

### G2M Research Multi-Vendor Webinar: Implementing NVMe<sup>™</sup> and NVMe-oF<sup>™</sup> for Cloud Service Providers



Thursday Nov 19, 2020



### Western Digital.

### Webinar Agenda

- **9:00-9:05** Ground Rules and Webinar Topic Introduction (G2M Research)
- **9:06-9:37** Sponsoring Vendor presentations on topic (8 minute each)
- **9:38-9:46** Key Question 1 (1-minute question; 2 minutes response per vendor)
- **9:47-9:47** Audience Survey 1 (1 minute)
- **9:48-9:56** Key Question 2 (1-minute question; 2 minutes response per vendor)
- **9:57-9:57** Audience Survey 2 (1 minutes)
- **9:58-10:06** Key Question 3 (1-minute question; 2 minutes response per vendor)
- **10:07-10:18** Audience Q&A (12 minutes)
- 10:19-10:20 Wrap-Up



### G2M Research Introduction and Ground Rules

Mike Heumann

Managing Partner, G2M Research





### **KIOXIA** Joel Dedrick

VP, General Manager www.kioxia.com





www.lightbitslabs.com



G2M Mike Heumann Managing Partner www.g2minc.com



#### Western Digital.

Mark Miquelon Director, Platform Partner Alliances <u>www.wdc.com</u>



### CSP Storage, Networking, and Compute Options – A Wide Variety of Choices

- CSPs focus on the delivery of one or more closely-coupled applications to an extremely large set of users.
- Most CSPs start out utilizing the cloud to host their business and data
  - Easy to scale up (or down), add resources
  - And no up-front investment in servers, etc.
- BUT as most CSPs scale, they start to explore building their own datacenters







### Why is The Cloud Problematic for CSPs?

- CSPs typically run complex applications across clusters of thousands (or more) servers
  - Usage patterns can/often are "bursty"
  - Responsiveness is very critical to user satisfaction
- <u>Cost</u>: Cloud compute/storage costs can be difficult to predict, often depend on hard-to-control factors
- <u>Performance</u>: Massively parallel applications are highly sensitive to network changes, VM migration, which are unpredictable in the cloud
- A CSP-owned datacenter solves these issues







### What Does a CSP Need In A Datacenter?

- <u>Performance</u>: Predictable, consistent performance is more important than peak performance
- <u>Efficient Implementation</u>: The most effective hyperscalers utilize custom HW solutions. OCP makes efficient HW solutions available to CSPs
- <u>Density</u>: Like efficiency, density impacts costs real estate is not cheap.
- <u>Cutting-Edge Technology</u>: The only way to stay ahead of the competition (and there is ALWAYS competition)







### Lightbits Labs

Josh Goldenhar Vice President, Product Marketing <u>www.lightbitslabs.com</u>

#### CSP Flash Deployment – 1 or More SSDs per server



NoSQL, In-memory, Distributed, Kafka, HPC workloads

#### They All Need:

- Low latency and high bandwidth
- Consistent response time no noisy neighbors
- Local flash (NVMe)

#### They All Suffer:

- Poor flash utilization
- No ability to share data
- No data protection when needed
- Applications tied to servers



#### **CSPs: Severe Flash Under-Utilization**



### 50-95% of flash is wasted!



### Customer Applications with Lightbits' Block NVMe/TCP



Higher uptime and utilization, faster recovery = lower TCO

#### Benefits:

- Local flash block performance
- Applications unaffected by drive failures
- Block volumes: any size and thin provisioned
- Volumes can be shared
- Server failure:
  - Instances launched anywhere
  - Recovery: seconds to minutes
  - Higher uptime, increased



#### Lightbits Central Values

- FLETIBILITY PERFORMAN **High Bandwidth** High IOPs ight**OS**™ Low latency ~ LOWER TCO
  - Standard TCP/IP networks
    - Rich data services
      - Varied configurations •

- Optimized for QLC NAND SSD
- Improved flash utilization
- Independent scaling of storage •

V

•

•

#### High Performance Software Defined Storage

Standard servers, NICs and SSDs, optional hardware accelerator



#### Flexibility: Ease of Adoption



lightbits

Light**OS**<sup>™</sup> CLUSTER

- Failure Domain aware
- Allows for failover between racks/domains
- Intelligent replication spans failure domains
- "Blast Radius" (failure domain) definable:
  - O Rack
  - O Row
  - O Power zone



LightOS enables QLC Flash:

- Up to 5 times endurance
- Aggregates writes for higher performance
- Allows for high capacity and high density at low cost while preserving high performance
- Use QLC where you could not locally

#### Reduce TCO





#### Cloud Native Applications with Lightbits' Block NVMe/TCP





#### Summary - Lightbits Differentiation

Utilize disaggregated NVMe SSDs as a remote low-latency pool at local flash speeds





## KIOXIA

### Kioxia

Joel Dedrick VP, General Manager, Networked Storage Software <u>www.kioxia.com</u>



### **G2M Webinar**

Nov 19, 2020





#### **Clearing Up some Confusion**



"CSPs" (Cloud Service Providers) come in a number of flavors:

#### (1) laaS, PaaS, e.g.:

- Amazon Web Services<sup>™</sup>
- Google Cloud Platform<sup>™</sup> service
- Microsoft Azure<sup>™</sup>

#### (2) B2B SaaS, e.g.:

- Salesforce.com<sup>™</sup>
- ServiceNow<sup>®</sup>

#### (3) B2C SaaS, e.g.:

- eBay<sup>™</sup>
- Uber<sup>™</sup>

#### Today, I'll focus on (1)

"Amazon Web Services" is a trademark of Amazon.com, Inc. or its affiliates in the United States and/or other countries. "Google Cloud Platform service" is a trademark of Google Inc. or its affiliates in the United States and/or other countries. "Microsoft Azure" is a trademark of Microsoft Inc. or its affiliates in the United States and/or other countries. "Salesforce.com" is a trademark of Salesforce.com, inc. or its affiliates in the United States and/or other countries. "ServiceNow" is a trademark of ServiceNow, Inc. or its affiliates in the United States and/or other countries. "eBay" is a trademark of eBay, Inc. or its affiliates in the United States and/or other countries. "Uber" is a trademark of Uber Technologies, Inc. or its affiliates in the United States and/or other countries.

#### **GCP<sup>™</sup> Storage Offerings:**



"Zonal/Regional SSD" (high performance block storage)<sup>1</sup>:

- Up to 100k IOPS, 1200 MB/S per instance
- 2 replicas, separate zones/regions
- \$340/TB/Month

"Filestore" (distributed filesystem)<sup>2</sup>:

- 15k/5k IOPS, 500/110 MB/S per client
- "Cloud Storage" (object store) <sup>3</sup>:
- + Several database services

<sup>1: &</sup>lt;u>https://cloud.google.com/compute/docs/disks/performance#performance\_factors</u> retrieved 11/17/20

<sup>2:</sup> https://cloud.google.com/filestore/docs/performance retrieved 11/17/20

<sup>3:</sup> https://cloud.google.com/storage/pricing retrieved 11/17/20

#### What Does it Take to Compete?



- ☑ Commodity hardware
- ☑ Cross-availability-zone data protection
- ☑ Multiple abstractions
  - Block
  - Shared File
  - Object
  - ...
- ☑ Automation (small SRE team)
- ☑ Software cost << media cost

#### AND

☑ Performance comparable to local NVMe flash ⇒ NVMe-oF

#### **KumoScale Control Plane Implementation**



Master Node



All KumoScale control services are now containerized, and managed by a private Kubernetes "micro-cluster"

KumoScale storage nodes are now also Kubernetes "worker" nodes

- Three nodes run K8s<sup>™</sup> Master & KumoScale Management services
- The rest can host distributed filesystems, object stores, etc.
- Storage services get *local*, very high performance NVMe volumes

Televisity Log Log DB G DB FEREFError (Log Log DB Man Providence) Log Log Log Communication (Log Log DB Man (Log Log DB ) (Log DB Man (Log Log DB ) (Log DB

Storage Node

# KIOXIA



### Western Digital.

### Western Digital

Mark Miquelon Director, Platforms Partner Alliances <u>www.wdc.com</u>

### Western Digital

### Implementing NVMe<sup>TM</sup> and NVMe-oF<sup>TM</sup> for Cloud Service Providers

Mark Miquelon Director, Platforms Partner Alliances 11.19.2020

### **NVMe and NVMe-oF Benefits**

**NVMe** is purpose-built for performance

- NVMe was created to enable consistently low latency storage access
- Takes advantage of ever-improving PCIe interface speeds
- Server architectures deploy high lane counts for SSD attach
- NVMe-oF brings the advantages of NVMe to the scale offered by a storage fabric
  - Disaggregation of storage from compute
  - Shared access
  - Eliminates PCIe limitations





Latency plo

ĊPU0 CPU1 CPU2



Cle<sup>®</sup> Roadman

### **Architectures Utilizing NVMe-oF**

#### Two ways to put NVMe SSDs onto a Fabric

#### Server Based



#### Limitations

- Server consumed by hosting NVMe SSDs
- Capacity likely out-weighs performance
- PCIe infrastructure tied to CPU roadmap cadence
- Qualified server types likely not well matched to every application you want to run

#### **Fabric Attached Devices**



Device

- Benefits
  - Server Infrastructure (CPU, DRAM, etc.) not used to serve NVMe namespaces

**F30** 

- Disaggregated storage optimizes scaling
- Performance matched to the Fabric
- Can deliver higher Density & lower Power



Data24



# Western Digital.

Western Digital, the Western Digital logo, and OpenFlex are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. The NVMe and NVMe-oF word marks are trademarks of NVM Express, Inc. All other marks are the property of their respective owners.



### Panel Questions and Audience Surveys

### Panel Question # 1

- There are a lot of factors that drive whether a CSP's compute and storage infrastructure should be cloud-based, on-premises, or both. When should a CSP start thinking about this question?
  - Lightbits Labs
  - Kioxia
  - Western Digital



### Audience Survey Question #1

 What is your organization's greatest concern when using cloud computing and storage? (select one):

Overall Cost:	6%
<ul> <li>Cost Predictability:</li> </ul>	17%
Performance:	11%
<ul> <li>Performance Predictability:</li> </ul>	11%
<ul> <li>Security/Data Privacy:</li> </ul>	33%
Other Concerns:	6%
<ul> <li>No concerns/no opinion:</li> </ul>	17%



### Panel Question #2

- What are some of the biggest hurdles/traps that CSPs should look for when considering implementing their own datacenter infrastructure?
  - Lightbits Labs
  - Kioxia
  - Western Digital



### Audience Survey Question #2

- When looking at building your own datacenter, what is your greatest concern? (select one):
  - CapEx/upfront costs:
    Scalability:
    Avoiding vendor lock-in:
    Hiring needed skillsets:
    Achieving the ROI expected:
    Other:
    No opinion:



### Panel Question # 3

- NVMe<sup>™</sup> and NVMe-oF<sup>™</sup> have helped accelerate a variety of applications. How can CSPs utilize these technologies to accelerate their application offerings?
  - Kioxia
  - Western Digital
  - Lightbits Labs



### Audience Q&A







### Thank You For Attending

