

Microsoft Details Project Olympus Servers

Written by Marco Attard
10 March 2017

Microsoft presents the latest iterations of [Project Olympus](#), its hardware take on Facebook's Open Compute Project (OCP), with designs based on Intel, AMD, Nvidia and, surprisingly, ARM-based processors.



Posited by the Windows maker as the "de facto open compute standard," Project Olympus is an open source server design aimed at advanced cloud workloads such as big data, machine learning and AI. The 2017 OCP Summit has the company show off Project Olympus running on Intel Xeon (Skylake and future updates featuring Intel FPGA or Intel-Nervana accelerators) and next-generation AMD "Naples" processors, marking AMD's increased push the server industry.

Interestingly, Microsoft is also working on a "long-term project" with Qualcomm and Cavium involving ARM-based cloud servers. Qualcomm and Cavium have designed motherboards for the Project Olympus form factor using ARM chips, namely Centriq 2400 and ThunerX2 ARM v8-A. Both are integrated systems-on-chips with PCIe, SATA and gigabits of ethernet.

Also presented are collaborations with Nvidia and Foxconn subsidiary Ingrasys in GPU-based servers. Nvidia has HGX-1, a hyperscale GPU accelerator chassis for AI designed to support up to 8 Pascal GPUs and NVLink multi-GPU interconnect technology. The chassis is scalable, and can support high bandwidth interconnectivity for up to 32 GPUs by connecting 4 HGX-1 units.

"This is a significant moment as we usher in a new era of open source hardware development with the OCP community," Microsoft concludes. "We intend for Project Olympus to provide a blueprint for future hardware development and collaboration at cloud speed."

Microsoft Details Project Olympus Servers

Written by Marco Attard
10 March 2017

Go [Ecosystem Momentum Positions Microsoft's Project Olympus as De Facto Open Compute Standard](#)