

G2M Research NVMe over ▶ Fabrics (NVMe-oF) Webinar

December 12, 2017

▶ Webinar Agenda

- 9:00-9:02 Introductory Remarks, Introduction of Panelists (Mike Heumann)
- 9:02-9:07 What is NVMe-oF, and Why is it Important (Howard Marks)
- 9:07-9:20 Panelist Question #1 (Howard plus panelists)
- 9:20-9:22 Survey Question #1 (Mike)
- 9:22-9:35 Panelist Question #2 (Howard plus panelists)
- 9:35-9:48 Panelist Question #3 (Howard plus panelists)
- 9:48-9:50 Survey Question #2 (Mike)
- 9:50-10:03 Panelist Question #4 (Howard plus panelists)
- 10:03-10:12 Audience Q&A
- 10:12-10:15 Closing remarks (Mike)

Panelists



Rob Davis
VP of Storage Technology
Mellanox Technologies

John Kim
Director of Storage Mktg
Mellanox Technologies




Eliot Rosen
Product Marketing Director
Broadcom

Frankie Fan
System Architect
Broadcom



Taufik Ma
CEO
Attala Systems

Bryan Cowger
VP of Sales/Marketing
Kazan Networks



Ahmet Houssein
VP Marketing
Solarflare Communications

Praveen Midha
Director of SW Engineering
Cavium



Hosts:

Howard Marks
DeepStorage, LLC

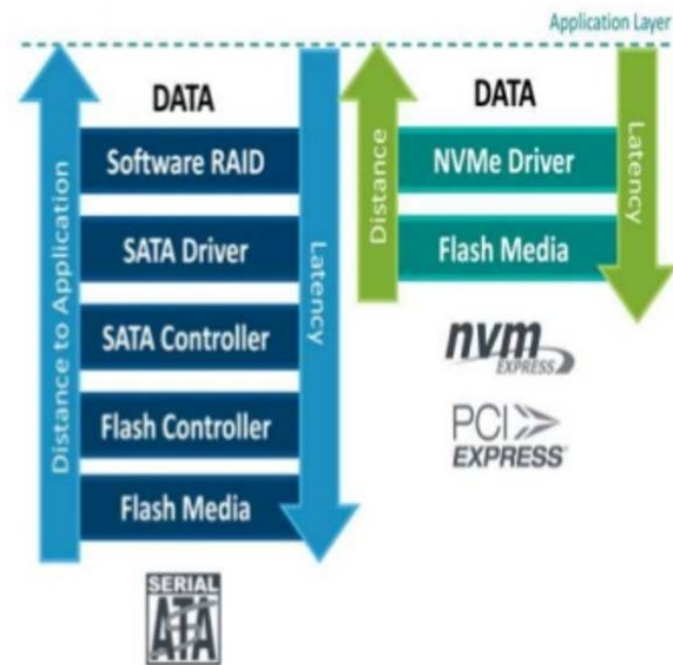


Mike Heumann
Managing Partner
G2M Research




▶ What is NVMe?

- ▶ A new software protocol for non-volatile memory access
- ▶ Gen1 and 2 PCI SSDs
 - ACHI (SATA command set)
 - Proprietary (Fusion-IO, Verident) with heavy software
- ▶ Lower compute overhead than SCSI
- ▶ 64K queues of 64K entries vs SCSI 1 queue of 32 entries




▶ NVMe = Lower Overhead & Latency



- ▶ NVMe replacing SAS/SATA from desktops (M.2) to the data center (U.2)

▶ NVMe Over Fabrics (NVMe-oF)

- ▶ Extends/encapsulates NVMe semantics over
 - Ethernet with RMDA
 - ROCE – RDMA over Converged Ethernet
 - iWARP – RDMA over TCP
 - Fibre Channel
 - Infiniband
 - TCP/IP in draft
- ▶ Adds name spaces and discovery
- ▶ 10-50 μ sec protocol and network overhead



▶ NVMe-oF Spreads Like The Flash

- ▶ iSCSI took 3-4 years from spec to Equallogic/Lefthand
- ▶ Now 18 months since NVMe-oF spec
- ▶ Mainline array vendors announce support:
 - Pure FlashArray//X
 - Kaminario Flex
 - HPE 3PAR
- ▶ Software based NVMe-oF killed DSSD's hardware based approach



▶ Panel Question #1

- ▶ Where is NVMe-oF being deployed today?
 - Mellanox
 - Cavium
 - Solarflare
 - Broadcom
 - Attala
 - Kazan

Audience Survey Question #1

- ▶ What are your organization's plans to deploy NVMe-oF?
(check one) → 71 responses
- Have already deployed NVMe-oF in our organization: 10%
 - Plan on deploying NVMe-oF in 2018-2020: 28%
 - Plan on deploying NVMe-oF, but after 2020: 6%
 - No current plans to deploy NVMe-oF in the next 5 years: 8%
 - Unsure of our plans to deploy NVMe-oF: 48%

▶ Panel Question #2

- ▶ What are the factors preventing or enabling the deployment of NVMe-oF as a replacement of existing storage protocols?
 - Attala
 - Broadcom
 - Kazan
 - Cavium
 - Mellanox
 - Solarflare

▶ Panel Question #3

- ▶ What do you think the primary applications will be that utilize NVMe-oF? When do you expect to see widespread deployment of these applications on NVMe-oF?
 - Kazan
 - Cavium
 - Attala
 - Solarflare
 - Broadcom
 - Mellanox

Audience Survey Question #2

- ▶ What are your primary concerns when considering deployment of NVMe-oF? (check all that apply → 51 responses)
- Don't need NVMe-oF's performance in the near future: 4%
 - Cost of NVMe-oF vs current storage networking solutions: 33%
 - Unsure of what use cases NVMe-oF makes sense for: 25%
 - NVMe-oF technology immaturity (standards, interoperability): 47%
 - NVMe-oF ecosystem immaturity (arrays, switches, management applications, etc.): 51%
 - Other: 14%

▶ Panel Question #4

- ▶ What is your perspective on NVMe-oTCP? How do you expect it's use case to differ (if at all) from that of RDMA-based NVMe-oF transports?
 - Solarflare
 - Mellanox
 - Broadcom
 - Attala
 - Kazan
 - Cavium

► Audience Q&A



► Questions from the Audience

- ▶ *Does RoCE use UDP as transport. if yes, how is data integrity addressed?*
 - It is provided by flow control/data integrity mechanisms within RoCE.
- ▶ *Do you just expect to see NVMe-oF being deployed with the higher end storage arrays, or do you expect it in the lower end arrays too?*
 - Expect it to show up in all arrays, as well as JBOFs.
- ▶ *To Solarflare presenter - you talked about NVMe-oTCP but finished with NVMe-oL2(Eth)?*
 - NVMe-oEthernet was meant to mean NVMe-oTCP.
- ▶ *Note: DC-TCP != vanilla TCP. Since with NVMe, storage is (finally) faster than networking, in-cast issues might get to a completely different level of magnitude.*
 - This has always been the challenge with networked storage (in one direction or another).



Thank You For Attending

- ▶ You will be sent a link to the recorded webinar within 48 hours

G2M
RESEARCH

DeepStorage

