

Zaius & Barreleye G2

Accelerator Ecosystem around Google / Rackspace 48V OpenPOWER Platform

OpenCAPI / Nvlink / PCIe Gen4

Adi Gangidi

Sr. Systems Design Engineer
Rackspace



Farther | Faster | Together

Agenda


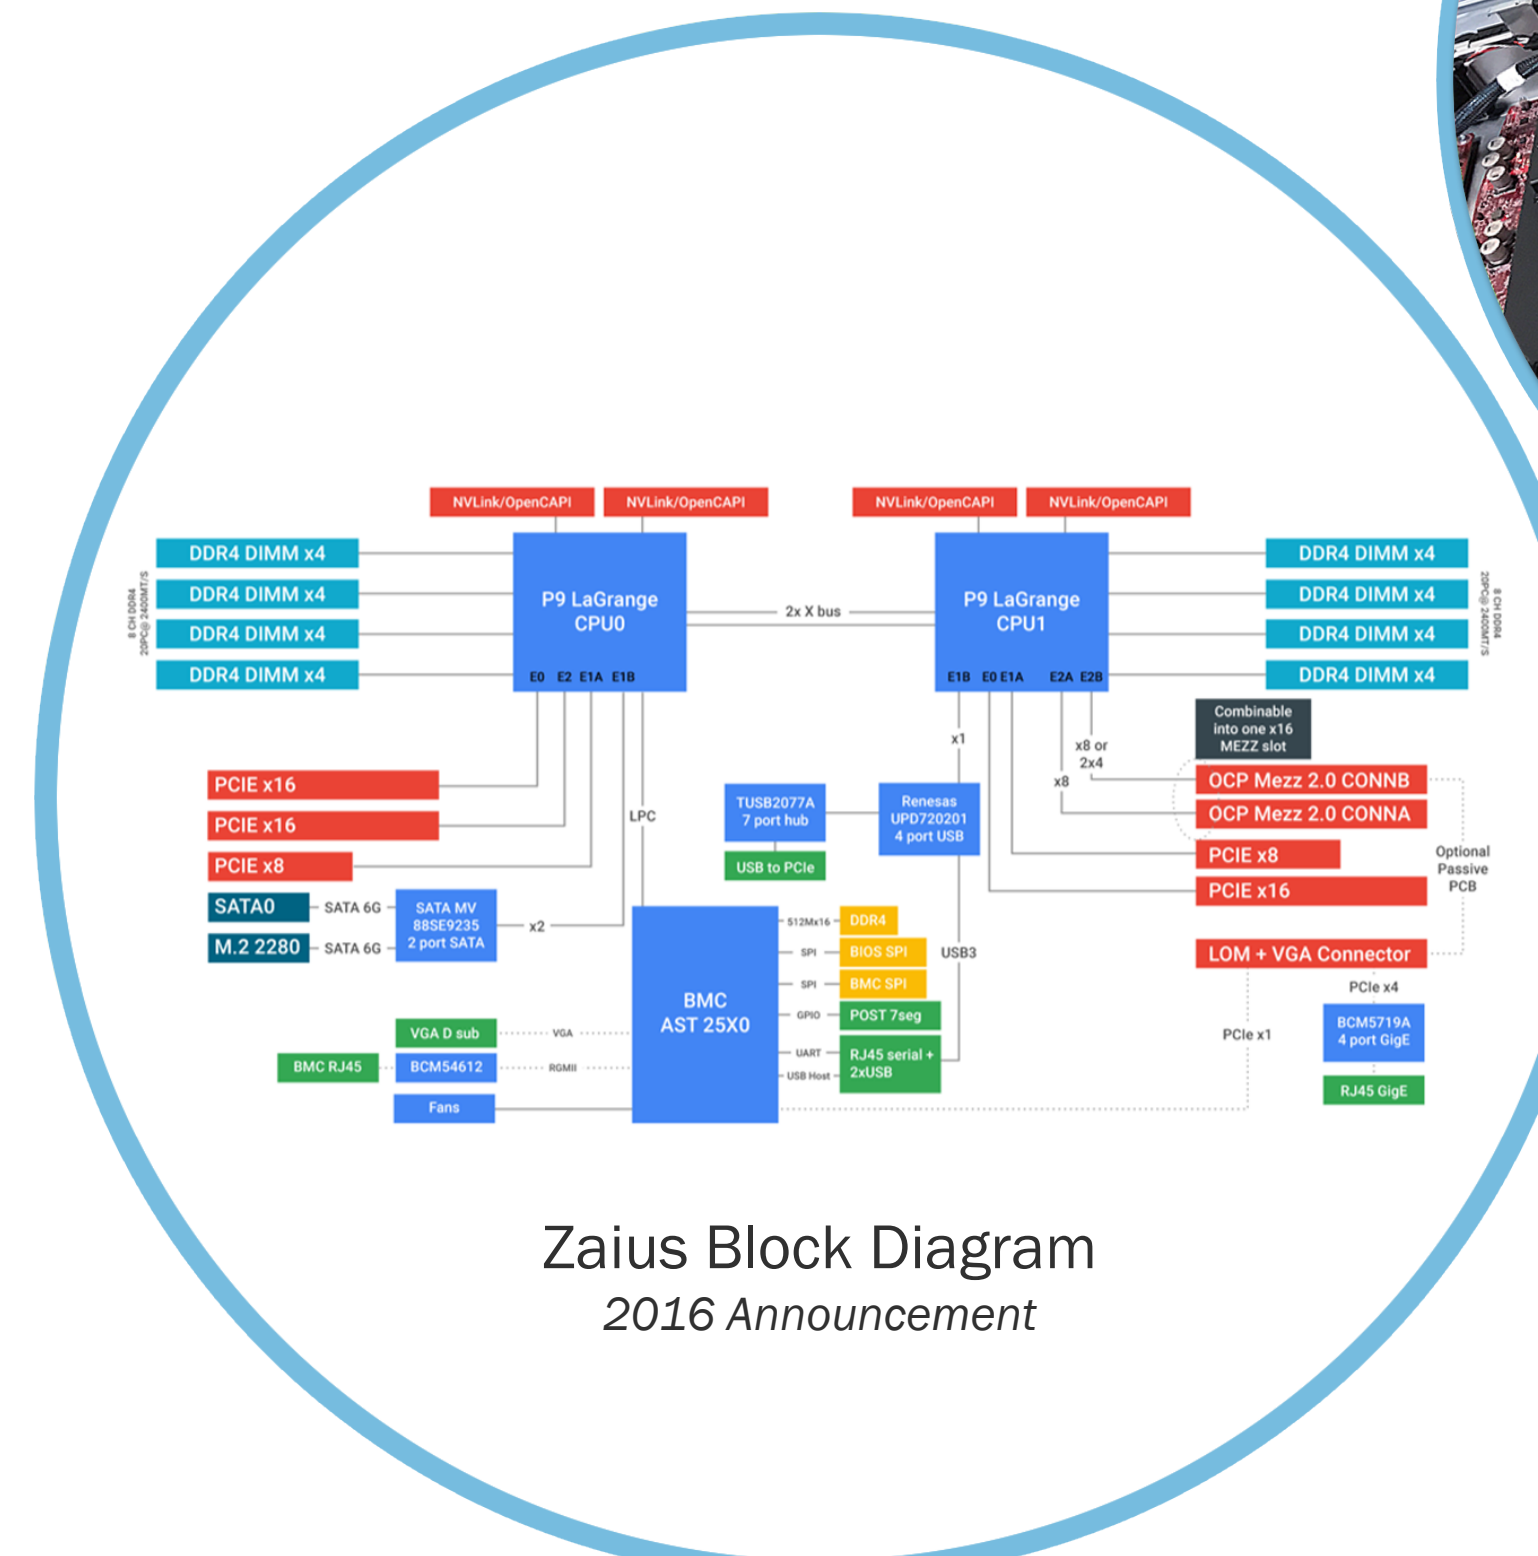
- Introduction | Recap
- Development Update
- Technology Update
- Accelerators
 - PCIe Accelerators
 - NVLink Accelerators
 - OpenCAPI Accelerators
 - Enablement Frameworks
- Samples



Introduction

Zaius + Barreleye G2




- 2016: Collaborative Effort Announced
 - OpenPOWER Foundation, Open Compute Project, OpenCAPI, IBM, Google & Rackspace
- 2017: Donated EVT, DVT Specifications to OCP Server Workgroup
 - Open/Emerging Technologies
- 2018: Development Update & IO/Accelerator Eco-system
 - Today's Announcements



Zaius Sled
2017 Show-case (EVT)

Family Tree

Zaius + Barreleye G2



ZAIUS Shelf & Sled

- Compact enclosure of Zaius MB
- 1.50U height
- Shorter than HH cards
- Compatible with 48V open rack v2 with deployment shelf

ZAIUS Motherboard

- 2 x POWER9 LaGrange
- 48V input
- Front IO & service access
- 80 Lanes of PCIe Gen4
- 32 Lanes of OpenCAPI / NVLink 2.0
- Open Source BMC & Host Firmware


BARRELEYE G2 Chassis

- Full-depth 48V open rack v2
- 2 0U chassis supports FHFL cards
- Hot swap storage bay (24 Drives)
- Hot swap fans and VGA access
- Wattage plan to support accelerators via Power board

DEVELOPMENT UPDATE

From Design...


...to Test



To Test...


...through EVT & DVT

OPENCAPI & NVLINK
PCIe GEN4




ZAIUS MOTHERBOARD


UNIVERSAL BACKPLANE
(SAS/SATA/NVME)



STORAGE EXPANDER




48V OPEN RACK
BUS BAR & FAN ADAPTER



PVT Entry:


The Green Board

OPENCAPI & NVLINK
PCIe GEN4




ZAIUS MOTHERBOARD


UNIVERSAL BACKPLANE
(SAS/SATA/NVME)



STORAGE EXPANDER



BUS BAR & FAN ADAPTER
48V OPEN RACK



What's New?

- 48V Motherboard
- Processor: Power9
- IO: PCIe Gen4 / OpenCAPI / NVLink 2.0

The Approach

- Solve for various workloads
- Provide future flexibility
- Decrease adoption time

IO QUALIFICATION			
NETWORKING	CX5 100GbE OCP Mellanox	CX4 25GbE OCP Mellanox	NetXtreme 25GbE Broadcom
RAID / HBA	9460-16i Broadcom	250S+ Nallatech	PEX 8734 Switch Board
GP GPU	PCIe Volta V100 NVIDIA	SXM2 Volta V100 NVIDIA	
OPENCAP	ADM-9V3 Alpha Data	Innova 2 Mellanox	Flash Storage Accelerator Molex
STORAGE	PM863A SATA Samsung	9200 Max NVMe Micron	PM963 NVMe Samsung
MEMORY	DDR4 2666 MHz 16GB 1R x4 Micron	DDR4 2666 MHz 8GB 1R x4 Micron	DDR4 2400 MHz 1R x4 Samsung

48V Rack

Power System

30 kW Power Shelf (Delta)


3 OU height, 4x 60A AC input

12x 3KW PSU Modules

6x ATS Modules (1 per 2 PSU)

10+2 Redundancy (If ATS Fails)

Rack Management Controller



6 PCS 6kW ATS





RMC

12 PCS 3kW PSU




48V Open Rack V2

Rack, Bus-bar & Clip

OCP 48V Rack
+48VDC bus bar (1 power zone)
30kW power shelf (3 OU, 36kW without redundancy)
17 Barreleye G2 Servers 41 OU total



Barreleye G2 48V Rack



48 V Bus Bar Clip

Amphenol / TE

Increase chamfer to improve Lining with Bus-bar

Cloud Evolution

The Road to 10x

Open Compute

2011: Rackspace joined OCP

OpenPOWER


2014: Rackspace joins OpenPOWER
2015: 1st OpenPOWER based OCP design
2017: Zaius / BG2 OCP Contribution

Accelerators

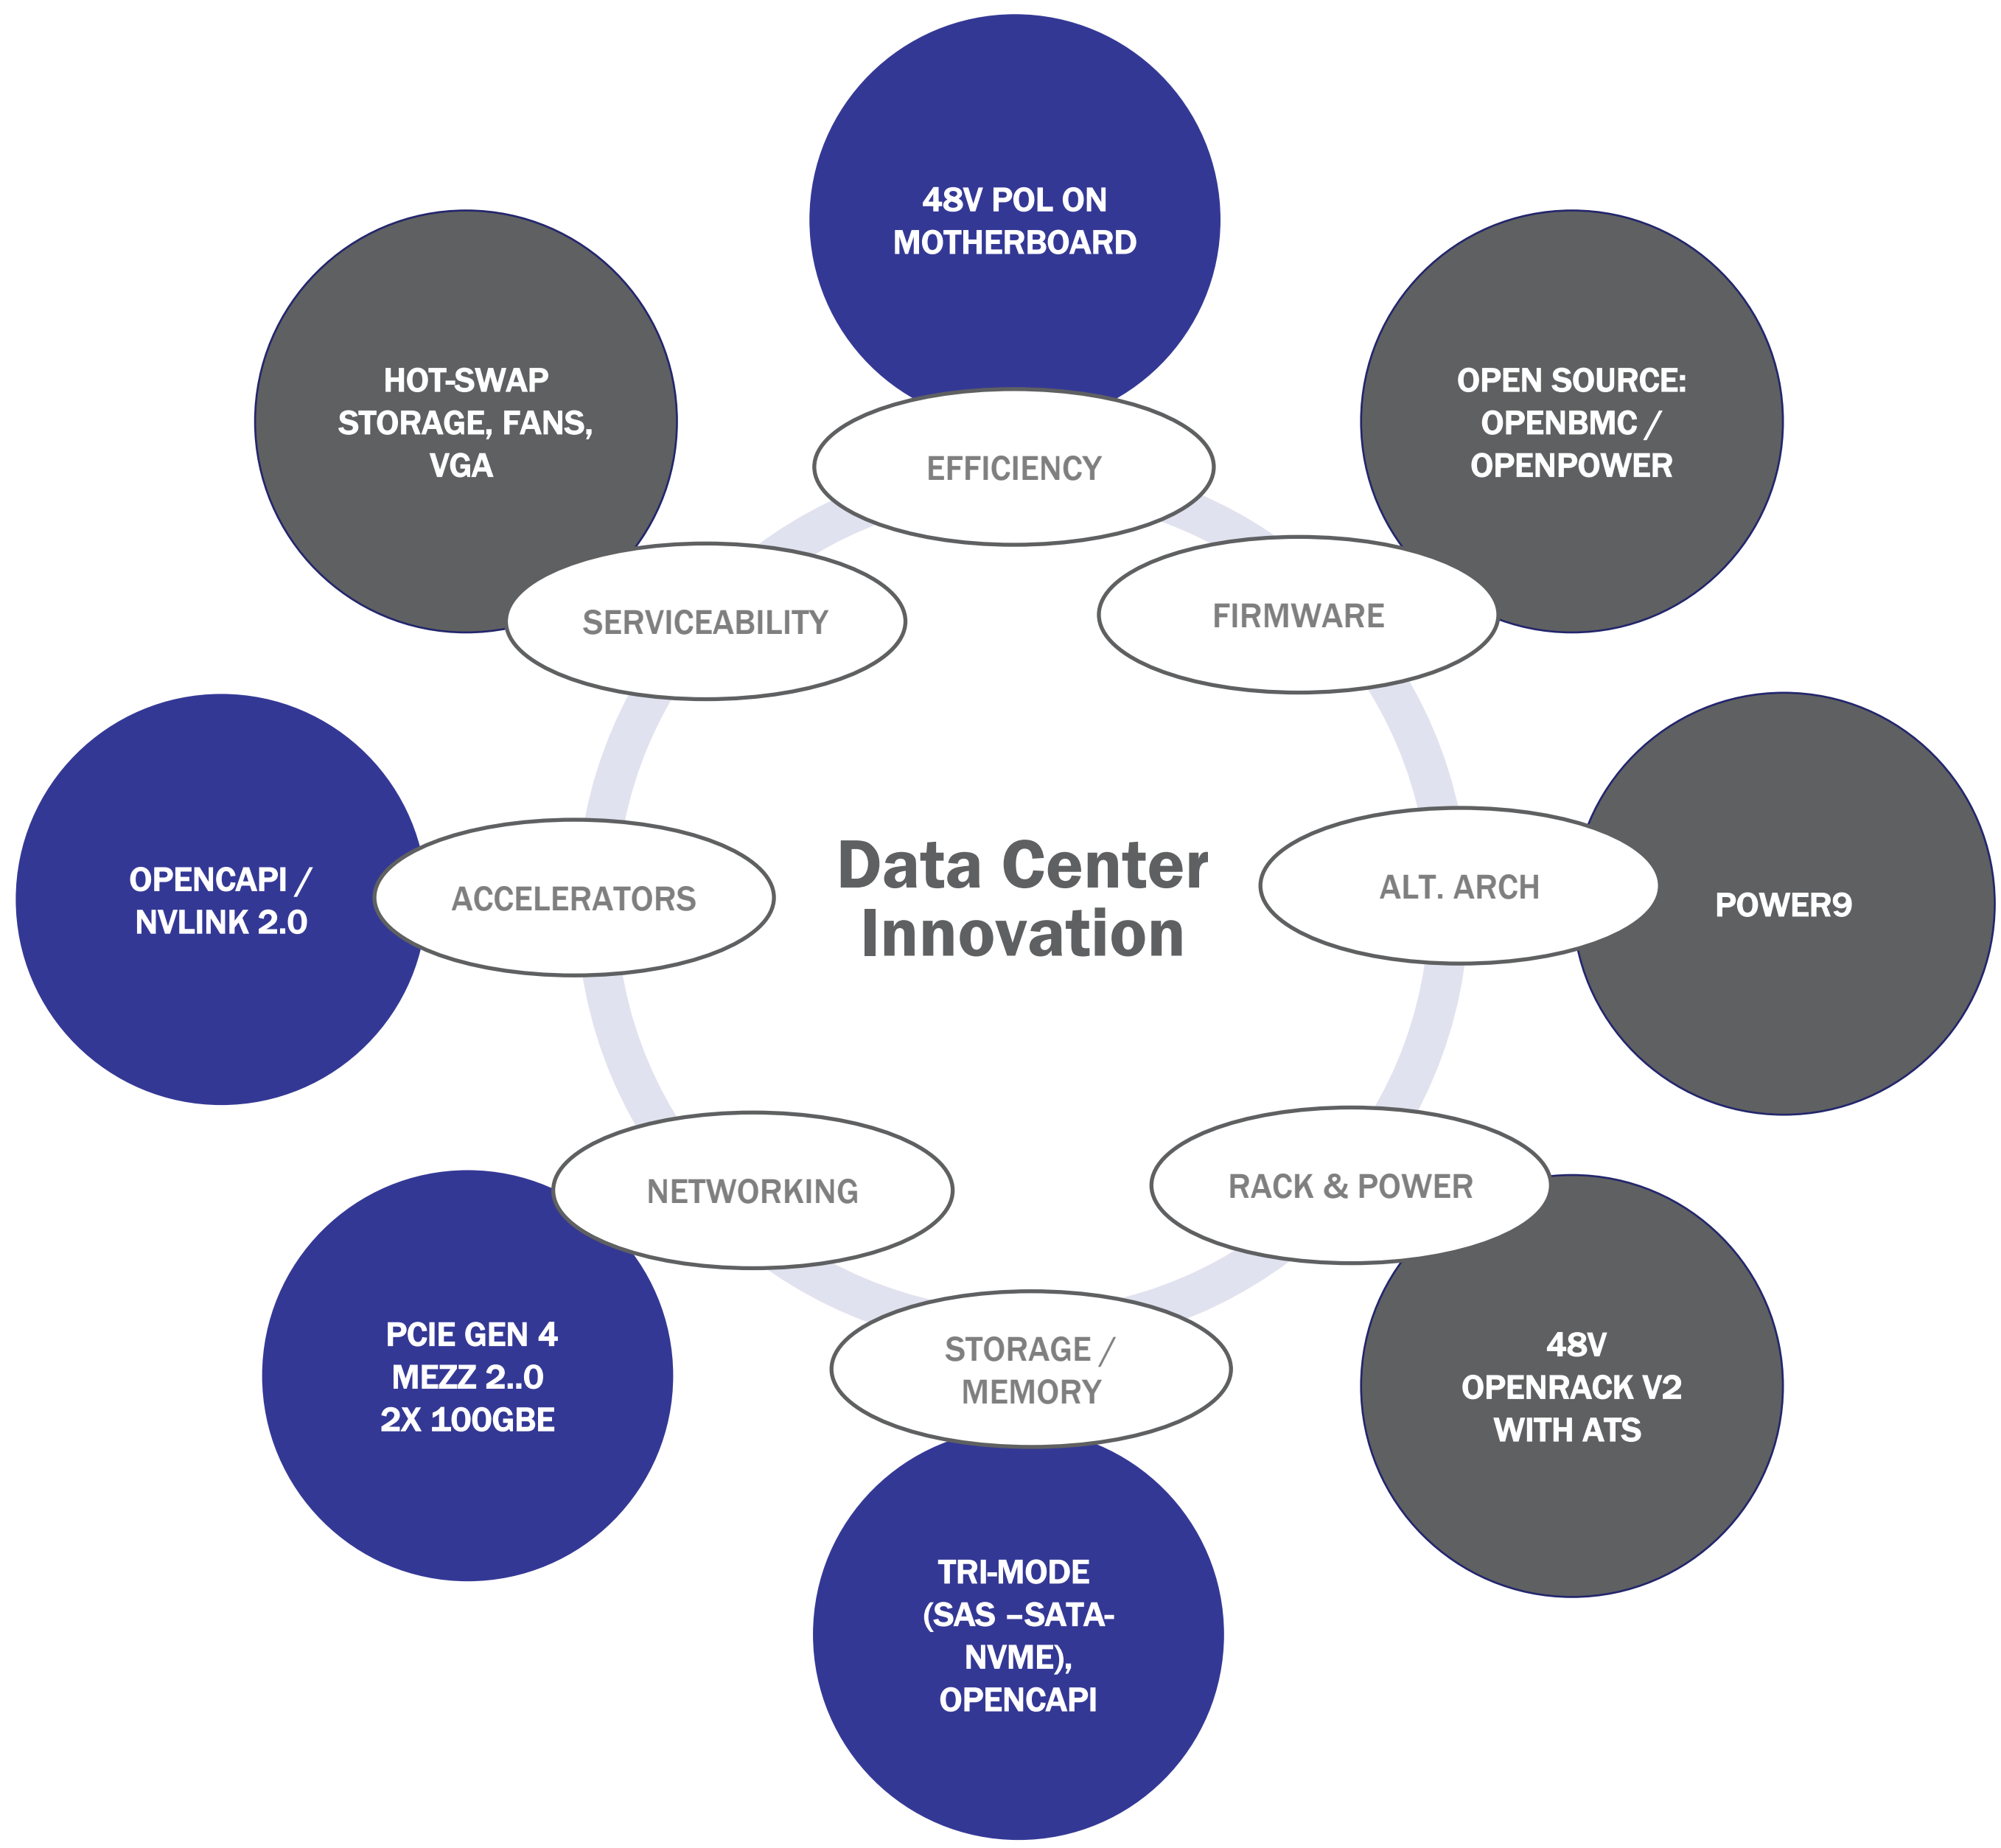
2017: Started qualifying Co-processors

Pilot

2018: Barreleye G2 Limited Access Beta




TECHNOLOGY UPDATE



PCIe G4/Mezz 2.0

Recap

- 6 PCIe G4 slots
- 2 x8 slots
- 3 x16 slots
- 1 x16 / x8 OCP Mezz 2.0 slot
- (x8 if using LOM+VGA mux card)




SLOT TYPE	QTY	TOTAL THROUGHPUT (UNI-DIR)
x16 Gen4	3	31.5 Gbytes / sec (252 Gbps)
x8 Gen4	2	15.75 Gbytes / sec (126 Gbps)
x16 Mez A/B Gen4 or x8 Mez A Gen4	1	31.5 Gbytes / sec (252 Gbps) 15.75 Gbytes / sec (126 Gbps)

PCIe G4/Mezz 2.0 Demo

Network & Storage Adapters


- GEN4 X16 NETWORK CARD
Mellanox CX5 Mezz 2.0
- *Gen3 Limitation: ~94 Gb/s*



Result: **187.7 Gb/s**

- GEN4 X8 STORAGE ADAPTER
Eideticom NVM Express Offload

Gen3 Limitation: ~6.8 GB/s



Result: **13.5 GB/s**


```
[ 3] local 10.0.0.22 port 44573 connected with 10.0.0.21 port 5001
[ 3] local 10.0.0.24 port 38176 connected with 10.0.0.23 port 5001
[ ID] Interval      Transfer    Bandwidth
[ 3]  0.0-10.0 sec  108.6 GBytes  93.6 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3]  0.0-10.0 sec  109.2 GBytes  94.1 Gbits/sec
```

```
ubuntu@ubuntu:~/NoLoad-Demos/fio$ sudo fio config.fio --filename=/dev/nvme0n1
simple: (g=0): rw=read, bs=4M-4M/4M-4M/4M-4M, ioengine=libaio, iodepth=32
...
fio-2.2.10
Starting 16 processes
^Cbs: 16 (f=16): [R(16)] [45.2% done] [13560MB/0KB/0KB /s] [3390/0/0 iops] [et
fio: terminating on signal 2

simple: (groupid=0, jobs=16): err= 0: pid=2626: Wed Feb 21 14:37:50 2018
read : io=178768MB, bw=13551MB/s, iops=3387, runt= 13192msec
```


PCIe G4 at Scale

Demo




```
me/ubuntu# lspci -vvvvvv | grep LnkSta | grep 16
0000:00:00.0 PCI bridge: IBM Device 04c1
0000:01:00.0 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0000:01:00.1 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0001:00:00.0 PCI bridge: IBM Device 04c1
0001:01:00.0 Ethernet controller: Mellanox Technologies MT27710 Family [ConnectX-4 Lx]
0001:01:00.1 Ethernet controller: Mellanox Technologies MT27710 Family [ConnectX-4 Lx]
0002:00:00.0 PCI bridge: IBM Device 04c1
0002:01:00.0 SATA controller: Marvell Technology Group Ltd. 88SE9235 PCIe 2.0 x2 4-port SATA 6 Gb/s Controller (rev 11)
0003:00:00.0 PCI bridge: IBM Device 04c1
0003:01:00.0 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0003:01:00.1 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0030:00:00.0 PCI bridge: IBM Device 04c1
0030:01:00.0 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0030:01:00.1 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0031:00:00.0 PCI bridge: IBM Device 04c1
0031:01:00.0 Ethernet controller: Mellanox Technologies MT27710 Family [ConnectX-4 Lx]
0031:01:00.1 Ethernet controller: Mellanox Technologies MT27710 Family [ConnectX-4 Lx]
0032:00:00.0 PCI bridge: IBM Device 04c1
0032:01:00.0 USB controller: Renesas Technology Corp. uPD720201 USB 3.0 Host Controller (rev 03)
0033:00:00.0 PCI bridge: IBM Device 04c1
0033:01:00.0 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
0033:01:00.1 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]
me/ubuntu# lspci
```

```
me/ubuntu# lspci -vvvvvv | grep LnkSta | grep 16
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk- DLActive+ BWMgmt- ABWMgmt+
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
LnkSta: Speed 16GT/s, Width x1, TrErr- Train- SlotClk- DLActive+ BWMgmt+ ABWMgmt+
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk- DLActive+ BWMgmt- ABWMgmt+
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk- DLActive+ BWMgmt- ABWMgmt+
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
LnkSta: Speed 16GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
```

PCIe G4 at Scale

Demo



```
[ 3] local 10.0.0.22 port 44573 connected with 10.0.0.21 port 5001
[ 3] local 10.0.0.24 port 38176 connected with 10.0.0.23 port 5001
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  108.6 GBytes 93.6 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  109.2 GBytes 94.1 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  107.4 GBytes 92.5 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  106.6 GBytes 91.8 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  110.4 GBytes 95.1 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  107.8 GBytes 92.9 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  106.9 GBytes 92.1 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  109.6 GBytes 94.4 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  57.0 GBytes 49.1 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  55.5 GBytes 47.8 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  56.1 GBytes 48.3 Gbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 3] 0.0-10.0 sec  56.8 GBytes 48.9 Gbits/sec
[2]+ Done iperf -c 10.0.0.1 & iperf -c 10.0.0.3
10.0.0.11 & iperf -c 10.0.0.13 & iperf -c 10.0.0.15 & iperf -c 10
root@ubuntu:/home/ubuntu#
```

1 Tbps
Network Bandwidth


Server-Server

- RECEIPE
- 2 x Barreleye G2 servers
- 8 “x16” PCIe Gen4 + 8 “x8” PCIe Gen4 NICs (Mellanox ConnectX5)
- 100 G Switch





PCIe G4 Switch Ecosystem

Last Step for G4 Adoption

GPU EXPANSION



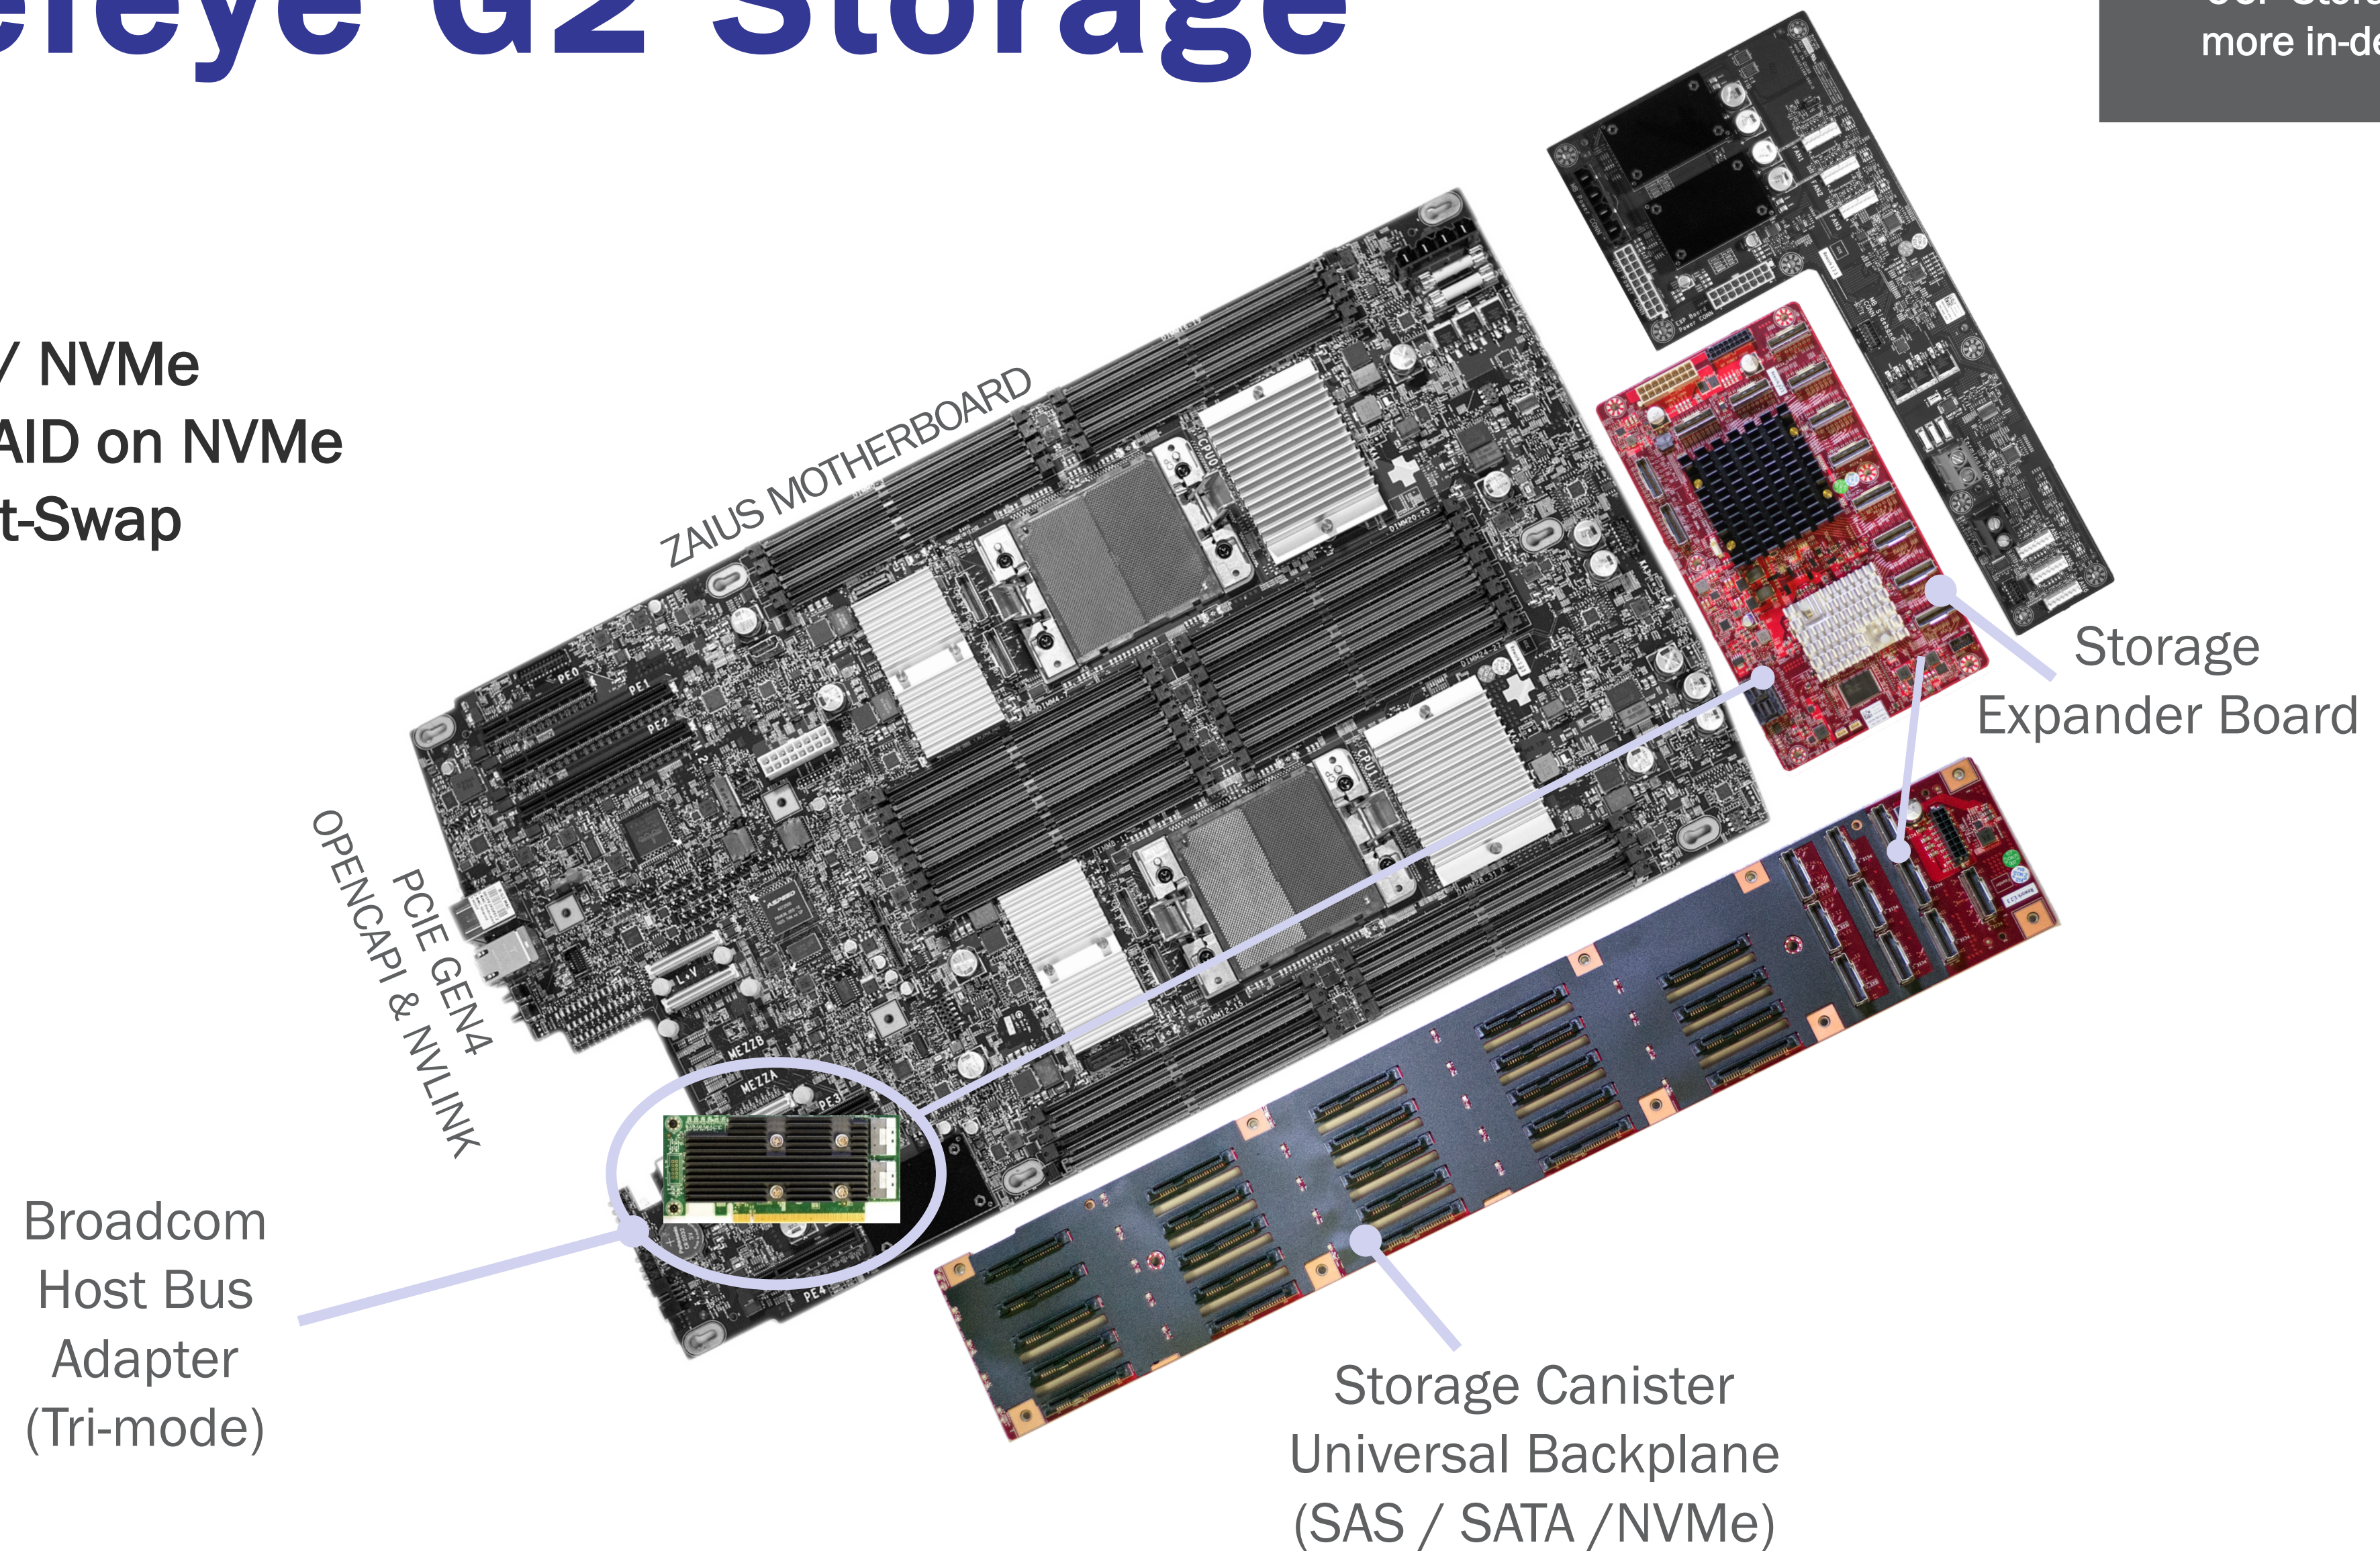
STORAGE EXPANSION



Barreleye G2 Storage

Tri-mode

- 24 Drives
- SAS / SATA / NVMe
- Hardware RAID on NVMe
- Tri-mode Hot-Swap




Interested?

Join me at 1:00pm at the OCP Storage Workshop for more in-depth information.

48 V Barreleye G2


Update




* Zaius sled & shelf use only motherboard. Additional PCBAs not used.

** Barreleye G2 chassis uses Zaius motherboard and additional PCBAs shown here.

OpenCAPI / NVLink Connector & Speeds




2 “Bricks” per CPU Socket





DESCRIPTION	QTY	TOTAL THROUGHPUT (UNI-DIR)
NVLink 2.0 / OpenCAPI 3.0 Brick	1	25 Gbytes / sec (200 Gbps)
Bricks per Socket	2	50 Gbytes / sec (400 Gbps)
Bricks per Server	4	100 Gbytes / sec (800 Gbps)

Zaius / BG2 OpenCAPI & NVLink Interface

- 2 “Bricks” per CPU Socket
- 24G Slimline SAS Connectors
- (25GByte per Brick)
- Follows SFF-8654 Standard



```
root@ubuntu:/home/fxn# ./OC_loopback.py
getscom(0x000000000901082c) = 2100ffff28107100
getscom(0x000000000901082d) = 2100ffff28107100
Training Passed on chip 0!
getscom(0x000000000901082c) = 2100ffff28107100
getscom(0x000000000901082d) = 2100ffff28107100
Training Passed on chip 1!
root@ubuntu:/home/fxn#
```



OpenCAPI at Scale: Update

Demo with 4x OpenCAPI Cards



```
Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — ssh fieldsaj@9.3.3.90...
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 000000004503A7E5 000000004503A7E5 000000001321E036 0.11461499
0.56902345 0.11380469
Total Cycles 0000000093B34A38 0000000093B34A38 00000000219F6388 0.20142273
1.00000000 0.20000000
Good Resp Total 00000000FE9F00BC 00000000FE9F00BC 0000000039F65480 22.22296857
110.32999199 22.06599840
Good Resp Load 00000000100B13A4E 00000000100B13A4E 0000000039F65620 22.22297794
110.33003851 22.06600770
Good Resp Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Total 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 0000000058257D82 0000000058257D82 000000001321D59D 0.11461492
0.56902674 0.11380535
Austin — root@zaiuslp11: ~ — ssh root@9.3.23.34 — 80x24

Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — ssh fieldsaj@9.3.3.90...
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 000000001EBFAFE3 000000001EBFAFE3 000000001321B009 0.11461172
0.56902330 0.11380466
Total Cycles 000000005073EDA6 000000005073EDA6 00000000219F78FF 0.20141801
1.00000000 0.20000000
Good Resp Total 000000008AB165EE 000000008AB165EE 0000000039F6784E 22.22243963
110.32995491 22.06599098
Good Resp Load 000000008CC3913C 000000008CC3913C 0000000039F6784E 22.22243981
110.32995582 22.06599116
Good Resp Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Total 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 0000000031E184E7 0000000031E184E7 000000001321D504 0.11461106
0.56902093 0.11380419
Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — ssh fieldsaj@9.3.3.90...

Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — ssh root@9.3.23.34 — 80x24
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 0000000045036365 0000000045036365 000000001321DE7E 0.11461238
0.56902842 0.11380568
Total Cycles 0000000093B2CD8F 0000000093B2CD8F 00000000219F7348 0.20141741
1.00000000 0.20000000
Good Resp Total 00000000FE9E14F6 00000000FE9E14F6 0000000039F65FD0 22.22228859
110.32953059 22.06590612
Good Resp Load 00000000100B030BE 00000000100B030BE 0000000039F65024 22.22219691
110.32907541 22.06581508
Good Resp Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Total 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 000000005825180C 000000005825180C 000000001321B4A7 0.11460806
0.56900772 0.11380154
Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — -bash — 80x24

Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — ssh fieldsaj@9.3.3.90...
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 000000001EBFE543 000000001EBFE543 000000001321E26F 0.11461642
0.56902059 0.11380412
Total Cycles 00000000507463C7 00000000507463C7 00000000219F86C6 0.20142567
1.00000000 0.20000000
Good Resp Total 000000008AB22F5A 000000008AB22F5A 0000000039F69FAA 22.22337629
110.33040916 22.06600183
Good Resp Load 000000008CC45A3E 000000008CC45A3E 0000000039F69F8E 22.22337547
110.33040507 22.06600101
Good Resp Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Total 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Loads 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
Retries - Store 0000000000000000 0000000000000000 0000000000000000 0.00000000
0.00000000 0.00000000
No cred cycles 0000000031E1C7AF 0000000031E1C7AF 000000001321E26C 0.11461593
0.56902346 0.11380469
Austin — root@zaiuslp11: ~/ocapi/cxl-tests-opencapi — -bash — 80x24
```


OpenCAPI at Scale: Update

Demo Details




- 4x OpenCAPI Cards
- *Signal Integrity*: 25.78125 Gbps (3x)
- *Round Trip Latency* < 80ns (~5x)
- *Bandwidth Test*: 88 GB/s (~700 Gbps)
- *Coherent*: No kernel Overhead
- Upstream Driver

**OpenCAPI:
Faster / Cooler PCIe**


ACCELERATOR FORM-FACTORS

Generic PCIe Accelerators

Coherence tunneled via PCIe




PERPENDICULAR MOUNT CARD




Perpendicular Mount Accelerators

With Additional Coherent Attach

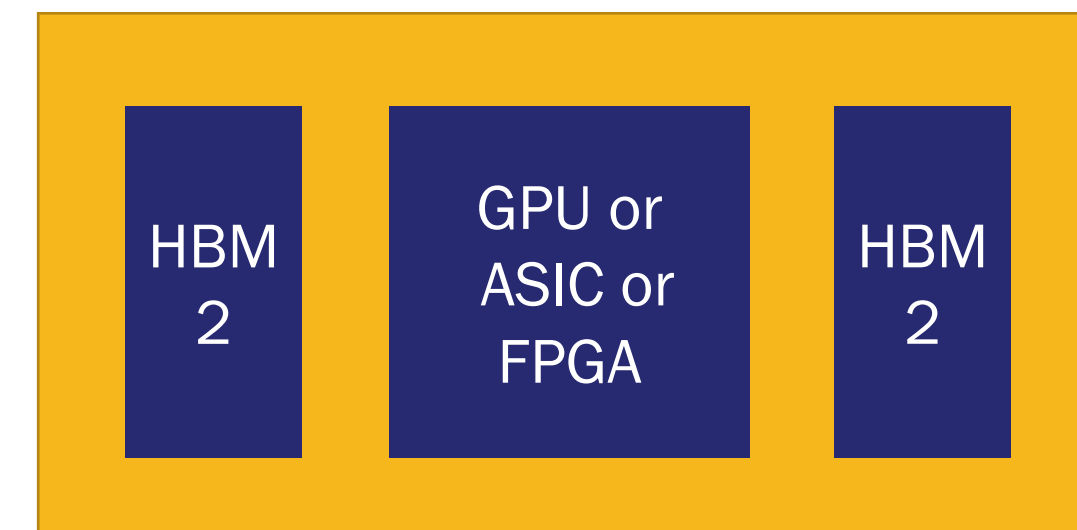



PERPENDICULAR MOUNT CARD WITH COHERENT ATTACH




Parallel Mount Accelerators

With Additional Coherent Attach



Top View of Parallel Mount card



PARALLEL MOUNT VIA MEZZ** WITH COHERENT ATTACH




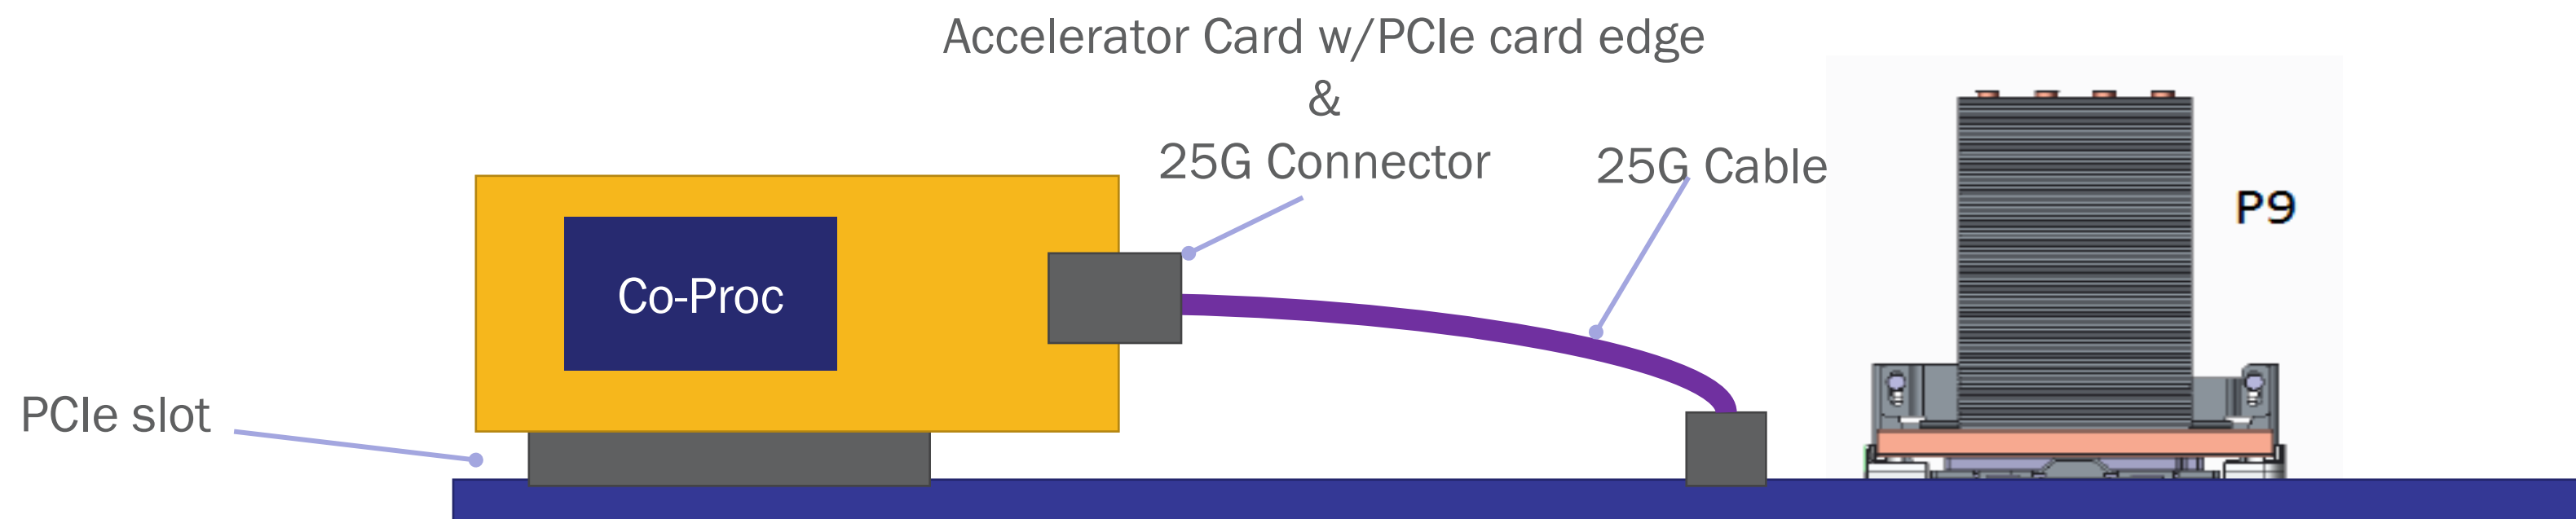
** Does not refer to OCP Mez

Accelerator Cards and Cable Mount Options*



Generic PCIe



Perpendicular Mount



Parallel Mount w/ Mez**



* Possible approach to co-processor card, mount, and cable designs.

** Does not refer to OCP Mez standard.

PCIe ACCELERATORS



VOLTA V100

NVIDIA



- 640 Tensor Cores , 250 W
- 16 GB HBM2
- Gen3 x16 Input
- FHFL card, fits with Riser
- 12 V power from fan board
- Deep learning



250S+ (Cabled)

Nallatech

- Gen4 x8 Host Connection
- FPGA Card with Cabled NVMe Storage
- Upto 12.8 TB
- Ideal for Coherent Memory Expansion




250S+ (On-Board)

Nallatech

- Gen4 x8 Host Connection
- 4x M.2 NVMe (x4)– On-Board
- Accelerated NoSQL Databases

NVLINK 2.0 ACCELERATORS



FUTURE

Easy to Re-spin for changes

Leading to JBoG implementations

Solves the CPU-GPU Bottleneck




TESLA V100 SXM2 NVIDIA

300 W , 120 TFLOPS
Tensor Performance

16 GB HBM2, 640 Tensor
Cores, 512 CUDA Cores

Ideal for GPU Analytics, Large
Deep Learning Models
(Kinectica, MapD , BlazingDB)



CPU-GPU NVLink 2.0 OpenPOWER


100 GB/s Bandwidth (Bidirectional)

3x Faster, Coherence

NVLink 2.0 on wire

2x SlimSAS Cables per GPU

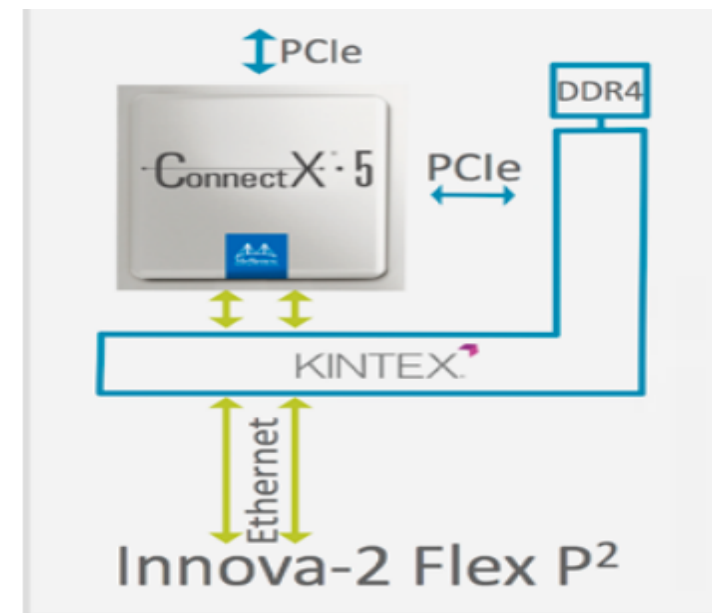

OPENCAPAPI ACCELERATORS



ADM 9V3

ALPHA DATA



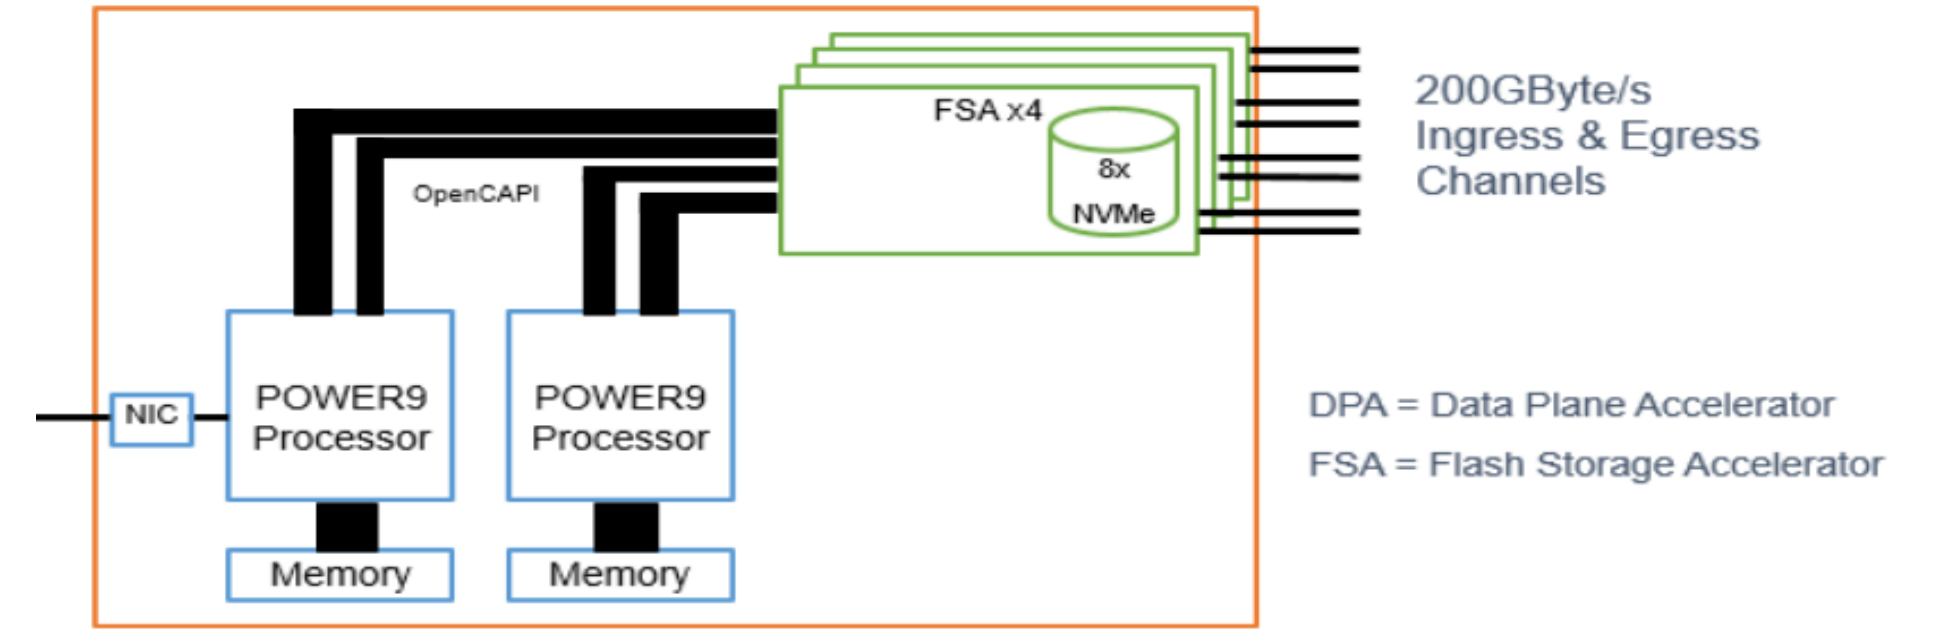
- x8 25G OpenCAPI
- Development Platform
- Offers Board Support Package
- HPC / Market Analysis,
- Smart NIC
- Inference



Innova 2 Flex

MELLANOX

- x8 25G OpenCAPI
- CX-5 + FPGA goodness
- Network Acceleration (NFV, Packet Classification)
- Security and Storage Acceleration

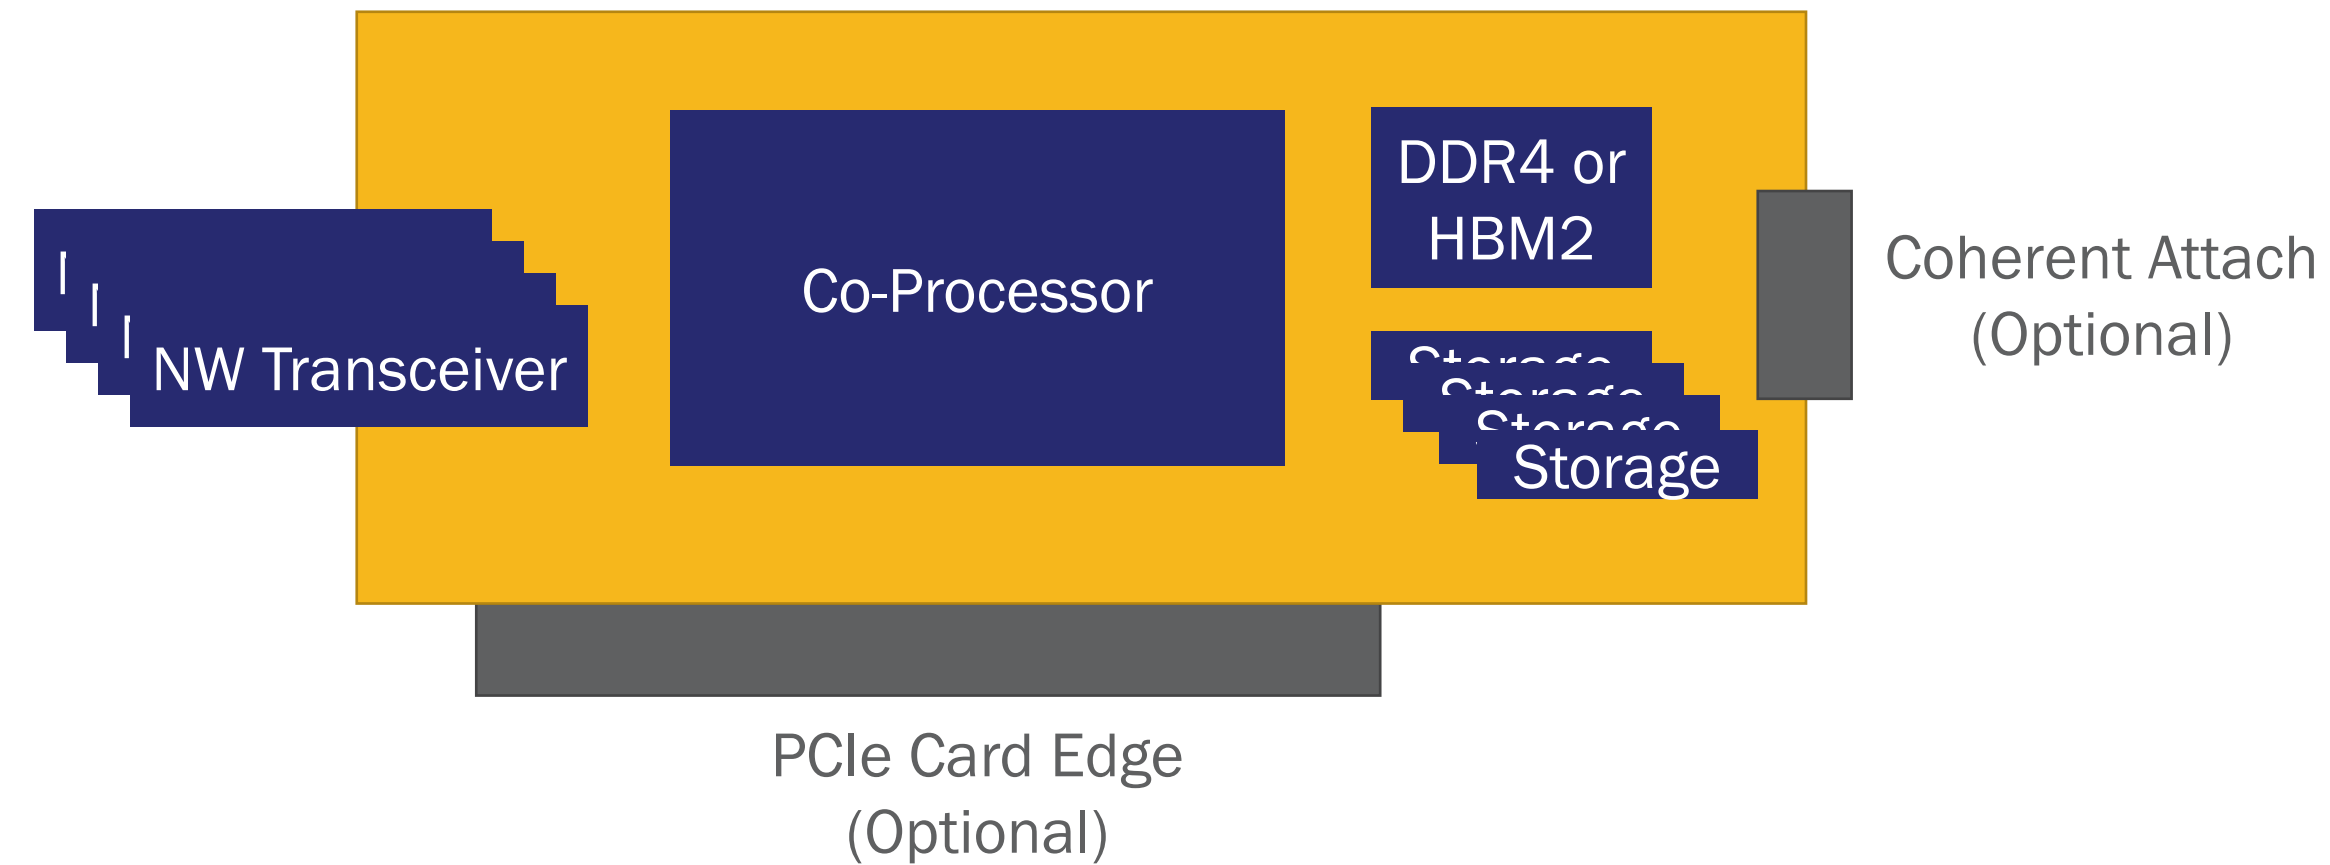


Flash Storage Accelerator

MOLEX

- x8 25 G OpenCAPI
- Coherently Attached with FPGA & ARM cores
- 8 NVMe (x4) via OpenCAPI
- Low latency Storage (Up to 16TB)
- Streaming (CDN), NoSQL, HPC Simulations, Hyperconverged Architecture (Like Microsoft Catapult)

Accelerator Composition



	Nallatech 250 S+	Alpha Data (ADM-9V3)	Mellanox Innova 2 Flex	Molex FSA
Data Input	PCIe Gen4 x8	x8 25G	x8 25G	x8 25 G
FPGA / Processor	Xilinx Kintex Ultrascale+	Xilinx Virtex Ultrascale +	Xilinx Kintex Ultrascale+ & Connect-X-5	Xilinx Zync Ultrascale+ (with ARM A53 4x 1.5GHz)
Number of LUTs	523K	394 K	522 K	523 K
Storage	12.8 TB	-	-	16 TB
Memory (DDR4)	4 GB	32 GB	8 GB	64 GB
Max Cards in G2 Chassis	2	4	4	4
Network	-	2 x 100GbE	2x 25GbE	2x 100 GbE
Power and Initialization	PCIe Gen4 x8	PCIe Gen4 x8	PCIe Gen4 x8	48V Power / Fan Board

ACCELERATOR ENABLEMENT



SNAP

OpenPOWER

Development
framework

Program in high-level
languages



PowerAI

IBM

Deep learning
made easy

Take advantage of
CPU-GPU NVLink 2.0



CAPI Flash

IBM

API driven, In memory
expansion

Redis, Cassandra,
Neo4J



Solution Providers


Ex: Eideticom, Burlywood

Make end-to-end solutions
or License IP

Others available via card
vendors



RACK & POWER SHELF



Amphenol


25G
CONNECTIVITY



ACCELERATORS



NETWORKING



MEMORY/
STORAGE

ACCELERATORS



molex



EIDETICOM

POWER9




VOLTAGE
REGULATOR




MOTHERBOARD, PCBA, CHASSIS, INTEGRATION

Collaboration at Scale



OPEN
Compute Project



Solving Bottlenecks . . .

. . . Easing Adoption

TOGETHER

Get on Board





Learn More

- Engineering Presentation, 2017 OCP Summit
<https://www.youtube.com/watch?v=2XJQvGX9yVE>
- OCP Community Preview
<http://files.opencompute.org/oc/public.php?service=files&t=4302505c9d59160ec03aaf4ea67b146c&download>
- OCP Server Mailing List
<http://lists.opencompute.org/mailman/listinfo/opencompute-server>
- OCP Server Mail Archives
<http://lists.opencompute.org/pipermail/opencompute-server/>

Participate

- Interested in consuming this hardware?
- Want to the platform available to wider audience?
- Are you an OpenCAPI & NVLink solution developer?
- Inquire about development samples

Reach out to:

adi.gangidi@racksapce.com

Design Package

- Available on OCP Server Wiki and GitHub
- Enhancements Coming Throughout 2018
- GitHub
<https://github.com/opencomputeproject/zaius-barreleye-g2> (don't forget to install Git LFS)
- OCP
http://www.opencompute.org/wiki/Server/Working#Open_Rack

Thank You