

OPEN POSSIBILITIES.

Zion Modular System Architecture Overview



OCP
GLOBAL
SUMMIT


NOVEMBER 9-10, 2021

Zion Modular System Architecture Overview

Hao Shen, Hardware Engineer, Meta

Michael Haken, Mechanical Engineer, Meta

Tyler Hart, RTP Engineer, Meta



OPEN POSSIBILITIES.



Meta Open Fleet

2011	2012	2013	2014	2015	2016	2017	2018	2019	2021			
 Data Center	 Triplet Rack	 Windmill (Intel)	 Knox	 Open Rack V2	 BluRay	 Wedge 100	 Backpack	 Wedge 100S	 Big Basin	 Big Basin V2	 Open Accelerator Module	 Delta Lake
 Battery Cabinet	 Freedom Servers	 Watermark (AMD)	 Winterfell	 Mezzanine Card V2	 Leopard	 Big Sur	 Lightning	 Bryce Canyon	 100G CWD4-OC	 Twin Lake	 Minipack	 Yosemite V3
 Spitfire Server (AMD)	 Mezzanine Card V1	 Open Rack V1	 Cold Storage	 Wedge	 Yosemite	 Yosemite V2	 OCP NIC3.0	 Minilake	 Wedge400			
 Power Supply		 Group Hug	 Micro Server (Panther)	 Honey Badger	 Six Pack	 Tioga Pass	 FAV3		 Minipack2			

OPEN POSSIBILITIES.

AI in Meta

AI is used extensively in Meta

Ranking

Content Understanding

Pattern Detection

Speech Recognition

Translation

Powerful AI Models need Powerful Hardware!

OPEN POSSIBILITIES.



SERVER




Zion System Overview

Zion is designed to support AI workload.



SERVER



Presented in 2019 OCP Summit

OPEN POSSIBILITIES.



Zion System Overview


2-socket Modular Server



2 or 4 Socket Host



OCP NIC + Storage Modules



Expander Box



Open Accelerator Module




8 Accelerator Shelf



SERVER

Zion system




OPEN POSSIBILITIES.

Zion System Overview

Zion is designed to support AI workload.



SERVER



Emerald Pools



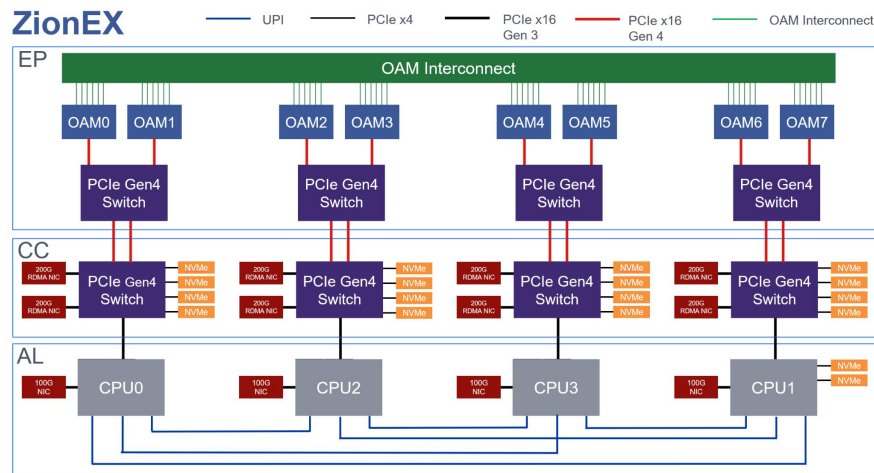
Clear Creek



Angels Landing



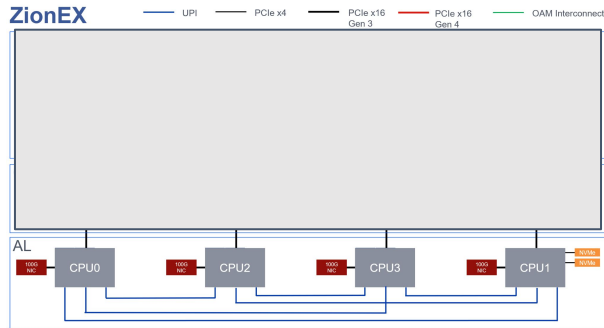

ZionEX



OPEN POSSIBILITIES.



Angels Landing

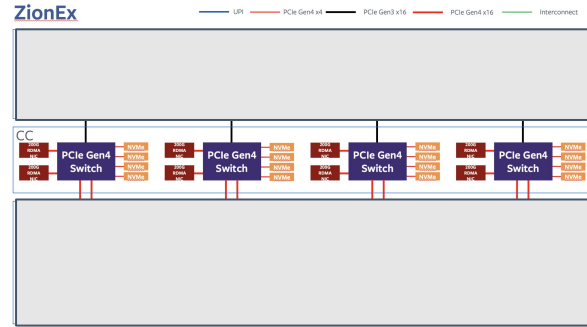


SERVER

- Up to 4 socket Intel Cooperlake CPUs
- 4x 100G OCP3.0 NICs
- 1.5TB DDR4 RAMs
- Fully connected UPI through backplane

OPEN POSSIBILITIES.

Clear Creek

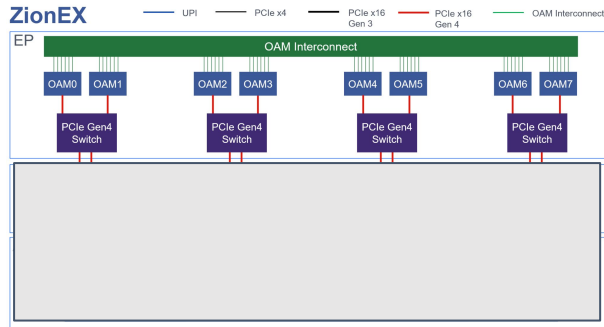


SERVER

- 4x PCIe Gen4 Switch
- 8x 200G NICs for scale out
- 16x E1.S/M.2 SSDs

OPEN POSSIBILITIES.

Emerald Pools




SERVER

- 8x Open Accelerator Modules
- OAM interconnections support high speed communications between accelerators


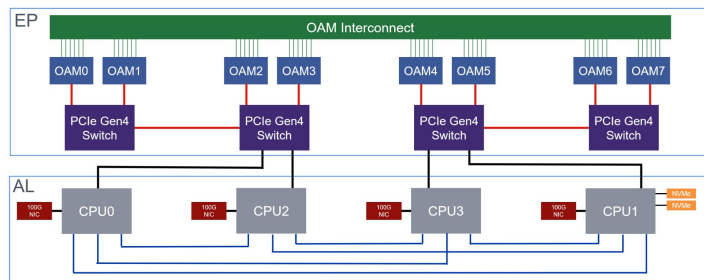
OPEN POSSIBILITIES.

Flexible Configurations

ZionEX



Zion2S/4S



SERVER

Modular system designs enable hardware to be tailored for each AI use case


OPEN POSSIBILITIES.



Designed for Field Service



SERVER

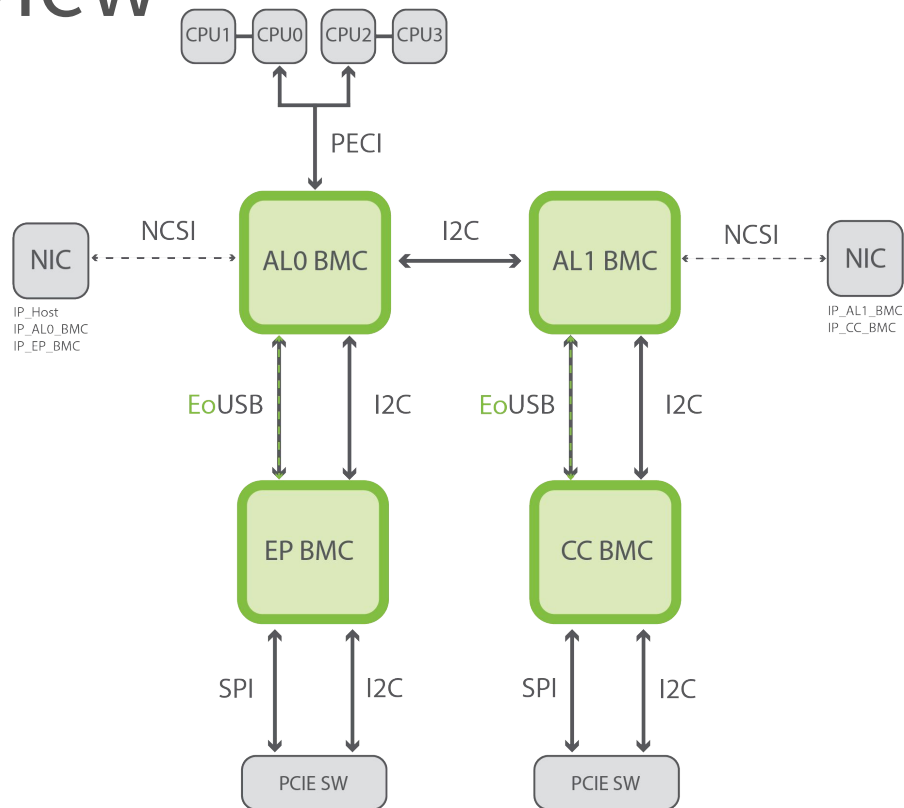


1. All field replaceable units (FRUs) with significant failure rates are accessible without removing cabling
2. PCIe cabling is routed from the back of each board, around the sides and to the front of each system
3. The OAMs are accessible from a sliding rail kit
4. CPUs, DIMMs, and storage modules are accessible on front accessible trays

OPEN POSSIBILITIES.



BMC Overview

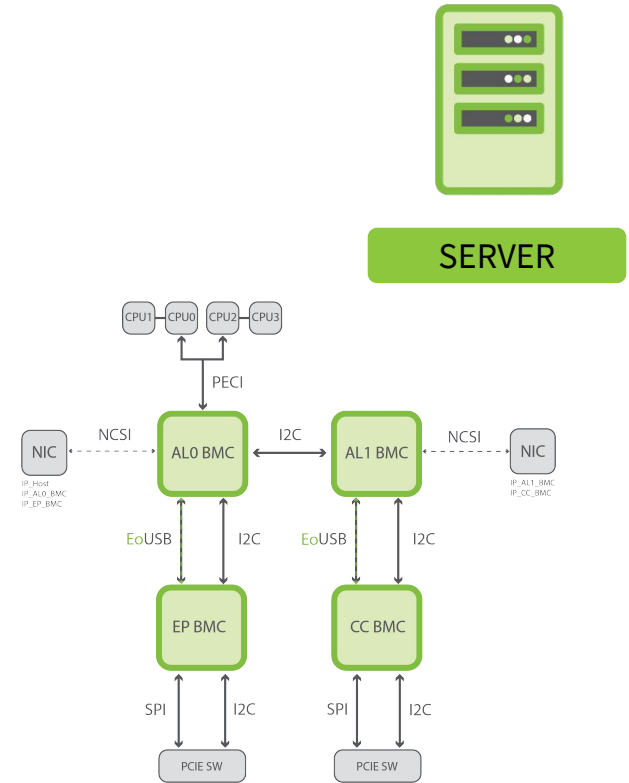


SERVER

OPEN POSSIBILITIES.

Crash Dump

1. CATERR/MSMI pin triggered, SEL created in AL0 BMC
2. AL0 BMC starts crash dump collection
3. MC Banks collected, sensors recorded
4. Crash dump log saved into BMC flash
5. Logging service extracts crash dump to database



OPEN POSSIBILITIES.

Call to Action

- Meta are contributing Zion system, Angels Landing, Clear Creek and Emerald Pools Specification to OCP Server/OAI group
- QCI will contribute the design collaterals soon.
- Zion System is already in MP stage.

Where to buy: <https://www.opencompute.org/products>

Project Wiki with latest specification : <http://www.opencompute.org/wiki/Server/OAI>

Mailing list: <http://lists.opencompute.org/mailman/listinfo/opencompute-server>

OPEN POSSIBILITIES.



Thank you!



NOVEMBER 9-10, 2021