

G2M Research Multi-Vendor Webinar: Composable Infrastructure – What is It, and Why is It Important?

Tuesday July 30, 2019



Webinar Agenda

- **9:00-9:05** Ground Rules and Webinar Topic Introduction (G2M Research)
- **9:06-9:26** Sponsoring Vendor presentations on topic (7 minute each)
- **9:27-9:28** Audience Survey 1 (2 minutes)
- **9:29-9:39** Key Question 1 (2-minute question; 2 minutes response per vendor)
- **9:40-9:41** Audience Survey 2 (2 minutes)
- **9:42-9:52** Key Question 2 (2-minute question; 3 minutes response per vendor)
- **9:53-9:54** Audience Survey 3 (2 minutes)
- **9:55-10:05** Key Question 3 (2-minute question; 3 minutes response per vendor)
- **10:06-10:18** Audience Q&A (13 minutes)
- 10:19-10:20 Wrap-Up





G2M Research Introduction and Ground Rules

Mike Heumann Managing Partner, G2M Research







Dave Montgomery Director, Data Center Systems (<u>www.wdc.com</u>)





Sumit Puri CEO and Co-Founder (www.liqid.com)



Alan Benjamin President and CEO (www.gigaio.com)

Host/Emcee:

Mike Heumann Managing Partner G2M Research www.g2minc.com



What is Composable Infrastructure, Anyway????

- Composable Infrastructure (CI) allows compute platforms (both servers and workstations) to share a variety of datacenter resources, such as:
 - General-Purpose Graphics Processing Units (GPGPUs)
 - FPGA and ASIC-based accelerator cards
 - Storage Elements (SSDs, Optane Memory, etc.)
- A network or fabric is utilized to connect the compute nodes to the resources. Possible networks/fabrics:
 - PCI Express (PCIe)
 InfiniBand
 - Ethernet
- A management system is utilized to set up and tear down connections between compute nodes and resources





Is Composable Infrastructure a New Concept?

- Like most things in Information Technology (IT), Composable Infrastructure (CI) has had predecessor concepts and systems.
- I/O Virtualization (2001-2013) was the most recent predecessor of CI
 - Focused on sharing of I/O resources (Ethernet NICs, Fibre Channel HBAs, InfiniBand HCAs)
 - Some vendors extended it to include GPGPUs and PCIe flash cards
 - Sharing was generally across PCIe or InfiniBand
- Composable Infrastructure has added several key elements:
 - Ability to add and delete resources from the overall resource pool
 - Ability to add and delete resources to compute elements without reboot
 - Ability to automate dynamic provisioning of resources



- In general, CI works where expensive resources are used intermittently by a given user/set of users, but which are critical to the task(s) of the user(s)
- Example: 4K or 8K digital video pre-production/mastering. Key resources: – GPGPUs: NVIDIA Titan RTX (\$2,499)
 - GPUs: NVIDIA GeForce RTX 2080 (\$820; need 2)
 - Storage: Intel Optane DC SSD 280GB (\$255; need 8)
 - Total cost per instance: \$6,200
- If you have a team of 100 users that each use the config 10% of the time:
 - Dedicated resources for each user: \$620,000
 - Shared resources using CI: \$62,000 (not including CI HW/SW costs)





Western Digital®

Western Digital

Dave Montgomery Director, Data Center Systems <u>www.wdc.com</u>

Western Digital's Open Composability Vision





Example: Flash Intensive Workload





Composability \rightarrow disaggregate hardware elements

Provides the advantages of Hardware Composed Infrastructure with no vendor lock-in



Disaggregate hardware components from the server so they can be efficiently pooled

Orchestrate virtual systems that can be optimally sized to the task



NVMe-over-Fabric Attached Devices



- No physical systems Only virtual systems Procured from separate suppliers
- Each device provides a resource that is offered over the network
- No established hierarchy CPU doesn't 'own' the GPU or the Memory
- All devices are peers on the network & they communicate with each other



Orchestrate Virtual Systems Using Fabric Attached Devices









Virtual System #4





OpenFlex[™] Composable Infrastructure

All-Flash Fabric Device and Enclosure

High-Capacity Fabric & High-Performance Compute Devices





OpenFlex D4000 — OpenFlex C2000 \rightarrow

OpenFlex F3000

High performance, low latency for fast data (flash) High capacity for big data *(disk)*

High performance compute (CPU, memory)

OpenFlex

NVMe-over-Fabric | Infrastructure Disaggregation | Software Composable



Getting Started: Open Composable Compliance Lab



www.opencomposable.com

Inquiries: OpenComposableAPI@wdc.com

Objectives

- Create interoperability across eco-system
- Enable customers to confidently purchase Fabric Attached Devices
- Enable composition / orchestration providers to focus on their value add
- Provide an environment for multiple companies to debug interoperability issues

Principals

- All Fabric Attached Devices are peers
- All Vendors benefit when customers can confidently purchase
- Provide an open environment to debug
- Collaborate with UNH-IOL



University of New Hampshire InterOperability Laboratory

- NVMe-oF Compliance Testing
- Ethernet Compliance Testing





Western Digital.





b the second sec

Liqid

Sumit Puri CEO and Co-Founder www.liqid.com

Data Center Efficiency



Liqid Enables Significant Improvement in Resource Utilization



Data Center Roadmap

Ligid chosen as 2018 Cool Vendor in Cloud Infrastructure

Gartner Cool Vendor 2018

"Composable infrastructures will provide I&O leaders with simple, flexible resource utilization and faster application deployment." - Gartner





PROPRIETARY & CONFIDENTIAL – 4/27/2016 2

PROPRIETARY & CONFIDENTIAL

Dynamic Infrastructure Platform



Use Case: Dynamic Bare Metal Cloud



Dynamic Bare Metal Cloud Infrastructure as a Service

Benefits of On-Premise with Flexibility of Cloud



ze	p3.8xlarge	p3.16xlarge
יט	4	8
יט	32	64
ry	244	488
r k	10 Gbps	25 Gbps
	660 700	6100 070



4-GPU		8-GPU	
4		8	
32		64	
256		512	
100 Gbps		100 Gbps	
001 667	600/	000 000	-

Future A.I. Market



20 **\$3.9** 22 **TN**

"Over the next 10 years, virtually every app, application and service will incorporate some level of AI," - Gartner

* "Forecast: The Business Value of Artificial Intelligence, Worldwide, 20172025.", Gartner, March 2018

Use Case: Composable A.I. Platform



Dynamic Resource Allocation for Each Stage of A.I. Workflow



Use Case: Media & Entertainment



Dynamic Reallocate Resources to Enable Improved Hardware Utilization

Benefits of Composable Disaggregated Infrastructure



Go-to-Market





Liqid GTM focused on Leveraged Partnership Model



Markets & Customers



AI & Deep Learning

- GPU Scale out
- Enable GPU Peer-2-Peer at scale
- GPU Dynamic Reallocation/Sharing

Dynamic Cloud

- CSP, ISP, Private Cloud
- Flexibility, Resource Utilization, TCO
- Bare Metal Cloud Product Offering



HPC & Clustering

- High Performance Computing
- Lowest Latency Interconnect
- Enables Massive Scale Out



5G Edge



- Utilization & Reduced Foot Print
- High Performance Edge Compute
- Flexibility and Ease of Scale Out

Thank You





GigalO

Alan Benjamin President and CEO <u>www.gigaio.com</u>

The Legacy Architecture is Being Forced Apart



Server

HUGE surge in data

- More storage
- Big Data Analytics & Databases
- Emergence of AI / DL / ML

End of Moore's Law

- Heterogeneous compute
- New faster, larger storage

"Tomorrow's advances will come from architectures optimized for specific workloads." - John Patterson, John Hennessy

Evolving into A Rack Scale Architecture



$G | G \land | O$

Rack Scale Systems and Disaggregation

- Drive better performance
- Enable more capable configurations
- Run converged workloads
- Reduce Costs Capex & Opex



"For all the value disaggregation brings, it is useless, unless we can reaggregate the pieces without getting killed in performance"



Rack Scale: Approach One – Bigger Servers



Bigger Server

- + Great performance
- Limited ability to scale the size / power in a single box
- Cost
- Utilization

Rack Scale: Approach Two – Use the Network



+ Scales well – through individual servers – storage

- Latency is a problem
 - Limits use beyond storage
- Complexity of SW & HW
- Proprietary
- Expensive

Rack Scale: Approach Three - PCIe Switching



A better "Big" server....

- + Very good performance, particularly latency
- + Works for storage and compute
- + Better sharing and utilization
- + Cost effective
- Fundamental limitations in building scale
- Still need to network the servers

FabreX is the Only Approach for Effective Rack Systems



FabreX[™]:

P2P Networking via non-transparent bridges

- + Great Performance latency AND bandwidth
- + Works for all types of storage and compute
- + Scales without limits; Resource Boxes and server based via NVMe-oF and GDR
- + Open-program platform via Redfish APIs
- + Easy Application / Framework / Container integration

Solves Both Scale AND Performance















Easy Integration into All Software Environments





The Network is *Finally* the Computer



Panel Discussion

Where is your organization in the process of investigating composable infrastructure (check one):

٠	We have deployed CI in production environments:	%
٠	We have deployed CI for limited "test" applications:	%
٠	We are engaged in a CI proof of concept evaluation:	%
•	We are talking to CI vendors:	%
•	We are identifying possible use cases for CI:	%
٠	We don't know enough about CI to know if it is applicable to us:	%
•	We don't see any application of CI for our organization:	%





- How does composable infrastructure today differ from the technologies that preceded it 10-15 years ago?
 - Western Digital
 - Liqid
 - GigalO



What use cases for composable infrastructure is your company interested in (check all that apply):

 Shared GPGPUs (computing): 	%
 Shared GPUs (graphics/video mastering): 	%
 Shared high-performance storage pools ("data lakes"): 	%
 Workflow-based shared storage pools: 	%
 Sharing FPGA or ASIC-based resources (AI, FinTech, etc.): 	%
Other:	%
None:	%





- What composable infrastructure use case have you seen the most interest in from you customers so far in 2019?
 - Liqid
 - GigalO
 - Western Digital



What are the greatest concerns that your organization has about utilizing composable infrastructure? (select all that apply):

 Overall maturity/lack of maturity of CI: 	%
 How much can we save? 	%
 Will CI impact resource performance? 	%
 Choosing the right CI network(s)/fabric(s): 	%
 Availability/support of CI by major server vendors: 	%
 Maturity/scale of CI go-to-market channels: 	%
Other:	%





How do you see composable infrastructure evolving over the next 5-10 years?

- GigalO
- Western Digital
- Liqid











Thank You For Attending

