Tb
64Tb
800 x 10GE
or
80 x 100GE
6400 x 10GE ports

With Current Hardware

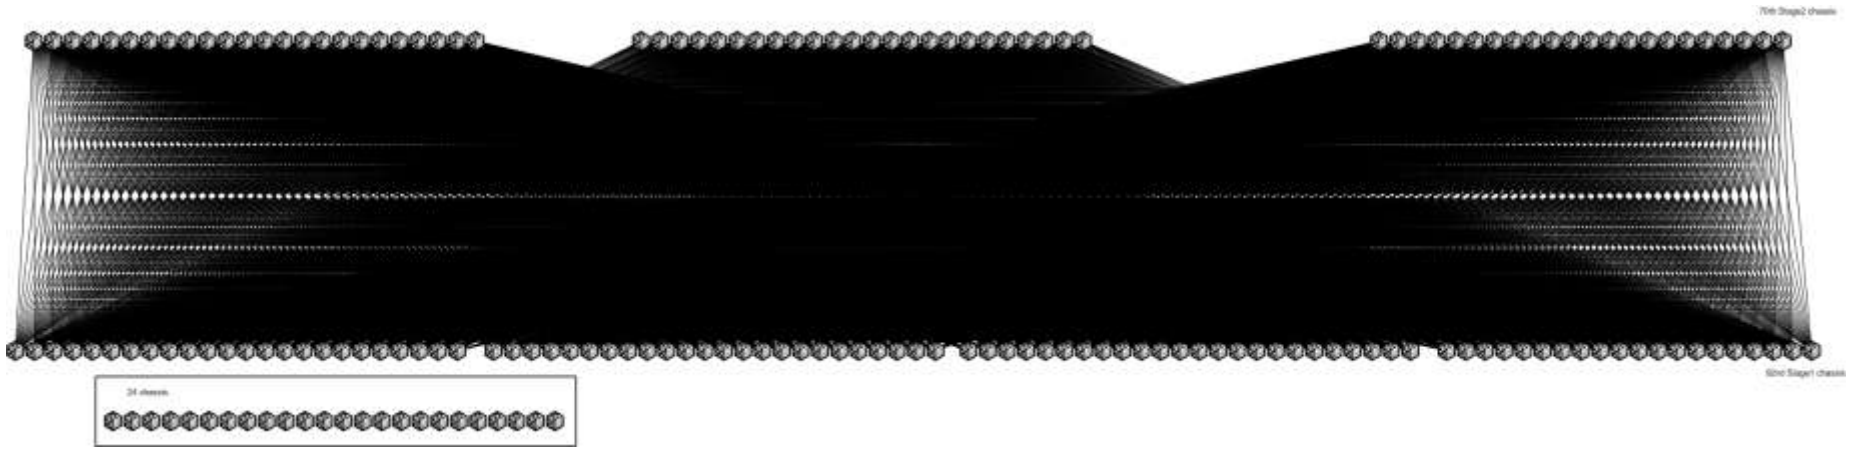

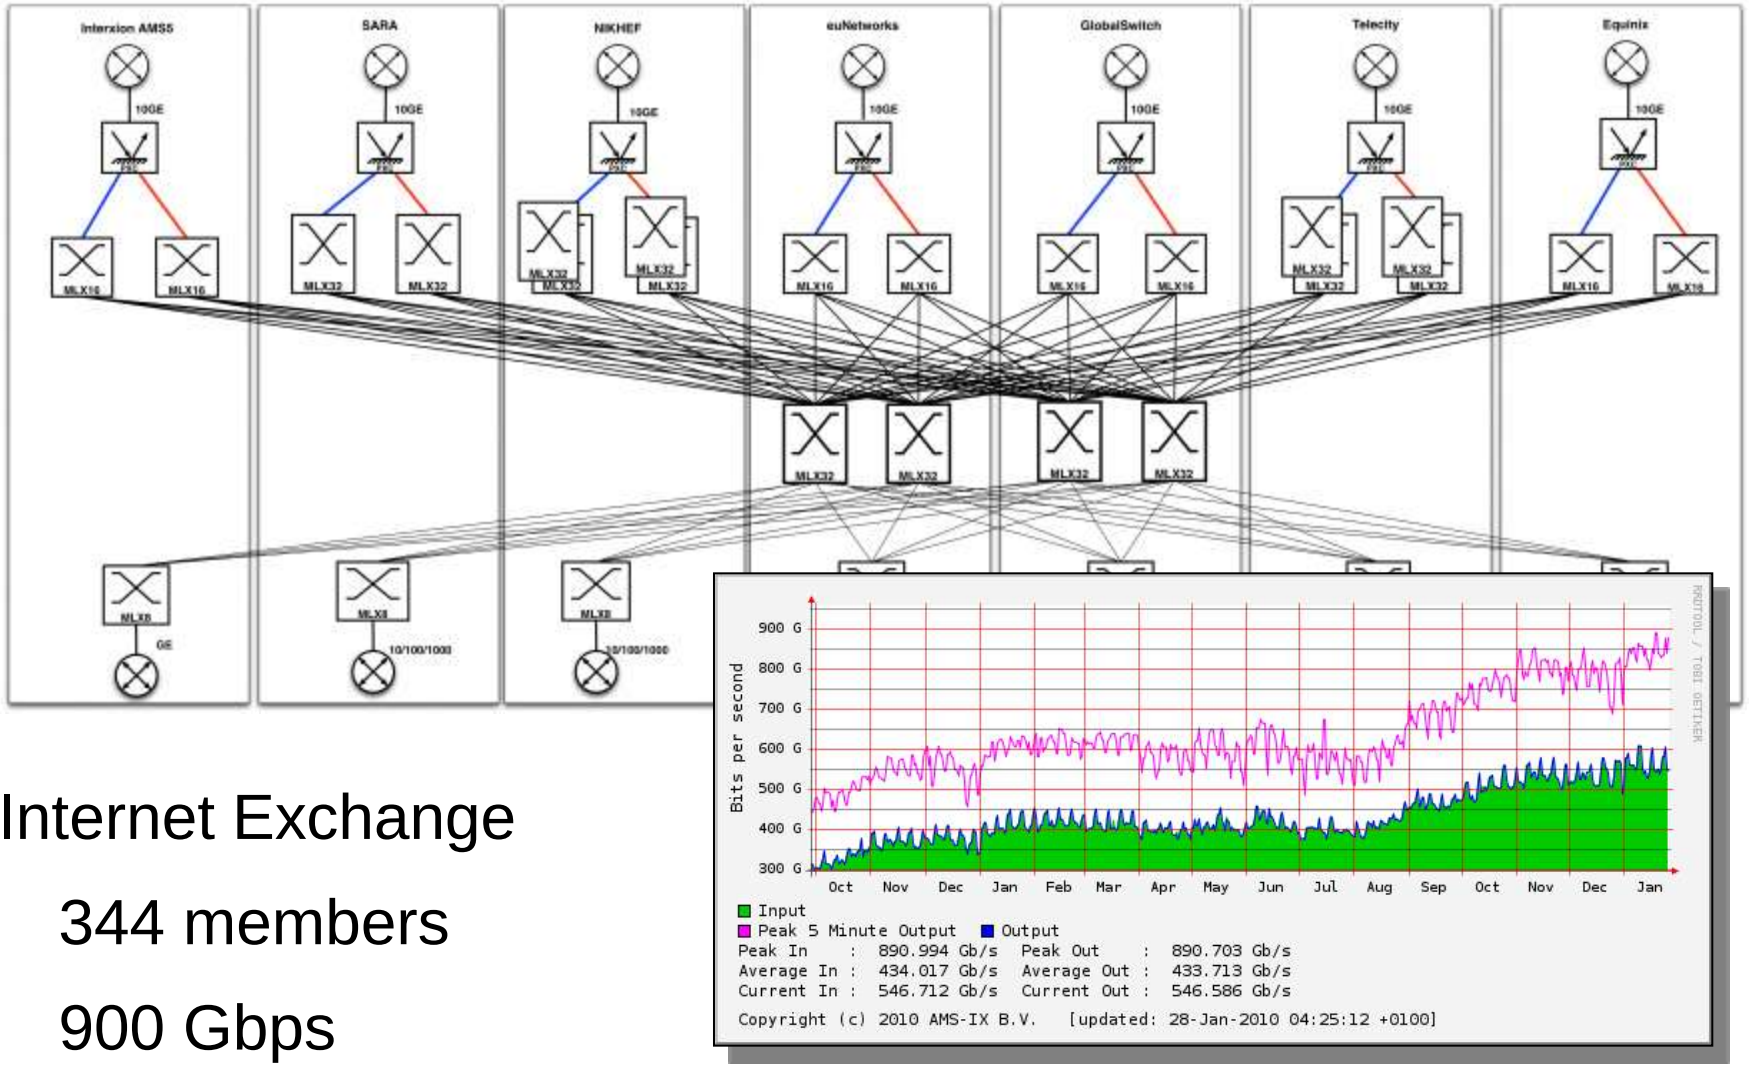

Internet Exchange

- 344 members
- 900 Gbps
- 8x10GE (and some 10x10GE) members
- 40/100GE aggregation needed

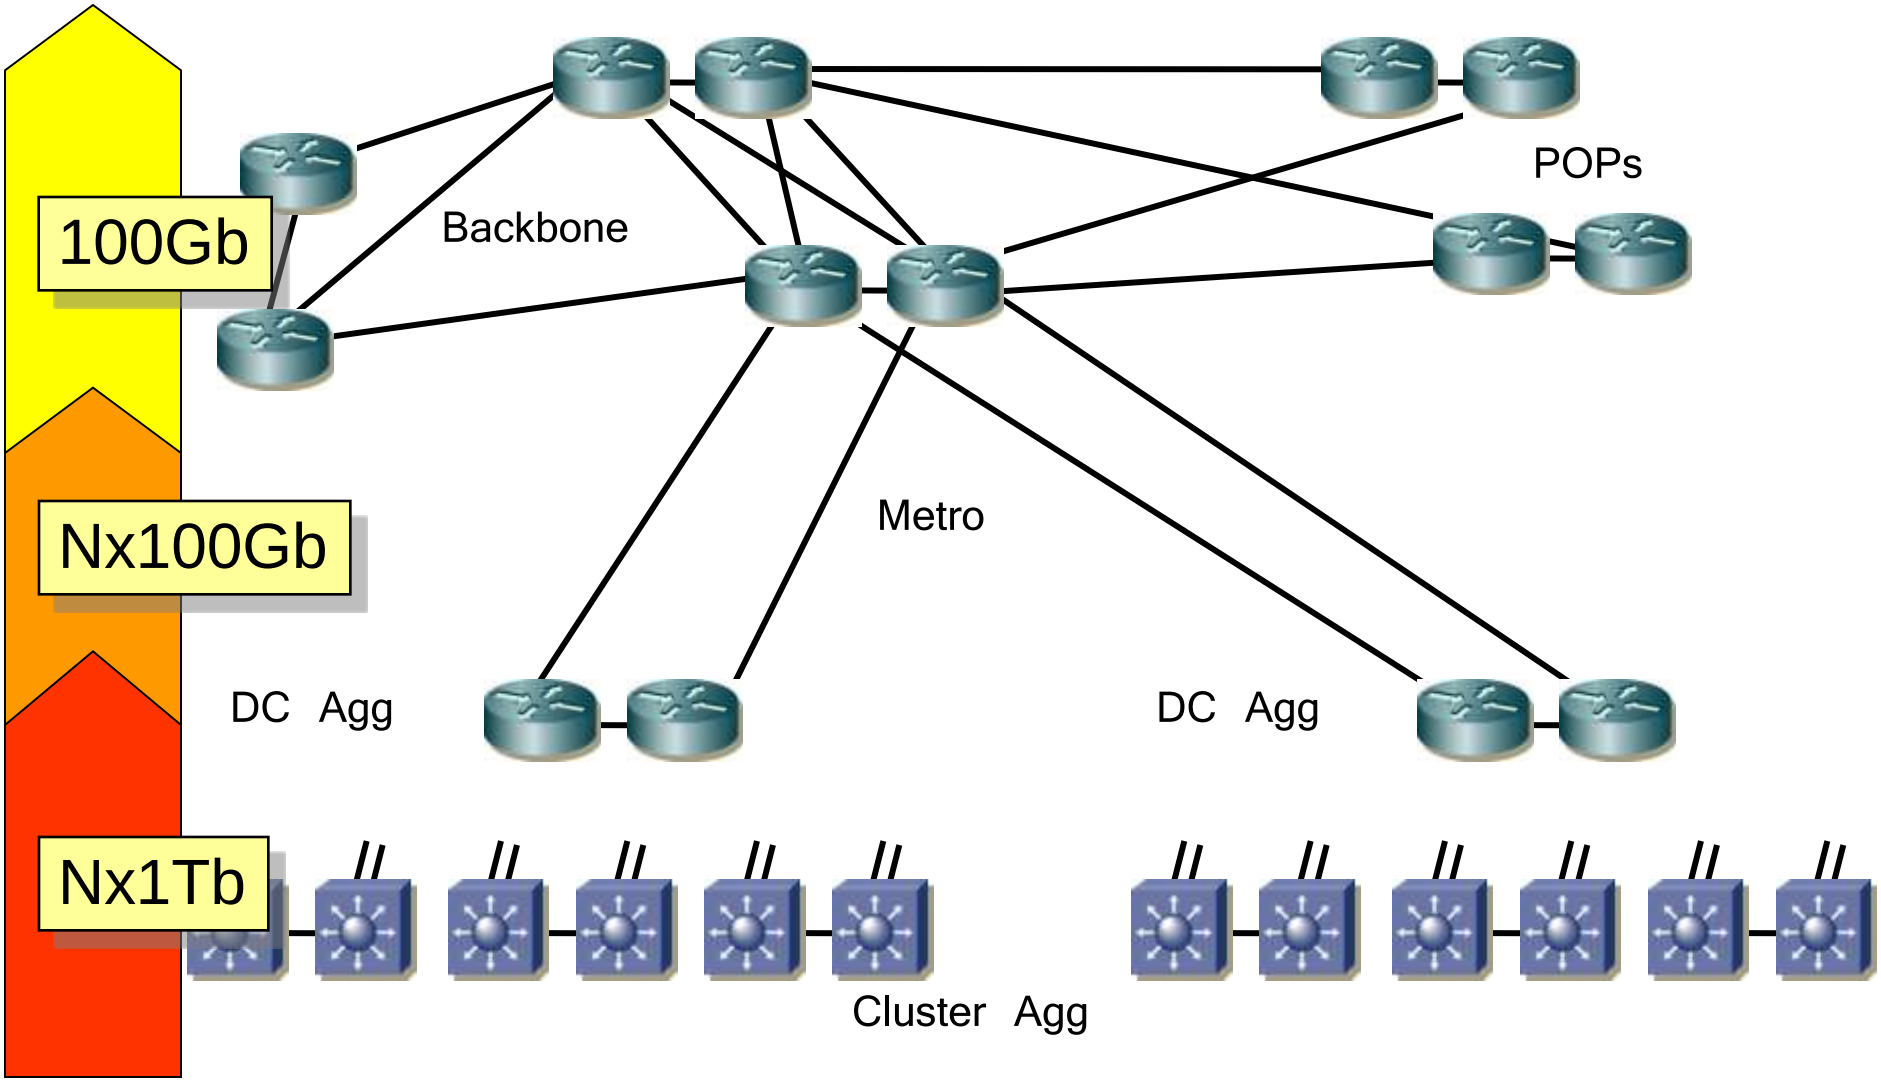
100Gb
Nx100Gb
Nx1Tb
Backbone
POPs
Metro
DC Agg
DC Agg
Cluster Agg

Thank You