Intro to Arm HPC Systems

Centre for Development of Advanced Computing (C-DAC) / National Supercomputing Mission (NSM)

Arm in HPC Course

Phil Ridley

phil.ridley@arm.com

2nd March 2021

Agenda

- Intro
- Overview of Arm
- Current HPC Deployments

Arm's HPC Field Engineering Team (HPC-FE)

Porting, Tuning, Training, and Enablement

- Application performance engineering
  - Get help optimizing for maximum performance
- System tuning
  - Tune HPC system parameters for your workloads
- Hackathons and tutorials
  - Education, mentoring, and hands-on events to help jumpstart HPC developers

arm

Introduction

70%

of the world's population uses Arm technology

Arm Technology Already Connects the World

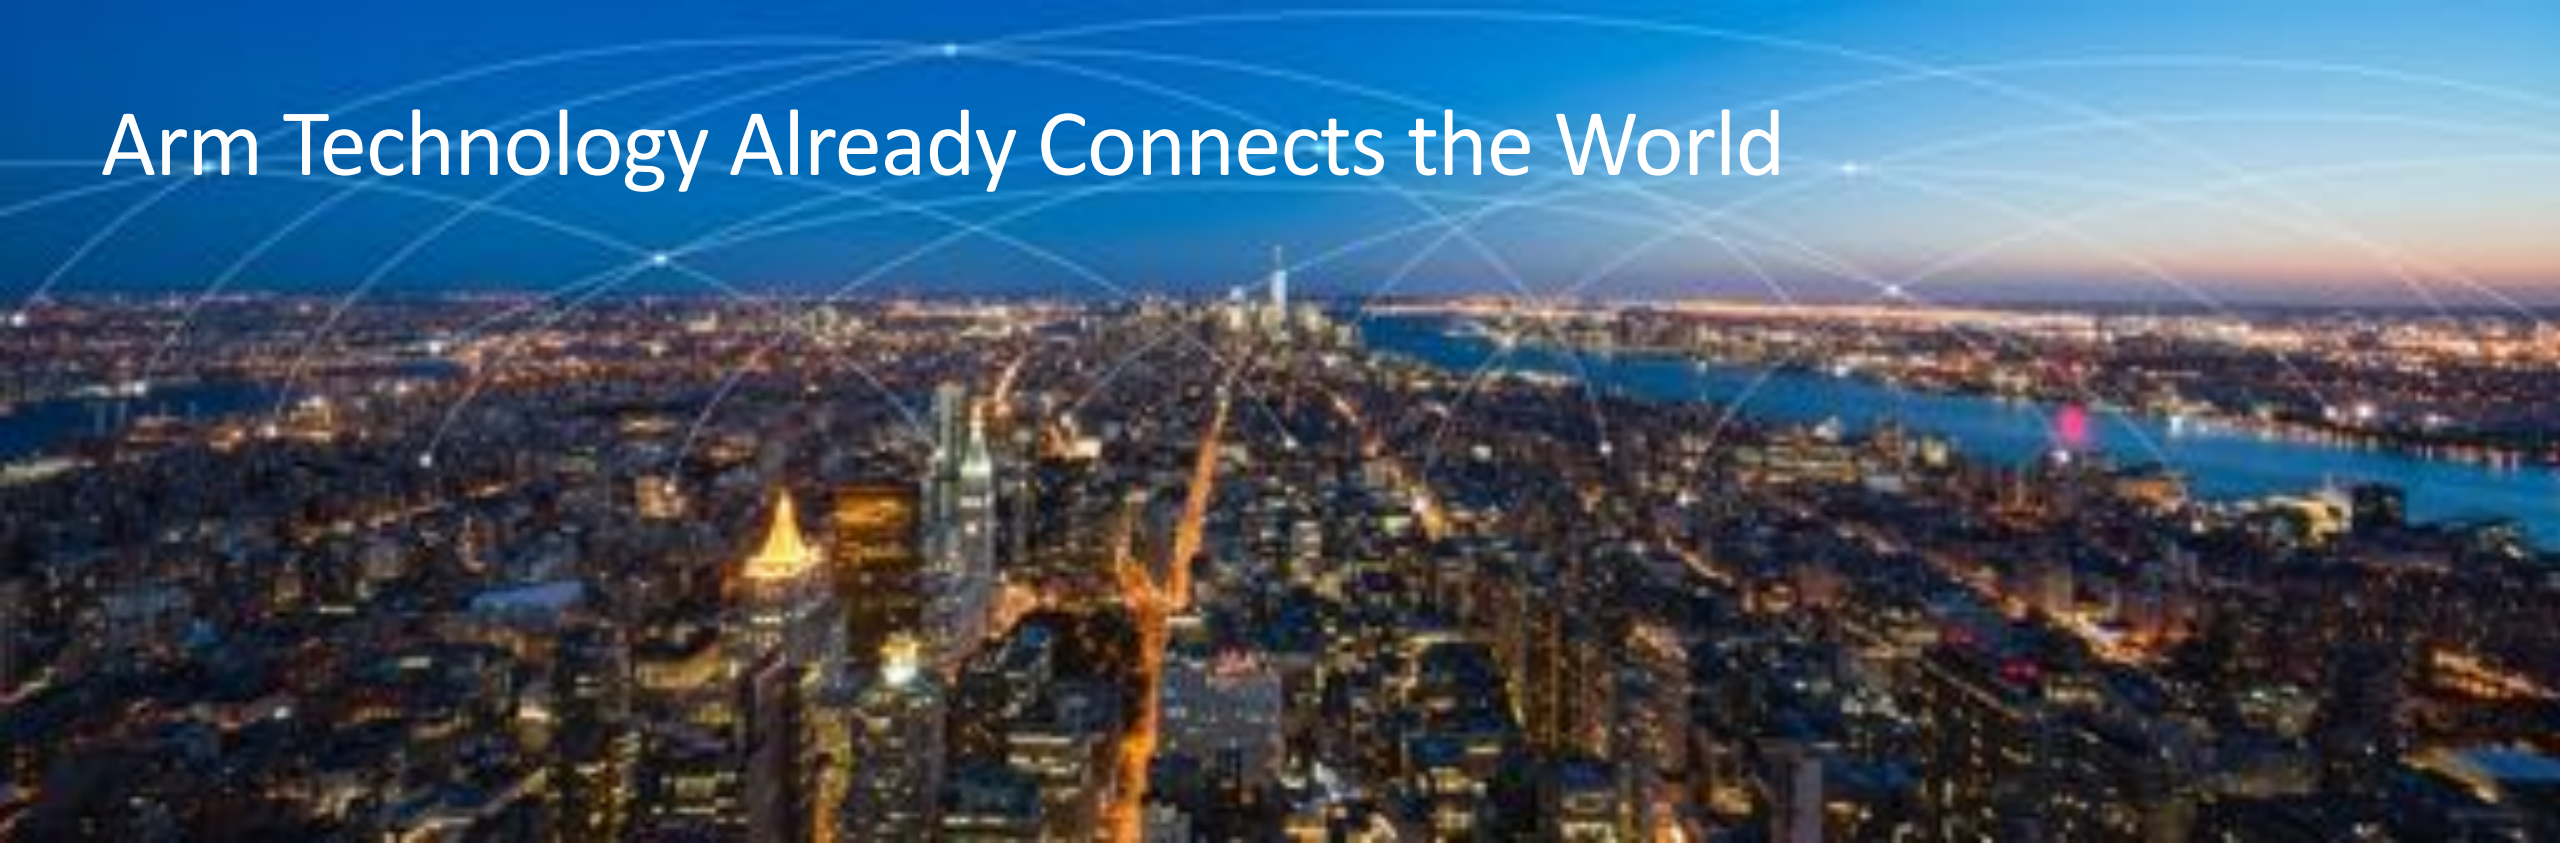

Arm is ubiquitous

21 billion chips sold by partners in 2017 alone

Mobile/Embedded/IoT/ Automotive/Server/GPUs

Partnership is key

We design IP, not manufacture chips

Partners build products for their target markets

Choice is good

One size is not always the best fit for all

HPC is a great fit for co-design and collaboration

arm

What is Arm?

ARMv8.0	ARMv8.1-A	ARMv8.2.-A	ARMv8.3-A	ARMv8.4-A	ARMv8.5-A	ARMv8.6-A
October 2011	December 2014	January 2016	October 2016	November 2017	September 2018	September 2019
• AArch64 execution state • A64 instruction set	• Atomic memory ops • Type2 hypervisor support	• Half-precision float • RAS support • Statistical profiling	• Pointer authentication • Nested virtualization • Complex float	• Improved crypto extensions • Improved virtualization • Signed and unsigned SDOT and UDOT instructions	• Memory Tagging Extension (MTE), • Branch Target Indicators (BTI) • RNGs	• GEMM • Bfloat16 • SIMD matrix manipulation instructions, BFDOT, BFMMLA, BFMLAL and BFCVT

CPU Engagement Models with Arm

Arm IP is the basic building block for extraordinary solutions.

Core License

- Partner licenses complete microarchitecture
- CPU differentiation via:
  - Configuration options
  - Wide implementation envelope with different process technologies

Arm IP

Architecture License

- Partner designs complete microarchitecture
- Clean room, scratch
- Maximum design freedom:
  - Directly address needs of the target market
- Arm architecture validation preserves software compatibility

Pizza Engagement Models

The basic building blocks for an extraordinary pizza!

Core License

Arm design

Partner IP

Yum!

IP = Irresistible Pizza

Architecture License

Recipe

Ingredients

Yum!

arm

arm NEOVERSE

The Cloud to Edge Infrastructure Foundation for a World of 1T Intelligent Devices

Arm Neoverse Platform Roadmap

In planning

IP Date	Production	Now	2021	2022+
Platform features	7nm PCIe Gen4, DDR4, HBM2 CCIX 1.0	7/5nm PCIe Gen5, DDR5, HBM2e CCIX 1.1	5nm PCIe Gen5, DDR5, HBM3 CCIX 2.0, CXL 2.0	5/3nm PCIe Gen5/6, DDR5, HBM3 CCIX next, CXL next
V-series Max, ST Perf, ML		**V1 Platform (Zeus)** N1+50% ST perf. uplift SVE 2x256b, bFloat16		**Poseidon Generation Platforms** +30% infra WL perf. ML/Vector uplift Greater core density
N-series Perf/W, Scale out, 5G	**N1 Platform** A72+60% ST perf. uplift		**N2 Platform (Perseus)** N1+40% ST perf. uplift SVE 2x128b, bFloat16	
E-series Data efficiency	**E1 Platform** A53+2.7x throughput			

Arm HPC R&D for the Future

European Processor Initiative

arm

Why Arm?

Infrastructure / HPC / Scientific Computing / ML

Hardware

- Flexibility: Allow vendors to differentiate
  - Speed and cost of development
- Provide different licensing
  - Core - Reference design (A78/N1)
  - Architecture - Design your own (Marvell®ThunderX2®, Fujitsu A64FX)
- Other hardware components
  - NoCs, GPUs, memory controllers
  - “Building blocks” design
- Architecture validation for correctness

Software

- All based on the same instruction set
  - Commonality between hardware
  - Reuse of software
- Comprehensive software ecosystem
  - Operating systems, compilers, libraries, tools
  - Not just vendor - third party too
- Large community
  - Everything from Android to HPC

arm

Current Arm-based HPC Processors

Processor Market

Marvell (Cavium)		CAVIUM THUNDERX		CAVIUM THUNDERX2			
Ampere (X-Gene)	applied micro X-Gene		applied micro X-Gene 2		AMPERE eMag$_{TM}$		AMPERE Altra$_{TM}$
Fujitsu						FUJITSU A64FX	
Huawei (HiSilicon)		Hi1612		1616		Kunpeng 920	Hi1630
Amazon (Annapurna)					aws Graviton		aws Graviton2
EPI / SiPearl							SIPEARL
Other				PHYTIUM	QUALCOMM Qualcomm Centriq Processor		

arm

Marvell ThunderX2 CN99XX

Architecture License

- Marvell's next generation 64-bit Arm processor
  - Taken from Broadcom Vulcan
- 32 cores @ 2.2 GHz (other SKUs available)
  - 4 Way SMT (up to 256 threads / node)
  - Fully out of order execution
  - 8 DDR4 Memory channels (~250 GB/s Dual socket)
    - Vs 6 on Skylake
- Available in dual SoC configurations
  - CCPI2 interconnect
  - 180-200w / socket
- No SVE vectorisation
  - 128-bit NEON vectorisation

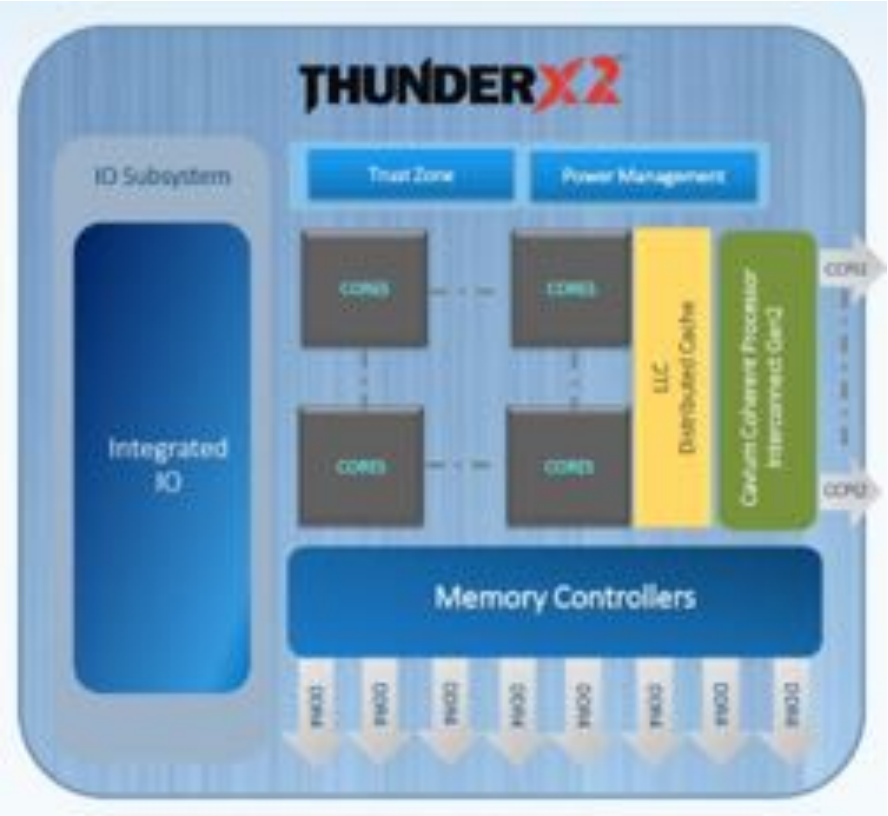

MARVELL®

arm

Marvell ThunderX2 Key Features

Architecture License

- Arm v8.1-A with 128-bit NEON
- 16nm CMOS FinFET
- 2.5GHz
- **Dual socket (32 cores / socket)**
  - 32KB L1$ per core
  - 256KB L2$ per core
  - 32MB L3$ shared
  - DDR4x8
  - 4 SMT per core, giving 256 cpus / node

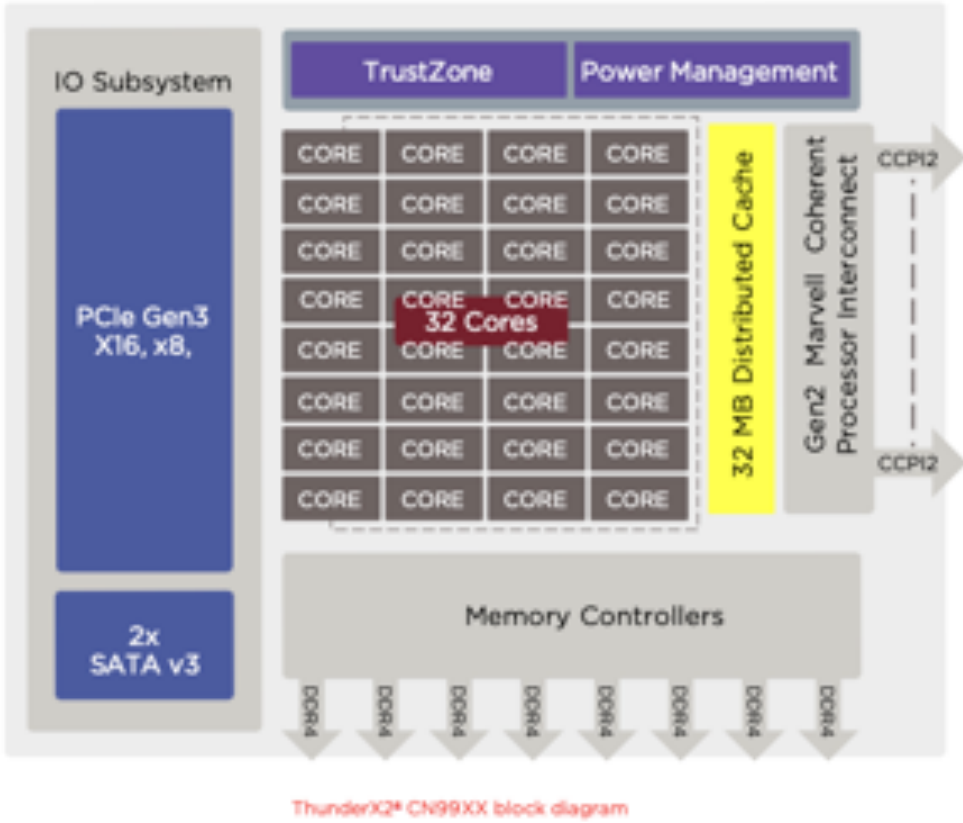

ThunderX2® CN99XX block diagram

https://www.marvell.com/content/dam/marvell/en/public-collateral/server-processors/marvell-server-processors-thunderx-cn99xx-product-brief-2019.pdf

Fujitsu A64fx Key Features

Architecture License

https://www.fujitsu.com/downloads/SUPER/a64fx/a64fx_datasheet.pdf

- Arm v8.2-A with 512-bit SVE
- Custom Fujitsu u-arch
- 7nm CMOS FinFET
- 2.2GHz, 2.0GHz, 1.8GHz
  - Constant clock: no turbo, no downclock
- **4 Core Memory Groups (CMGs)**
  - 12 cores (13 in the FX1000)
  - 64KB L1$ per core
    - 256b cache line
  - 8MB L2$ shared between all cores
    - 256b cache line
  - Zero L3$
  - 8 GB HBM at 256GB/s

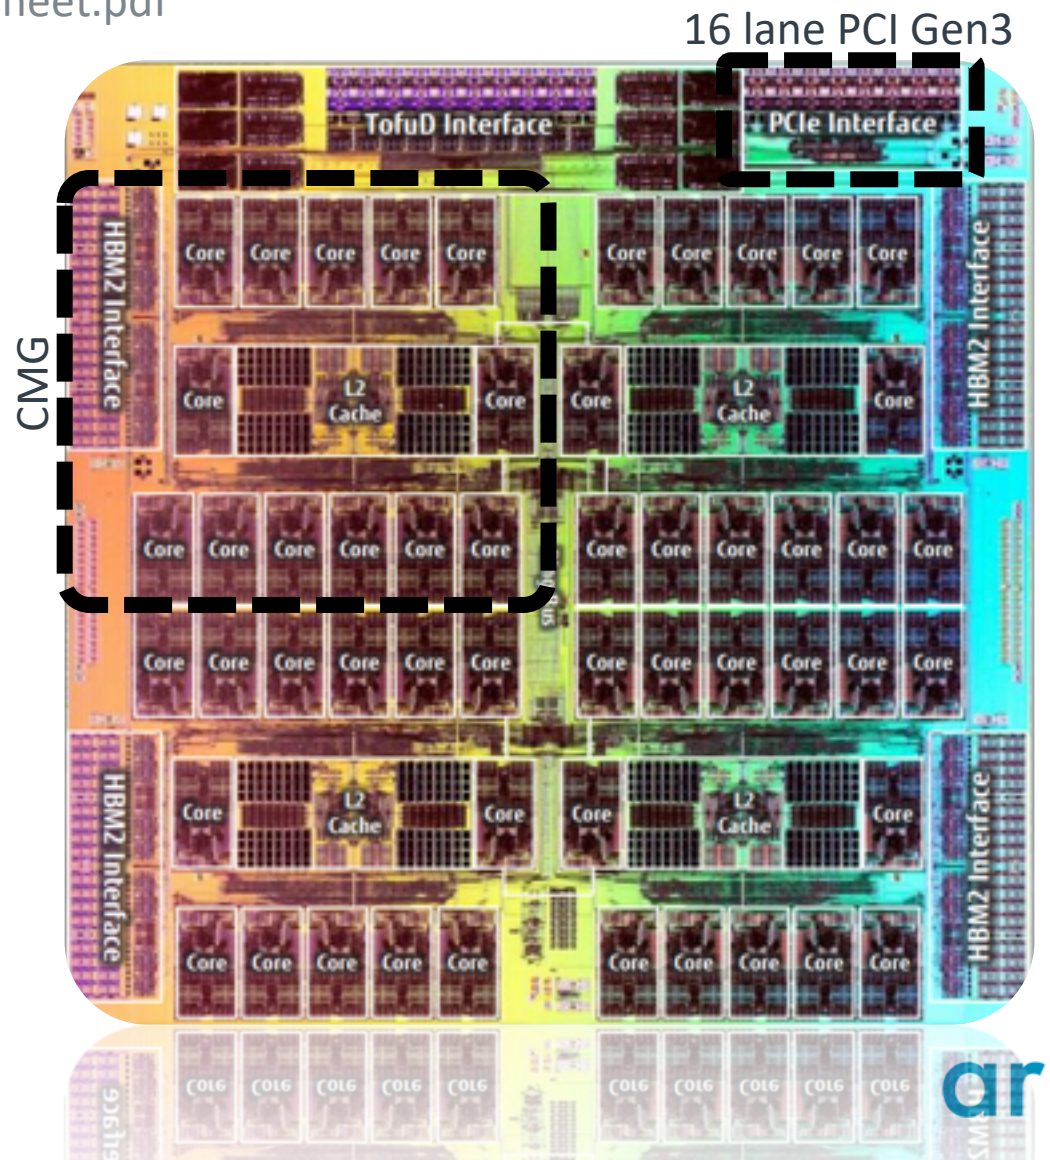

AWS Graviton2

Core License

- Designed by Annapurna Labs
  - Known as the Amazon Graviton
  - For use in the AWS EC2 M6g, R6g C6g and T4g instances
- 64 Core Socket at 2.5 GHz
  - Based on Arm Neoverse N1 (core license)
- Non-NUMA (single socket, 1-NUMA node)
- Memory: Up to 512GB
- Enhanced Network Bandwidth: Up to 25 Gbit/s
- EBS Bandwidth: Up to 18.5 Gbit/s

Marvell ThunderX2 / Fujitsu A64fx / AWS Graviton2/ Ampere Altra

	Arm ISA	Vector size	Cores/ socket	Socket(s)	L1	L2	L3	Clock	Mem
Marvell ThunderX2	Armv8.1-A	128	32	2	32K	256K	32MB	2.5GHz	DDR4x8
Fujitsu A64fx	Armv8.2-A+SVE	512	48	1	64K	8MB	-	1.8GHz	HBM2x8
AWS Graviton2	Armv8.2-A (Neoverse N1)	128	64	1	64K	1MB	32MB	2.5GHz	DDR3x8
Ampere Altra	Armv8.2-A (Neoverse N1	128	80	2	64K	1MB	32MB	3.3GHz	DDR4x8

Current Deployments

Deployments: HPE Astra at Sandia

Mapping performance to real-world mission applications

- HPE Apollo 70
- #281 on Top500 (Nov 2020)
  - 1.76 PFLOPs Rmax (2.2 PFLOPs Rpeak)
- Marvell ThunderX2 processors
  - 28-core @ 2.0 Ghz
  - 332 TB aggregate memory capacity
  - 885 TB/s of aggregate memory bandwidth
- 2 592 HPE Apollo 70 nodes
  - 145,152 cores
- Mellanox EDR InfiniBand
- OS: RedHat

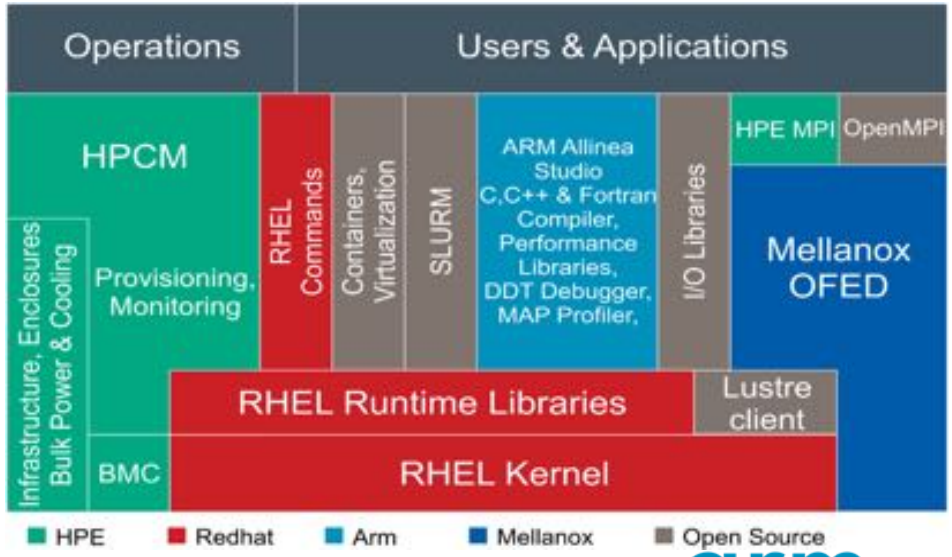

Deployment: CEA

- **Atos BullSequana X1310**
  - *A Mont-Blanc 3 design*

MONT-BLANC

- 292 blades
  - 2 sockets
  - Marvell ThunderX2
  - 32 cores, 2.2 GHz
  - 8 channels, DDR4 2666, 256 GB
  - Infiniband EDR

- Peak performance: 329 TFlops

Atos

arm

Deployment: Minho Advanced Computing Centre (MACC)

- Fujitsu PRIMEHPC FX700
- Operational early 2022
- 10 Petaflops

EuroHPC

Home About Participate News Events Media Documents Contact

Deucalion: a new EuroHPC world-class green supercomputer in Portugal

Published on 5 February 2021

The procurement contract for a new EuroHPC petascale supercomputer has been signed by the European High-Performance Computing Joint Undertaking (EuroHPC JU), the Portuguese Foundation for Science and Technology (FCT) , the hosting entity, and Fujitsu, the selected vendor.

The new HPC system, to be located in Portugal, will be called "Deucalion". It will be a petascale supercomputer, capable of peak performance of 10 Petaflops or 10 million billion calculation per second. The machine will use ARM-based technology – the Fujitsu A64FX CPU currently used by Fugaku, the fastest supercomputer in the world at the moment.

Deucalion will be installed in the recently created **Minho Advanced Computing Centre (MACC) of the Portuguese Foundation for Science and Technology (FCT)** which is **one of the four operational advanced computing centres in Portugal** coordinated by FCT. Deucalion will provide a unique and innovative context to demonstrate green computing principles in Europe and worldwide by leveraging a fully sustainable infrastructure, co-funded by national science funds, European structural funds allocated to the North of Portugal and the Portuguese Fund to Support Innovation and the Energy Efficiency Fund (the latter within the remit of the Ministry for the Environment and Climate Action).

Isambard

- **21,504** Marvell ThunderX2 cores (168n x 2s x 32c), 2.5GHz
  - Armv8.1-A, NEON 128-bit vectors
- Cray XC50
- High-speed **Aries** interconnect
- Cray Linux Environment 7 (Suse SLES 15)

- **3,456** a64fx cores, 1.8GHz (72 nodes)
  - Armv8.2-A+SVE (512-bit vectors)
  - 72 TB/s memory bandwidth
  - 202 TFLOP/s double precision
- HPE Apollo 80
- 100Gbps Mellanox InfiniBand ConnectX-6 EDR
- RHEL 8.2

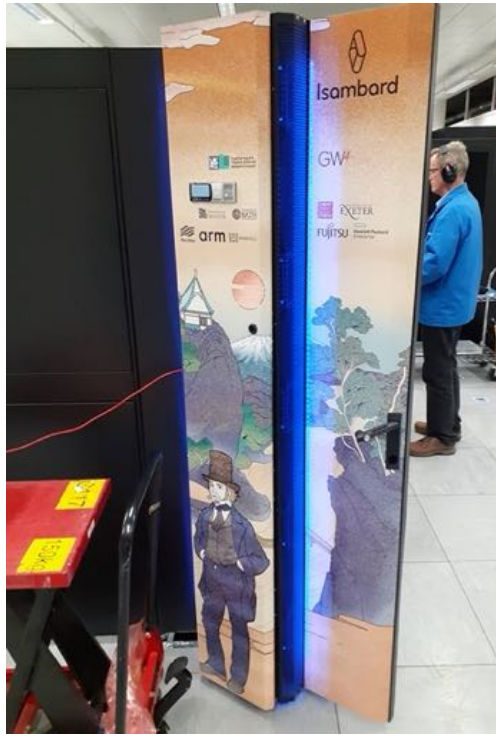

Fugaku Supercomputer

- Biggest Arm based deployment
  - Number 1 on Top500
  - 415.53 PFLOPS
  - 7.6M Arm cores (no accelerators)
- Energy Consumption
  - Designed to be low
  - ~150 W / node *[1]
- Delivered early to assist with COVID-19
  - National and international projects
  - Open Science

NEWS

Japan Captures TOP500 Crown with Arm-Powered Supercomputer

June 22, 2020

FRANKFURT, Germany; BERKELEY, Calif.; and KNOXVILLE, Tenn.—The 55th edition of the TOP500 saw some significant additions to the list, spearheaded by a new number one system from Japan. The latest rankings also reflect a steady growth in aggregate performance and power efficiency.

TOP500 The List.

HPCwire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

Home

Arm Technology Powers the World's Fastest Supercomputer

June 22, 2020

June 22, 2020 – At the International Supercomputing Conference (ISC) it was announced that the Fugaku supercomputer, a system jointly developed by RIKEN and Fujitsu Limited, and based on Arm technology, was awarded the number one spot of the TOP500 list. Having been crowned the world's most efficient supercomputer on the Green500 list in November 2019, Fugaku was today also given top honors on the HPCG list, a ranking of benchmarks across real-world applications, and the HPL-AI, which rates performance on tasks used in artificial intelligence applications.

arm

EPI: SiPearl Rhea (1st Gen)

- New initiative as part of EuroHPC
  - Drive for European technology for HPC
- Mixture of new technologies
  - Arm general purpose cores (Zeus - N2)
  - Accelerators: RISC-V, FPGA
  - Memory: DDR 4/5, HBM
  - Connectivity: PCIe G5, CCIX
- Targeted for key European markets
  - Automotive
  - HPC
  - AI / ML

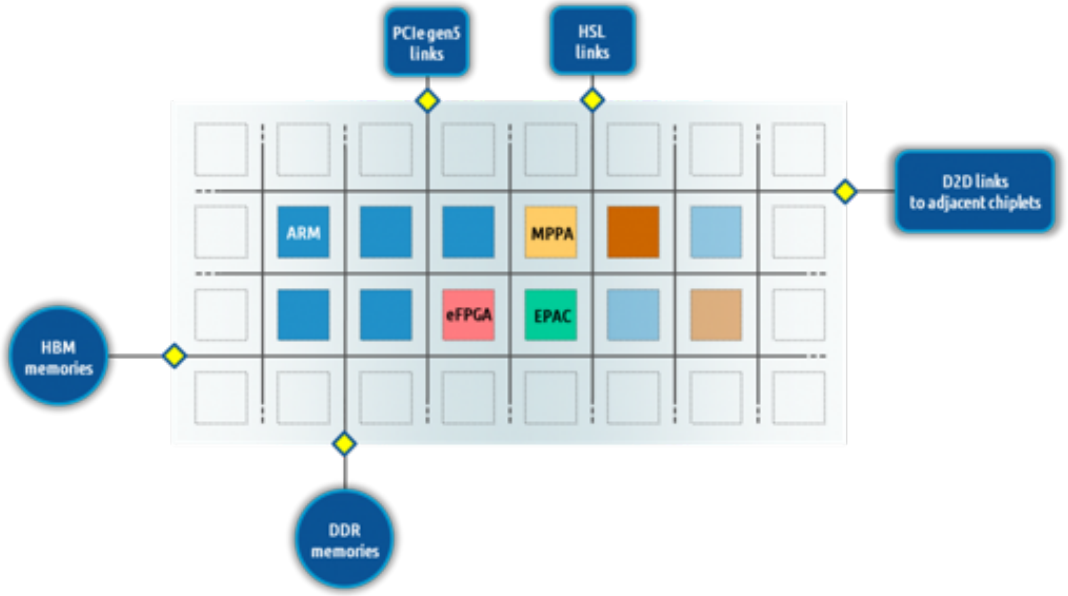

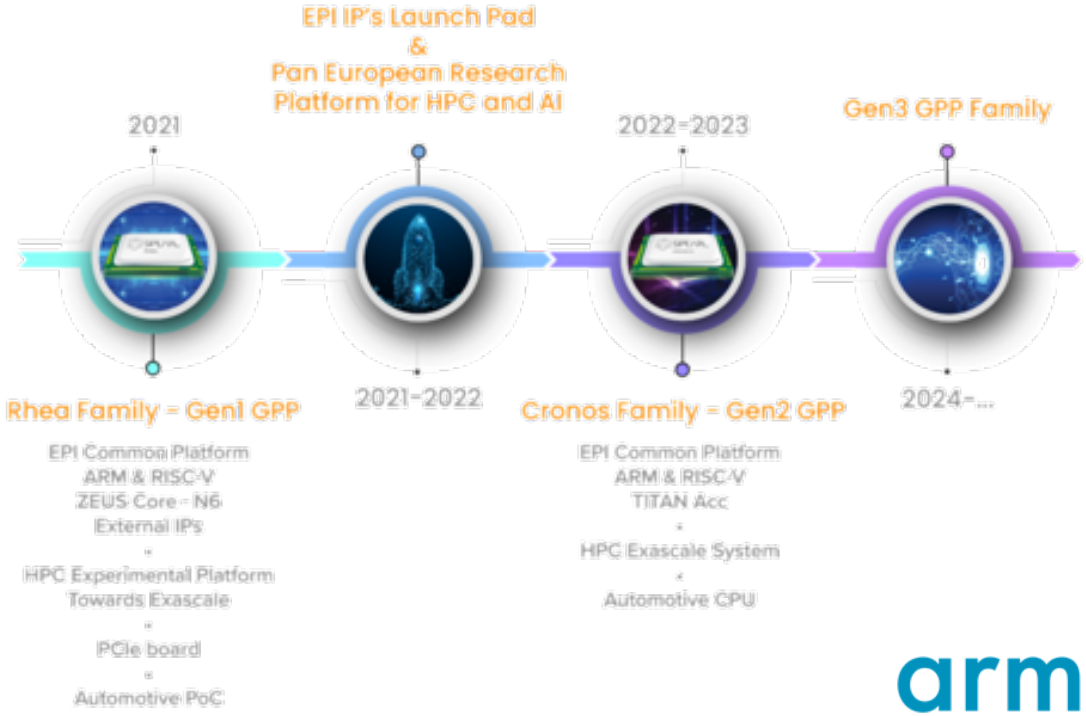

GPU acceleration on Arm

- NVIDIA support for Ampere, Fujitsu, HPE, Marvell
- Deployed at Oak Ridge and Sandia National Laboratories, the University of Bristol, Riken, and CINECA
- CUDA 11, OpenCL, PGI Compiler support, NVIDIA HPC SDK

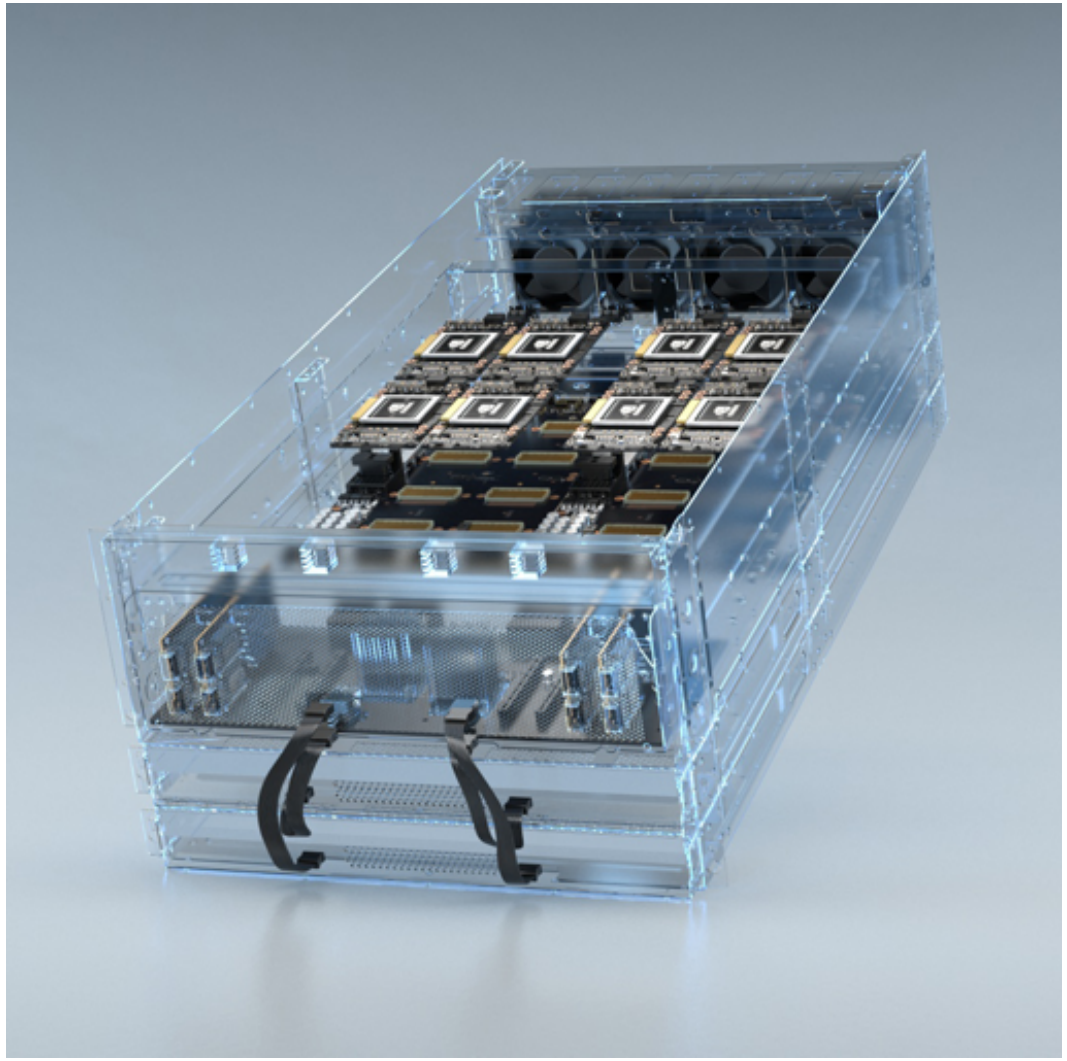

Arm-NVIDIA server. Image Credit: NVIDIA

The Cloud

Open access to server class Arm

HPC wire

AMPERE.

arm

Thank You
Danke
Gracias
谢谢
ありがとう
Asante
Merci
감사합니다
धन्यवाद
Kiitos
شكرًا
ধন্যবাদ
תודה