

OPEN POSSIBILITIES.

Meta AMD 1-socket server:
Scaling Performance Efficiency



OCP
GLOBAL
SUMMIT

NOVEMBER 9-10, 2021


Meta AMD 1-socket server: Scaling Performance Efficiency

Damien Chong, Hardware TLM, Meta
Bharath Muthiah, Technical Sourcing Manager, Meta
Jay Kirkland, CVP, Server Platform Solutions Eng, AMD

OPEN POSSIBILITIES.



SERVER



OPEN
PLATINUM™

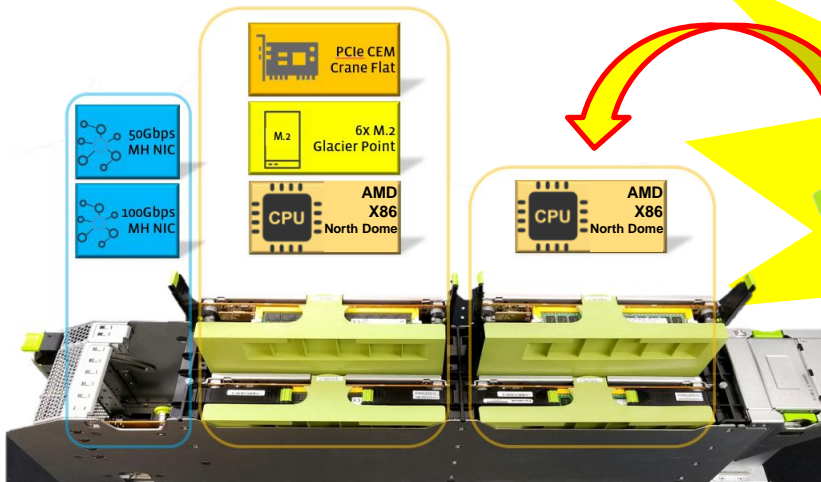


Yosemite V2 1-socket platform

- Yosemite V2 platform is CPU-agnostic
- North Dome AMD 1-socket server as alternative CPU module option



SERVER



North Dome
AMD CPU module



OPEN POSSIBILITIES.

Introducing North Dome

Meta 1-socket server based on AMD SP3 socket

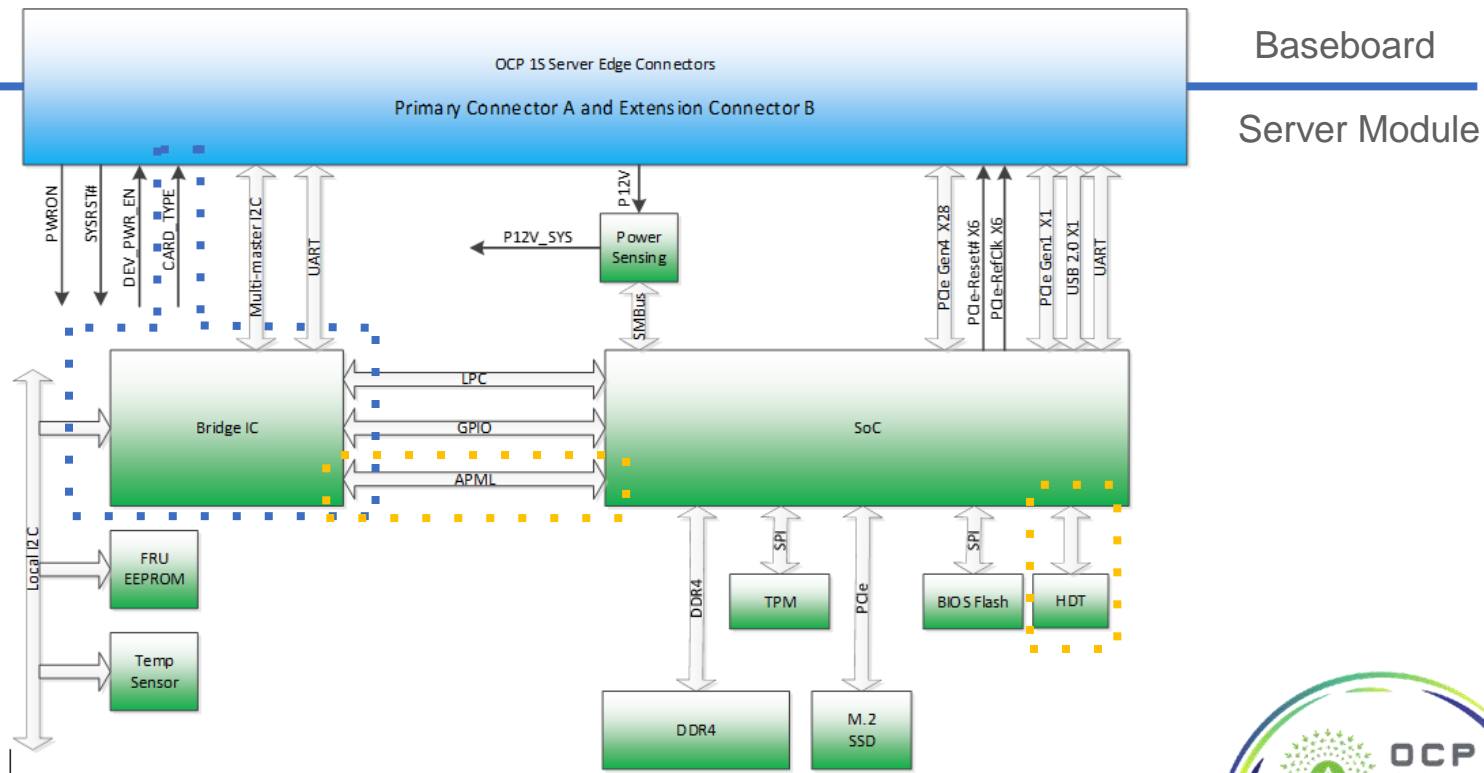


On-board Features	Details
CPU	AMD SP3 socket (3 rd Gen EPYC™ “Zen3” qualified)
DDR	Six channel DDR4 up to 3200MT/s
SSD	One 2280 NVMe x4 One 22110 NVMe x4

Expansion Features	Details
PCIe	Twenty-nine PCIe lanes expansion (x16 + x8 + x4 + x1)
USB	One USB2.0
I2C	One pair
UART	One pair

OPEN POSSIBILITIES.

North Dome Block Diagram



OPEN POSSIBILITIES.

Single Socket Compute Rack

- Optimized for performance efficiency
- CPU SKU customized for Performance/Watt
- Modular platform designs to enable compute, accelerator or flash configurations

OPEN POSSIBILITIES.

Blanking Panel			
TOR switch			
Blanking Panel			
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
Power Shelf			
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
Power Shelf			
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2
40U Yosemite V2	40U Yosemite V2	40U Yosemite V2	40U Yosemite V2



North Dome CPU Requirements

- Scale chiplet arch to lower TDP's for Yosemite
- Fit into existing YV2 platform
- Performance optimized
- Granular power efficiency enablement




SERVER

OPEN POSSIBILITIES.



AMD EPYC™ “Zen 3” Core Architecture



8 CORE
COMPLEX

Lower Latency

Significant Accelerator for HPC and Cloud Applications

2X*

Direct Access L3 Cache

Better Performance for Large Virtual Machines

2X*

Throughput for AI Inference

Doubling of INT8 Pipeline and Faster Floating Point

19%*

Core Complex IPC Uplift

Performance on SPECrate® 2017 at Same Power

OPEN POSSIBILITIES.

*When compared to EPYC™ “Zen 2” Core Architecture



Chiplet Architecture Enables Broad Scaling

Scalability by Design


- “Zen 3” architecture spans up-to 64c @ 280W
- High performance & throughput

Meta Challenge

- Best-in-class power efficiency
- Tight power envelope @ 95W
- Enough per-thread performance to meet latency

AMD created a custom device to meet Meta’s needs

- Built from standard 3rd Gen AMD EPYC™ components
- Provides 36c high-performance CPUs @ 95W



OPEN POSSIBILITIES.

Joint Investment & Optimization – Available To All Through OCP

Platform Design Optimizes Power Efficiency

- Tuning of power management algorithms for Facebook
- Board level optimizations to reduce power consumption

Performance Tuning

- Enhancing tool chain to increase performance
- OS and BIOS level optimizations

Joint Workload Optimization

- Optimizations for Facebook workloads on AMD
- Focus areas include AI/ML, web-tier, and database



OPEN POSSIBILITIES.

Call to Action

- Spec contribution is in progress and draft version available on Server wiki
- Design collaterals will be contributed by ODM shortly
- Multiple customizations on new chiplets architecture enable perf/watt optimized at low TDP
- Join us to improve performance/power/cost efficiency of datacenter server

Project Wiki with latest specification : <https://www.opencompute.org/wiki/Server/Working>

OPEN POSSIBILITIES.



Thank you!



NOVEMBER 9-10, 2021