US Patent Pending

Ultra-high Performance Parallel Big Data Transfer Software for Data-intensive Science

Chin Fang, Ph.D. , Founder
fangchin@zettar.com
650-644-9722

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Scale-up or Scale-out?

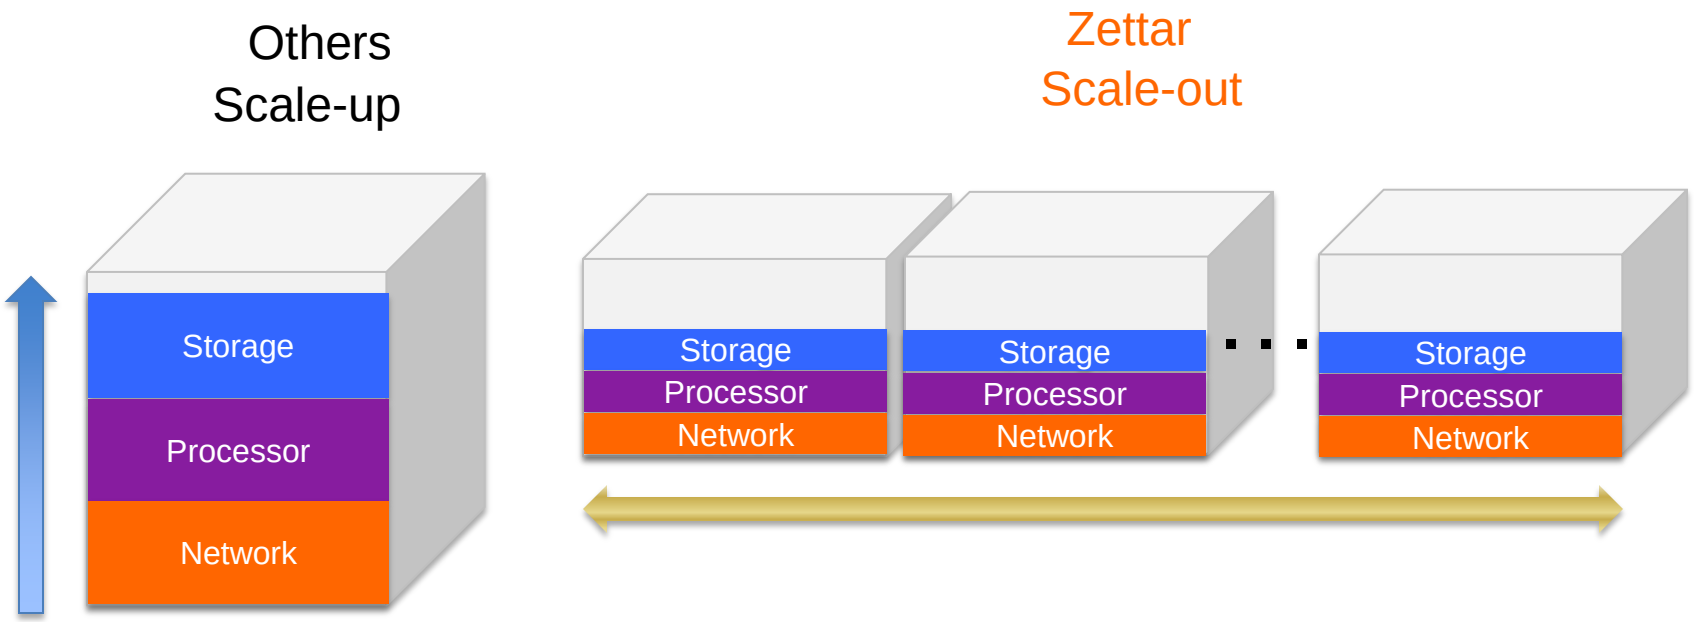

1. Too much data to store? Scale-out storage
2. Too much data to analyze? Scale-out compute
3. Too much data to transfer? Why not scale-out transfer?

Bounded Throughput, SPOF, Low In-Transit Security

Scalable Throughput, HA, High In-Transit Security

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

ZX – Boundlessly Scalable Data Transfer Throughput

US Patent Pending

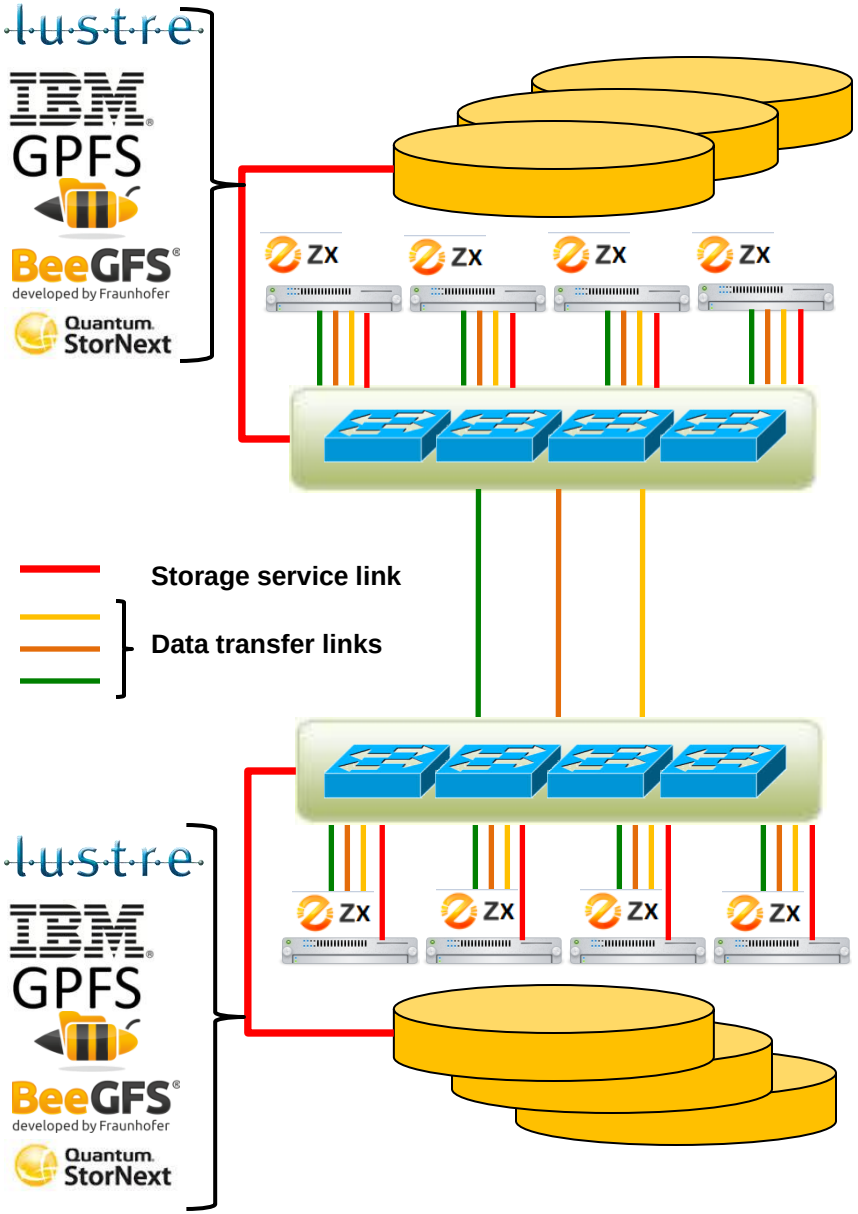

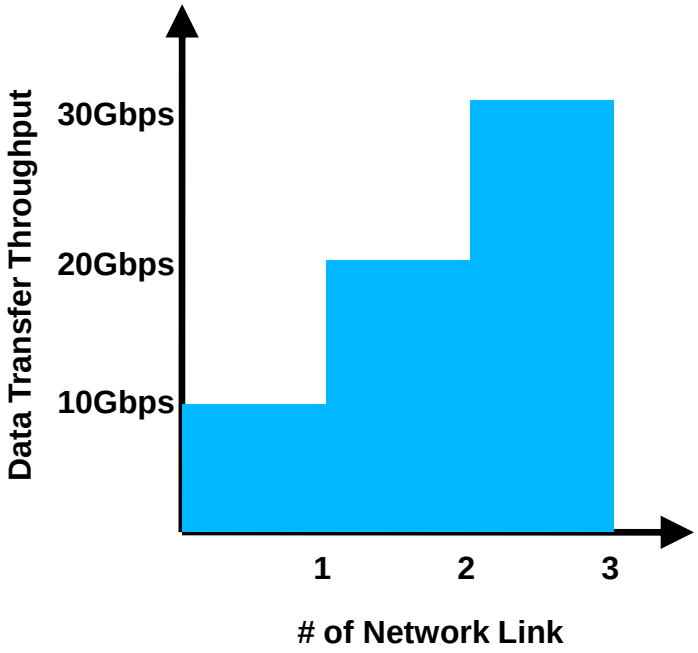

SSD, eSATA, USB 3, Thunderbolt, Network Speeds (MiB/s)

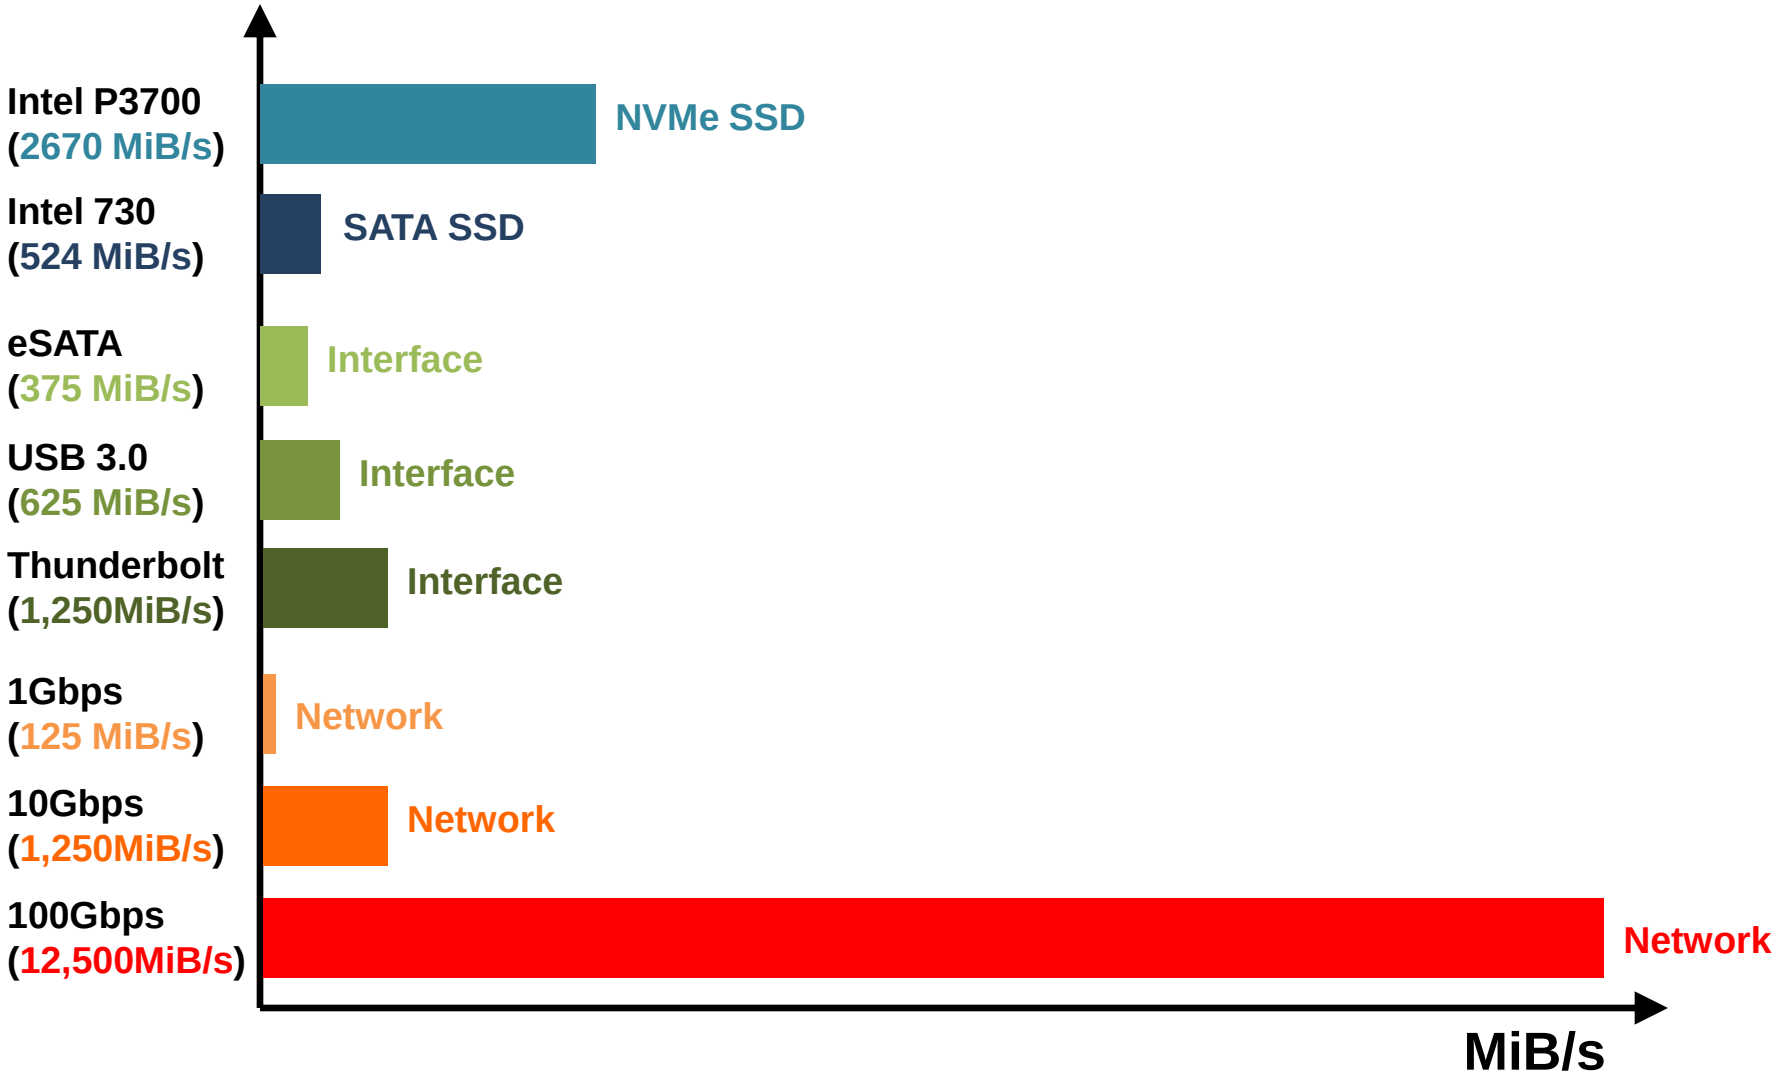

Fast Data Transfers Demand More Than Networks - I

All 7 pairs are 10GbE connected, why the throughput differences?

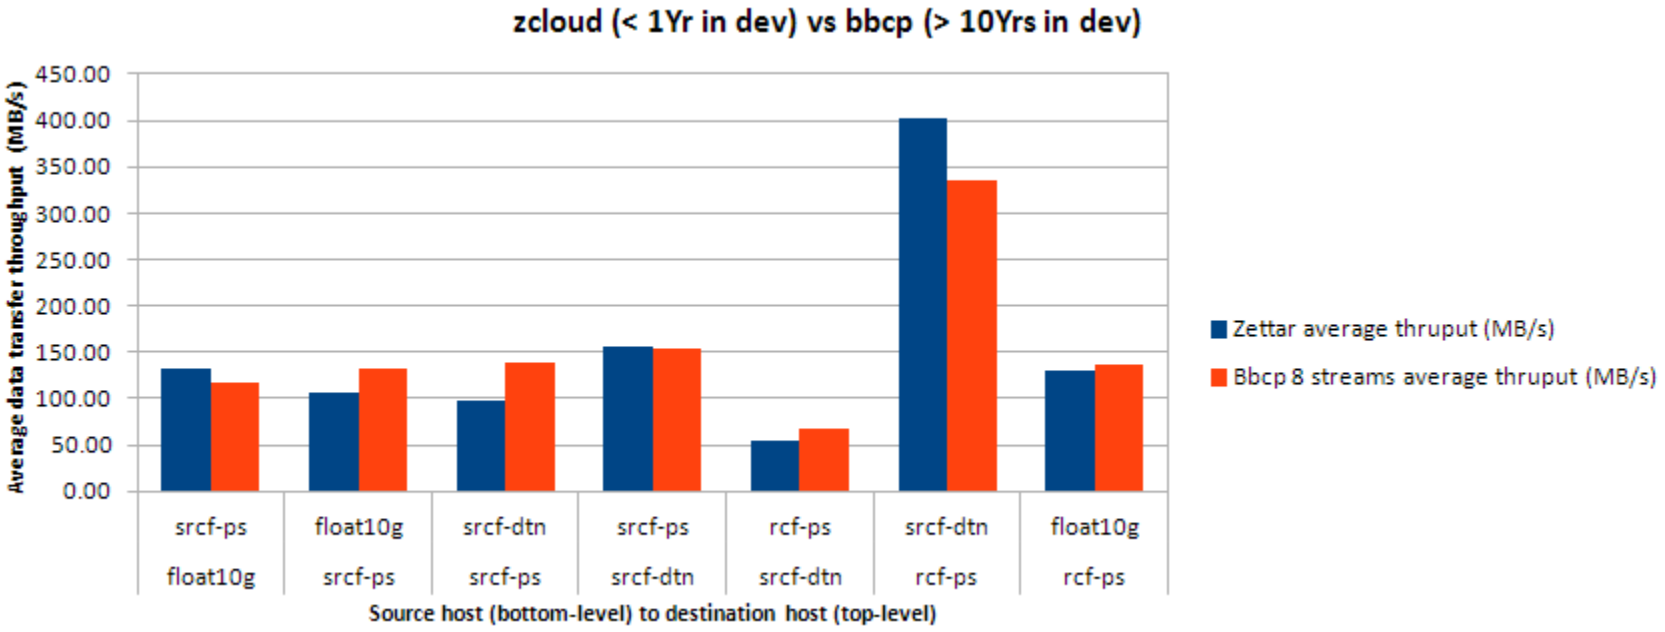

Data source: Dr. Reese, Stanford Research Computing and Dr, Chin Fang, Zettar Inc. 2013

Fast Data Transfers Demand More Than Networks - II

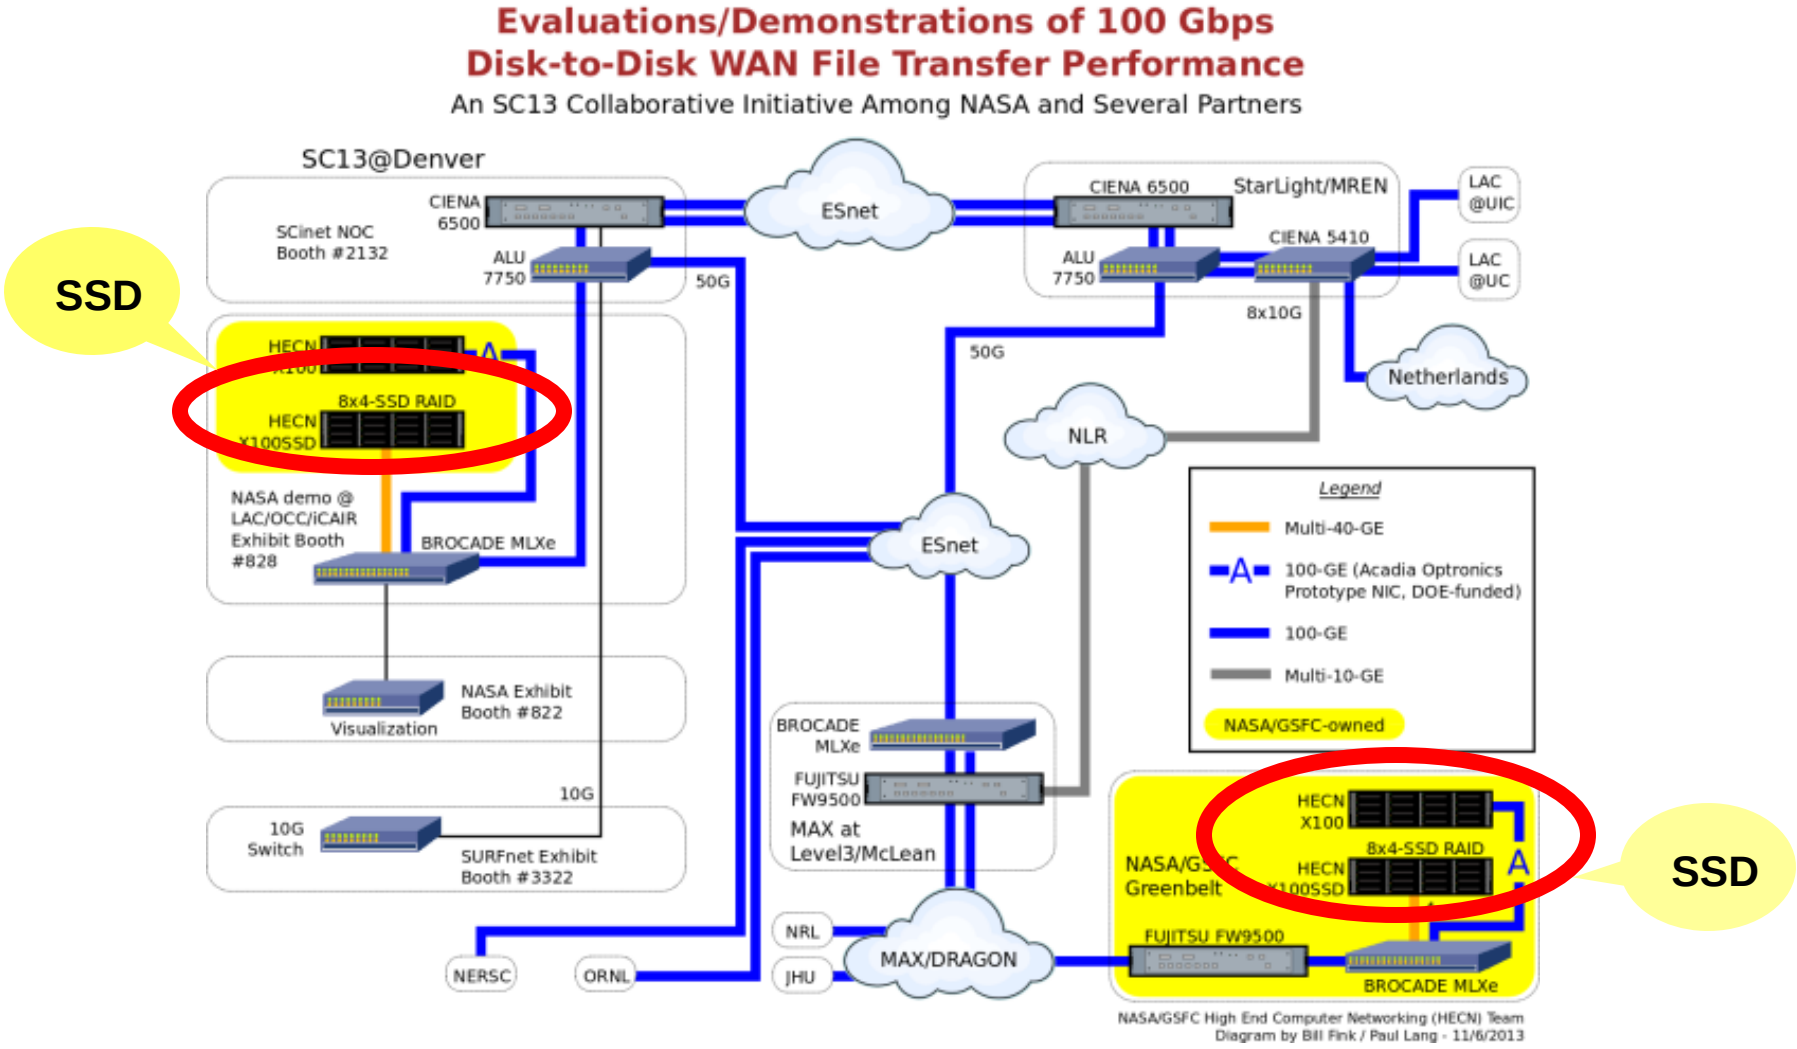

Verizon EVPL

Ethernet Access

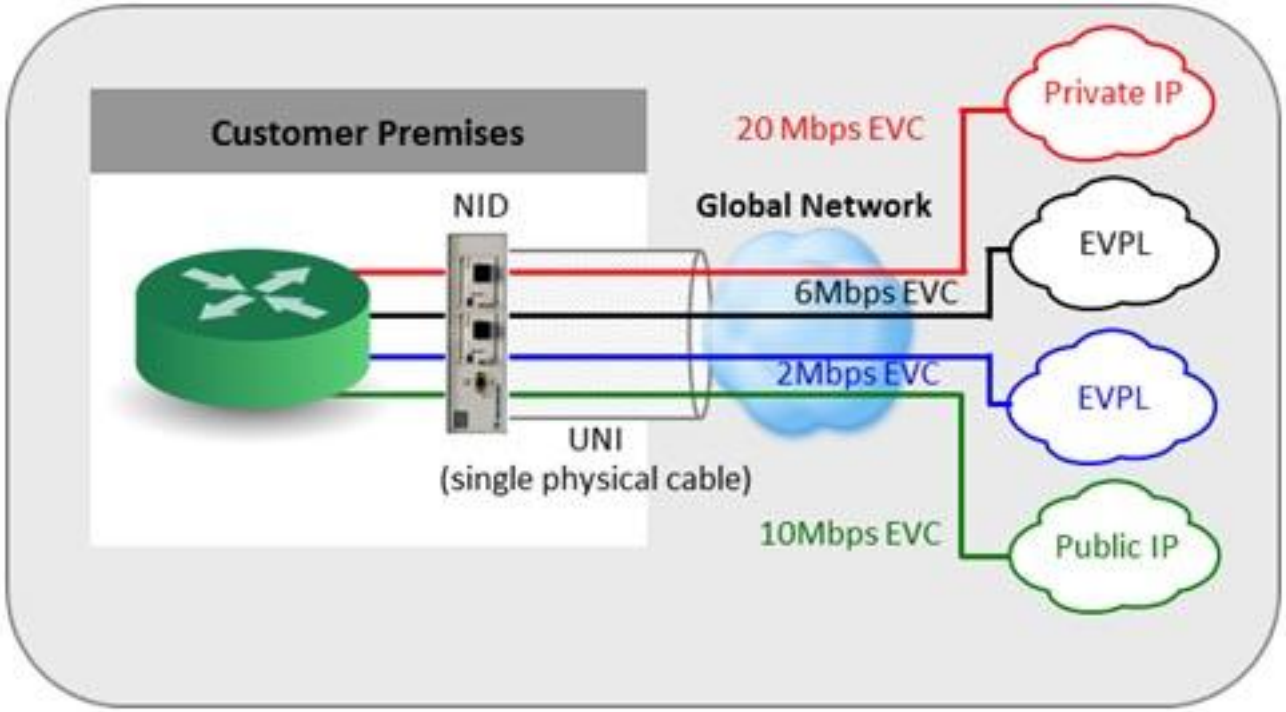

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Data Transfer Patterns

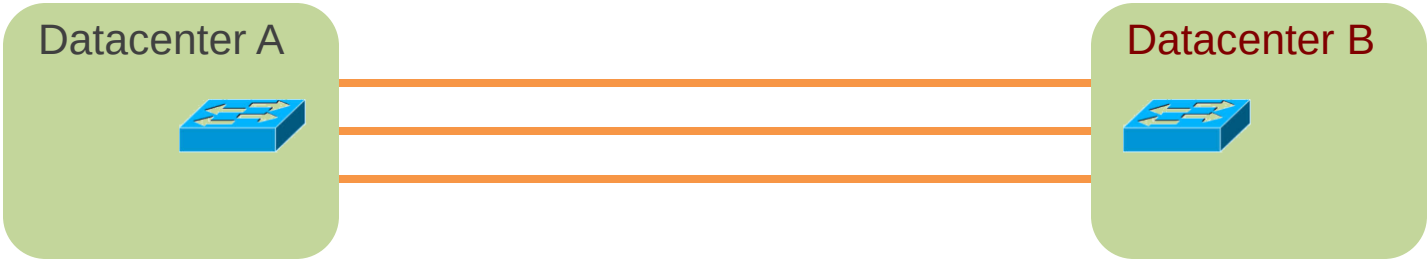

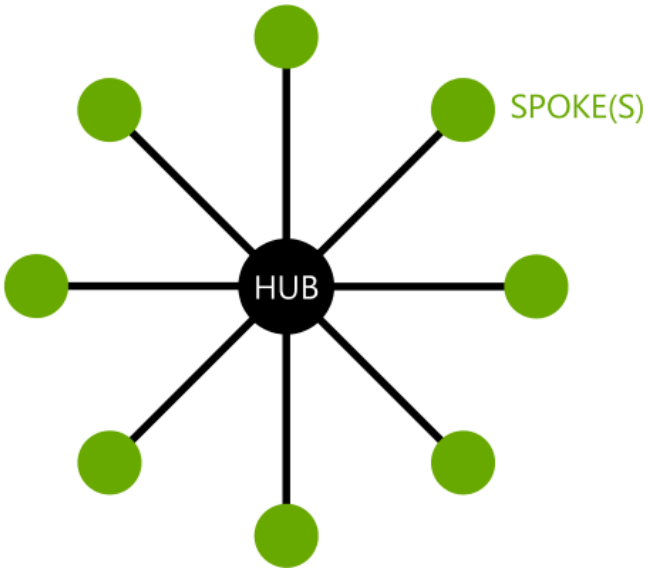

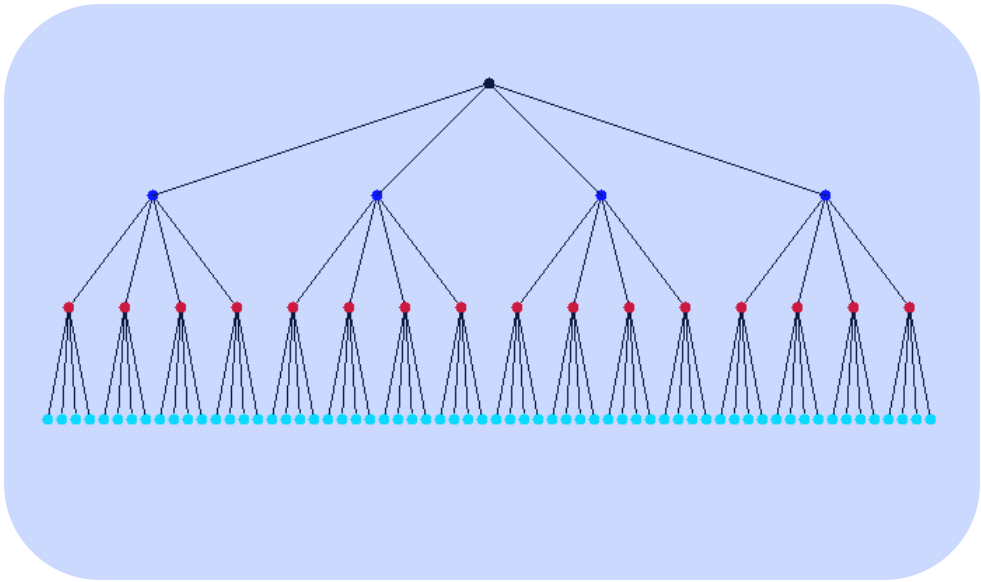

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

ZX – User-mode P2P-based Multipath Data Transfers

US Patent Pending

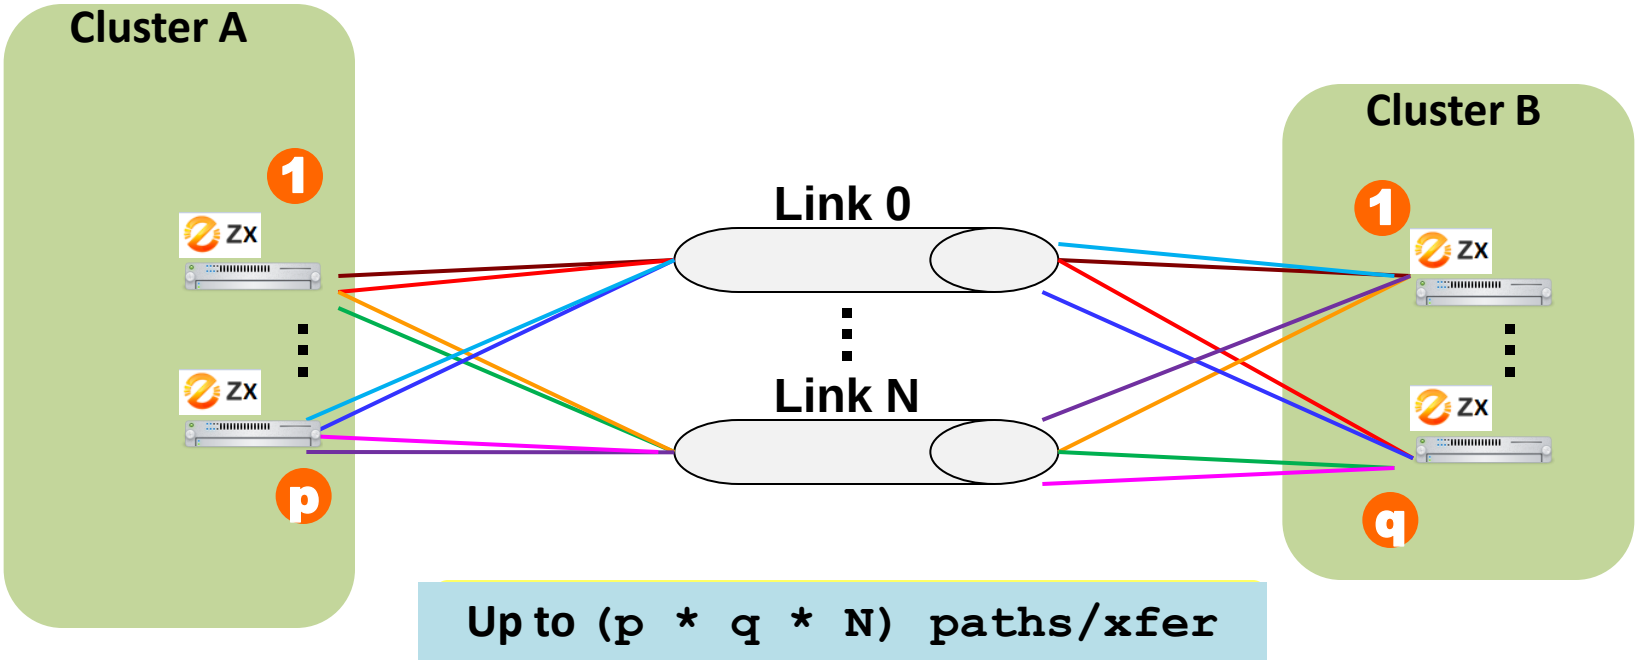

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Engineering Trade-offs To Meet SC14 Deadline

1. KISS to trade for short-term scalability & performance
2. Routing complexity to trade for scalability & performance
3. Demanding programming model to trade for product longevity

Never compromised: ease-of-use + good OOTB performance

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Design And Implementation Challenges

1. Distributed, highly concurrent, *asynchronous* programming model
2. Resources for testing and verification. Debugging too!
3. Human factors, e.g. communication

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Supercomputing 2014 Live Demo Setup

The Same Setup, Now Hosted In SLAC Building 50

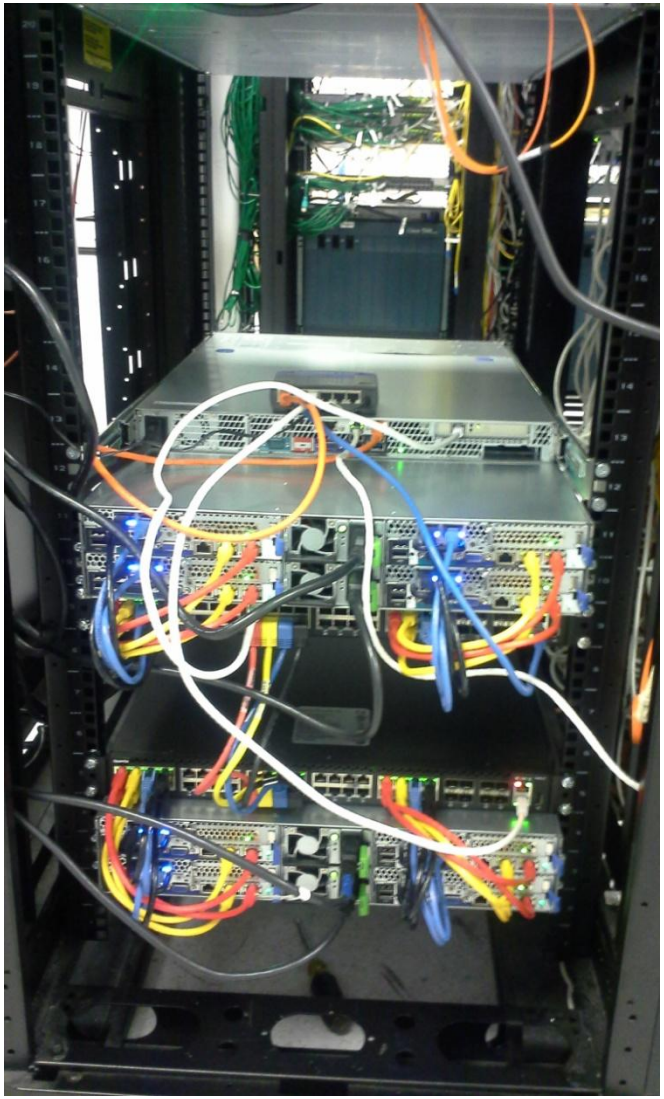

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Pilot Deployment At SLAC

1. Campus data transfer to start
2. LOSF cases?
3. Mix-sized and large file cases?

Lets use existing hardware as much as possible

Recommended

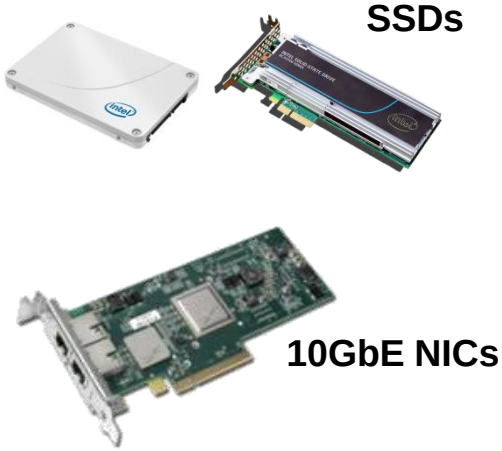

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

Upcoming SLAC Publications

1. **SLAC-TN-15-001**: Using NVMe Gen3 PCIe SSD cards in high-density servers for high-performance big data transfers over multiple network channels. *Published on Feb. 24, 2015*

2. A study of the Linux XFS and parallel file system overheads on the data transfer performance using high-density servers equipped with NVMe PCIe SSD devices.

3. More to come…

SC15: Surpass The Best LOSF Throughput Over 100Gbps

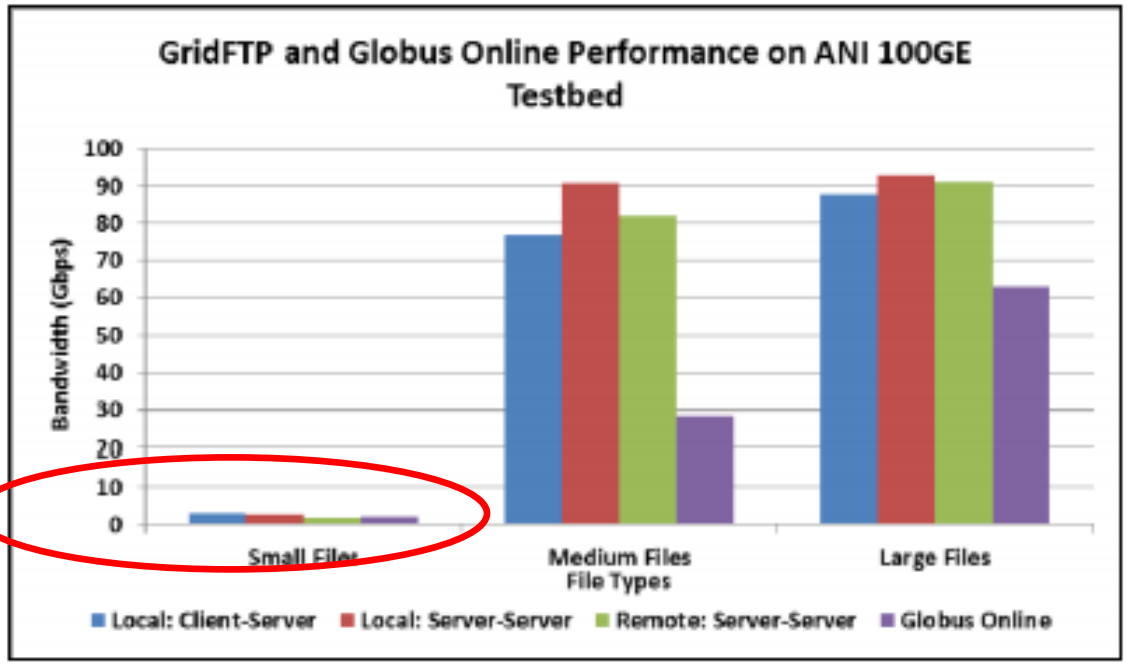

Figure 6: GridFTP & Globus Online performance comparison

4Gbps!

Source: Optimizing Large Data Transfers over 100Gbps Wide Area Networks, 2013

SC15: Transferring Big Data ~100 Gbps With A 2U Server

Picture credit: Caltech

QuantaPlex T41SP-2U (4-Node)

- Intel Haswell CPUs
- 10Gbps NICs
- NVMe SSDs

INTERNET 2

100Gbps

ESnet
Energy Sciences Network

100Gbps

US Patent Pending

Ultra-high Performance Parallel Big Data Transfer Software for Data-intensive Science

Chin Fang, Ph.D. , Founder
fangchin@zettar.com
650-644-9722

Agenda

Problem solved & motivation of the design

Unique capabilities & a key to high-speed data transfers

Use cases

Architecture overview

Engineering trade-offs to meet SC14 deadline

Design and implementation challenges

A quick review of the SC14 setup

SLAC pilot deployment discussions

Upcoming SC15 excitements ☺

Q&A

High-Speed TCP Data Transfers - Background

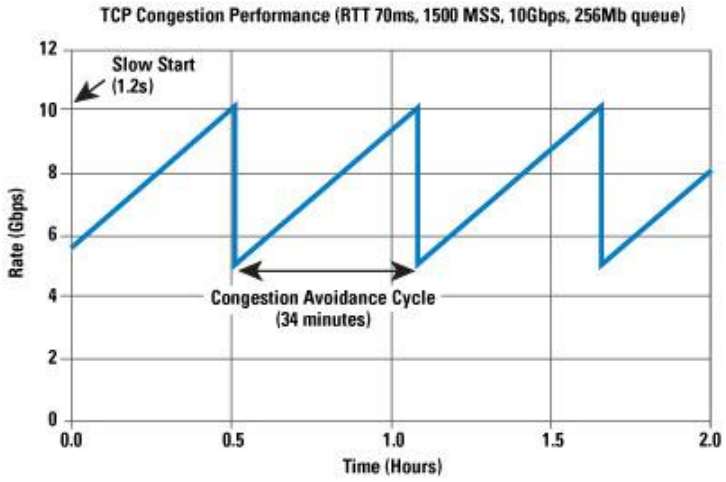

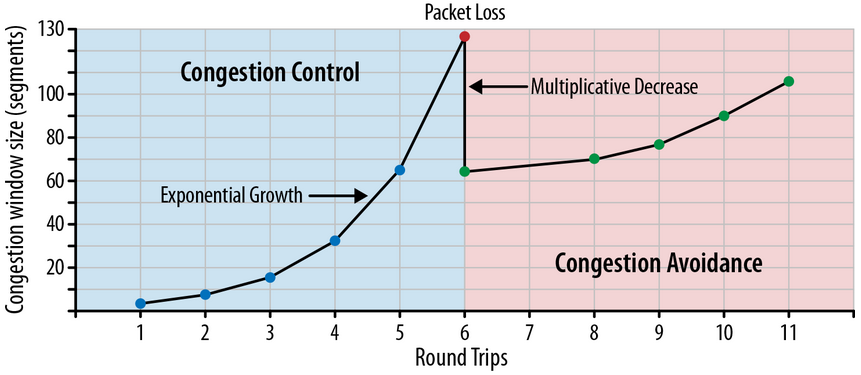

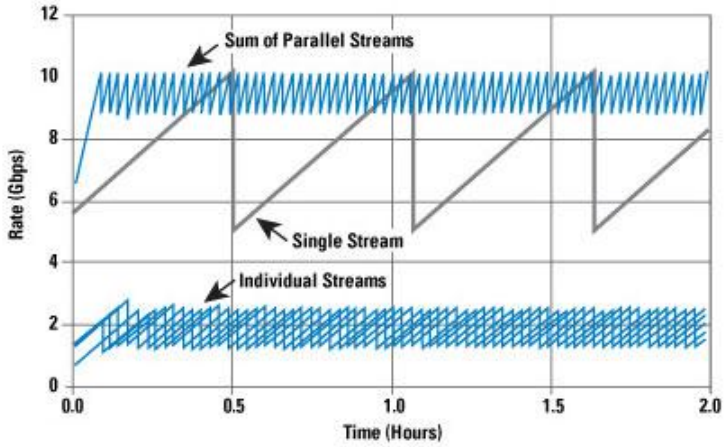

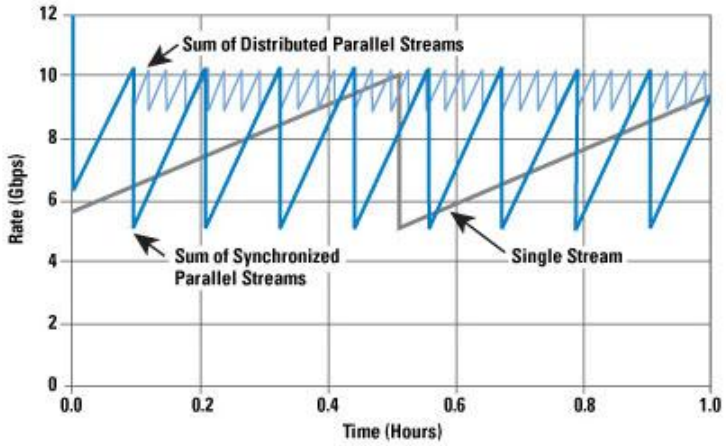

Picture credit: Geoff Huston, APNIC & Ilya Grigorik, Google Inc.

Computer Memory Hierarchy Diagram

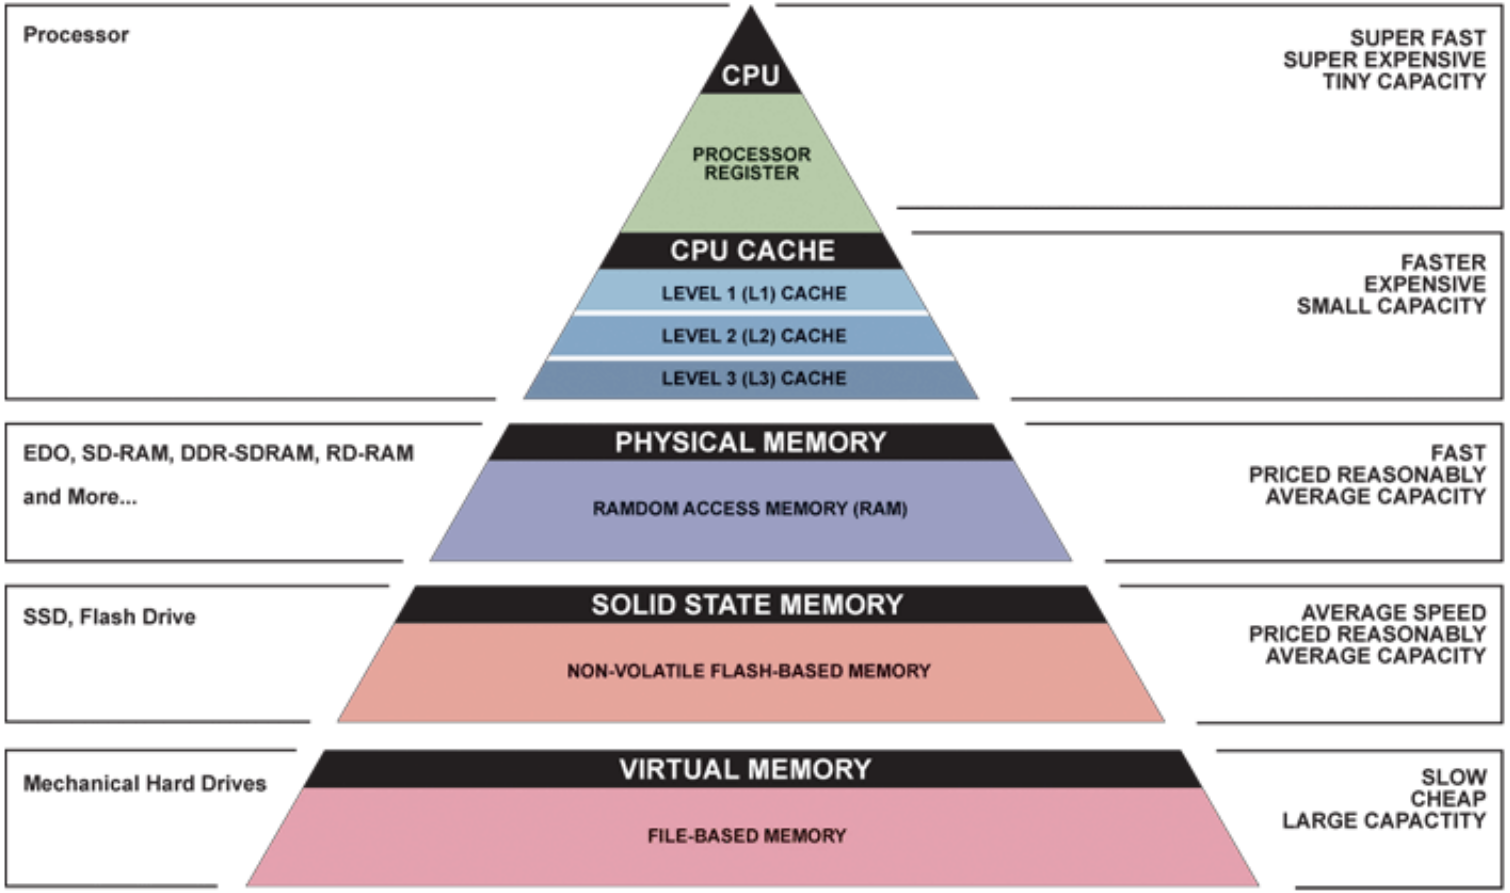

▲ Simplified Computer Memory Hierarchy
Illustration: Ryan J. Leng

Apply Storage Auto-Tiering To High-speed Data Transfers

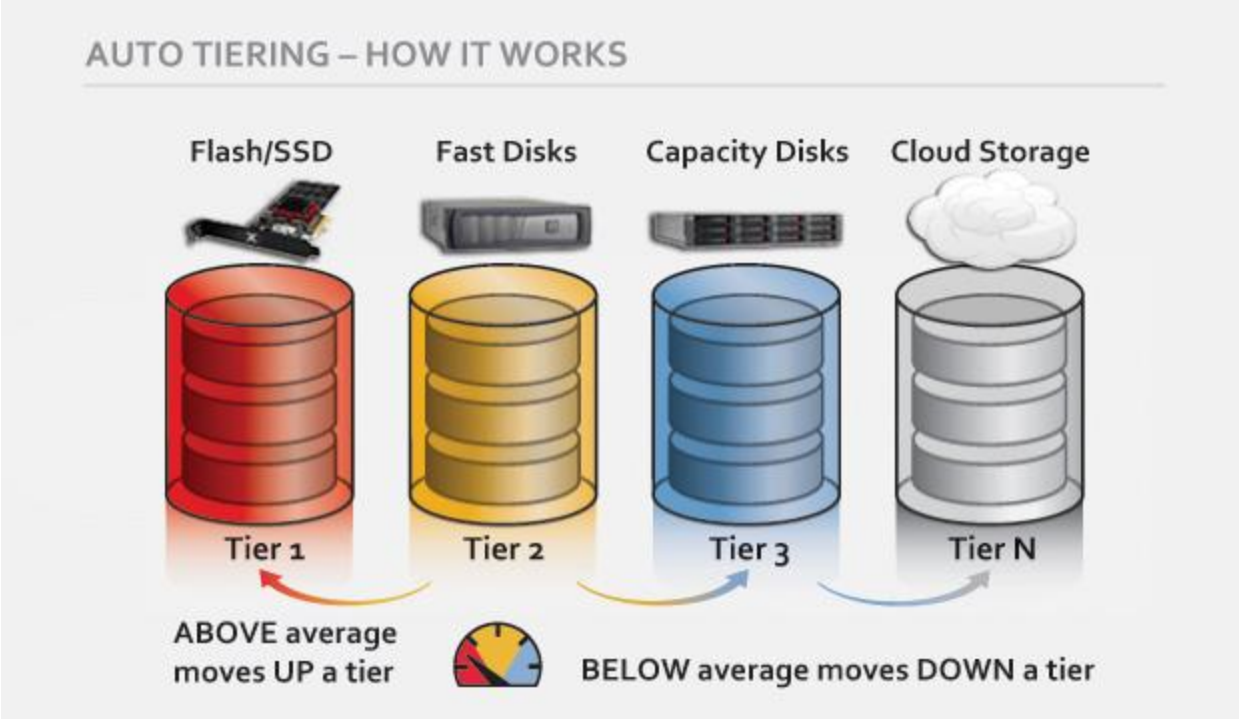

Picture credit: DataCore Software