ARM: Power-efficient Compute for HPC

CERN OCP HPC Conference

Darren Cepulis
darren.cepulis@arm.com

The Architecture for the Digital World®

ARM®

Over 12B ARM Based Chips Shipped in 2014

	2013	2014	Growth	
Mobile	5.1bn	5.4bn	300m	
Embedded	2.9bn	4.1bn	1.2b	
Enterprise	1.8bn	1.9bn	100m	
Home	0.6bn	0.6bn	-	

ARM CPU Core Unit Shipments

ARM Architecture: Licensing Overview

350+ customers

1000+ overall licenses

1B smartphones sold in 2013

50B chips shipped

121 new licenses sold in 2013

200+ Cortex-M licenses, over 160 companies

10B ARM based chips shipped in 2013

620 licenses sold in past five years

50 ARM v8-A licenses

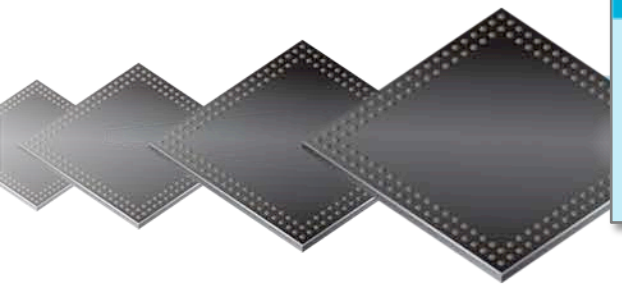

Hundreds of optimized system-on-chip solutions

ARM®

Extensible Architecture for Heterogeneous Multi-core Solutions

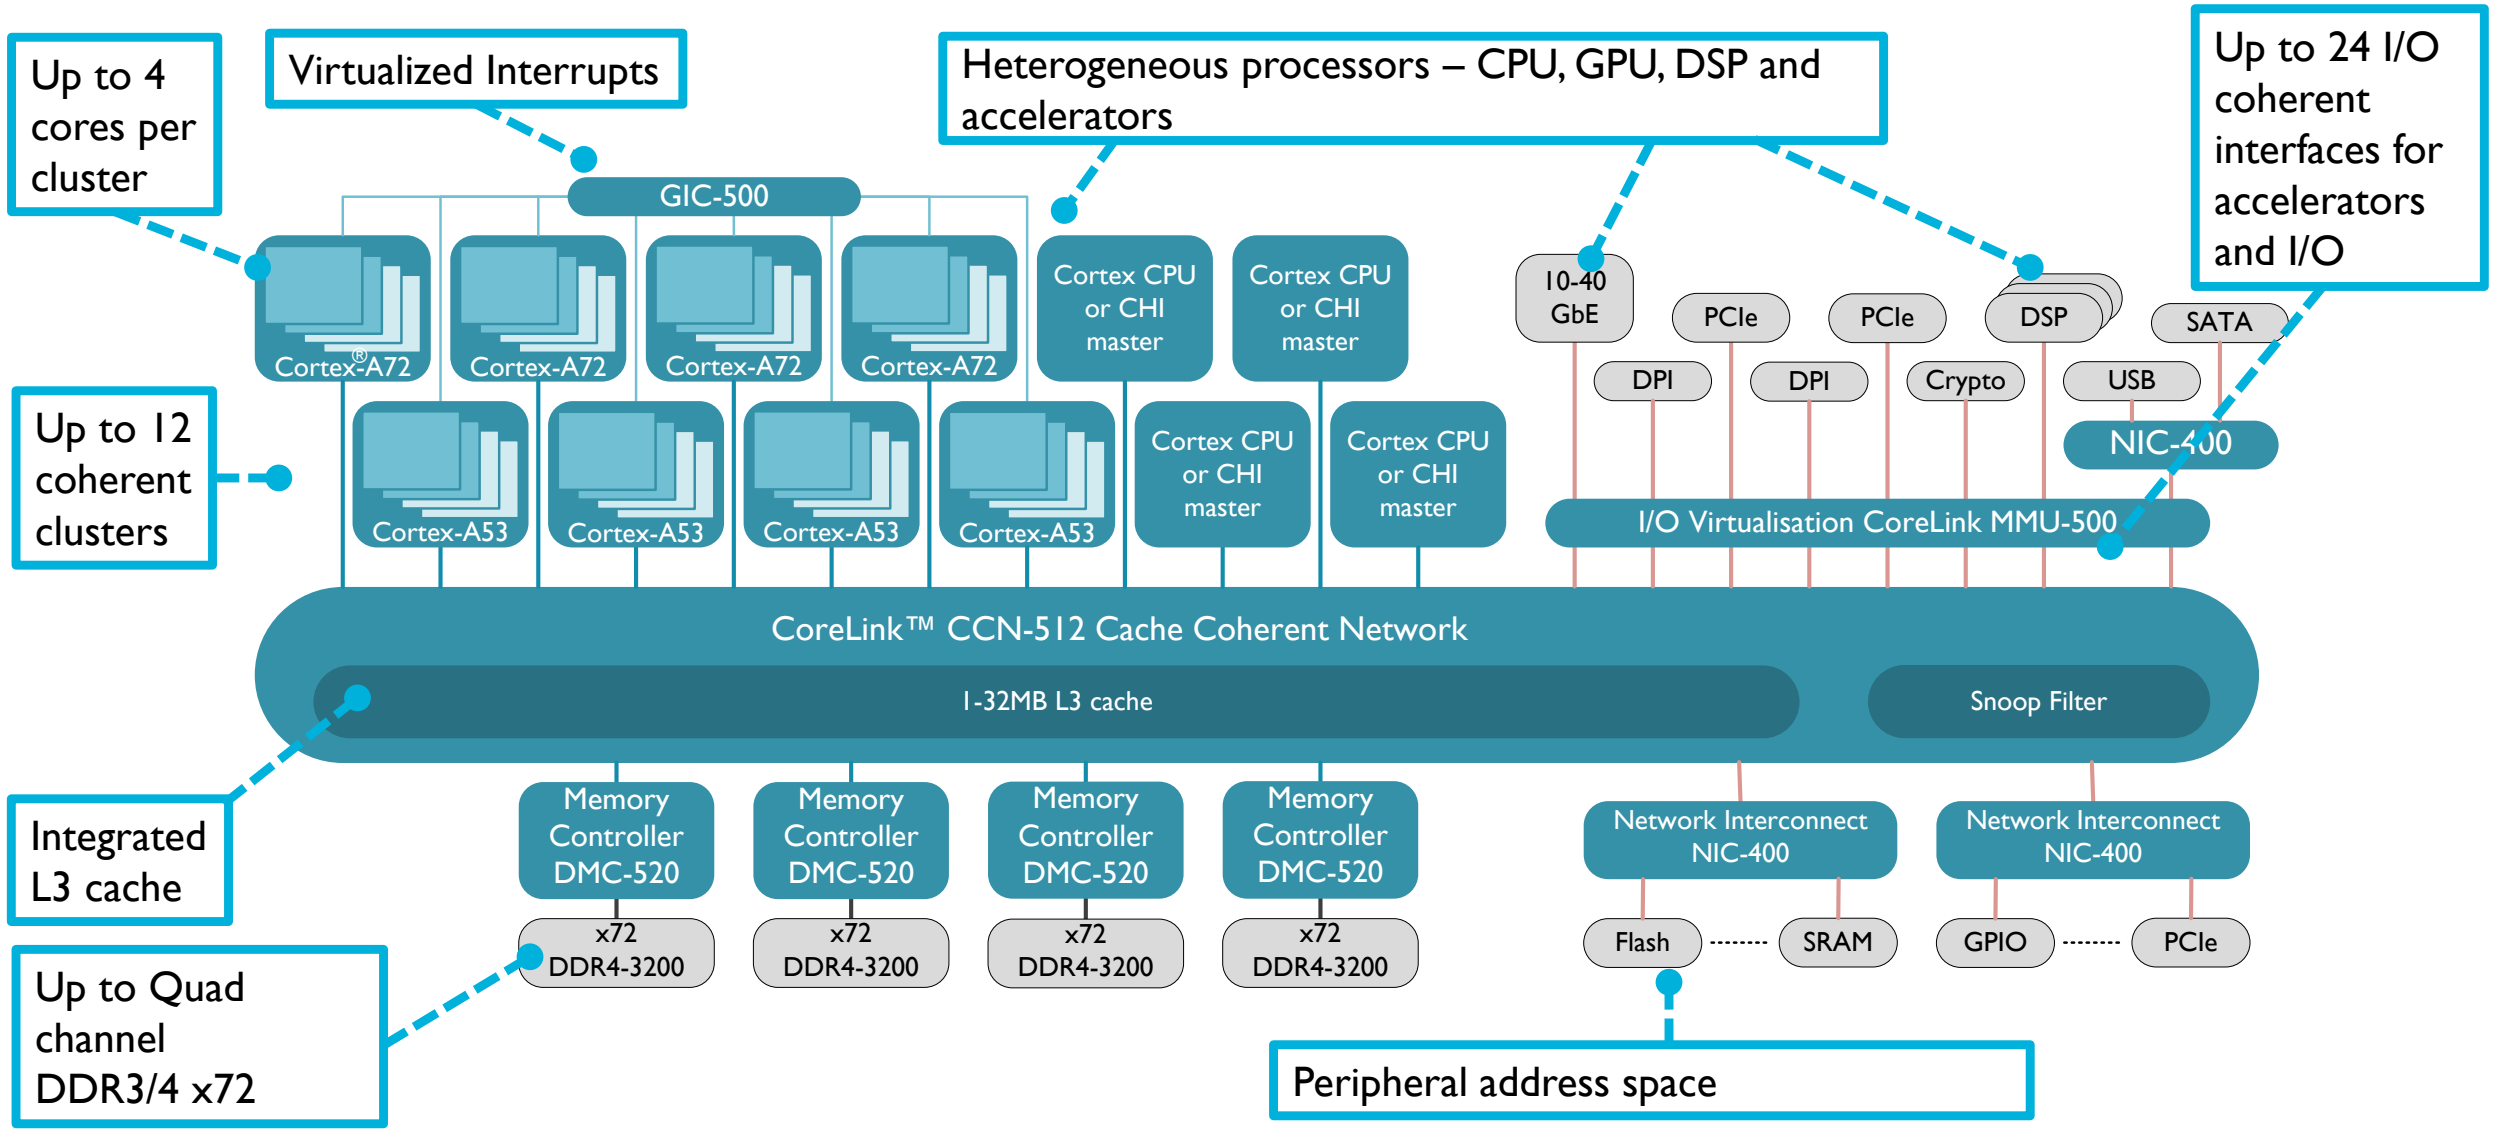

ARM in Datacenter/HPC Compute

- ARM-based SOC's enable compelling solutions:
  - Optimized performance/watt, performance/RU
  - Workload optimized balance of CPU, memory, cache, and IO
  - Application specific HW accelerators
  - Heterogeneous compute (DPS's, FPGA's, GPU's, etc).
  - Comprehensive SW Ecosystem
- Built to existing datacenter standards
  - ARM platforms use standard motherboard form-factors or chassis/rack form-factors
  - Standard interconnects (PCIe, RapidIO, etc.)
  - Standard platform firmware abstractions for deployment and management

Compelling TCO

Performance/Watt/$/RU

Flexibility and Workload Optimization

Choice and Innovation

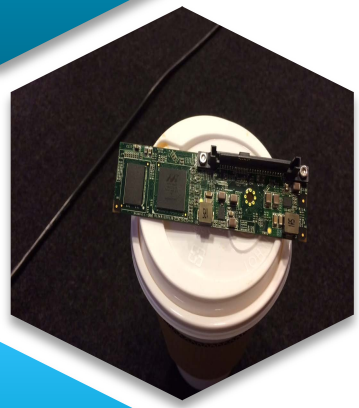

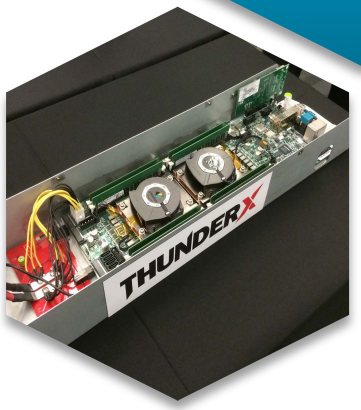

Standards based Ecosystem

Platform Management and Deployment

ARM®

ARM: All the Pieces for HPC

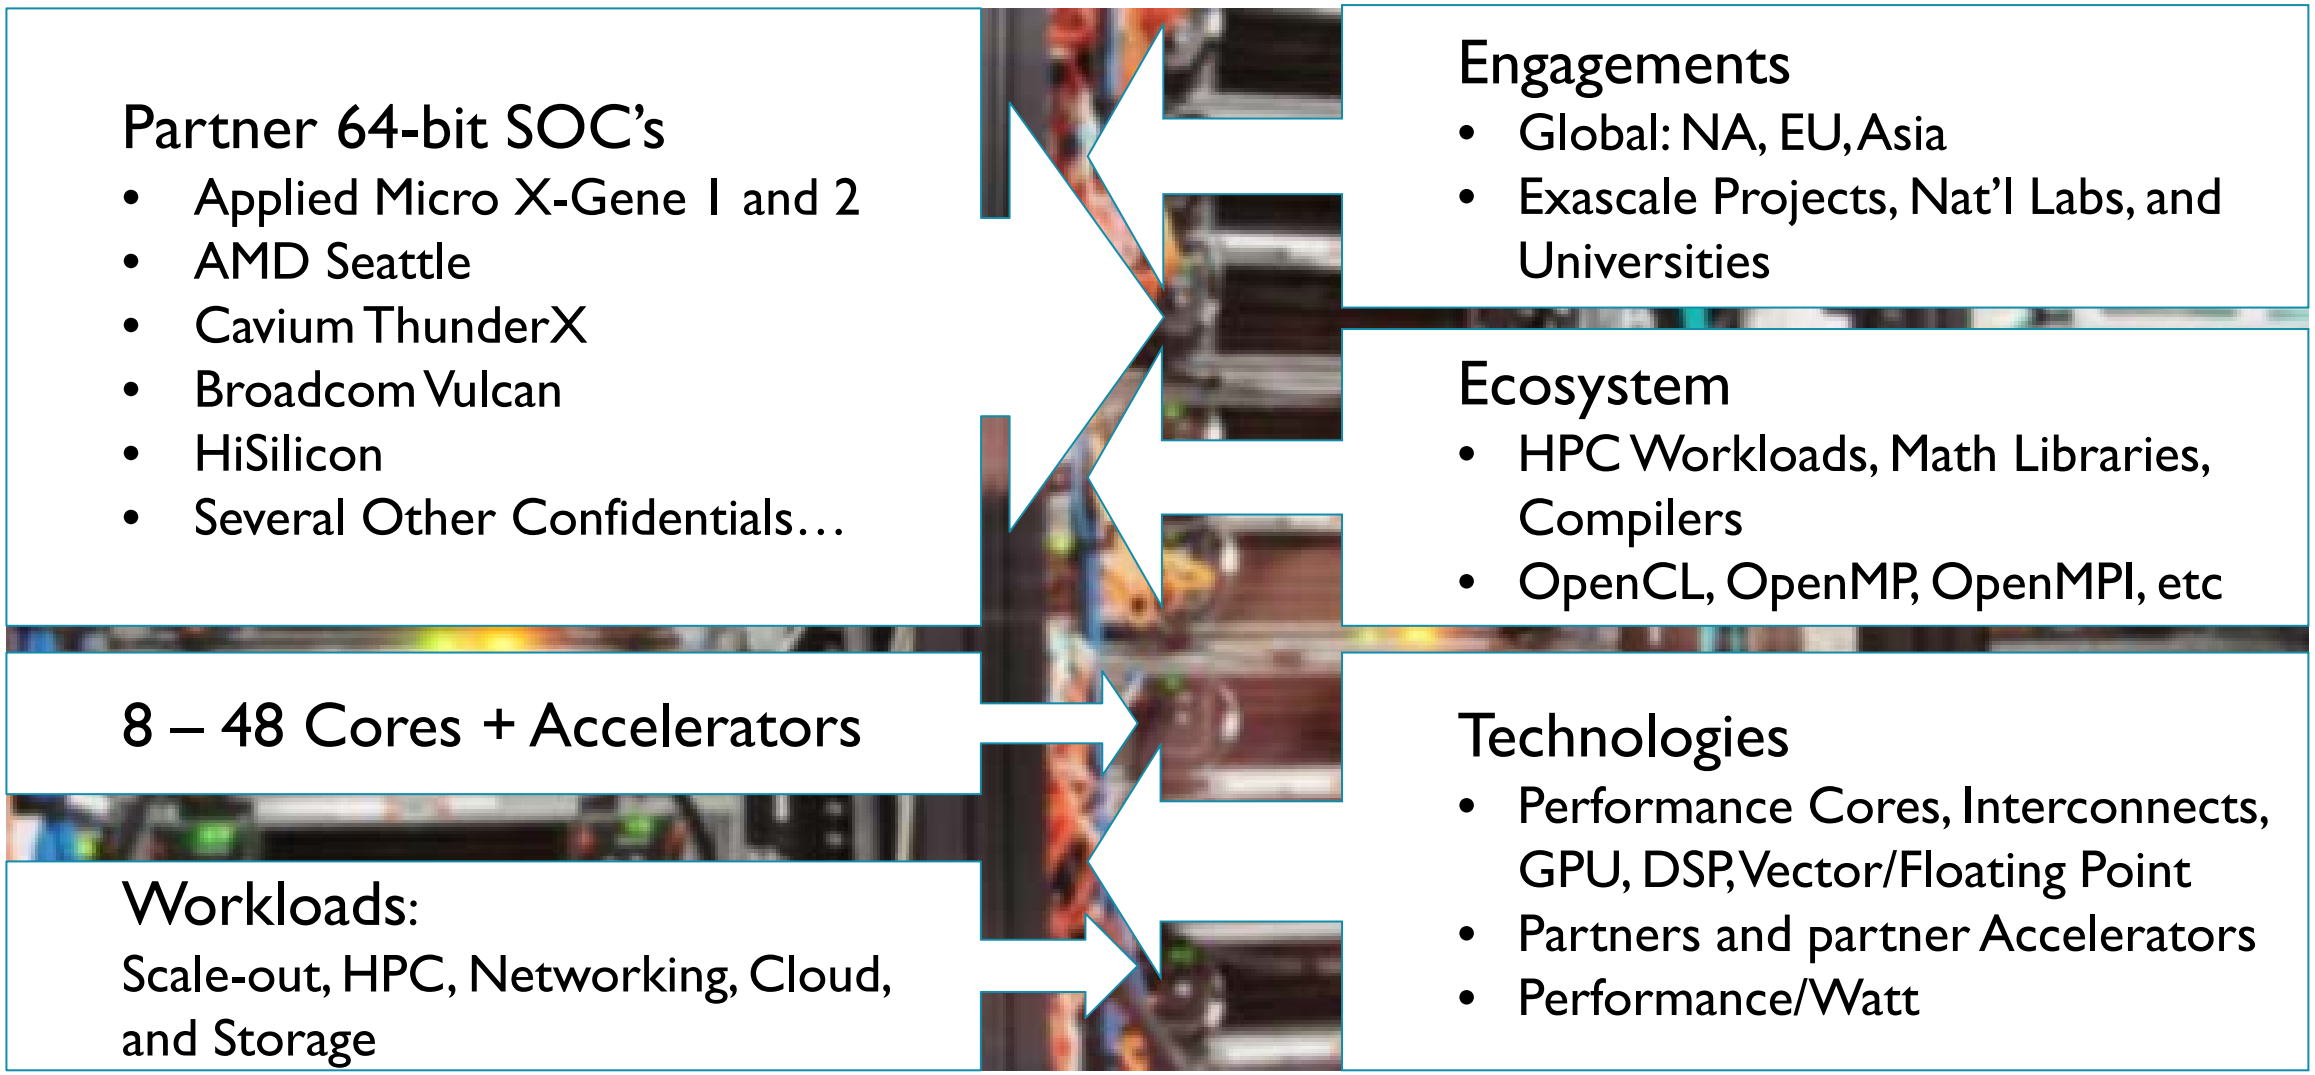

ARM HPC Software Ecosystem Overview

	Compilers C/Fortran	Analysis Tools	Math Libraries	Parallelism Libraries	Parallel File Systems	Cluster Management and Test SW
X86 Ecosystem Compare =>	ICC, Pathscale, PGI, NAG, GCC, Proprietary	Parallel Studio Allinea, Rogue Wave Cluster Studio XE, IBM Toolkit, Vtune	MKL (BLAS), BLIS, libATLAS, Eigen BLAS, ACML, FFTW	TBB, OpenMP, SequenceL, DistrParallelism: OpenMPI/PGAS	Lustre, Panasas OpenSFS, HDFS, Ceph, IBM GPFS	HP CMU, IBM Platform LSF, Adaptive Computing (Moab) Altair PBS Works
Supercomputer Vendors & Labs	Internal or proprietary Compilers	Vendor or Customer SW Stack	Internal or proprietary	Vendor or Customer SW Stack	Vendor SW Stack	Vendor or Customer SW Stack
Enterprise Volume HPC (Commercial SW)	GCC, LLVM Pathscale EKOPath, NAG	Allinea DDT Rogue Wave ARM DS-5 + PTP	Pathscale BLAS NAG Numerical	ARM OpenMP GCC Libgomp, Libtbb OpenMPI	Lustre Client/Server IBM GPFS Panasas	IBM LSF supported, HP CMU supported, Altair PBS and Moab
HPC Labs, Universities (Open Source SW)	GCC & Clang, Gfortran LLVM	TAU, Perf, HPCToolkit PTP	BLIS, FFTW libATLAS, EigenBLAS OpenBLAS	GCC, OpenMP LLVM run-time PGAS	OpenSFS (Lustre) Lustre Client/Server HDFS, CEPH, Gluster	SLURM OpenLava (LSF fork)

Maximizing Throughput Density: per mm^2, per Watt

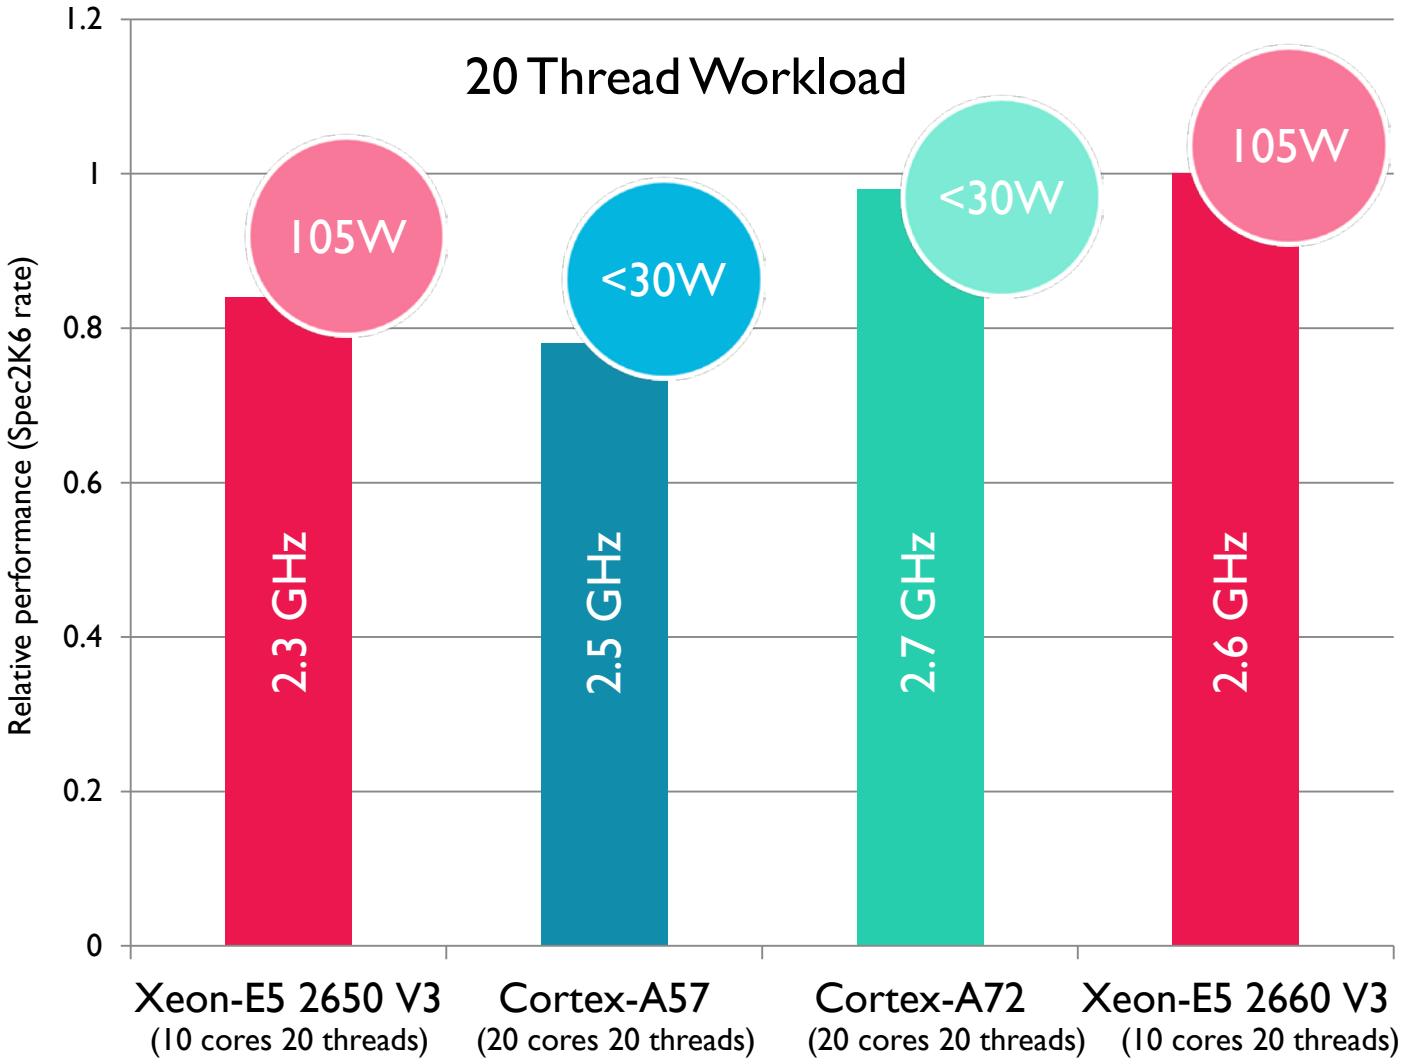

ARM Solution Benefits:

- Less than 1/3rd the power for equivalent performance
- Allows more specialized computing or significantly greater thread density in the same power budget

Comparison for equivalent number of threads

- Platforms used:
  - Xeon-E5 2660 v3 10C20T platform (measured)
  - Xeon-E5 2650 v3 10C20T platform (measured)
  - Gcc compiler v4.9 with –o3 flag
- Estimated result on example 20C ARM Cortex platforms with CCN-508, 28MB total L2+L3 cache
  - per-core measurements on RTL with relevant memory system
  - Gcc compiler v4.9 with –o3 flag
  - Scaled to 20T based on modelled and empirical results
  - Power estimated in 16nm based on ARM internal implementations for entire CPU+ interconnect

ARM®

Cortex-A72: Ideal for dense compute environments

Single Cortex-A72 core [2]
~1.15mm²

Cortex-A72 is <20 % size

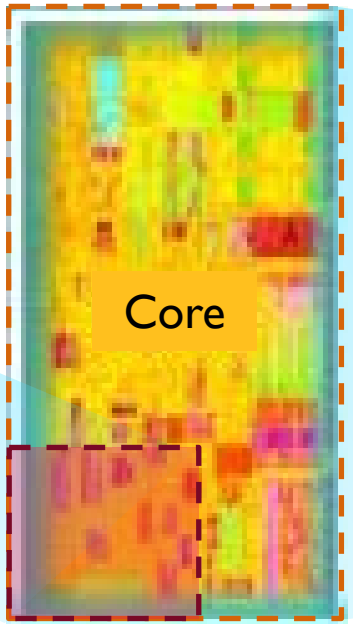

Single Broadwell CPU + 256K[1] L2
~8mm²

Cortex-A72 MP4 + 2MB L2[3]
~8mm²

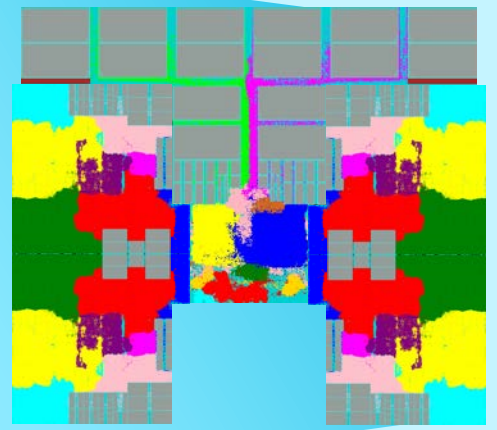

A quad core Cortex-A72 with 8x L2 cache RAM is the same size

[1]Source: Estimated from die-shot image provided by Intel at IDF 2014.
[2/3]Source: ARM trial implementations on TSMC 16FF+, using ARM Artisan libraries

Cortex-A72: More performance within constrained envelopes

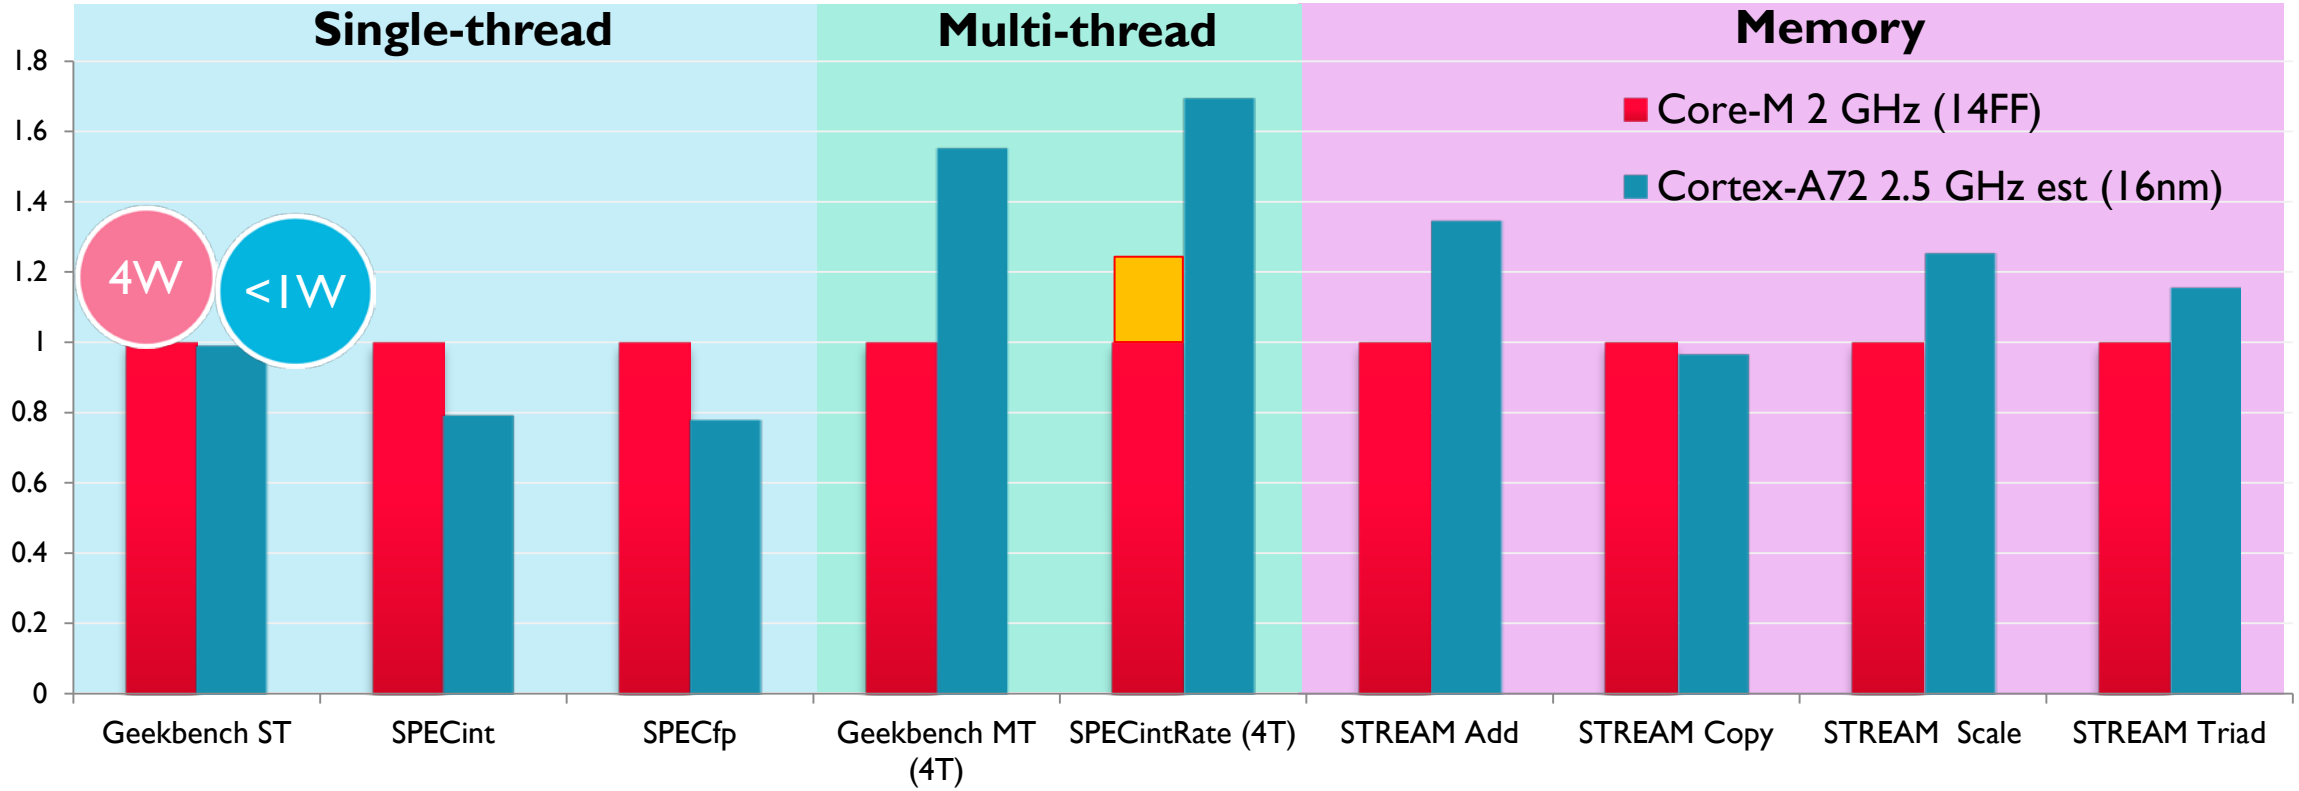

- Intel workloads measured on Dell Venue Pro II. SPEC benchmarks measured using gcc compiler v4.9 with –o3 flag.
- Cortex-A72 measured on RTL with realistic memory system with the same compiler settings
- Multi-threaded workloads use 2C4T Core-M CPU and estimated on 4C Cortex-A72 configuration w/2MB L2 cache.

Upside (scaled based on Xeon scaling) if core isn't thermally throttled

ARM®

ARM Partner SoC solutions c.2015

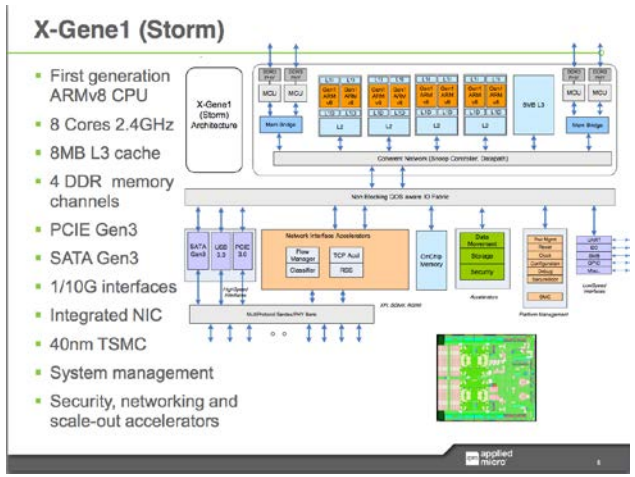

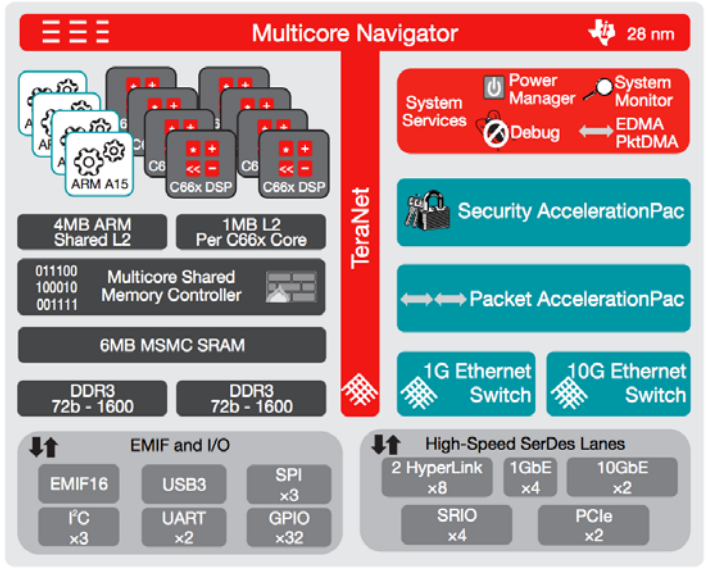

▲ TI's KeyStone 66AK2H14 SoC

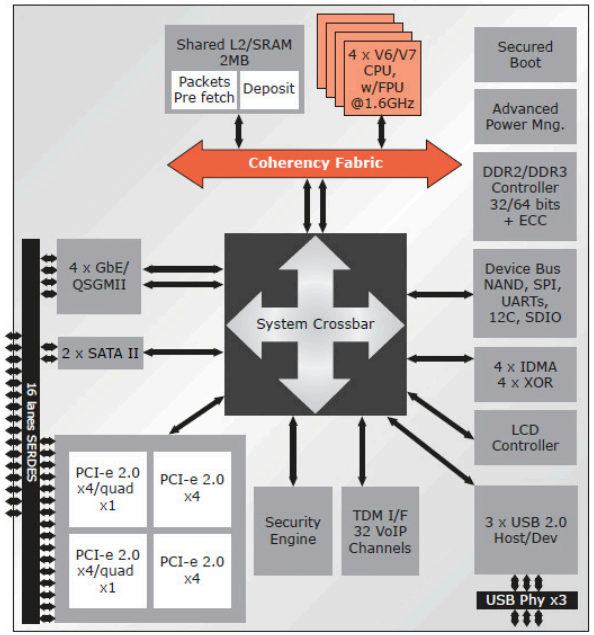

ARMADA XP Block Diagram

"SEATTLE" SOC OVERVIEW

Power Efficient Cores
- Up to Eight ARM Cortex-A57 cores
- Up to 4MB shared L2 cache total

Cache Coherent Network
- Full cache coherency
- 8MB L3 cache
- SMMU: I/O address mapping and protection

High Performance, Flexible Memory
- Two 64-bit DDR3/4 channels with ECC
- Two DIMMs/channel up to 1866Mhz
- SODIMM, UDIMM, RDIMM support
- Up to 128GB per CPU

Highly Integrated I/O
- 8x SATA 3 (6Gb/s) ports
- Two 10GBASE-KR Ethernet ports
- 8 lanes PCI-Express® Gen 3, supports x8, x4, x2

System Control Processor
- TrustZone® technology for enhanced security
- Dedicated 1GbE system management port (RGMII)
- SPI, UART, I2C interfaces

Cryptographic Coprocessor
- Separate Cryptographic algorithm engine for offloading encryption, decryption, compression, decompression computations

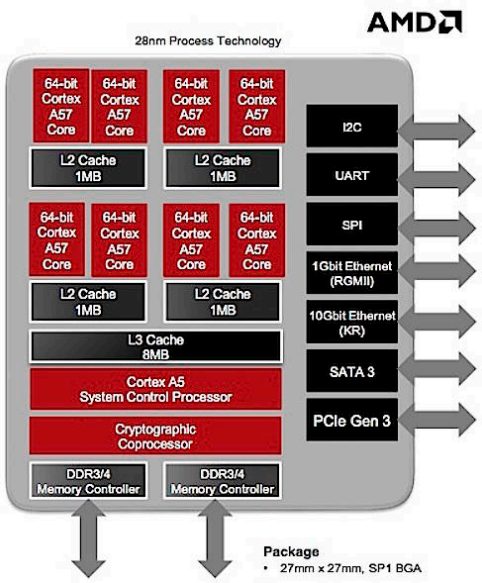

THUNDERX CAVIUM

Family of Workload Optimized Processors

- Up to 48 custom ARMv8 cores @ 2.5GHz
- 78K-I cache & 32K-D cache, 16MB L2
- 1S and 2S configuration
- Up to 4x72 bit DDR3/4 Memory Controllers
  - 1 TB system memory in 2S config
- Family Specific I/O's
- Standards based low latency Ethernet fabric
- virtSOC™: Virtualization from Core to I/O
- Family Specific Accelerators
- 4 Workload Optimized Processor Families:
  - ThunderX_CP: Compute servers
  - ThunderX_ST: Storage Servers
  - ThunderX_NT: Network/Telco Servers
  - ThunderX_SC: Secure Servers

More ARM 64-bit solutions on the horizon…

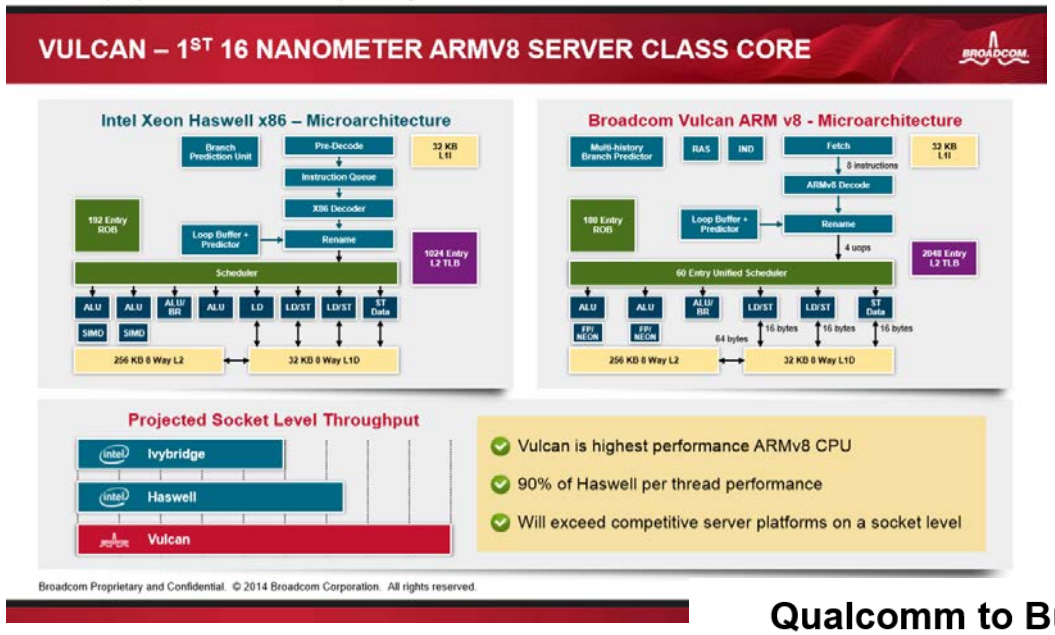

Source: Broadcom Presentation at IDC HPC USER FORUM APRIL

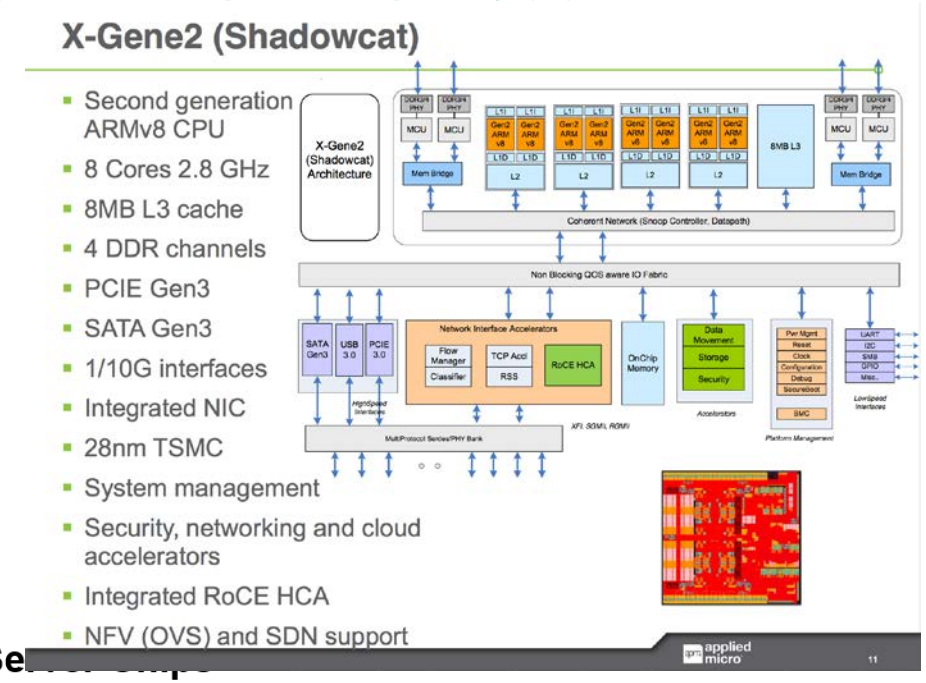

Qualcomm to Build ARM-Based Se

By Jeffrey Burt | Posted 2014-11-19 Email Print

Share 2 Tweet 19 Google + 1 Share 1 Like 11 Recommend 11

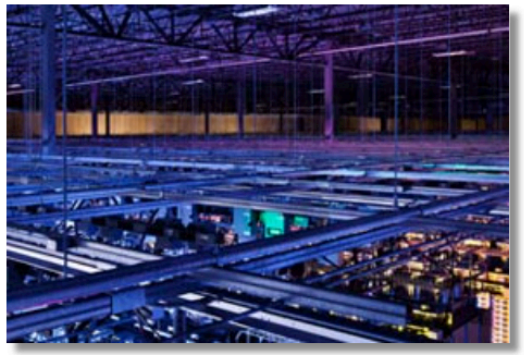

CEO Steve Mollenkopf gave few details, but Qualcomm will present a challenge to both smaller ARM server chip makers and dominant player Intel.

Qualcomm, the world's top mobile chip maker, is ready t get into the crowded ARM-based server chip business.

At the company's annual analyst day Nov. 19 in New York CEO Steve Mollenkopf said company engineers have bee working on the technology "for some time. Now we ar going to have a big product that goes into the server."

The *Wall Street Journal* was the first to report on Qualcomm's move.

ARMv8-A Infrastructure Ecosystem Building Momentum

ARM

Empowering Enterprise Software Developers

- Multiple options for software developers on ARM.

PayPal – Real-time Data Analysis

Application logs
Data center metrics
Server machine data
Metadata
Social media data.

3M events/sec, 25Tb/hour

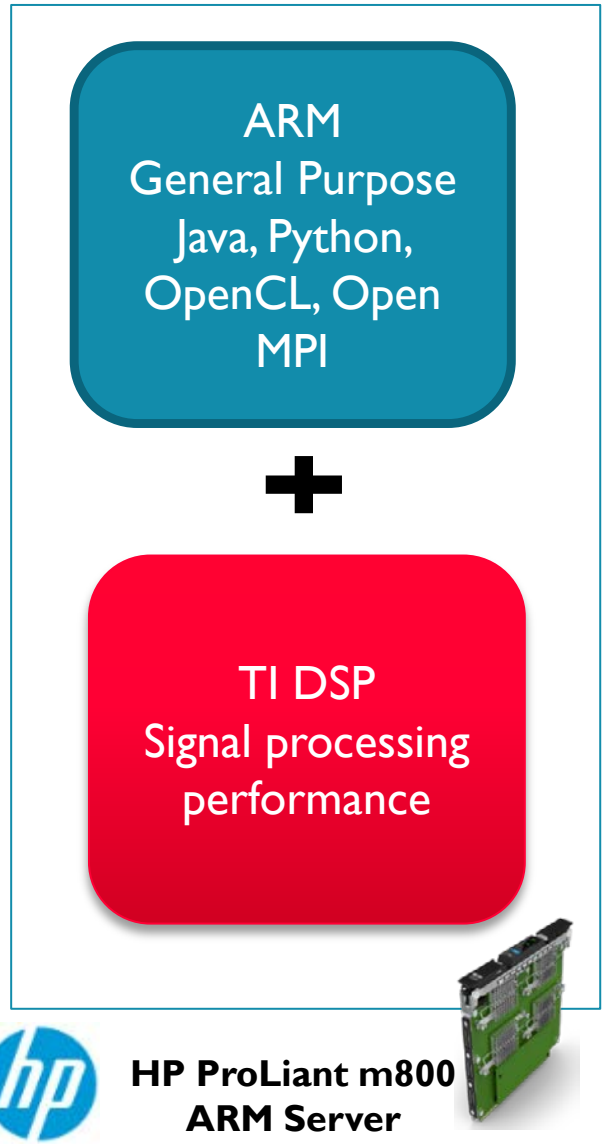

- Leverages HPC and Networking/Telco data-plane Technology
- **55W per cartridge (measured) => 11.2GFlops/watt**
  - Top of Green500.ORG SuperComputer list is 4.4GFlops/W
- SOC Specific Advantages
  - Right Sized General Purpose Compute
  - TI DSP cores for high performance signal processing performance
  - Low latency response times
  - An integrated, high-performance I/O fabric

Questions