



OCP 2019 Global Summit

Luca Atzori
CERN IT-CF

HEPiX
16/10/2019



Open Compute Project

- Organization that shares designs of data center products among companies
- Mission: *design and enable the delivery of the most efficient server, storage and data center hardware designs for scalable computing*
- Started by Facebook in 2011; other founders are Intel, Rackspace ¹
- Includes now ~195 members: Alibaba, AMD, ARM, Asus, Flex, Google, HPE, Huawei, IBM, Lenovo, LinkedIn, Mellanox, Microsoft, Nokia, Quanta, WD...



OPEN
Compute Project®

Source: www.opencompute.org

• ¹Also including: Goldman Sachs and Andy Bechtolsheim

Open19

- Specification that defines a cross-industry common server form factor
- Goal: *create a flexible and economic data center and edge solution for operators of all sizes*
- Started by LinkedIn in 2016, together with HPE and VaporIO
- Currently 25+ members: Flex, GE, Packet, Amphenol, Asus, Inspur, Mellanox, Molex, Supermicro, Wiwynn...



Source: www.open19.org

OCP Global Summit

- Yearly summit @ San Jose (CA) Convention Center
- 3600 participants from 727 organizations, 42 countries ¹
- 2 high-level keynotes, 13 Engineering Workshop tracks with 135 presentations
- Advanced Cooling, Data Center Facility, HPC & GPU/FPGA, Hardware Management, Networking – HW & SW, NIC 3.0 Development, OSF/Security, Rack & Power, Server, Storage, Telco & openEDGE



¹ <https://www.opencompute.org/news/by-the-numbers-2019-ocp-global-summit>

Keynotes (Microsoft)


Hardware innovations for data growth challenges at cloud-scale

Project Denali

- EDSFF Prototype
- Reduces costs of the non NAND part of the SSDs

Project Zipline

- Global Datasphere: 33 ZB in 2018, 175 ZB in 2025 ¹
- New compression standard: Algo + SW + HW
- Open sourced!



<https://github.com/opencompute/Project-Zipline>

¹<https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf>

Keynotes (Intel)

Taking OCP to the next level-workload optimization

High-Density Cloud-Optimized Platform


- 2U, 4S Xeon Scalable, 48 DDR4/Apache Pass slots
- Also 2/4/8S Cooper Lake Platforms announced

OCP Cards support new AI Accelerators

- Training: optimized memory and interconnect
- Inference: 10nm process, M.2 form factor

CXL: New CPU to Device Interconnect Standard

- Memory coherent, high-speed
- Optimized stack w/ x16 PCI Express Gen 5 (32 GT/s)
- Expected in 2021



Intel® NNP-L 1000 for Training

OCP NIC 3.0 specs (Intel)

<http://www.youtube.com/watch?v=QfL4wvGScVI>

- PCIe Gen5, fully OCP compliant
- Two form factors
 - Small x16 ASIC-based NICs
 - Large x32 Smart/programmable NICs
- PCIe bifurcation for multi-host support, greater density in server
- Hot swappable! (if platform supports it)
- 1GbE, 10GbE and 25GbE now available (XXV710 series)

<https://www.opencompute.org/wiki/Server/Mezz>



Source: Intel®

Experience Data Together (Seagate)

http://www.youtube.com/watch?v=uh3bH_FdFL0

- By the end of 2025, 80% of the enterprise bytes will continue to be HDD
- Enterprise data 13.6 ZB by 2025 (70% vs 30% of user)

HAMR

- 20+ TB HAMR drives in 2020
- no increase in drive temperature
- media heating is in nanoseconds
- 40+ TB in 2025?

SSD: pushing on QLC, Persistent Memory Regions (Logs, Journals)

Project Olympus 3U PCIe Expansion Server Deep Dive (ZT Systems)

<https://www.youtube.com/watch?v=fvT4grJdAE8>

- Microsoft leading provider for hyperscale data centers
- Project Olympus

3U PCIe Expansion Server with 3x 1kW PSUs




ZT Systems: ZT-XPO200-3UA

Source: ZT Systems

5 x FHFL, Double-wide
300W x16 PCIe Cards
+
1 x FHHL Single-wide
75W x16 PCIe Card

AMD® SKU



Source: ZT Systems



ZT Systems: ZT-XPO200-3UN

6 x FHFL, Double-wide
300W x16 PCIe Cards
+
1 x FHHL Single-wide
75W x16 PCIe Card

12 x FHFL, Single-wide
75W x16 PCIe Cards
+
1 x FHHL Single-wide
75W x16 PCIe Card

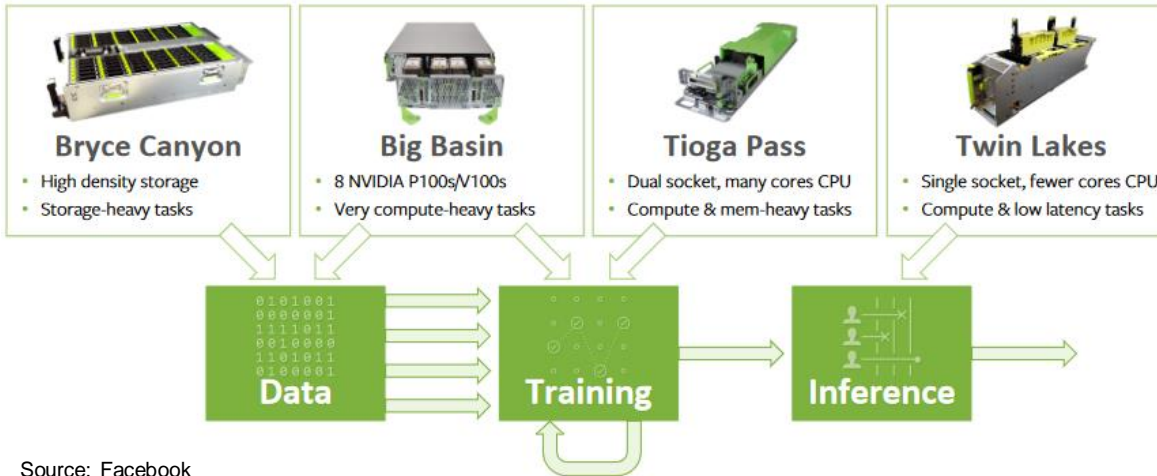
Intel® SKU

Facebook AI Infrastructure (Facebook)

<https://www.youtube.com/watch?v=MYICesArTWk>

Frameworks:

- **PyTorch** both for prototyping and deployment
- In terms of backend, they rely on **ONNX** and the **Glow** compiler



New: **King's Canyon**

- M.2 form factor Inference ASIC
- DRAM, 12W TDP
- inside a carrier card in a Yosemite v2 chassis


Source: Facebook

Memory Storage Innovations Fueling the Public Cloud (Samsung)

<https://www.youtube.com/watch?v=pf00ulyaB3g>

Flash Native Server (e.g. NF1)

- small size (30.5 x 110mm)
- hot plug
- high capacity (16+ TB)
- front panel cooling



Source: Samsung

Smart SSDs

- PM983F prototype w/ Xilinx FPGA
- hot plug
- 3-4x boost for IO bound workloads
- PCIe format now, U.2 end of the year



Source: Samsung

Tools and Process for creating a Redfish profile (DMTF)

https://www.youtube.com/watch?v=20mtQs3o7_M

- jsonschemavalidator.net
- Redfish-Interop-Validator (github.com/DMTF)
- Latest inter-op specs: <https://www.dmtf.org/dsp/DSP0272>

OpenBMC - Project Update (openBMC)


https://www.youtube.com/watch?v=qhp_uKeEcul

- Release 2.7 coming in August, 2.8 coming in February 2020
- Currently in 2.6: KVM over IP, Remote Media, IN-band Updates, SNMP Traps, IPv6, KCS, PECL, IPMB, IPMI2.0, fan control, LDAP, Certificate Management...

Keynotes (Linkedin)

Enabling High Density Computing Through Open19 Technology


Open19 specifications are now contributed to OCP



Source: Open19

Blind mate
400w power cables

Blind mate
100G data cables



Source: Open19

Keynotes (Schneider Electric)

- Computing to the edge, fully integrated modular/micro data center
- Cloud based management tool, powered by AI: downtime avoidance, data normalization, automatic updates, pay as you grow

Keynotes (Inspur)

- SR-AI Rack: Resource pooling for AI compute
- Three OCP Platforms:
 - NF8260M5 Flexible High Density 2U 4-Socket Serve
 - 8GPU JBOG (XM2)
 - 4-Socket Project Olympus 80cores + 16GPUs box



Open19 Platform Deep Dive (Open19)

<https://www.youtube.com/watch?v=uCNQWu89rK4>

- **Brick Cage:** 12RU and 8RU options, 2RU modularity, Snap-on rear opening
- **Power cables:** 400W per server – 8/12 servers per cable
- **Data cables:** 100G per server – 8/12 servers per cable
- **Bricks Form Factor:** Brick ($\frac{1}{2}$ wide 1RU), Double High Half Width (2RU), Double Wide Brick (1RU), Double High Brick (2RU)
- **Power Shelf:** 1RU, 19.2kW (6 x 3.2 kW), full AC/DR range, per-server protection and monitoring
- **Future plans:** different cage form factors, new server models, automated data center deployment model, optical snap-on data cabling, liquid cooling...
- **Some numbers:** Infra 40-60min per rack, 96 servers: 10 minutes per rack, Total ~100min per rack (2 technicians), 2 days for a full cluster – 1536 Servers

Conclusion

- Nice event, great location, interesting talks!
- Many interesting products at the booths in the Expo Hall
- Can we benefit from it?
 - OCP and Open19 modularity and flexibility are key features
 - Many more available products nowadays

Questions?

Disclaimer

All product names, trademarks and registered trademarks are property of their respective owners. All company, product and service names used in this presentation are for identification purposes only. Use of these names, trademarks and brands does not imply endorsement.



Backup (I)

OCP Accelerator Module (OAM) System: An Open Accelerator Infrastructure Project (Facebook, Baidu, Microsoft)

Video <http://www.youtube.com/watch?v=kIHLNDqdVjY>

Open19 Technology (Linkedin)

Video <https://youtu.be/tTAtwLffXFI>

Leverage OCP Design Advantages on EIA 19" GPU Server (Wywinn)

Video <https://youtu.be/p1we2BwlPkw>

OCP AI Accelerator Module for Deep Learning Training with Intel Nervana Neural Network Processor (NNP) (Intel)

Backup (II)

OCP Design for Huawei Cloud Infrastructure (Huawei)

Breaking Barriers in AI: New Hardware, Standard Platforms (Intel)

Video <http://www.youtube.com/watch?v=wTOfFYU8YXc>

F16: the next-generation fabric (Facebook)

Video <https://youtu.be/OXmrSaEcoUM>

Slides <https://146a55aca6f00848c565-a7635525d40ac1c70300198708936b4e.ssl.cf1.rackcdn.com/images/eb47e324d4bdb8980b0dfd491097ca996bd875d1.pdf>