

Facebook Reveals More OCP Servers

Written by Marco Attard
10 March 2017

Facebook builds on its Open Compute Project (OCP) with four new servers-- Bryce Canyon, Yosemite v2, Tioga Pass and Big Basin, all designed to serve different functions for the social network.



As part of the OCP initiative, all server designs are freely available, and interested parties can take them, modify them and have them built by OEMs, be it Chinese companies such as Quanta or big vendors of the HPE variety. Facebook says a number of major companies are involved with OCP, including Google, Microsoft, Apple, AT&T and Verizon, among others.

Bryce Canyon is a refresh of Open Vault, the Facebook storage server design. It is a storage chassis able to hold 72HDDs in four OU, an HDD density 20% higher than Open Vault. A modular design allows multiple configurations, and it supports more powerful processors and a larger memory footprint. Cooling comes from larger 92mm fans able to simultaneously cool the front 3 rows of HDDs, and the design is compatible with the Open Rack v2 standard.

Yosemite v2 is, as one might gather from the name, an update on Yosemite, a multi-node compute platform. Designed for high-density scale-out datacentre use, it consists of a multi-node compute platform holding four 1S cards. It is Open Rack v2-compatible, and uses a new 4 OU vCubby chassis with each cubby holding either four 1S server cards or two servers plus two device cards. Interestingly it supports hor service, meaning servers continue to operate even if a card is pulled out of the chassis.

On the server platform side Facebook has Tioga Pass, a dual-CPU update on Leopard. It has a dual-socket motherboard and supports single- and double-sided designs, with the double-sided design offering DIMMs on both PCB sides for maximum memory. The onboard mSATA connector gets an update to an M.2 slot, and the chassis is Open Rack v2-compatible.

Facebook Reveals More OCP Servers

Written by Marco Attard
10 March 2017

The final server is Big Basin, a GPU server designed for "larger and deeper" neural networks. Designed as a JBOG (Just a Bunch of GPUs) to allow complete disaggregation of CPU compute from GPUs, it does not have compute and networking built-in, requiring an external server head node. This allows one to connect OCP servers as a separate building blocks from the Big Basin unit, scaling each block independently as new CPUs and GPUs are released. In total the Big Basin server holds 8 GPUs (specifically Nvidia Tesla P100 GPU accelerators), and the design is Open Rack v2-compatible and occupies 3 OU of space.

Facebook says all these servers find good use-- users of the social network watch 100 million hours of video daily, post over 95 million photos and videos to Facebook and Instagram, and around 400 million people use voice and video chat on messenger.

Go [The End-to-End Refresh of Our Server Hardware Fleet](#)