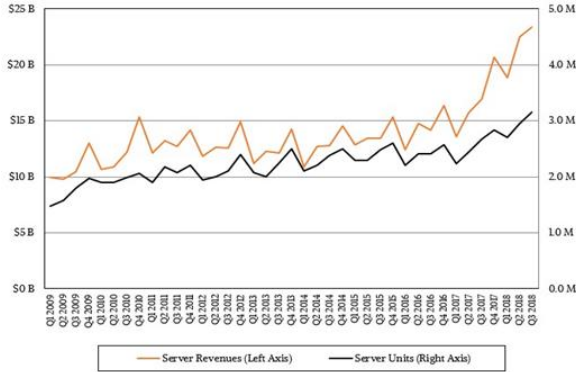


HEPiX TechWatch WG: Server

Chris Hollowell, Brookhaven National Laboratory
Michele Michelotto, Università e INFN, Padova (IT)
On behalf of the HEPiX Technology Watch Server Market Subgroup

HOW Workshop
18-22 March 2019
Jefferson Lab, Newport News

Server Market




Global server market revenue up substantially

- \$23.37 billion in Q3 2018 with 3.16 million servers shipped
- Primarily due to hyperscale cloud provider purchases
- May grow to \$100 billion/year industry

Revenue outpacing units sold

- Increased cost of components including CPU, GPUs/FPGAs, memory and solid state storage




Total server shipments per quarter now almost double what they were for an entire year ~20 years ago (1997)

- Vast majority of revenue from x86 servers
- Other server platforms only sold \$1.6 billion in Q3 2018: ~6.8%

New ARM Neoverse server CPU platform hopes to change the downward trend for non-x86 servers

Server Market (Cont.)



ODMs (Original Design Manufacturers) combined have a larger server sales revenue than any single tier one vendor

These servers are often purchased by hyperscale cloud providers Some examples ODMs include Foxconn, Quanta, and Inventec

Approximately one in four servers shipped now going to hyperscale cloud providers

Dell, HPE and Lenovo, respectively, leading tier one server vendors in terms of sales revenue

AMD-based server shipments still a small portion of the total market However was 0% prior to EPYC Naples' release

AMD anticipates capturing additional server market share after the 2019 release of EPYC Rome CPUs

Hyperscale cloud providers currently most heavily driving AMD server sales

Server Market - Environment and Cooling

Data centers used ~3% of the world's total electricity output in 2017
Estimated to jump to ~20% by 2025

US DoE recommending EPEAT environmental certification for servers
"Electronic Product Environmental Assessment Tool"

<https://greenelectronicscouncil.org/epeat/epeat-overview/>

Both Dell and HP participating in the program, among others

Rates products on 51 environmental criteria, including recycled content, toxin reduction, energy efficiency, etc.




Figure 1.8: Approximate distribution of peak power usage by hardware subsystem in a modern data center using late 2017 generation servers. The figure assumes two-socket x86 servers and 12 DIMMs per server, and an average utilization of 80%.

Traditionally, there's been reluctance and little need for data centers to adopt liquid cooling technologies, although they're considerably more efficient than air-based cooling
Impediments for adoption including cost and increased maintenance complexity

However, expect to see more demand for both rear-door heat-exchange and direct liquid cooling in 2019

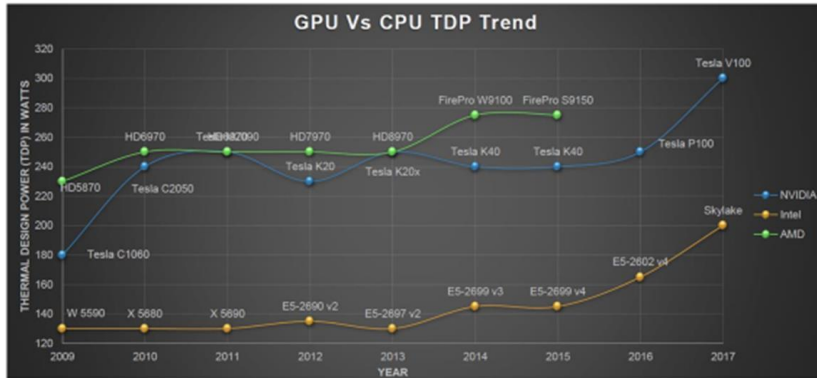
High demand for co-processors (GPU, FPGAs, etc.) increasing the power density of servers

Increasing use of "hyperconverged infrastructure" - combining separate network, compute and storage elements into a single higher-density systems


Cooling

- Common density 10-12 kW/rack max to 12-20 kW/rack
- Today a 42U rack full of 2S nodes (80 server) can draw 40 kW
- Full air cooling is not enough

TDP trends



Source: Alibaba.com - Immersion cooling for Green Computing - OCP2018



Datacom Equipment Power Trends and Cooling Applications --ASHRAE


Liquid Cooling Landscape

Close Coupled Cooling

- In-Row, In-Rack, Rear Door Exchangers

Close Coupled Cooling

- Air is still the only mean to cool board/chip
- Limited by existing chip maximum temps
- Requires additional fans



DLC (Cold Plates)

- Positive Pressure, Negative Pressure

Cold Plates

- Individual Heatsinks
- Board-Specific
- Individual Chip Fluid-cooling



Immersion Cooling

- Single Phase, Two-Phase

Immersion

- Specialty fluids (\$\$\$)
- Can require separate cooling coil
- Orientation sensitive




Open 19

- Aims at support data center design, open, economical and customizable
- Similar to Facebook OCP Open Compute Project that is more addressed to major internet companies
- Open19 focused on optimizing data center of any size including small edge platforms, unlike OCP
- Define a common, open standard server form factor, using rack, cages and pre-defined network and power



Open 19

- Standard 19" 4-post rack, Brick cage
- Brick (B), Double Wide Brick (DWB), Double High Brick (DHB), Double High & Wide Brick (DHWB)
- Power shelf: 12v distribution, OTS power modules with any AC or DC inputs
- Optional Battery Backup Unit (BBU)
- Networking switch (ToR)
- Snap-on power cables – up to 400w per brick, linear growth with size
- Snap-on data cables – up to 100G per brick, linear growth with size



Open Computing Project



- Server compute node designs included one for Intel processors and one for AMD processors. In 2013, Calxeda contributed a design with ARM processors
- Several subprojects:
 - HPC: heterogeneous computing networking and fabric platform for multi-node processor
 - Networking: Network hardware and software
 - Rack and Tower: rack standards, integrating the DC infrastructure
 - Open System Firmware: OS Hw platform initialization
 - Data Center Facility: maximizes mechanical performance and thermal and electrical efficiency



Left - Open Compute V2 Server



Right - Open Compute V2 Drive Tray, 2nd lower tray extended

References

- <https://www.nextplatform.com/2018/12/14/the-vital-engines-of-commerce/>
- <https://www.nextplatform.com/2019/01/30/amd-nails-its-epyc-server-targets-for-2018/>
- <https://www.anandtech.com/show/13959/arm-announces-neoverse-n1-platform>
- <https://www.crn.com/arm-takes-on-intel-with-neoverse-platforms-for-edge-cloud-and-5g>
- <https://rcpmag.com/blogs/scott-bekker/2018/12/server-market-hits-record-high.aspx>
- <https://technology.ihs.com/608048/the-canary-in-the-coal-mine-for-liquid-cooling-is-server-shipments-with-coprocessors>
- <https://www.amazon.com/gp/feature.html?ie=UTF8&docId=1000405461>
- <https://www.epeat.net/profile/us-department-of-energy/>
- <https://www.forbes.com/sites/forbestechcouncil/2017/12/15/why-energy-is-a-big-and-rapidly-growing-problem-for-data-centers/#19be27a55a30>
- <https://data-economy.com/data-centres-world-will-consume-1-5-earths-power-2025/>
- <https://www.datacenterknowledge.com/power-and-cooling/data-center-power-and-cooling-trends-watch-2019>
- https://www.theregister.co.uk/2017/05/24/open19_foundation_launch/
- <https://www.datacenterknowledge.com/edge-computing/why-open19-being-designed-edge-data-center-standard>
- <https://www.opencompute.org/projects>
- https://en.wikipedia.org/wiki/Open_Compute_Project
- <https://www.opencompute.org/files/Immersion-Cooling-for-Green-Computing-V1.0.pdf>