

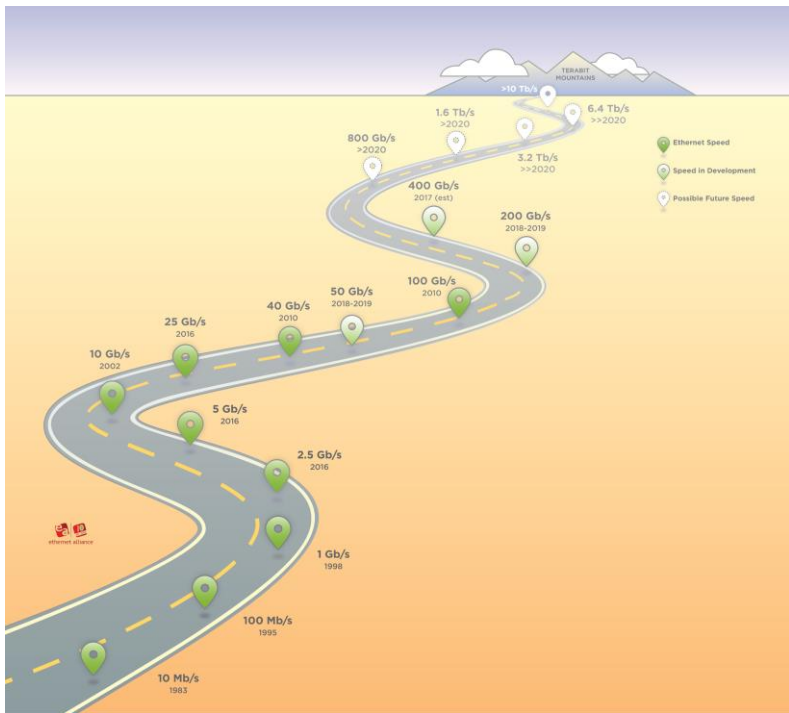
Ethernet Networked Flash Storage

J Michel Metz, Ph.D

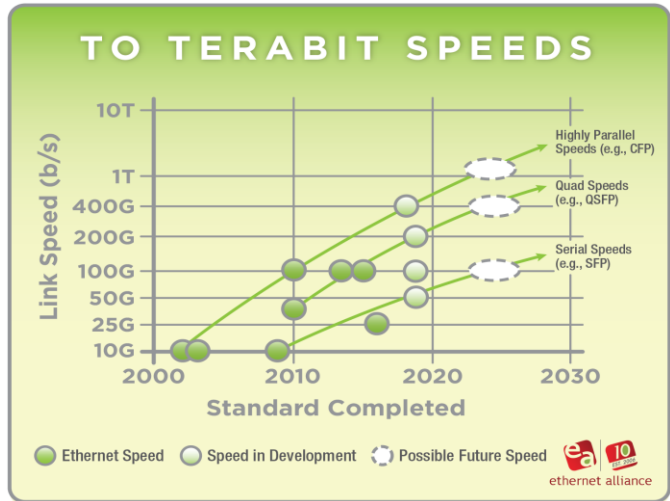
R&D Engineer - Advanced Storage

Cisco Systems, Inc.

Why Ethernet?



Source: <http://www.ethernetalliance.org/roadmap/>



- General Purpose network
- Massive capacity and roadmap
- Seemingly unlimited flexibility
- Broadest range of deployment options and architectures
- Ubiquitous and well-understood
- Easy-to-find expertise


Why (Maybe) *Not* Ethernet?

- Natural trade-off
 - Dedicated storage networks v. General Purpose storage networks
- “Just because you can, doesn’t always mean you should.”
- Choosing the wrong Ethernet storage networking architecture can “box” you in

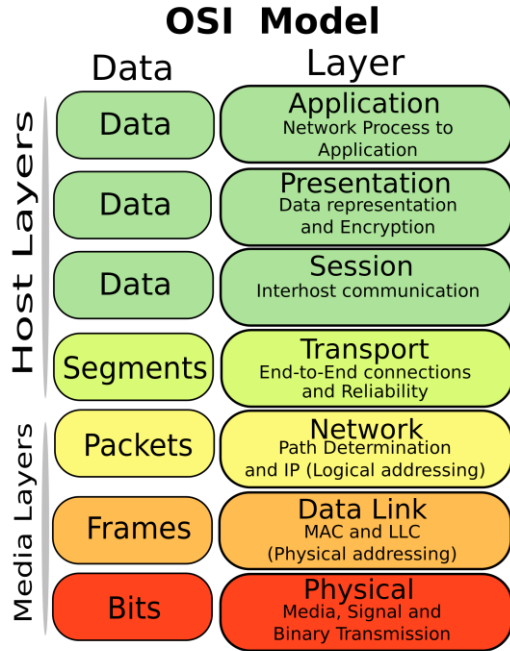


Types of Ethernet Storage Networks

- Closer to the wire?
- Closer to the software?
- What's the difference?
 - Congestion
 - Oversubscription
 - Latency concerns



L2 v L4

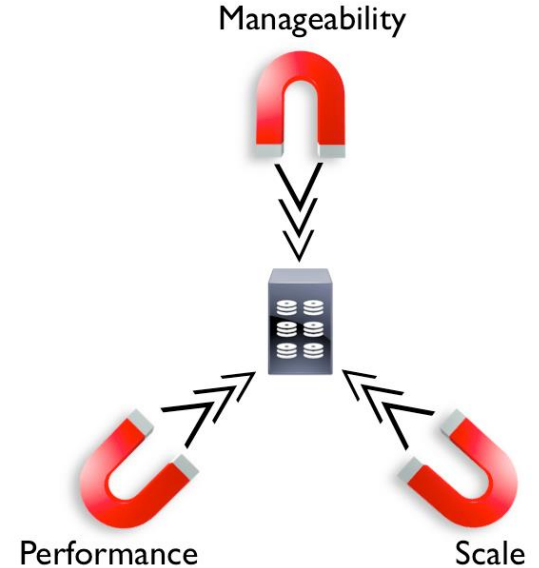


- OSI model:
 - Layer 2
 - For *deterministic* storage network traffic
 - RoCE, FCoE, Lossless iSCSI
 - Layer 4
 - For *non-deterministic* storage network traffic
 - iSCSI, iWARP, NFS, SMB, Object

Source: <https://commons.wikimedia.org/wiki/File:Osi-model-jb.svg>

Understand the Storage Forces

- There are forces pulling on your storage *and* your storage network
 - Growth over time
 - Pressures on budgeting



Source: <https://jmetz.com/2016/07/storage-forces/>

Top 5 Things To Know About Ethernet and Flash

- 1) **Not all storage protocols are created equal**
 - Each are designed to solve a problem, and they are not interchangeable
 - “IP” storage is designed for lossy networking environments, high oversubscription, and lower priority on latency
 - Consistency models vary (e.g., NFS, pNFS, Object)
 - Hyperconvergence models are deliberately restricted on scalability to manage these factors
 - “Lossless” storage (e.g., FCoE, RoCE, SMB Direct) is designed for low fan-in ratios, low oversubscription, high-predictability



Top 5 Things To Know About Ethernet and Flash

- 1) Not all storage protocols are created equal
- 2) **Identify the shift in latency burden**
 - Higher up the stack you go, the more latency you will suffer
 - Lower the stack you go, the more rigid your architectural choices
 - E.g., putting a 500us software stack ontop of a 10us NVMe media makes the software faster, but is it overkill?



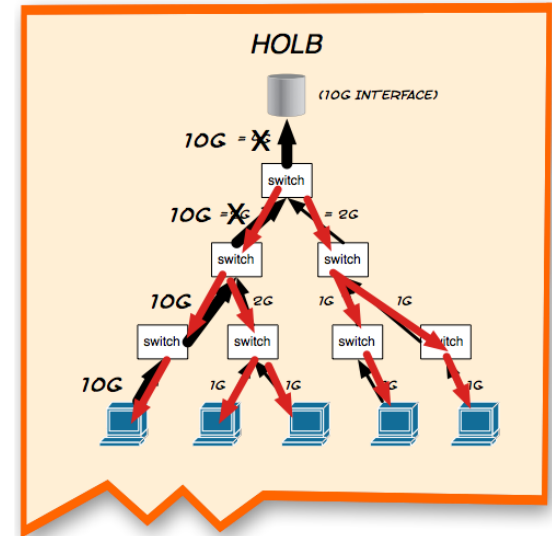
Top 5 Things To Know About Ethernet and Flash

- 1) Not all storage protocols are created equal
- 2) Identify the shift in latency burden
- 3) **Understand your *criteria for failure***
 - At what point will your preferred method break? Scale?
Performance? Manageability?



Top 5 Things To Know About Ethernet and Flash


- 1) Not all storage protocols are created equal
- 2) Identify the shift in latency burden
- 3) Understand your *criteria for failure*
- 4) **Do the Upper and Lower layers align?**
 - Are you robbing L4 by changing L2 (e.g., lossless storage networks with lossy protocols)?
 - E.g., Lossless iSCSI can be the worst of both worlds, especially with high throughput of flash devices



Source: <http://blogs.cisco.com/datacenter/the-napkin-dialogues-lossless-iscsi>

Top 5 Things To Know About Ethernet and Flash

- 1) Not all storage protocols are created equal
- 2) Identify the shift in latency burden
- 3) Understand your *criteria for failure*
- 4) Do the Upper and Lower layers align?
- 5) **Consider evolutionary growth patterns of storage needs and network development**
 - Will your thirst for storage outgrow your network's capabilities?
 - Will your network evolution create storage ghettos?
 - Are you putting workloads together that match with your storage networks?



Source: <https://jmetz.com/2016/07/storage-forces/>