X-Gene® in High Performance Computing

Rob Reiner
Director of Marketing
rreiner@apm.com

ARM HPC User Group Session @ SC16
November 14, 2016

X-Gene Driving ARM in Data Centers

OCTOBER 2009
Dedicated processor team begins 64b development

2009

2010

MARCH 2010
World's 1st ARMv8 Server Architecture Licensee

APRIL 2012
X-Gene FPGA release enables ecosystem partners

2012

2013

MARCH 2013
X-Gene silicon samples to customers

OCTOBER 2014
HP's m400 product launch

2014

OCTOBER 2014
X-Gene 2 wins ARM TechCon "Best in Show"

2014

2014

SPRING 2014
X-Gene 2 samples to customers

NOVEMBER 2015
X-Gene 3 Announced

2015

Spring/Summer 2015
Cirrascale RM1905D Announced

2015

NOVEMBER 2015
E4 RK003 Announced

Eurotech Aurora Hive Announced

2H'2016
Platform Announcements

2016

Target Customer Types

OEM ODM Networking

Hyperscale HPC

Target Deployments

Mainstream Web Tier Search Storage

Memcache Enterprise

ARM Platforms Powered by X-Gene

wistron GIGABYTE MiTAC kontron Cirrascale E4 COMPUTER ENGINEERING EuroTech China unicom Hewlett Packard Enterprise uniserver

Data Center

HP ProLiant m400

Tier 1 China Hyperscale end users

Gigabyte R120-P30 & MP30

ODM Platofrms

Storage

HPE StoreVirtual3200 Enterprise Storage Platform

Wistron X5 OCP Platform

Mudan Storage Platform

Networking

Tier1 Datacenter Switch

Kontron SYMKLOUD T4010

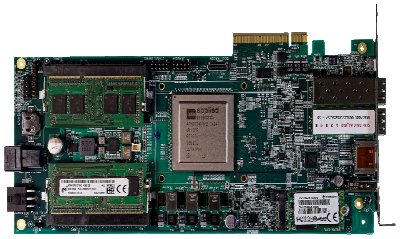

Programmable PCIE Accelerator Card

High-Performance Computing

E4 ARKA RK003

Cirrascale RM2916

Eurotech Aurora "Hive"

applied micro

X-Gene + GPU : Solid Platform for GPU-Optimized Apps

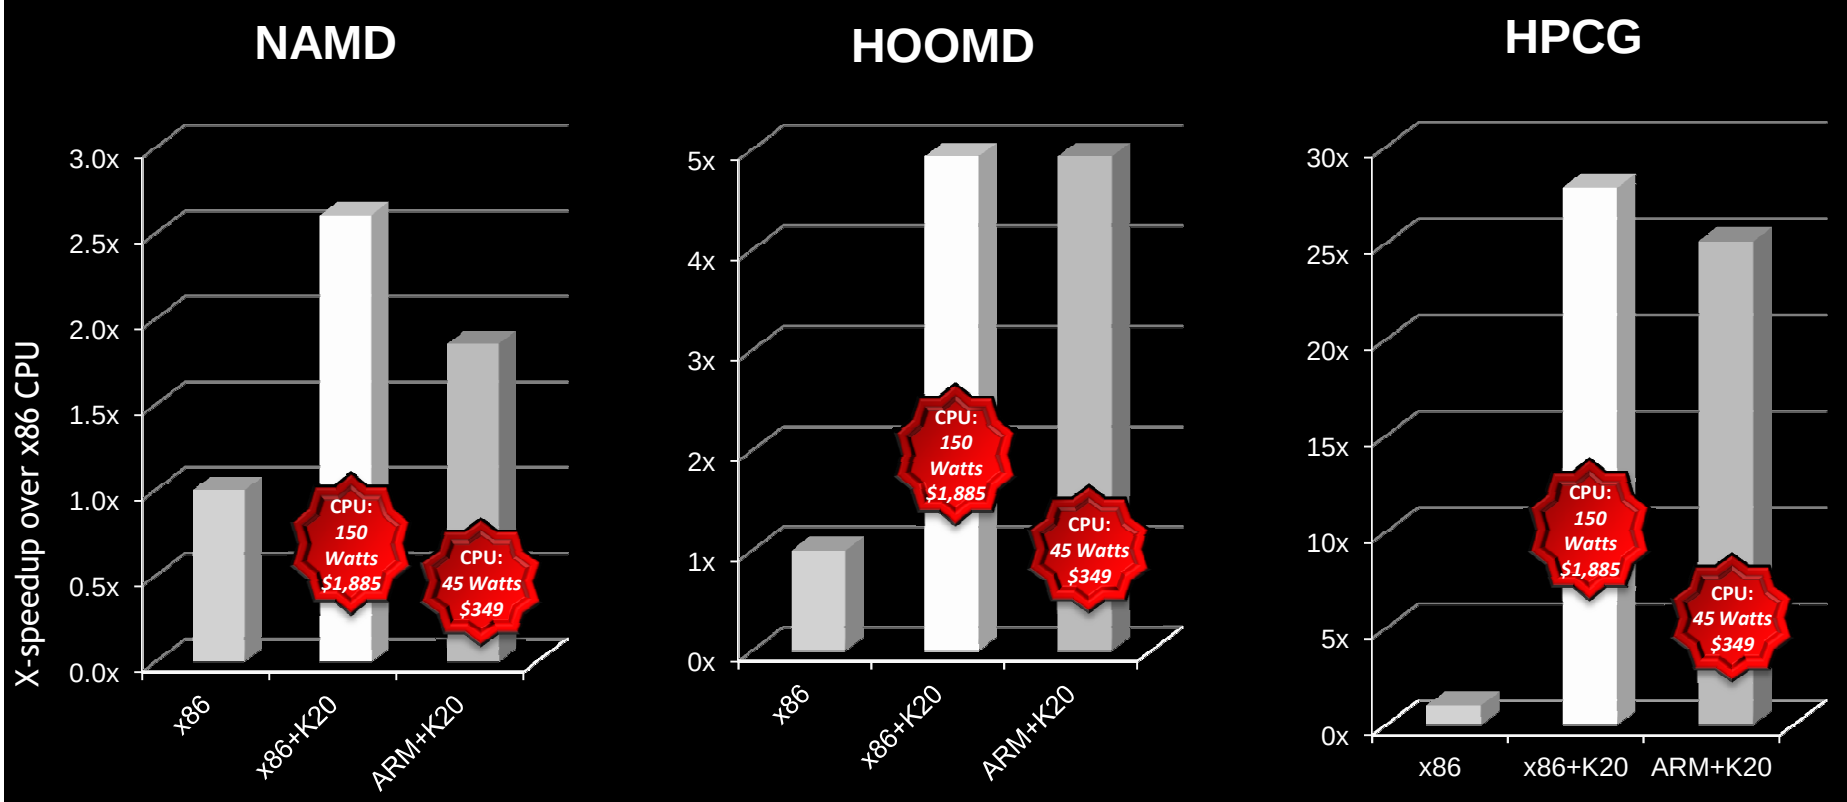

Application Workload Profile

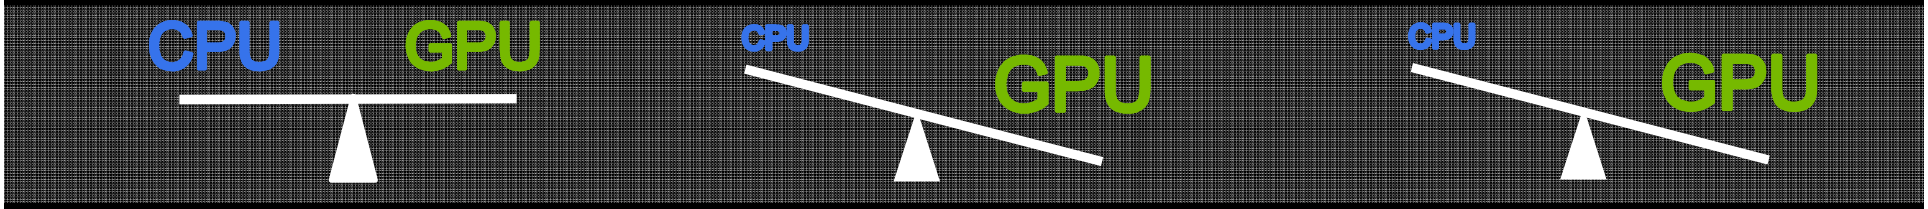

ARM: Applied Micro X-GENE CPU, X86: Xeon E5-2697, X86*: E5-2687W @ 3.10GHz

Source: nVidia

applied micro

X-GENE & X86 NVIDIA K40 BENCHMARK : HPCG

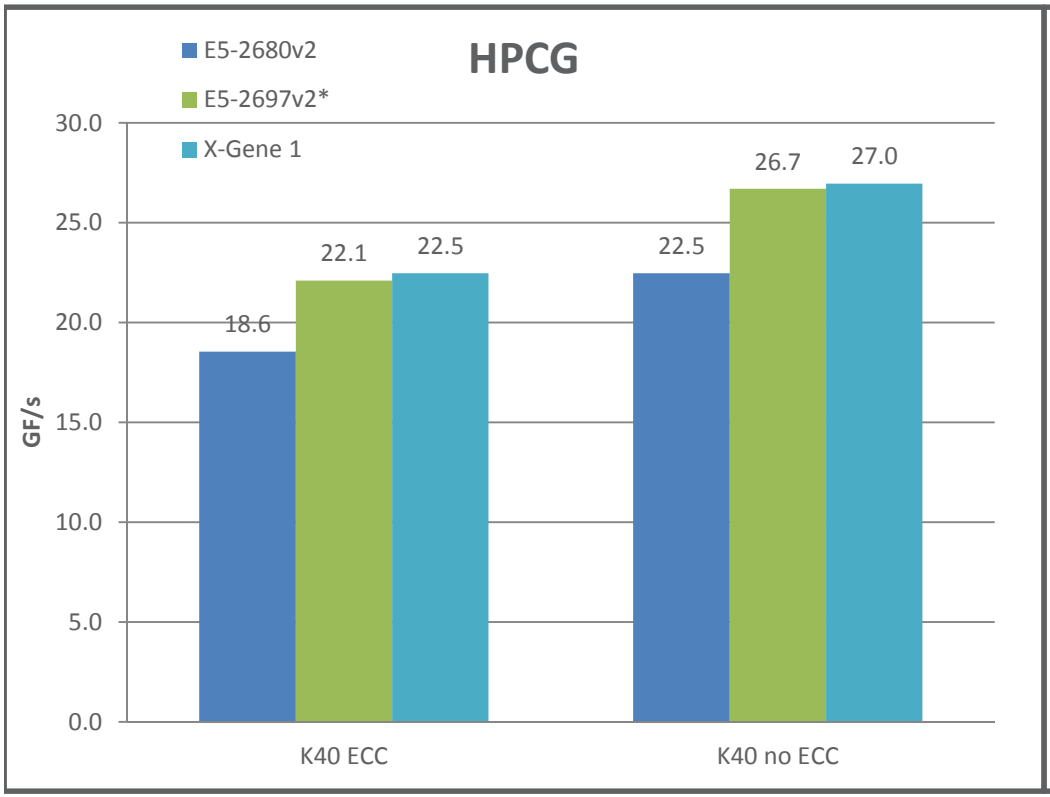

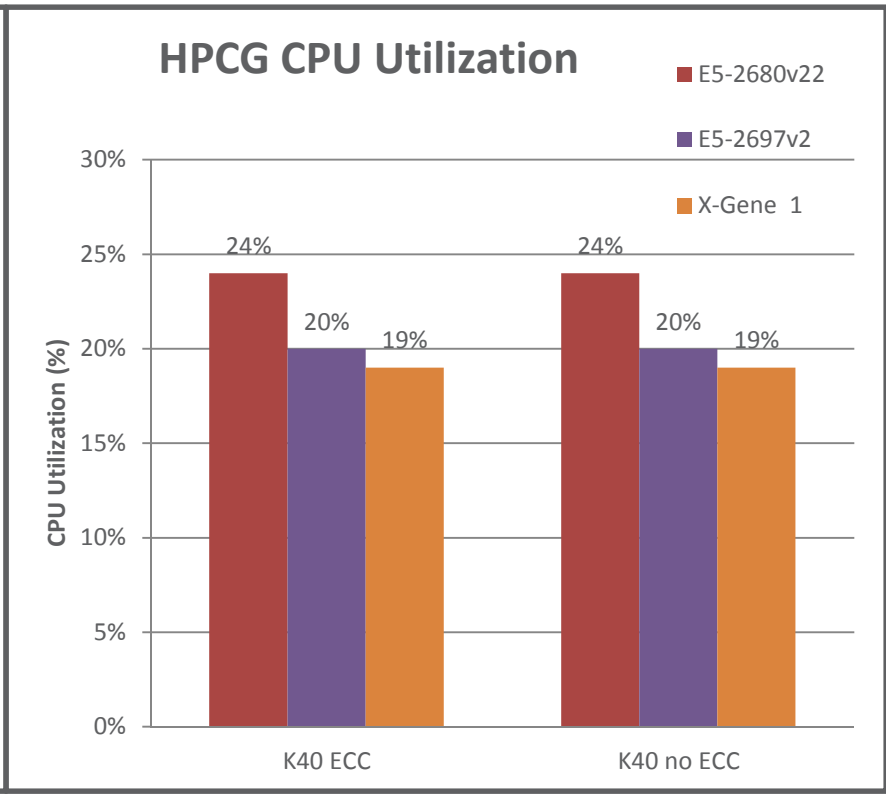

Source: APM, Nvidia

X-Gene: Applied Micro Gigabyte MP30 Platform 8-core @ 2.4GHz, 32GB DDR3-1600, Ubuntu 15.04
X86: Xeon E5-2680v2 10C/20T @2.8GHz Turbo/HT Enabled, 64GB DDR3-1600, CentOS 6.6, Xeon E5-2697v2 12C/24T @2.7GHz
CUDA 6.5

X-Gene + GPU : Solid Platform for GPU-Optimized Apps

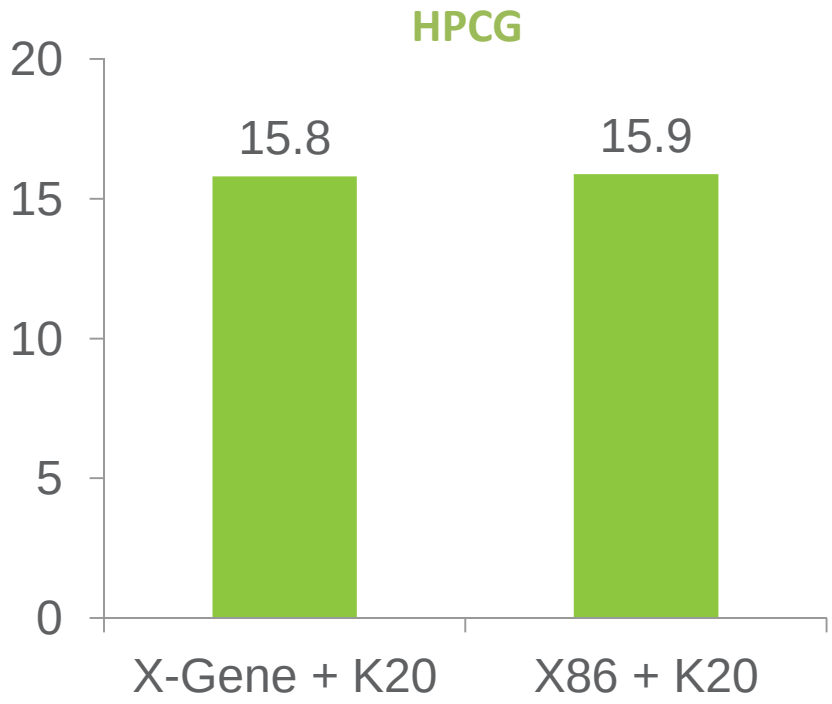

Source: nVidia

Cirrascale RM1905D

Cirrascale

NVIDIA TESLA

X-Gene: AppliedMicro X-Gene 1 ARM CPU

x86: Xeon™ E5-2697v2 CPU

K20: nVidia Tesla K20 GPU

E4 Arka HPC Platform based on X-Gene

E4 ARKA series the first ARMv8 plafroms dedicated to HPC

E4 COMPUTER ENGINEERING

E4 RK003

FEATURES	
Form Factor	2U
CPU	1x APM X-Gene 8 cores
GPU up to	2x NVIDIA Kepler® K20, K40, K80
Memory	Up to 128 GB RAM DDR3
Network	2x 10 GbE SFP+, 1x Infiniband FDR QSFP
Storage	4x SATA 3.0
Expansion slots	2x PCI-E 3.0 x8 (in x16)

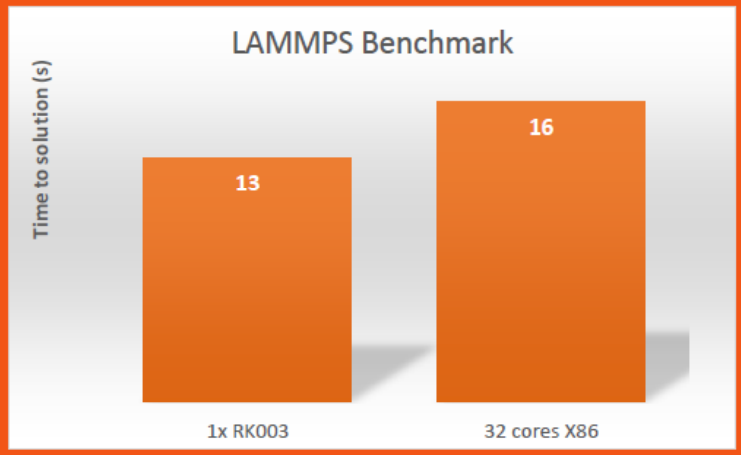

ARM applied micro NVIDIA

Source: E4

X-Gene HPC Platforms

Eurotech Aurora Hive

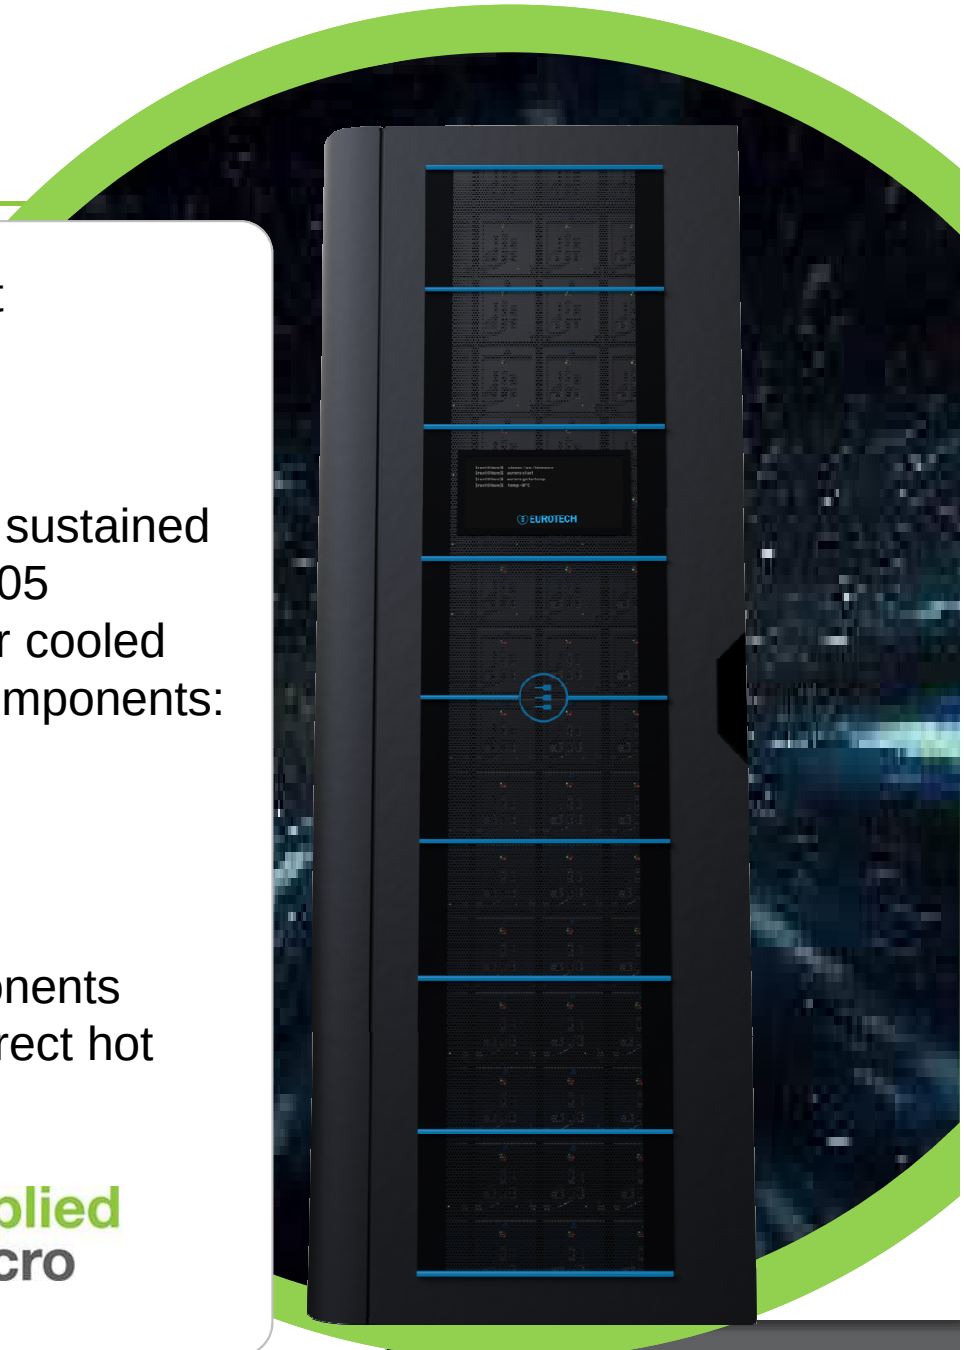

- **Most dense** architecture in the market
  - 1.5 PFlop/s DP peak per rack
  - 5 Pflop/s SP peak per rack
- **Highest energy efficiency.**
  - At system level > 6 GFlops / Watt sustained
  - At datacenter level with PUE of 1.05
- **Distinctive flexibility**. Direct hot water cooled with no constraints for the choice of components:
  - Intel or ARM CPU
  - NVIDIA or AMD accelerators
  - Intel coprocessors
  - Any other PCIe card
- **Direct hot water cooling** of all components
- **First and only** system with multiple direct hot water cooled GPGPUs

X-Gene® 3

- Third Generation ARMv8 Processor
- Quantum Jump in ARM Server Landscape
- 6x the performance versus current X-Gene family
- 8 64-bit DDR4 Memory Controllers with ECC
- 1TB Memory per Socket
- 30% Lower Power per Core
- Addresses mainstream work-loads used by Tier 1 hyperscale datacenter customers
- Targets data analytics, HPC, machine learning, web/caching servers, web search

Sampling Q1 2017

X-Gene® 3 Overview

Processor Subsystem

- 32 ARM™ v8 64-bit CPU cores at up to 3 GHz
- 32 KB I-cache/D-cache
- Shared 256 KB L2 cache per 2 cores

Memory

- 32 MB globally shared L3 cache
- Eight DDR4-2667 channels with ECC and RAS
- Up to 16 DIMMs and 1TB/socket

System Resources

- GICv3 Interrupt Controller
- Full IO Virtualization
- Enterprise server-class RAS

Connectivity

- 42 lanes of PCIe Gen 3 with 8 controllers
- 4 x SATA Gen 3 ports
- 2 x USB 2.0 ports
- 1 x RGMII port

Functionality

- ARM SBSA Level 3
  - – EL3, secure memory and secure boot support
- Advanced Power Management with ACPI/PSCI
- Full SBBR compliance with UEFI

Performance

- **SPECint_rate ~550; TDP: 125W**

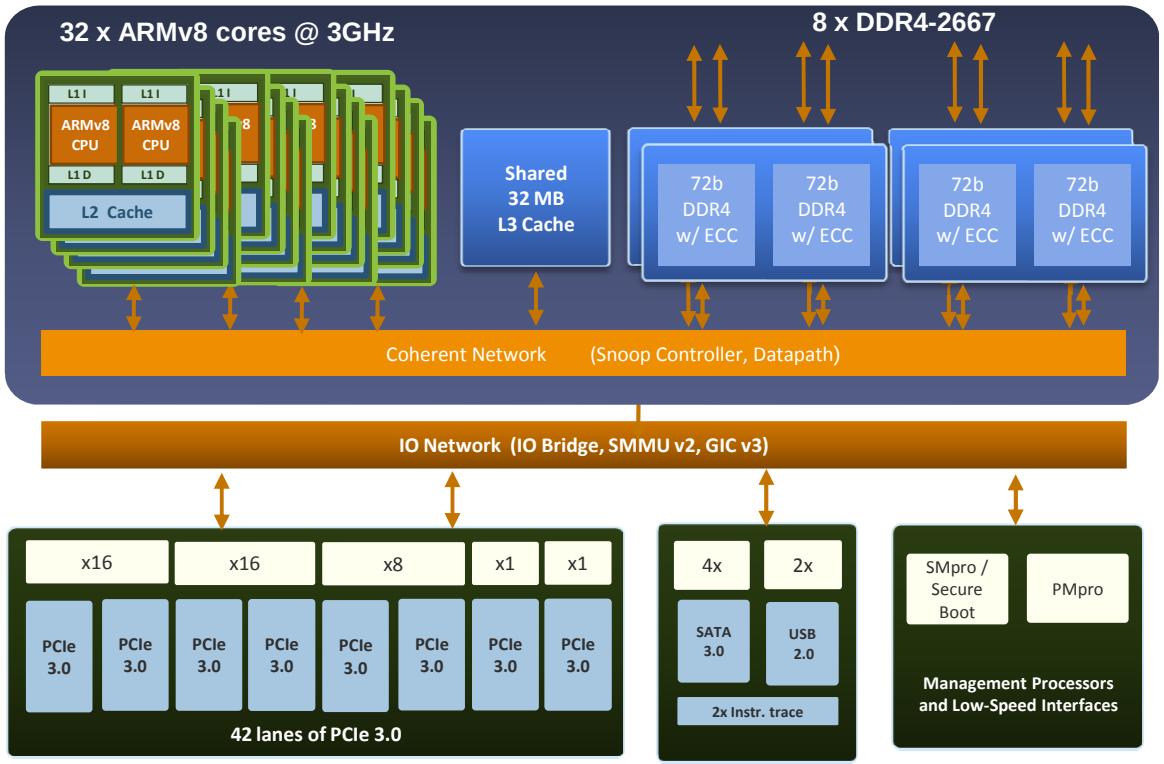

Conclusion

- X-Gene 1/X-Gene 2 are in production today
  - Wide Range of HPC Platforms and Deployments
- Higher Performance/$ and Performance/Watt reduces the overall TCO
- X-Gene 1 and X-Gene 2 pave the way for Third Generation X-Gene 3
- X-Gene 3 Delivers Quantum Jump in ARM Server Landscape

Thank You!