lightbits

NVMe/TCP is the Best Way to Disaggregate Flash

Muli Ben-Yehuda
Co-Founder & CTO

Flash Memory Summit

A little bit about Lightbits Labs and me

- Lightbits is a well-funded stealth mode startup with offices in Israel and San Jose, CA
- Doing cool things with NVMe and NVMe-oF
- Inventors of NVMe/TCP

- Me: Lightbits Labs co-founder & CTO
- Operating systems, hypervisor, clouds, high-speed networking and storage

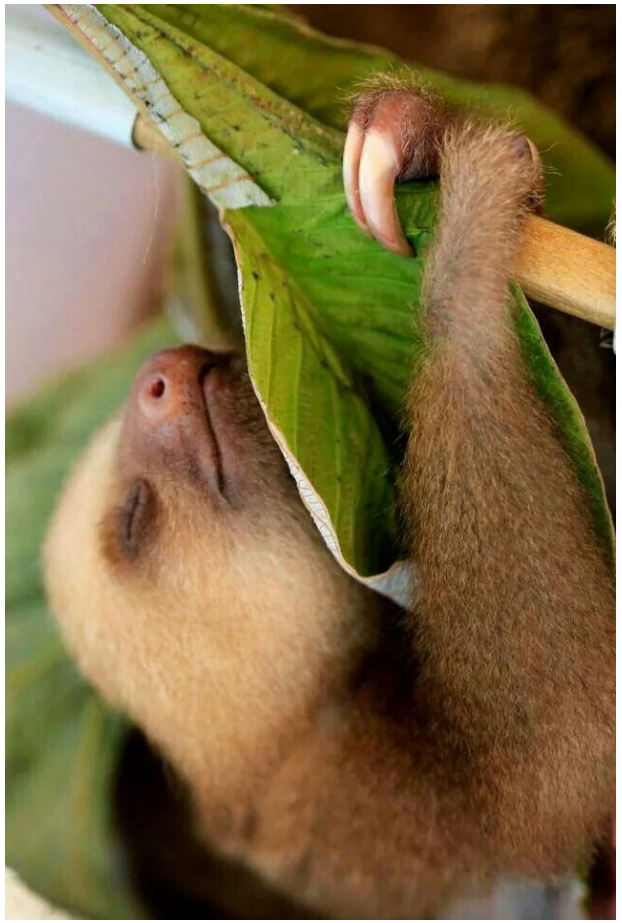

Disaggre-what?

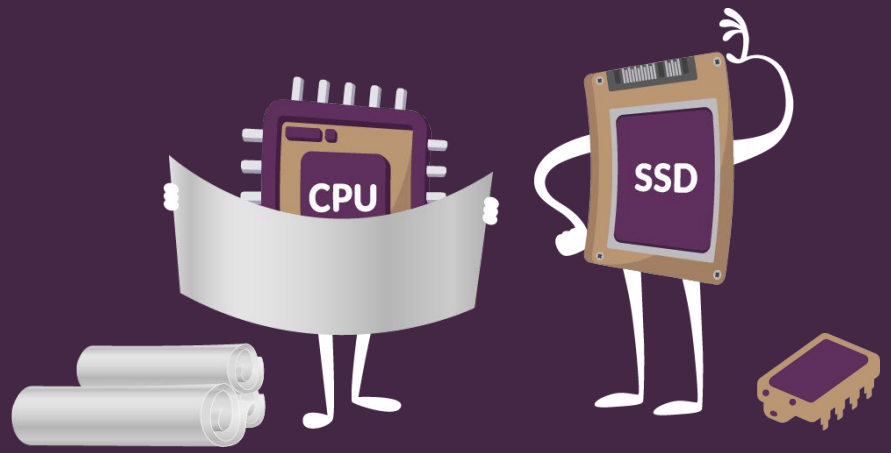

From direct-attached to a disaggregated cloud

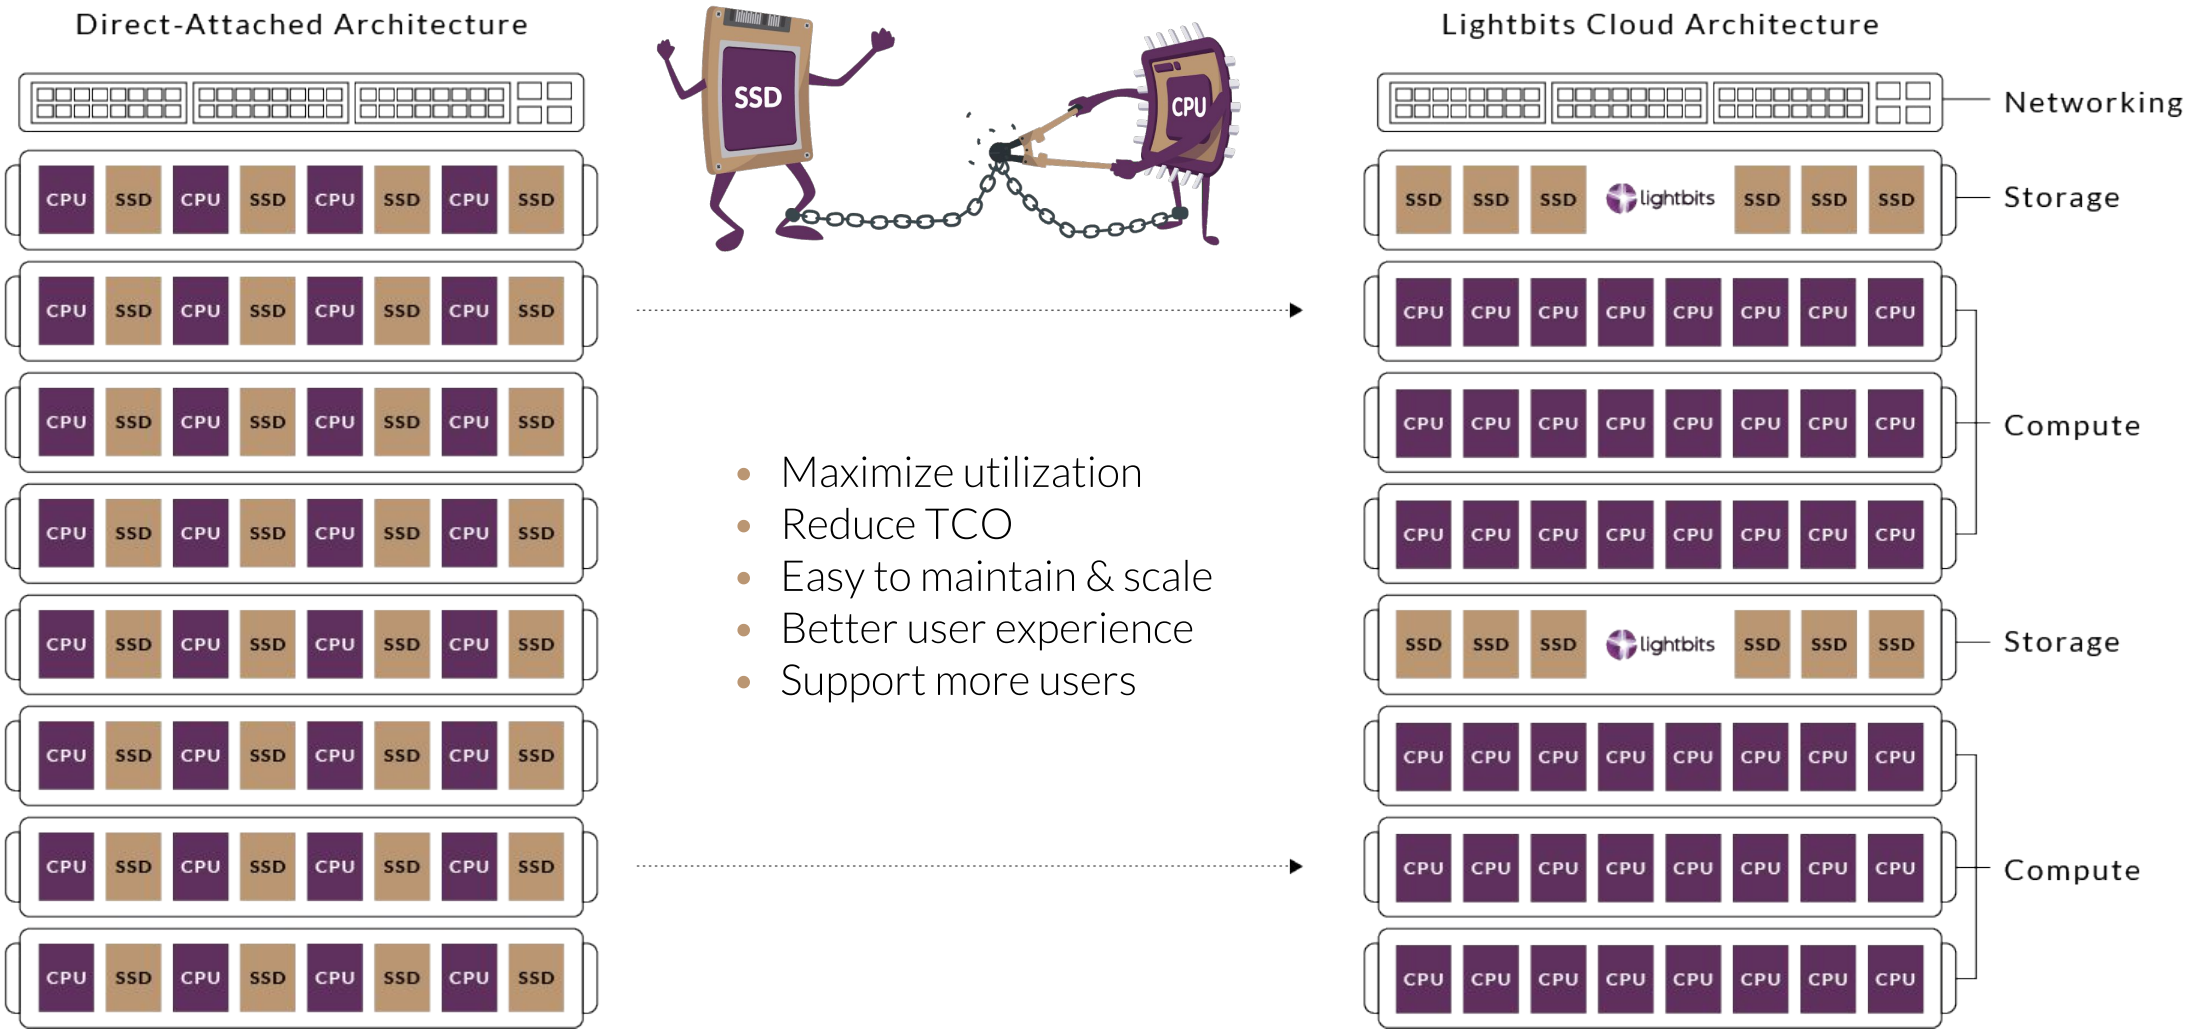

- Maximize utilization
- Reduce TCO
- Easy to maintain & scale
- Better user experience
- Support more users

NVMe over Fabrics

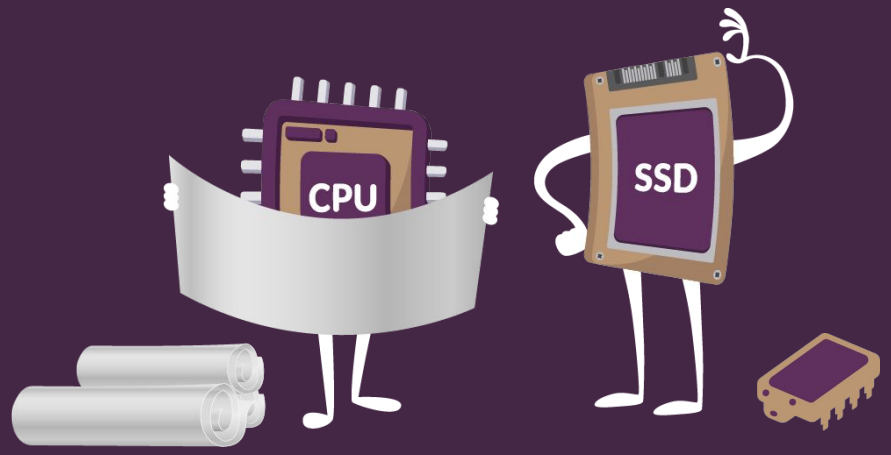

NVMe over Fabrics

- PCIe is a network (transport)
- Can we do NVMe over other transports?
- Retain NVMe efficiency and performance over network fabrics
- Eliminate unnecessary protocol translations
- Enable low-latency and high IOPS **remote** NVMe storage solutions

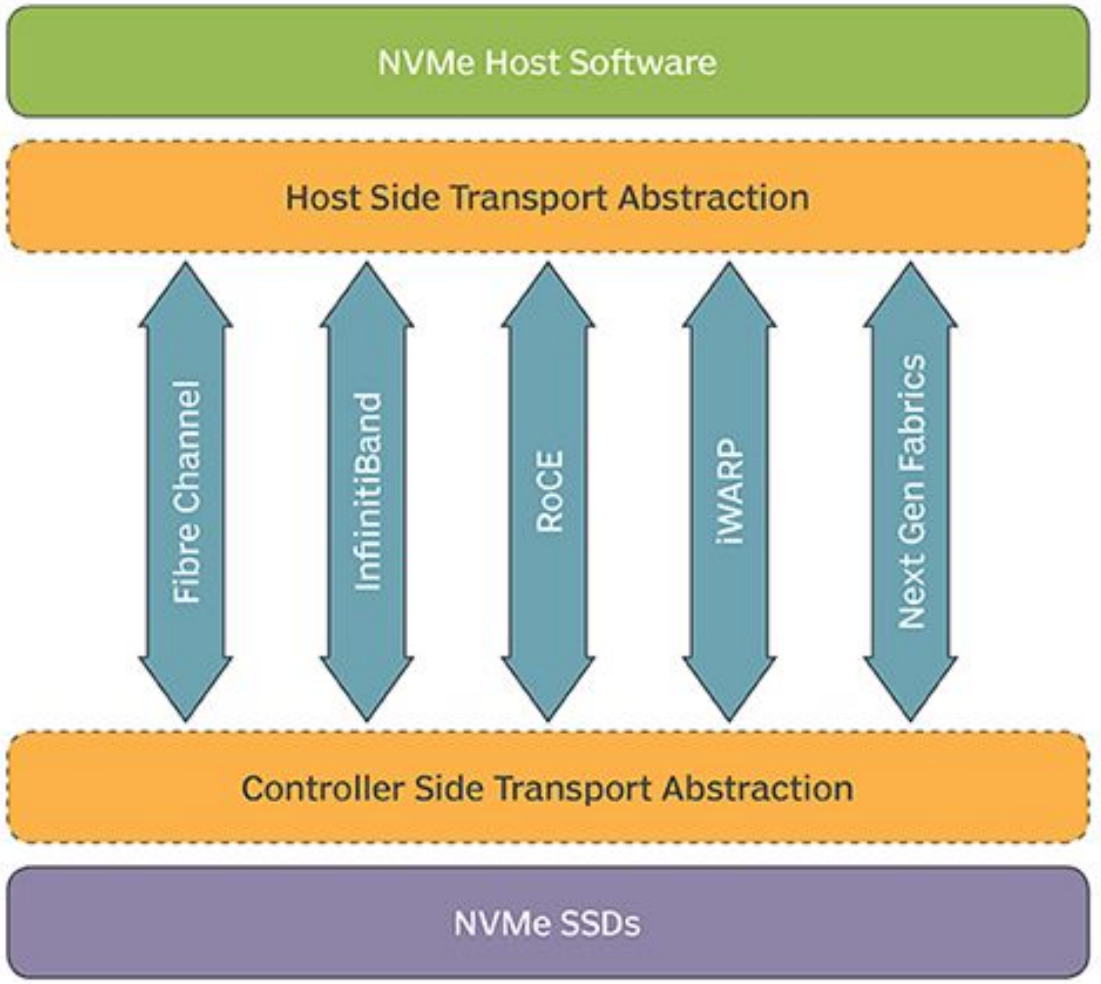

Flash Memory Summit

Spot the missing protocol...

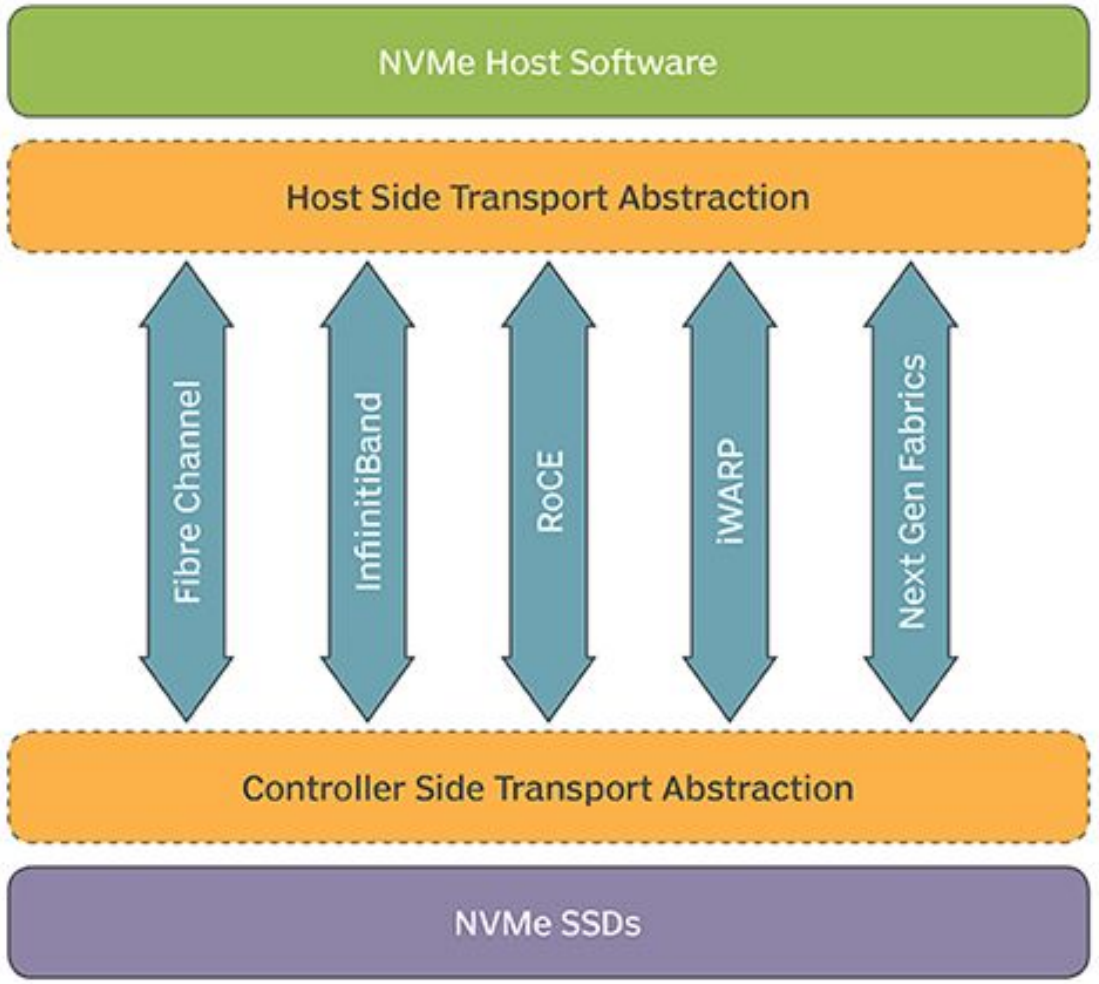

Enter TCP/IP

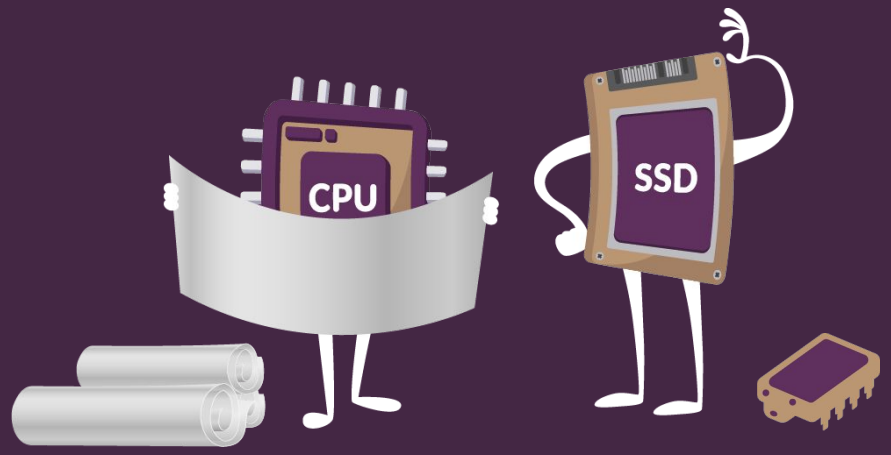

TCP/IP model	Protocols and services	OSI model
Application	HTTP, FTTP, Telnet, NTP, DHCP, PING	Application
		Presentation
		Session
Transport	TCP, UDP	Transport
Network	IP, ARP, ICMP, IGMP	Network
Network Interface	Ethernet	Data Link
		Physical

Santa Clara, CA
August 2017


why?

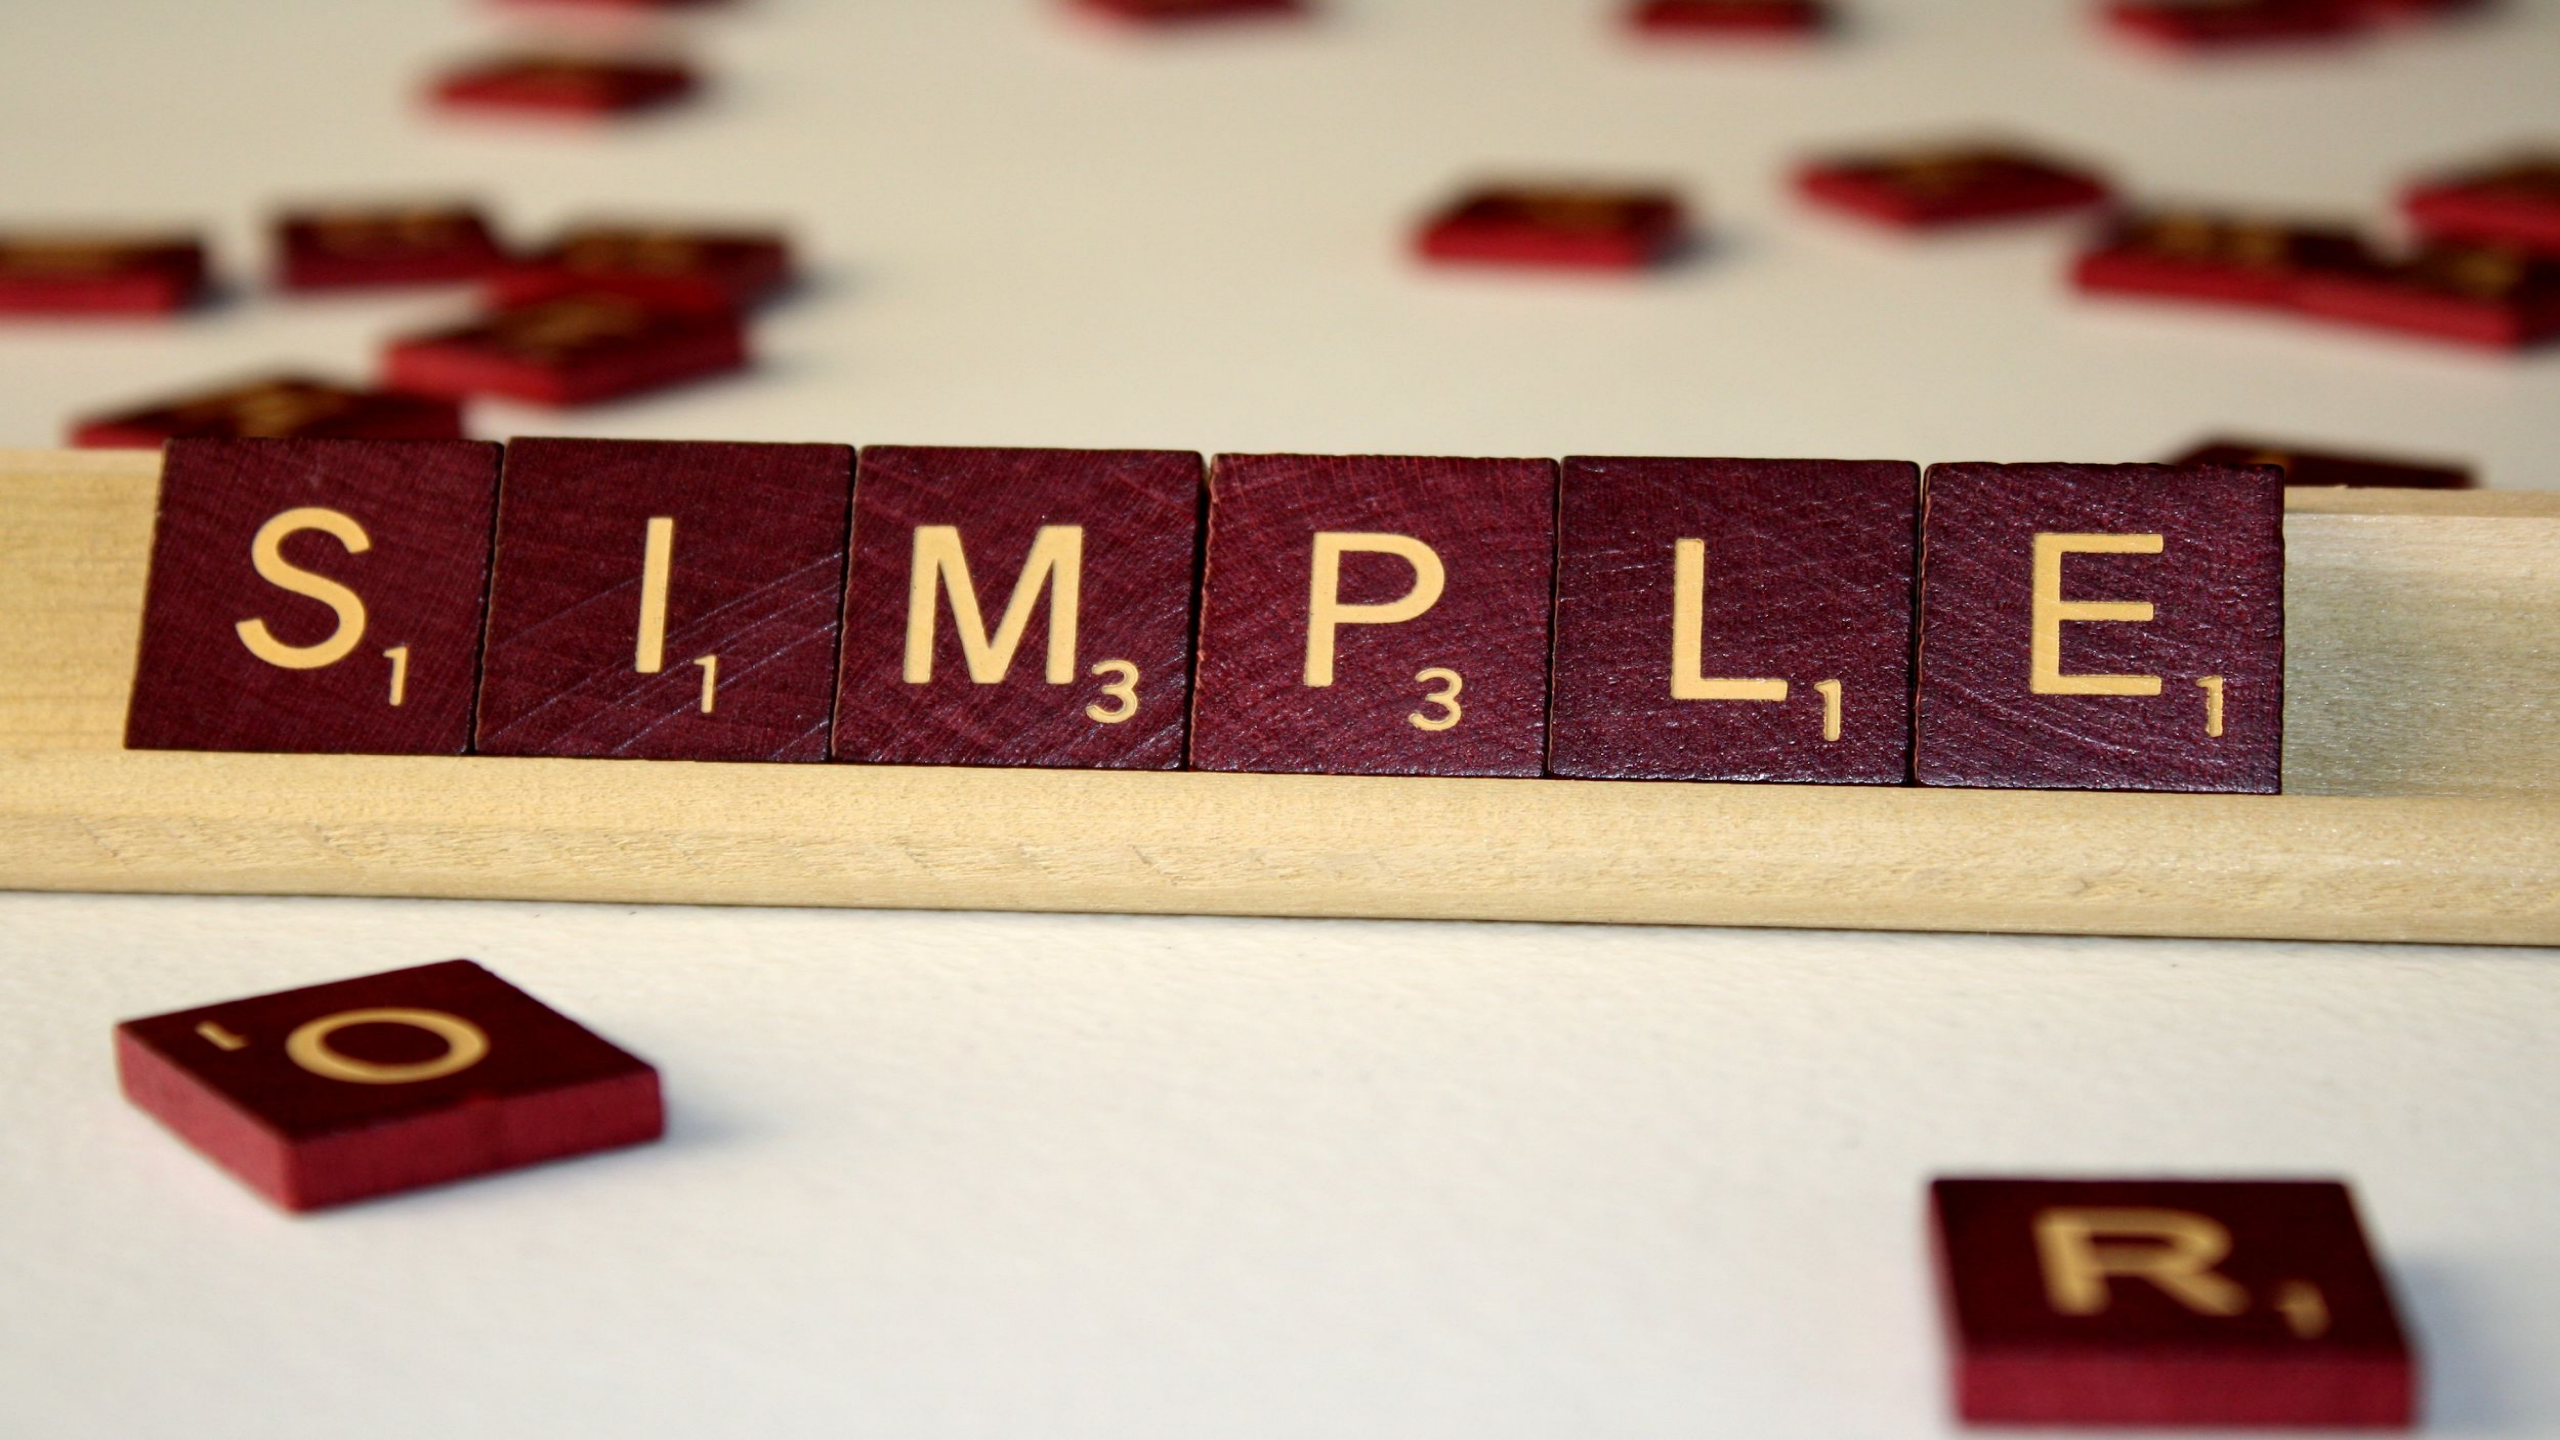
SIMPLE

ubiquitous (adjective)

1. Being everywhere at once: omnipresent.

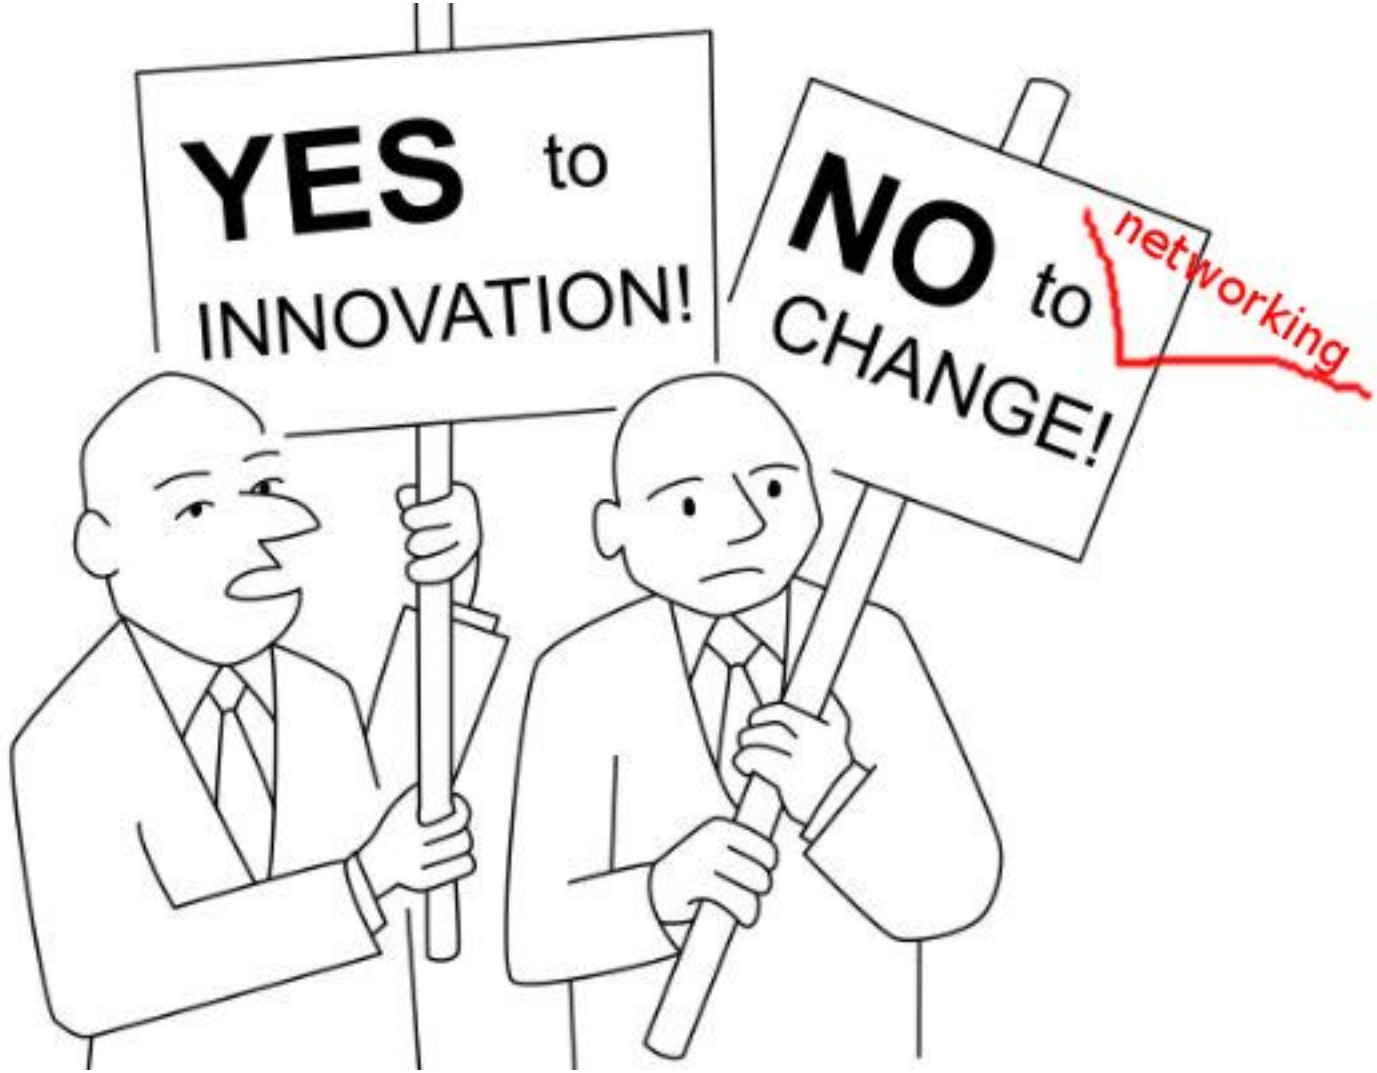
YES to INNOVATION!
NO to CHANGE!
networking


Innovation
Process
Research
Efficiency
Vision
30%
15%
7%
48 %
Investment
Strategy
Marketing
Development
Teamwork
Analysis
Management
Partner

ESSENCE

meaning, definition, explanation...

Flash Memory Summit

NVMe/TCP in a nutshell

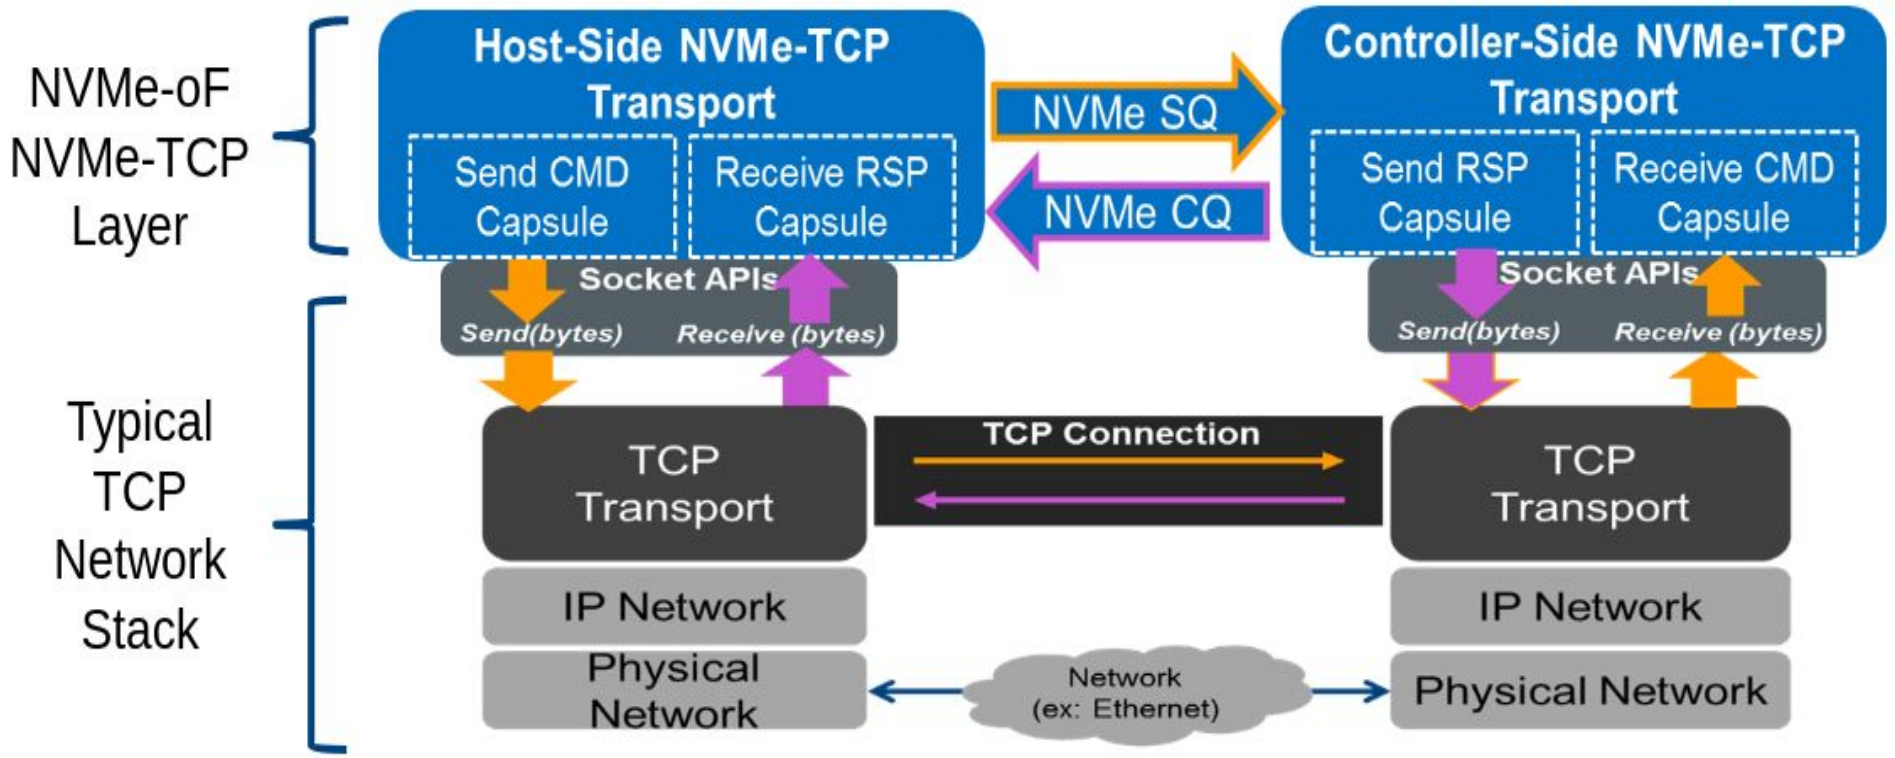

- A TCP/IP transport binding for NVMe over Fabrics
- NVMe-OF Commands sent over standard TCP/IP sockets
- Each NVMe queue pair mapped to a TCP connection
- TCP provides a reliable transport layer for NVMe queueing model

Is this the real life?
Is this just fantasy?

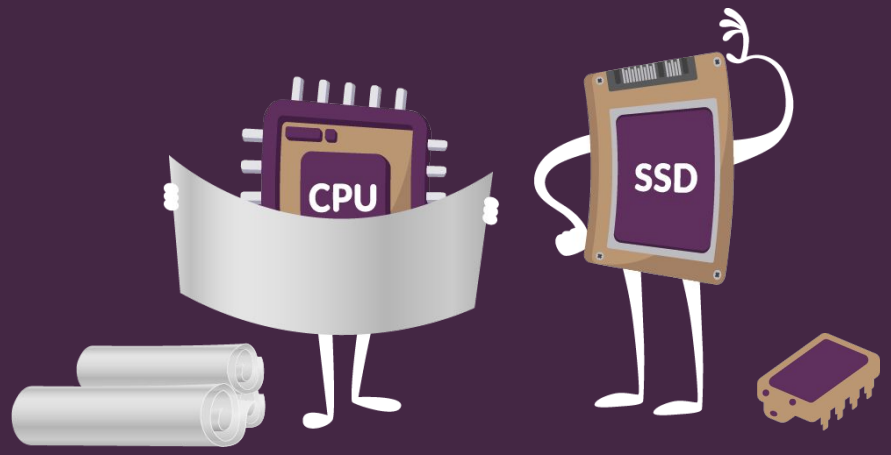

NVMe over TCP Standardization

Standardizing TCP/IP transport binding, adding to NVMe-oF spec alongside RDMA & FC

- Expect standard ratification in 2H 2018
- Supports remote NVMe SSDs with minimal additional latency compared to local SSDs
- Same NVMe model: sub-systems, controllers namespaces, admin queues, data queues

- **Lightbits is leading new TCP/IP transport**
  - Developed pre-standard client available to NVMe members
  - Tested preliminary implementation with several customers & partners
  - Key contributor to standard and Linux upstreaming process

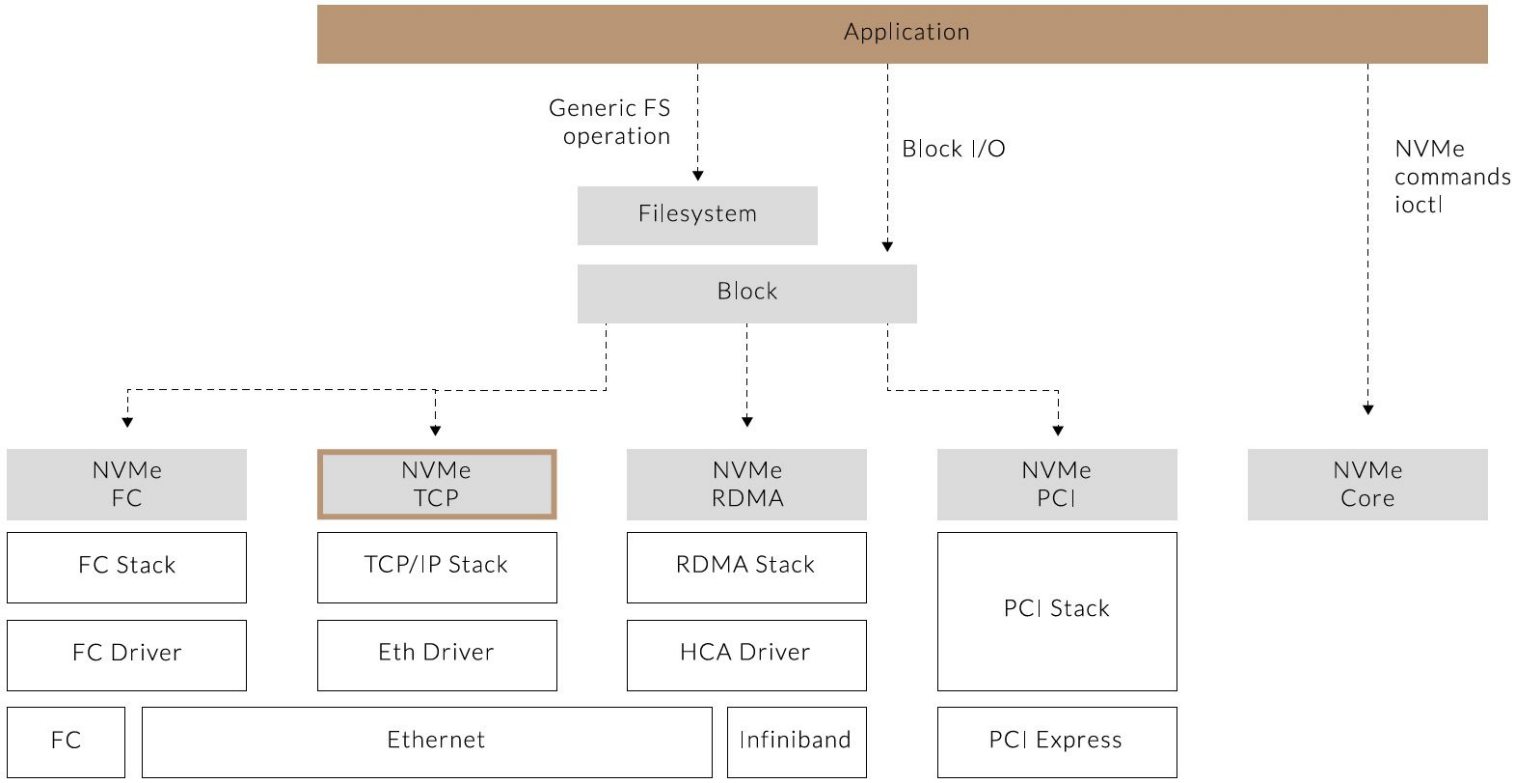

lightbits

Pre-Standard

Pre-Release

Comparing DAS vs. NVMe/TCP

IOPS, average and 99.99% latencies

lightbits

Comparing DAS vs. NVMe/TCP

IOPS, average and 99.99% latencies

Pre-Standard

Pre-Release

Come see the demo live at Intel booth #745!

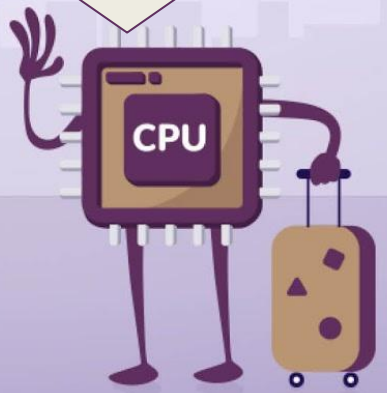

lightbits

NVMe/TCP unleashes disaggregation at DAS performance

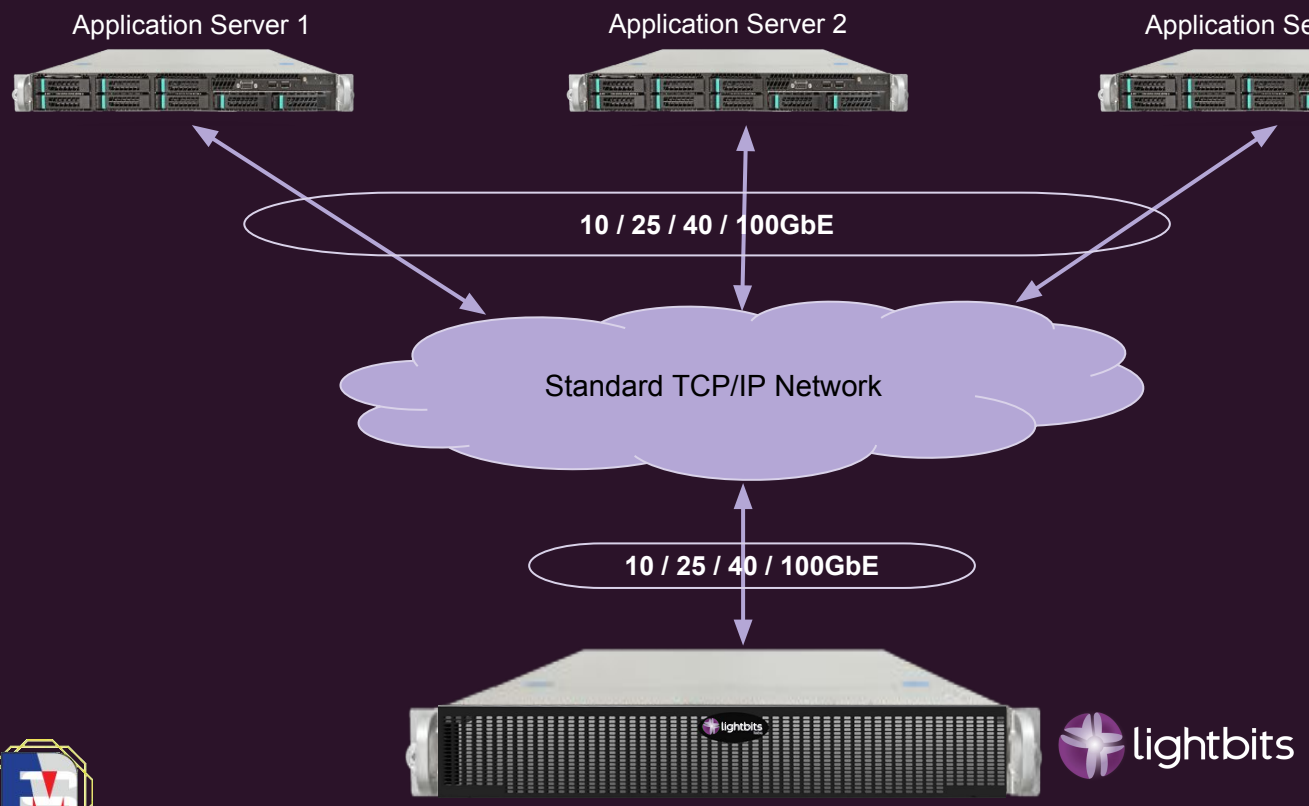

- Hyperscale data center deployment
- Multiple Application Server Live Demonstration
- Low latency and High performance as DAS
- Standard TCP/IP Network Infrastructure
- No modification in Application server Software

- High IOPS for serving many Application Servers
- Thin Provisioned Storage
- In-line Hardware Accelerated Data Reduction at line rate speed
- Optane™ ready - Supports variety of SSD technologies

lightbits

High Performance, Thinly Provisioned, Data Reduction

DAS PERFORMANCE

- Low Latency on NVMe/TCP
- Scalable IOPS for serving many application servers
- In-line acceleration

THIN PROVISIONING

- Provision per use only
- Buy only what you need
- Consume flash at the time of writes

DATA REDUCTION

- Maximize Flash Utilization
- No performance compromise

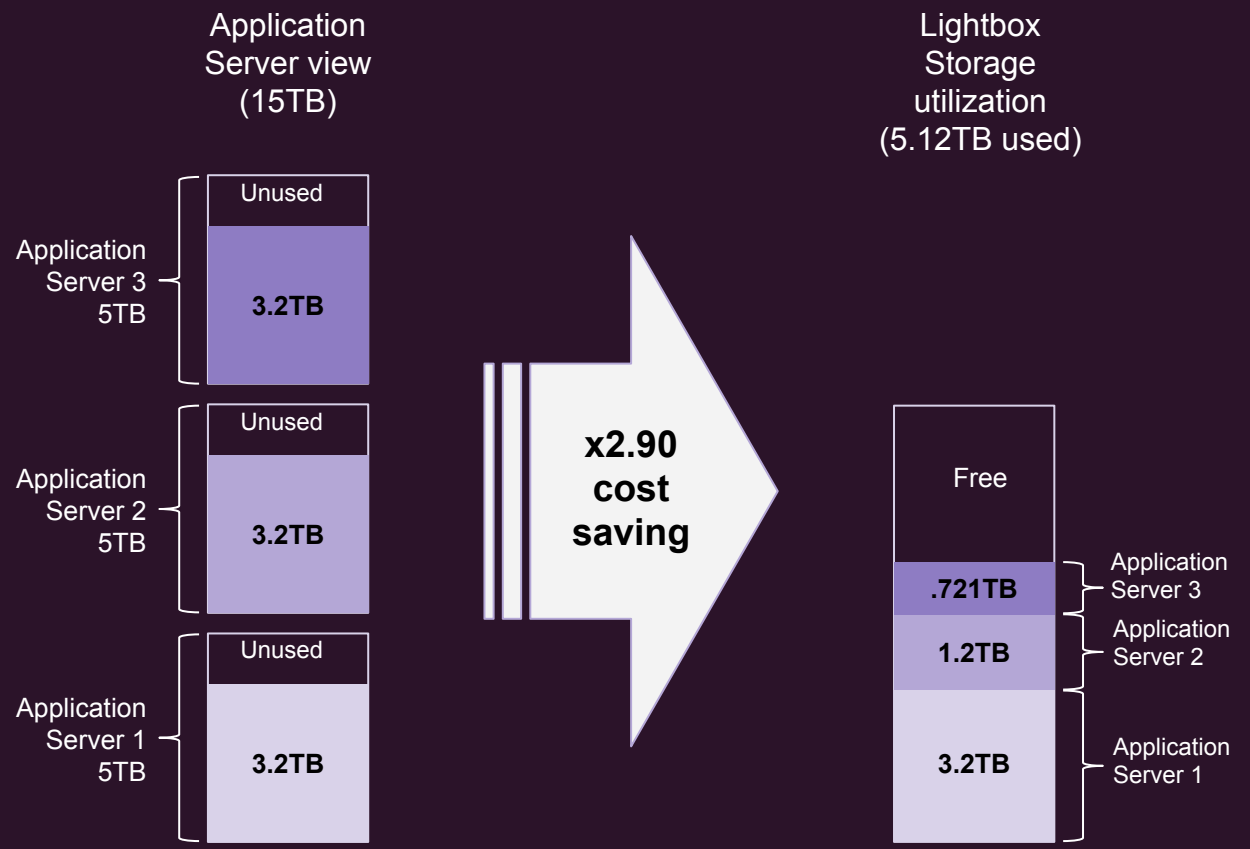

The Demo Dashboard

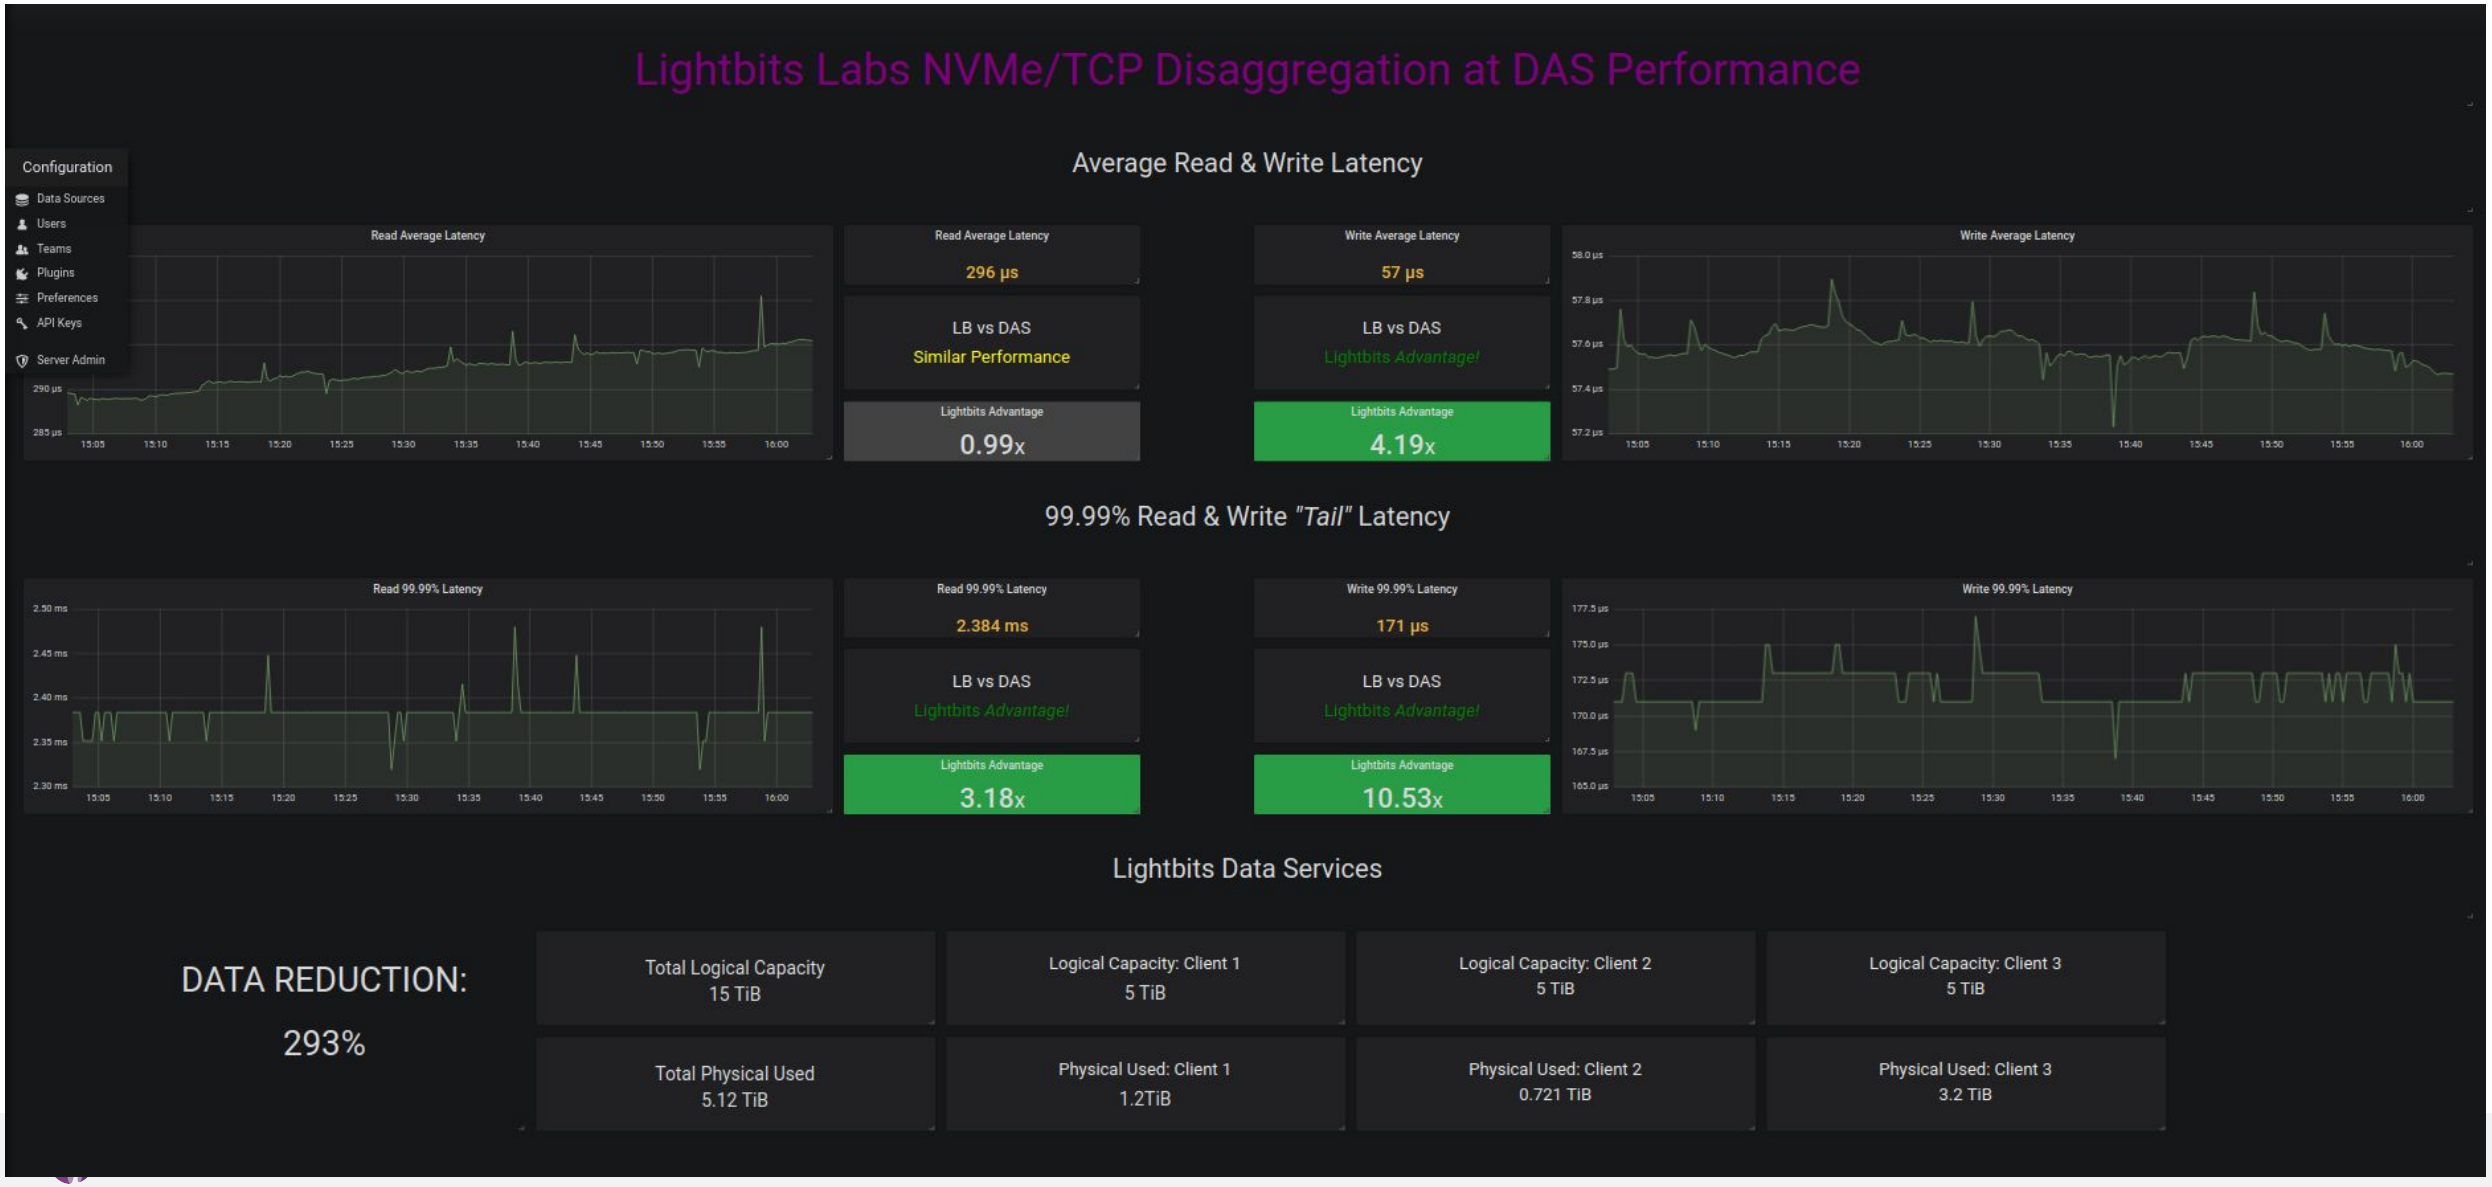

NVMe/TCP Standard

Lightbits NVMe/TCP disaggregation party

Lightbits NVMe/TCP technology enables hyperscalers to move from inefficient direct-attached storage to a shared flash model where compute and storage are scaled independently

NVMe/TCP is an open standard that enables flash disaggregation without compromising performance and without requiring any changes to compute clients or networking infrastructure.

Flash Memory Summit
Flash Memory Summit 2018
Santa Clara, CA

Thank you!