

**Hewlett Packard
Enterprise**

Further Updates From the Coalface of Arm in HPC

Andy Warner, Distinguished Technologist

Arm Research Summit | 16 Sept 2019 | Austin, TX

HPE and ARM: Long-standing Partnership



2012: Redstone

- Calxeda low power CPUs
- 288 nodes in 4U
- ARMv7 32bit
- 4 cores, 1.4GHz

2014: Moonshot


- Calxeda, TI, Applied Micro
- 45 XGene cartridges in 4U
- ARMv8, 64 bit
- 8 cores, 2.4 GHz

2016: "The Machine" (prototype)

- Broadcom Vulcan CPU
- 160 TiB of addressable memory
- Gen-Z fabric
- Fabric Attached Memory
- Integrated fabric optics

2017: Commanche

- Cavium TX2 Early Access
- Four 2P nodes in 2U
- 32 cores, 2.2GHz



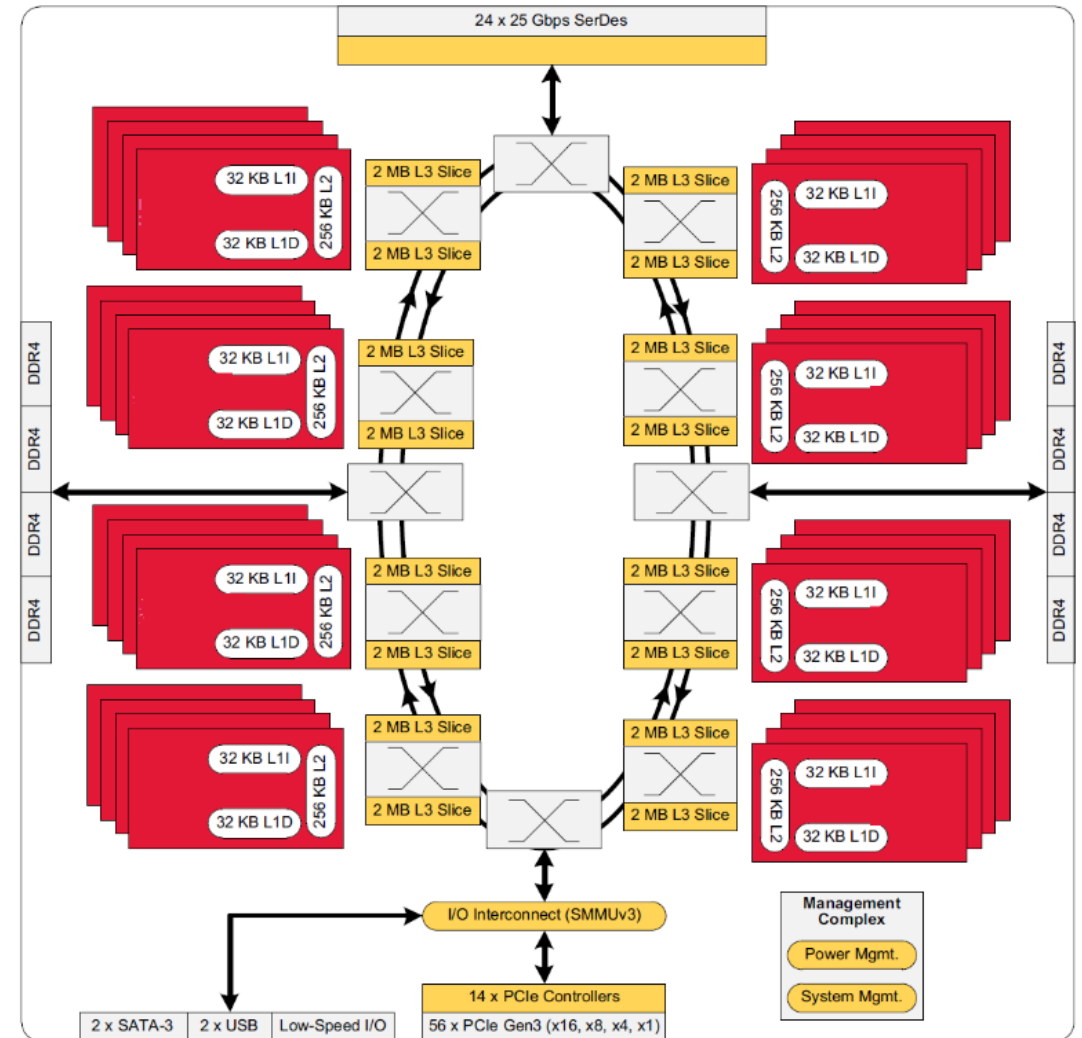
2018: Apollo70

- 28c & 32c SKUs offered
- Astra - Top500 system
- CatalystUK

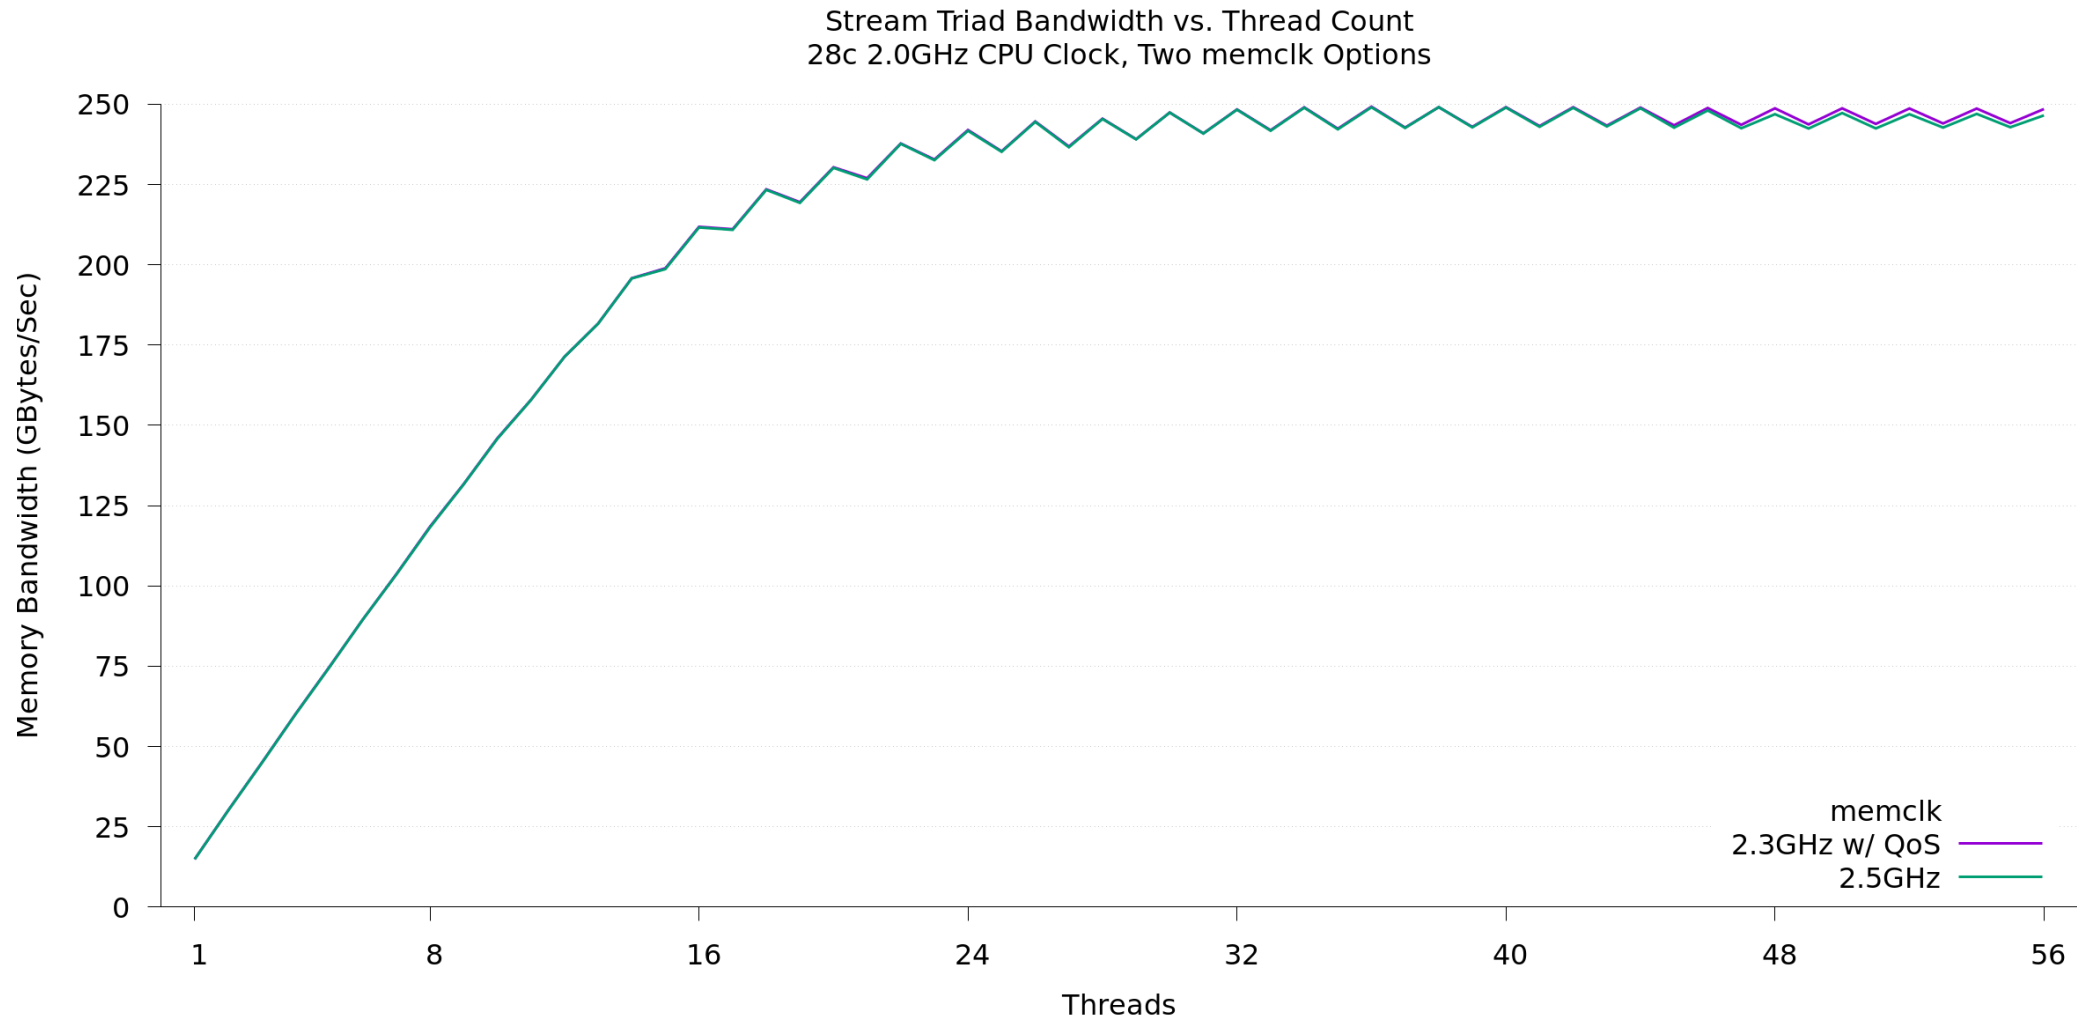


Marvell TX2 L3 Throughput

- Single configurable clock used throughout the L3 subsystem (“memclk”)
- Memclk not *directly* related to DIMM speeds
- 2.5GHz used throughout early deployments providing maximum performance with DDR4-2667 memory
- QoS capability designed into silicon to give small messages from GCU to local CPUs priority over cache lines
- Applications routinely decomposed to maximize NUMA-awareness, minimize the inter-socket cache-line traffic, increasing the proportion of NAKs
- Extensive testing demonstrated that memclk could be reduced to 2.3GHz while maintaining maximum memory bandwidth when QoS is enabled
- Power savings and corresponding thermal margin can be redirected to CPU cores, or banked
- 2.3GHz (with QoS enabled) is now the maximum memclk across all TX2 SKUs




Marvell TX2 L3 Throughput




Technology Demonstration: Liquid-Cooled Apollo 70

- CoolIT cold plates
- Unmodified A2K enclosure, firmware & software
- Displayed at ISC in Frankfurt




Technology Demonstration: Liquid-Cooled Apollo 70

CPU Clock Freq, Power & Die Temp vs. Time
28 Core, 2.0GHz, Turbo Enabled, HPL (866 GFLOPS)




Technology Demonstration: Liquid-Cooled Apollo 70

CPU Clock Freq, Power & Die Temp vs. Time
32 Core, 2.2GHz, Turbo Enabled, HPL (1,078 GFLOPS)



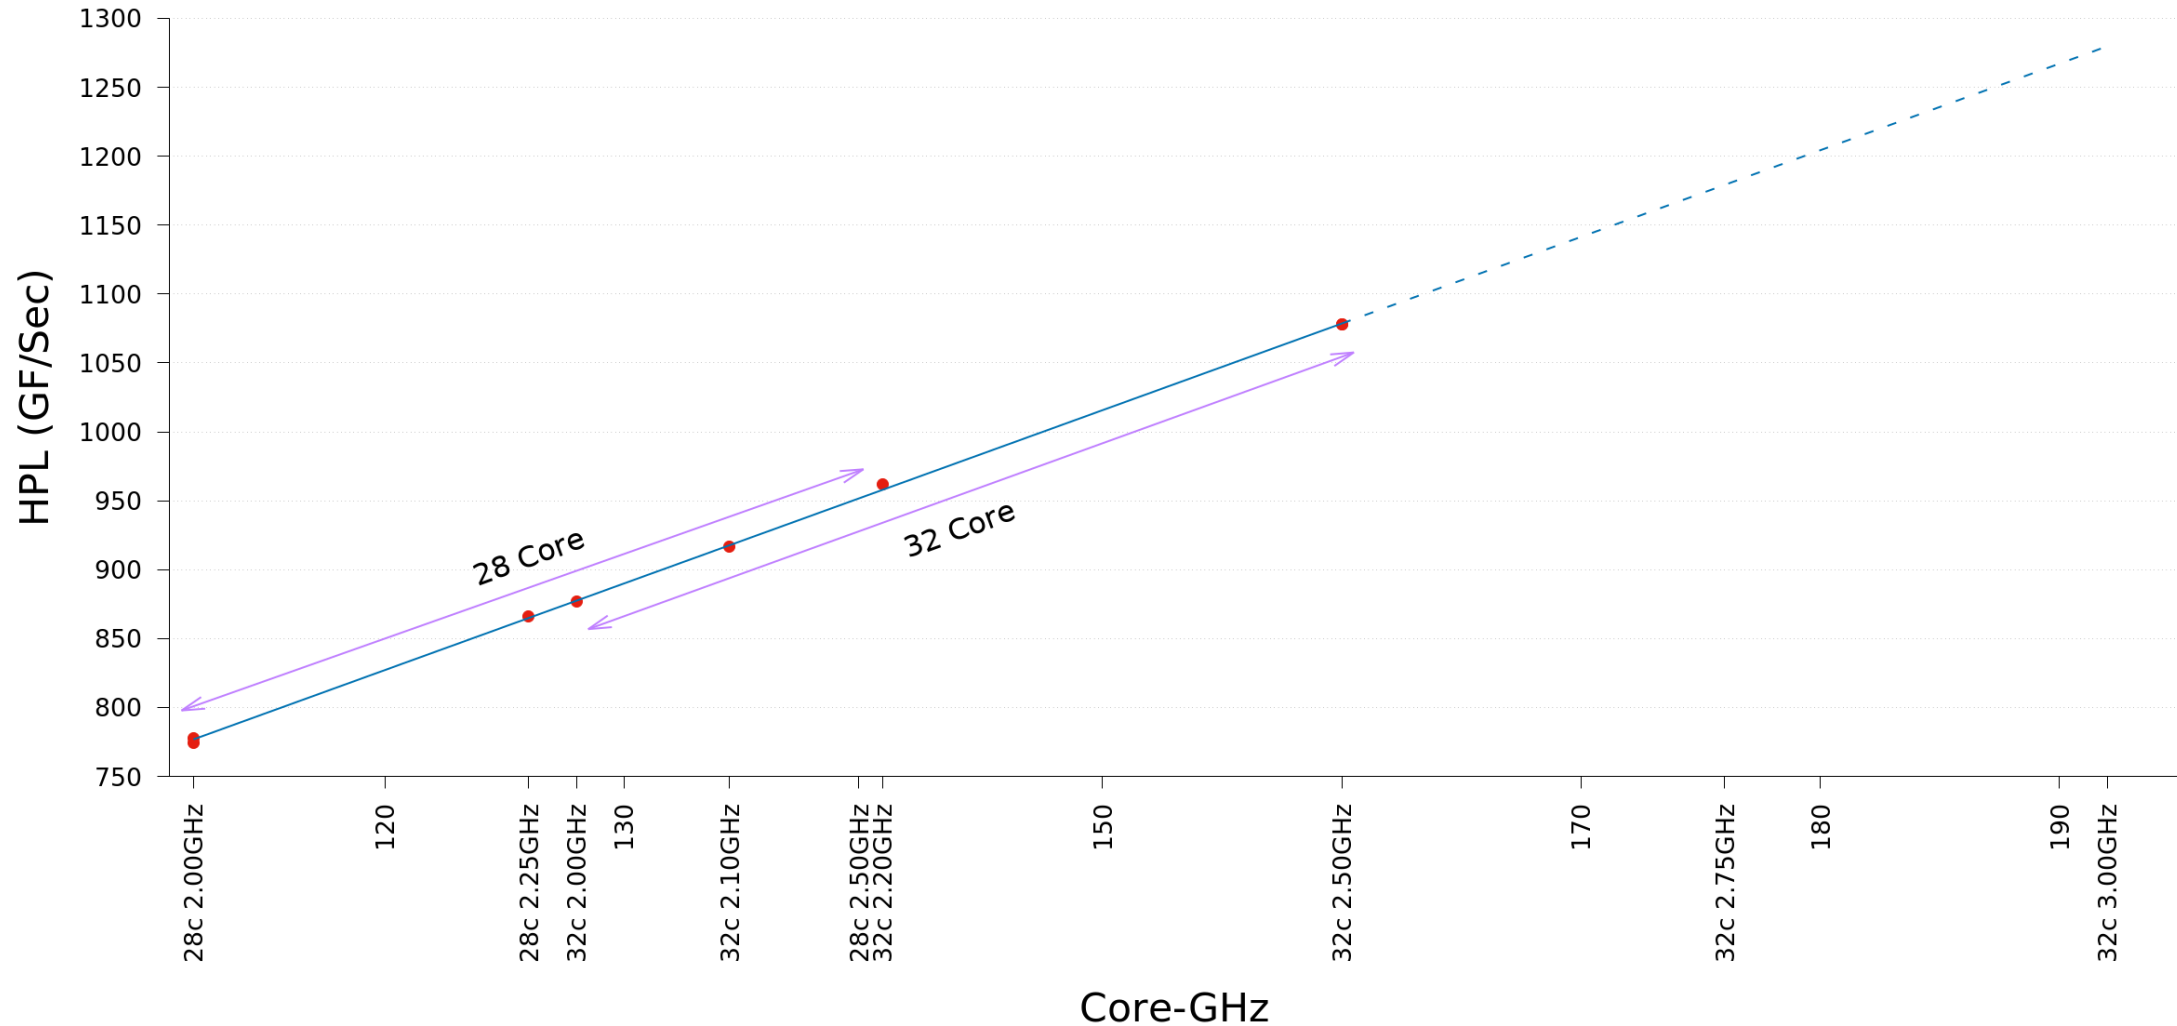
Technology Demonstration: Liquid-Cooled Apollo 70

CPU Clock Freq, Power & Die Temp vs. Time
32 Core, 2.5GHz, Turbo N/A, HPL (1,078 GFLOPS)



Technology Demonstration: Liquid-Cooled Apollo 70

TX2 Single Node (2 Socket) HPL Performance vs. Multiple Core Counts & Clock Speeds



Direct liquid cooling enables full theoretical performance

Technology Demonstration: Liquid-Cooled Apollo 70

Conclusions:

- DLC enables continuous operation at max turbo speeds
- 28c 2.0GHz (150W) SKU performance throttled by power cap
- 32c 2.2GHz (180W) SKU performs identically with turbo enabled to experimental 2.5GHz SKU due to continuous turbo capability
 - Significant headroom available (50°C thermal, 25-45W electrical power)
- Fans @ 40% throughout test (minimum speed for unmodified firmware), could reduce further in production

Weak Ordering Exposes Latent Software Issues

- In the main, delivering an HPC software ecosystem has proven remarkably boring*. However, weak ordering is doing what it so often does...
- For performance, MPI libraries often poll for messages or status. One recent bug in HPE MPI was uncovered on Astra using the IMB-MPI1 Exchange benchmark. At least 72 nodes required to reproduce. Resolution: additional barrier operation required to ensure message length read from work queue correctly in the case of small messages. This code path has previously supported billions of core-hours (on platforms with strong ordering) of operation without error.
- Linux Kernel is not immune, despite all the exposure and testing it has already received. High core and node count of HPC systems works to exacerbate latent problems. For example:

```
[172322.746745] [<ffff0000082cd5cc>] dcache_readdir+0x9c/0x170
[172322.752306] [<ffff0000082b4fd8>] iterate_dir+0x150/0x1b8
[172322.757691] [<ffff0000082b5780>] SyS_getdents64+0x98/0x170
```

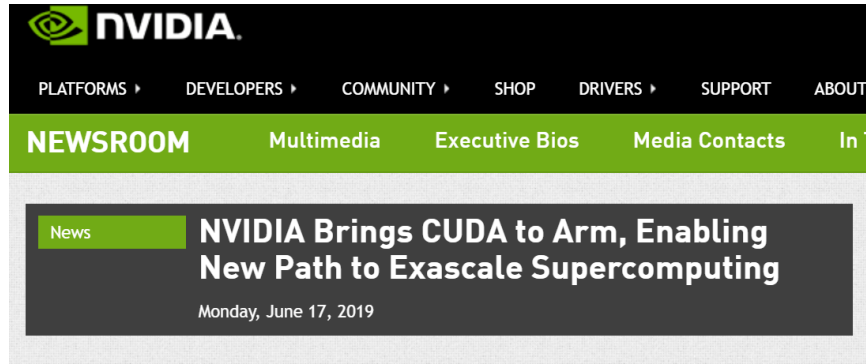
a.k.a. Redhat BZ 1702057 - Kernel panic on job cleanup, related to SyS_getdents64

This is actually a series of fundamental problems, which appear to be benign or vanishingly rare on strongly ordered, or low core count, processors.

See: <https://lore.kernel.org/linux-fsdevel/20190910215357.GH1131@ZenIV.linux.org.uk/>

- Similar issues surfaced in 2011 with the iPad2, which had a dual-core Cortex-A9 CPU.

NVIDIA GPU Update



Global HPC Leaders Join to Support New Platform

International Supercomputing Conference -- NVIDIA today announced its support for Arm CPUs, providing the high performance computing industry a new path to build extremely energy-efficient, AI-enabled exascale supercomputers.

HPE is working closely with NVIDIA to add their GPUs to the Arm HPC ecosystem

- First platform is Apollo70 with 2 GPUs per node.
- PCIe Gen3 x16 links connecting each GPU to a TX2
- Two nodes and four GPUs per 2U enclosure
- Look for **much** more information in November at SC19 in Denver, CO.

