

OPEN POSSIBILITIES.

Intel Ponte Vecchio Compute Accelerator OAM Product and System



NOVEMBER 9-10, 2021

Intel Ponte Vecchio Compute Accelerator OAM Product and System

Server/OAI

Sarosh Irani, Product Line Director
Song Kok Hang, Principal Engineer
Intel Corporation




OPEN POSSIBILITIES.



Ponte Vecchio (PVC)

General Compute Accelerator



Silicon Current Status

> 45 TFLOPS

FP32 Throughput

> 50 TBps

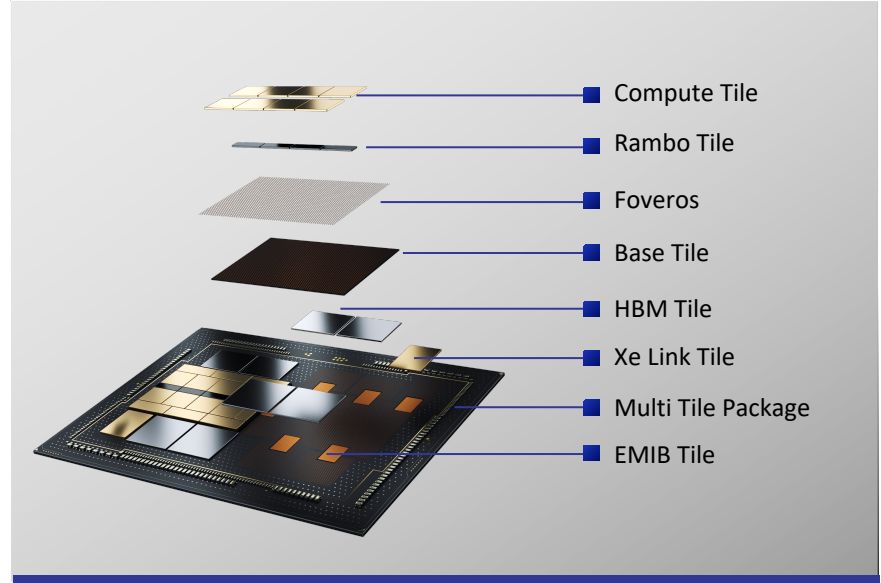

Cache Bandwidth

> 8 Tbps

Communication Bandwidth

OPEN POSSIBILITIES.


Ponte Vecchio Construction



OPEN POSSIBILITIES.

Intel Ponte Vecchio OAM Spec




- OAM v1.1 Compliant
- Air Cooled
- Liquid Cooled
- High Speed SerDes for Xe^e Links



OPEN POSSIBILITIES.

X^e Link for Scalability


Enabling a high number of coherent and unified accelerators



OPEN POSSIBILITIES.

4 PVC OAM Baseboard


- **4 GPU (OAM) Carrier Base Board with X^e-Links built in to enable high performance multi-GPU communication within the sub-system**
- **Half Width sled, to enable high density HPC designs**
 - Connect to CPU Motherboard over PCIe Gen5
 - Support for air and liquid cooled PVC OAM SKUs
 - Can support liquid cooled or air-cooled systems (1U or 2/3U chassis)
 - Support for UP or DP systems
 - High voltage power input (48V to 54V) for high efficiency power delivery
- **High Speed all to all connected X^e-Link between PVC modules**
- **Scale-out thru CPU host node**
- **Designed to enable system flexibility and faster OEM time to market**




OPEN POSSIBILITIES.

Accelerated Compute Systems

Ponte Vecchio
OAM




Ponte Vecchio
x4 Subsystem
with X^e Links



Ponte Vecchio x4 Subsystem
with X^e Links

+ 2S Sapphire Rapids



OPEN POSSIBILITIES.

Intel Accelerator UBB: 8 PVC OAM Baseboard


- **8 GPU (OAM) UBB with Xe-Links built in to enable high performance multi-GPU communication within the sub-system and across multi-node**

- **OCP UBB, to enable high performance AI/HPC designs**

- Connect to CPU Motherboard over PCIe Gen5
- Can support liquid cooled or air-cooled systems (3/4U UBB Tray)
- Support for aggregated or disaggregated systems
- High voltage power input (48V to 54V) for high efficiency power delivery

- **Fully Connected Topology with all to all connected Xe-Link between PVC OAM**

- **Designed to enable system flexibility and faster OEM time to market to support both PVC and a next generation Gaudi AI processor**



OPEN POSSIBILITIES.


Call to Action

Get involved in the project:

OCP Server Project: <https://www.opencompute.org/projects/server>

OAI Subgroup: <https://www.opencompute.org/wiki/server/OAI>

OAI Mailing List: <https://oc-all.goup.io/g/OCP-OAI>



OPEN
ACCELERATOR
INFRASTRUCTURE

OPEN POSSIBILITIES.



Thank you!



NOVEMBER 9-10, 2021