G2M RESEARCH

G2M Research Multi-Vendor Webinar #1: NVMe Scale-Out Storage Best Practices

- July 24, 2018

Sponsored By:

Webinar Agenda

9:00-9:05	Ground Rules and Webinar Topic Introduction (G2M Research)
9:06-9:20	Sponsoring Vendor presentations on topic (5 minute each)
9:21-9:22	Audience Survey #1 (2 minutes)
9:23-9:36	Key Question #1 (2-minute question; 4 minutes response per vendor)
9:37-9:50	Key Question #2 (2-minute question; 4 minutes response per vendor)
9:51-9:52	Audience Survey #2 (2 minutes)
9:53-10:06	Key Question #3 (2-minute question; 4 minutes response per vendor)
10:07-10:14	Audience Q&A (8 minutes)
10:14-10:15	Wrap-Up

Panelists

Josh Goldenhaar
VP Products
Excelero
www.excelero.com

Scott Schweitzer
Director, Technology Evangelist
Solarflare Communications
www.solarflare.com

Dan Liddle
VP of Marketing
Newisys
www.newisys.com

Host/Emcee:

Mike Heumann
Managing Partner
G2M Research
www.g2minc.com

Storage Arrays vs Storage Appliances

- What is a storage array?
  - An integrated system sold by the manufacturer as a complete solution
  - Entire solution is supported directly by the manufacturer (“one throat to choke”)
- What is a storage appliance?
  - An industry-standard server that runs 3rd party storage management software
  - Often integrated by resellers or SIs
  - Reseller/SI (with help from storage management software vendor) provides first line of support
  - Was often thought of as “the poor man’s storage solution” or “just better than DAS”

How Have Flash, NVMe, and 100Gb Ethernet Changed the Storage Landscape?

- Flash has significantly reduced the storage media latency
  - Reduced the need for huge RAM buffers, and eliminated the need to short-stroke hard disk drives
  - NVMe has further accelerated this trend
  - Does this reduce the importance of array controllers?
- Increases in Ethernet speed has impacted the advantages of Fibre Channel for SANs
  - 100GbE plus RoCE reduces latency and jitter
- Result: More storage array controllers are now based on X86 architectures, and less on ASICs, FPGAs, or other custom hardware

vs

nvm EXPRESS®

25 50 100 ALLIANCE

x86

G2M RESEARCH

Why Are These Factors Important?

- ***Cost/GB***
  - Storage appliances have been cheaper ($/GB), trading lower cost for lower performance/size
  - Flash and “thousands of node” scale-out software have largely eliminated those tradeoffs

- ***Data Management and Protection***
  - Backup/recovery of complex apps has always been difficult, even for expensive arrays
  - Apps (mostly) take care of this problem today

- ***Support***
  - As array vendors move to standard server H/W, they become S/W vendors – little advantage over appliances

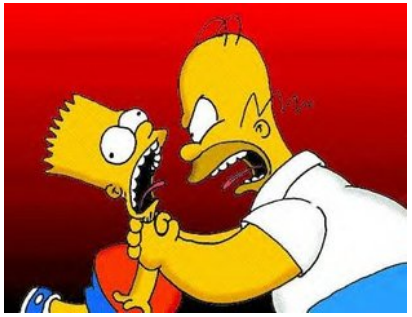

- Why pay significantly more to get nearly the same performance and scalability?

G2M RESEARCH

Excelero Introduction

- Josh Goldenhaar, VP Products

Excelero

Scale-out Server SAN

Software-Defined Storage

What is NVMesh server SAN?

NVMesh allows unmodified applications to utilize pooled NVMe storage devices across a network at local speeds and latencies.

Distributed NVMe storage resources are pooled with the ability to create arbitrary, dynamic block volumes that can be utilized by any host running the NVMesh block client.

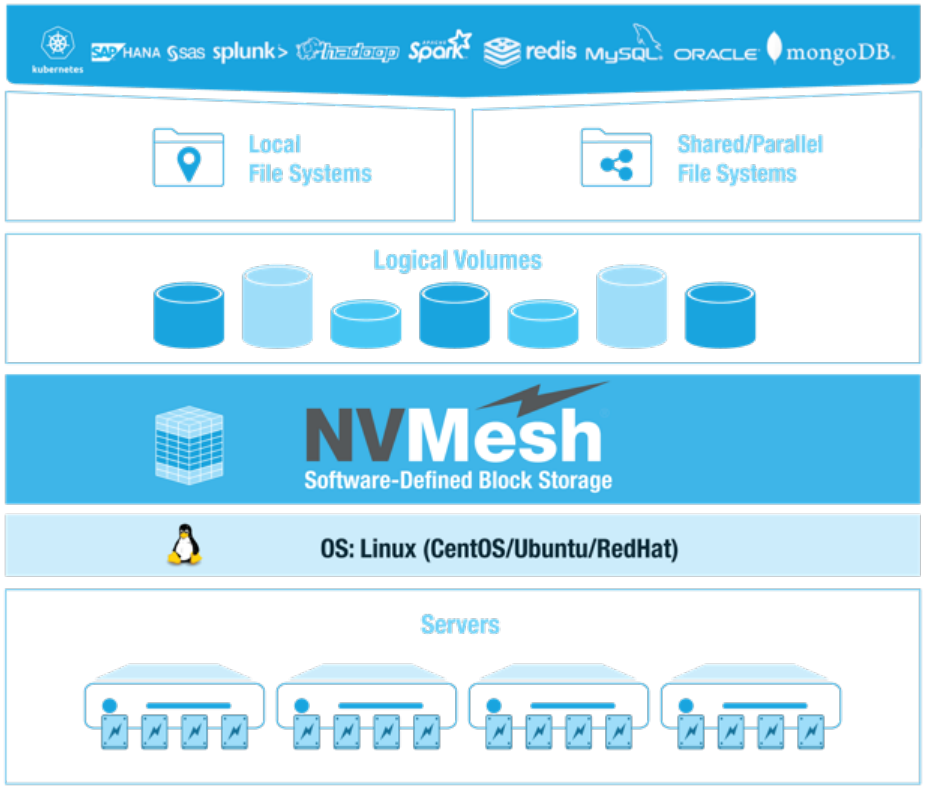

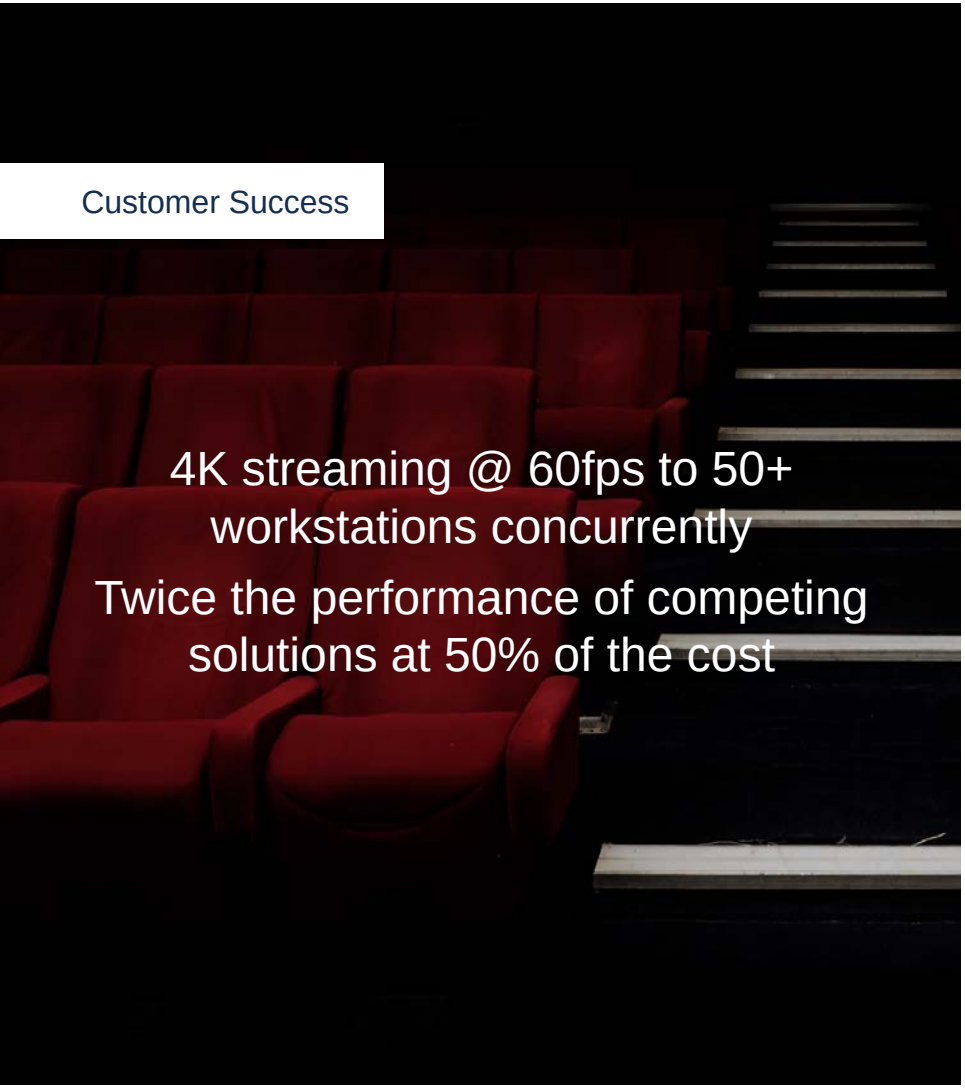

technicolor

Any-K Storage

Use Case

- Any-K production & post-production
- Streaming, editing and reel presentation to simultaneously active workstations

Problem

- 4K & 8K video is changing requirements for storage performance and capacity
- Workstations need high BW & low latency performance for streaming, editing and reel presentation
- Customers need future-proof solutions that will support 10k

Solution

With just 6 standard servers powered by 3D NAND, NVMesh enabled 4K streaming @ 60fps to 50+ workstations concurrently

The solution gives twice the performance of competing solutions at 50% of the cost and requires a much lower data center footprint

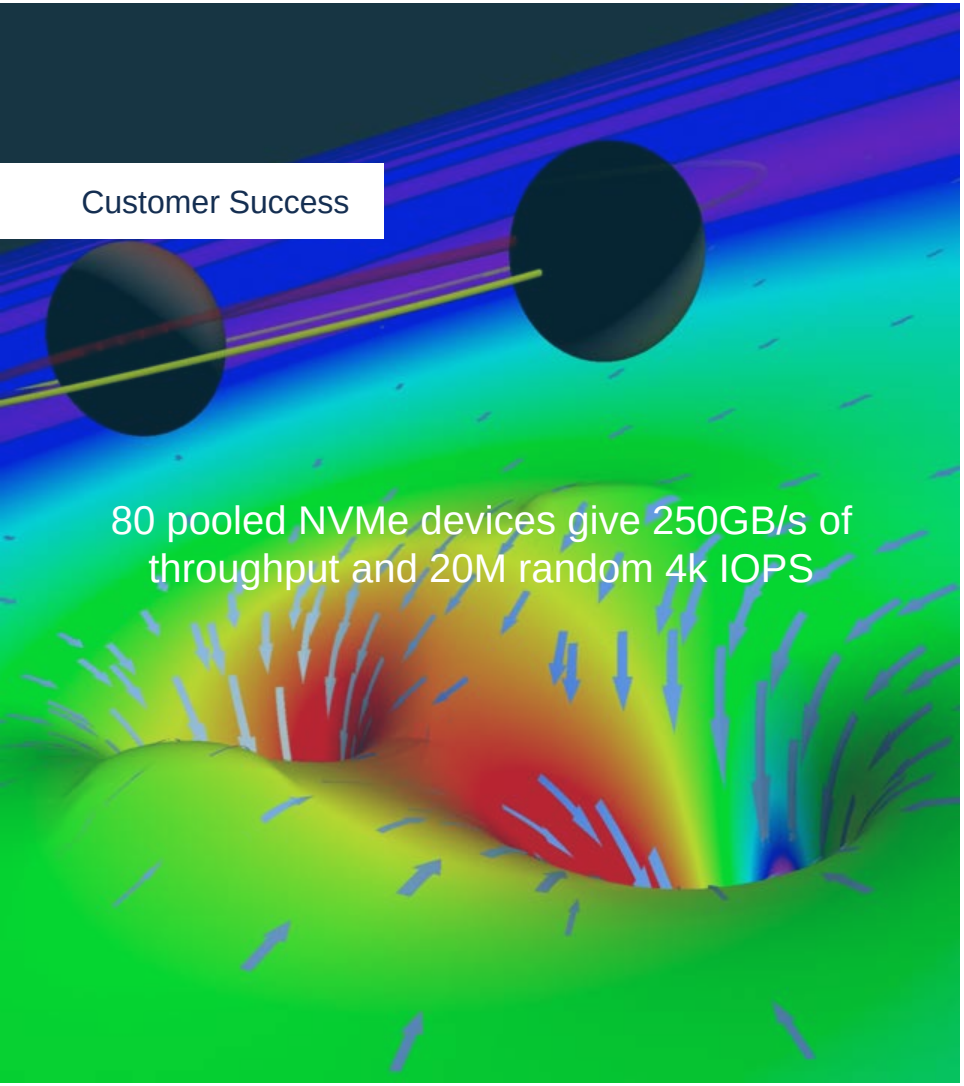

SciNet compute canada | calcul canada

Pooling NVMe Within GPFS NSDs enables new Science use cases

Use Case

- Large-scale modeling, simulation, analysis and visualization
- Visualizes supercomputer simulation data on 100s of compute nodes

Problem

- Finish check pointing faster and start running the job.
- Achieve performance of 250GB/Sec at the lowest price

Lenovo Mellanox TECHNOLOGIES Excelero

Solution

unified pool of high-performance flash distributed retaining the speeds and latencies of directly-attached media

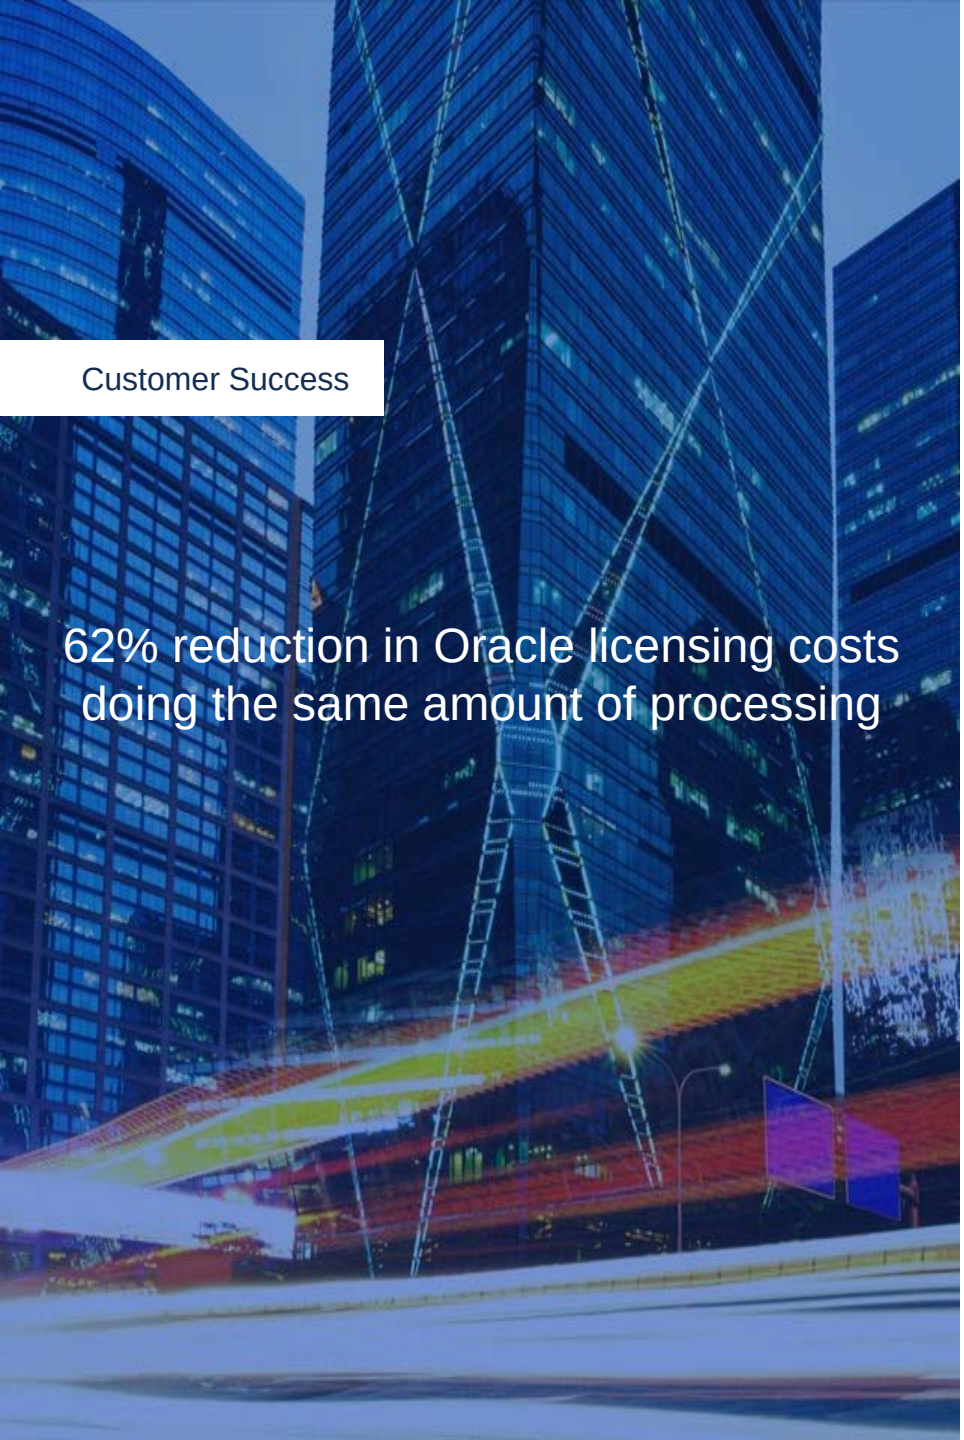

Customer Success

62% reduction in Oracle licensing costs doing the same amount of processing

CMA

High throughput & low latency for Oracle databases

Use Case

- Cluster performance is a key priority for large Oracle environments, especially for customers with tables that size up to billions rows and serve thousands of users driving analytics against that data

Problem

- Traditional AFA storage solutions give latencies of slightly less than 1 millisecond; customers often get thousands of queries a day and some of these can take over an hour to run.

Solution

NVMesh offers latencies of single digit microseconds. This significantly speeds up database scans and queries

Running Oracle database on NVMesh enables customers to reduce the number of licenses needed for their clusters.

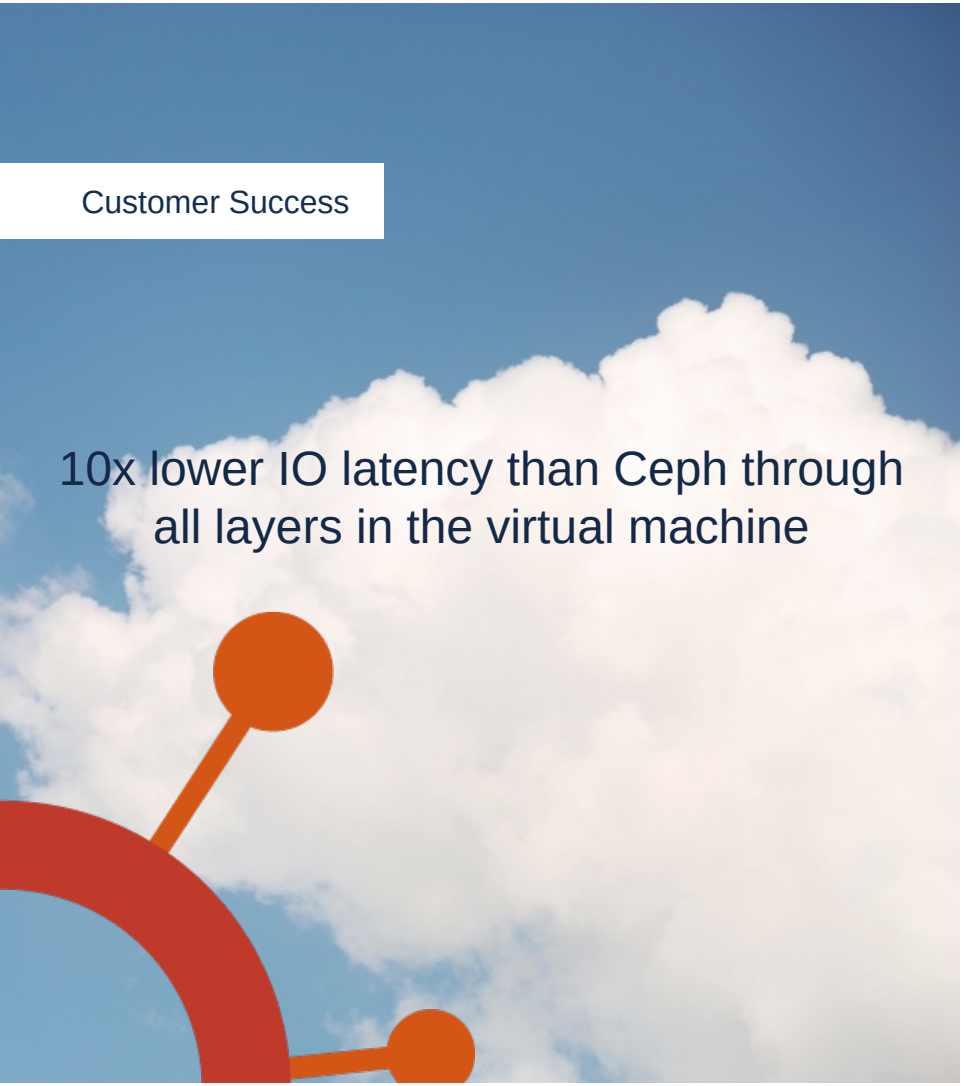

teuto.net

Exceptional performance with low-latency block storage for OpenStack

Use Case

- New workloads such as machine learning and artificial intelligence are creating new opportunities for service providers but also bring unseen performance requirements for OpenStack Clouds.

Problem

- Ceph is a popular storage solution for OpenStack, but the layered architecture has a big impact on the performance and makes the solution less attractive for new low-latency workloads.

Solution

TeutoStack Cloud runs NVMesh on SuperMicro servers from Boston Limited along with Mellanox ConnectX-4 and ConnectX-5 NICs running 2x25Gb Ethernet to each client node. It uses OpenStack Cinder to manage storage though the NVMesh Cinder driver.

Excelero

Thank you!

G2M RESEARCH

Solarflare Introduction

- Scott Schweitzer, Technology Evangelist

NVMe-oF Options – TCP is Lowest Cost; Ubiquitous; Scalable

Fabric Type	Latency	Lossless	Deployability	Scalability	Deployment Footprint	Cost
Ethernet (TCP)	**Low**	**Yes**	**No Forklift (Extends current infrastructure)**	**High**	**Ubiquitous (98%)**	**Lowest**
Fibre Channel	Low	Yes	No Forklift (Extends current infrastructure)	Medium	Specialized Storage (FORTUNE 500)	High
Ethernet (RoCE- RDMA)	Low	Yes *(Config option)*	Forklift Upgrade (DCB/Converged Ethernet)	Medium	Negligible* (2%)	Medium
Infiniband	Very Low	Yes	Special Network	Limited	Limited (Scientific HPC)	High

Large & Fast Growing Market: NVMe-oF TCP offers the most attractive Solution

* – Microsoft Azure Cloud

SOLARFLARE®

NVMe™ Feature Roadmap

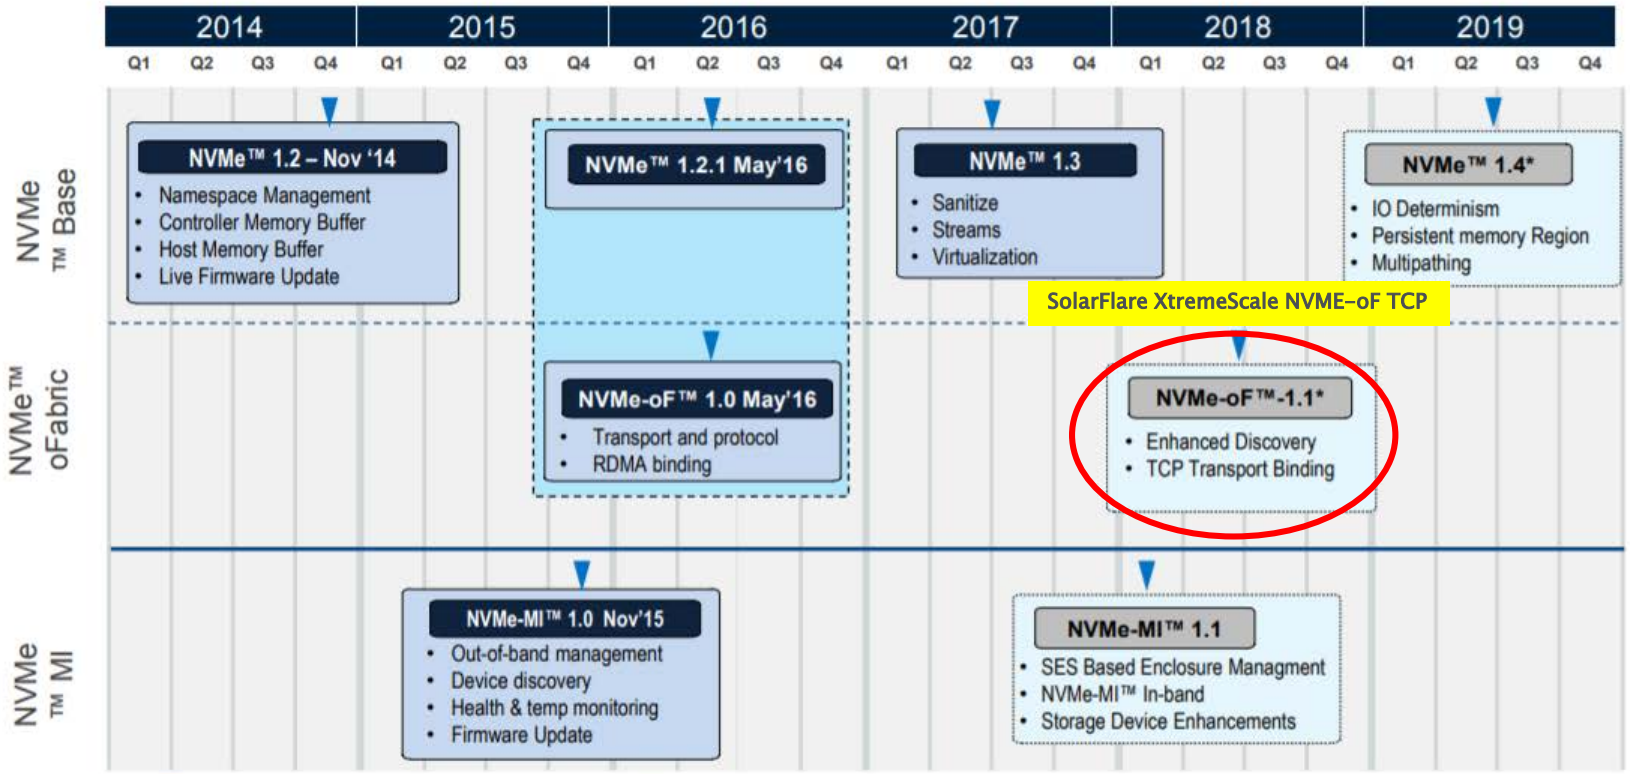

7

NVMe-oF TCP Equivalent To NVMe-oF RoCE

Benchmarking with Tier 1 OEM (Pre-standard drivers)

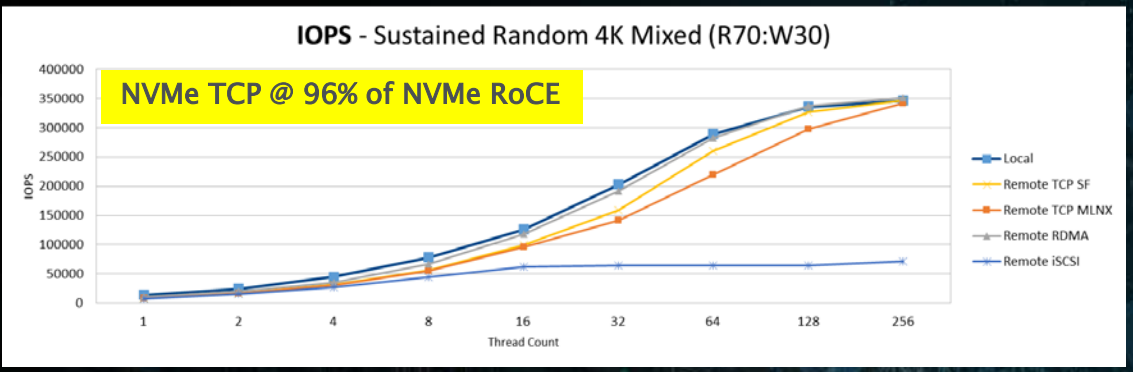

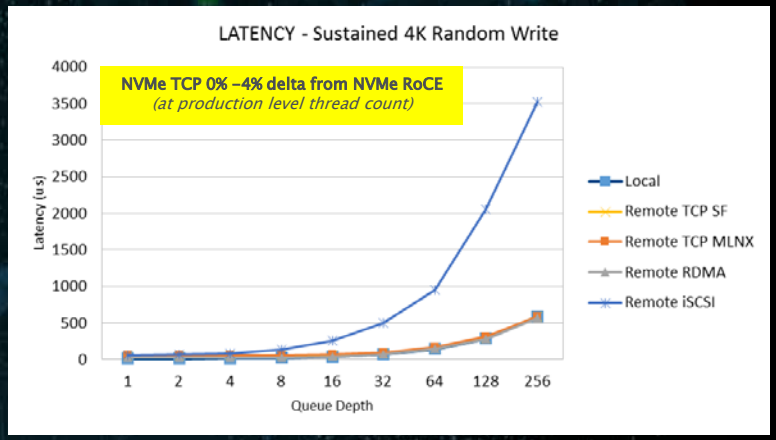

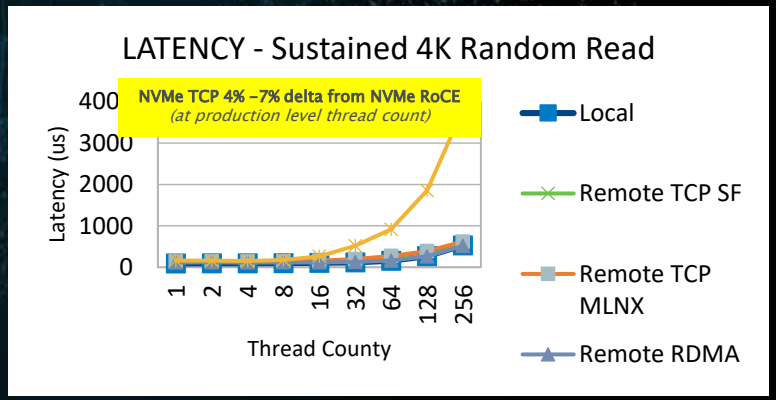

NVMe-oF TCP: Solarflare Engagement To-Date

	Q1'17	Q2'17	Q3'17	Q4'17	Q1'18
NVM Express		• Org Membership • TCP working group contributor			
Linux Community	• Solarflare submits NVMe-TCP plug-in RFC				
OEM			• Demonstrated TCP as a viable transport with Tier-1 OEM#1		• Validated enhanced code with Tier 1 OEM#1 • Validated enhanced code with Tier 1 OEM#2
Solarflare Internal	• PoC Code	• Feature add • Debug	• Scalability • Port to RHEL 7.4	• Scalability • Port to RHEL 7.4	

Summary

- Solarflare NVMe-oF TCP delivers equivalent performance to NVMe-oF RDMA without requiring a “forklift” infrastructure upgrade to Converged Ethernet adapters and switches— also known as Data Center Bridging.
- Workloads such as OLTP, real-time analytics and file sharing benefit from NVMe architecture.
- **Solarflare Linux kernel drivers for initiator and target mode NVMe-oF TCP are available now for PoCs.**

G2M RESEARCH

Newisys Introduction

- Dan Liddle, VP of Marketing

SANMINA

Communications • Computing & Storage • Medical Systems • Defense & Aerospace • Multimedia • Clean Technology • Industrial • Automotive

NEWISYS®

What we make, makes a difference™

NEWISYS Background

SANMINA

Newisys is wholly owned by Sanmina Corporation

- Newisys has been designing high performance/high-density server and storage for 25+ years for OEMs and CSPs (Cloud Service Provider)

Award Winning Compute and Storage Platforms

- All-Flash & Hybrid Storage and Servers
- NVMe Flash-based Storage, Servers and NVMeoF Platforms
- Primary & Cold Storage Servers and Expansion Arrays

Key Market "Firsts"

- 4RU: x60, x84, x96 bay JBODs;
- 2RU: x60 bay storage server;
- NVMe PCIe JBOF; NVMe-over-Fabric storage

Focus on product innovation that delivers

- Performance, capacity, cost leadership, power/cooling footprint

15 Year Track-record of innovation and technology excellence

NVMe Scale-out Storage Best Practices

SANMINA

This is an ongoing & evolving market so the roadmap matters

- Rapidly changing as new applications are enabled by these tremendous latency, throughput, and overall performance advances
- Still maturing and adding services and functionality
- Prices coming down over time, further expanding uses, which would have previously been cost-prohibitive

Offering a broad-range of NVMe platforms really matters

- Supports numerous third-party, purpose-built software applications
- Allows alignment of the underlying hardware platform configurations with applications for the greatest TCO efficiency
- Fewer NVMe vendors to manage

Multi-vendor interoperability is paramount (e.g., drives, cards, software, etc.)

- Provides performance tailoring and fine-tuning based on desired characteristics
- Avoids vendor lock-in, reduces system cost, and provides supply flexibility

NVM Express® (NVMe™) Platform Overview

SANMINA

NSS1160G-2N (1U8/16)
NVMe Server

NDS-2244 (2U24)
NVMe High-performance JBOF

NSS2247G-2N/4N (2U24)
High-performance Quad Server

NDS22482F (2U24)
NVMe over Fabrics JBOF

UXXXXX
U.2 NVMe SSD M.2 Carrier

NSSXXXX (2U56)
NVMe High-performance Dual Server

Achieving record breaking performance, new levels of density with a superior high availability design

Industry's broadest NVMe product portfolio

NEWISYS

SANMINA

Newisys Division of Sanmina

Addressing the storage needs of business

Communications • Computing & Storage • Medical Systems • Defense & Aerospace • Multimedia • Clean Technology • Industrial • Automotive

NEWISYS®

What we make, makes a difference™

Panel Question #1

- For many people, scale-out flash storage (SOFS) is synonymous with software-defined storage, and those people equate today's SOFS solutions with those from ten years ago. How would you say that today's offerings differ from those of ten years ago, from the standpoint of the user?
  - Excelero
  - Solarflare
  - Newisys

Audience Survey Question #1

- What would be the key factors for you to consider deploying an SOFS solution in your organization versus an all-flash array? (check all that apply):
  - Cost: %
  - Scalability: %
  - Performance: %
  - Avoiding Vendor Lock-In: %
  - Ability to mix heterogeneous hardware in a single solution: %
  - Ability to get effective support: %

Panel Question #2

- The term “scale-out” is being used by everyone in our industry today, from cloud service providers and software as a service providers (“scale infinitely”) to networking vendors, server vendors, and application vendors. What does “scale-out” mean for storage solutions, and what are the key factors that need to be considered when developing SOFS solutions?
  - Solarflare
  - Newisys
  - Excelero

Audience Survey Question #2

- How important do you think that support for NVMe flash media is in an SOFS solution? (check one):
  - Extremely important – NVMe will become ubiquitous in enterprise SOFS solutions: %
  - Important – NVMe flash-based arrays will be used in performance-focused SOFS solutions: %
  - Not Important – Most storage users can't take advantage of flash performance, let alone NVMe flash performance: %
  - Unsure – I don't have the information to answer: %
  - No Opinion: %

Panel Question #3

- What do you see as the “killer apps” for SOFS solutions, and what are the best practices for maximizing the value of SOFS for these use cases?
  - Newisys
  - Excelero
  - Solarflare

Audience Q&A

G2M RESEARCH

Thank You For Attending

G2M

RESEARCH