An Overview of Fugaku Supercomputer

Yutaka Ishikawa
Leader, Flagship2020 Project
RIKEN Center for Computational Science (R-CCS)

Russian Supercomputing Days 2019, 2019/09/24, 11:10-11:35

About RIKEN

- National institute for basic research in Japan covering physical to biological sciences
- Foundation in 1917
- 14 centers in 10 sites across Japan

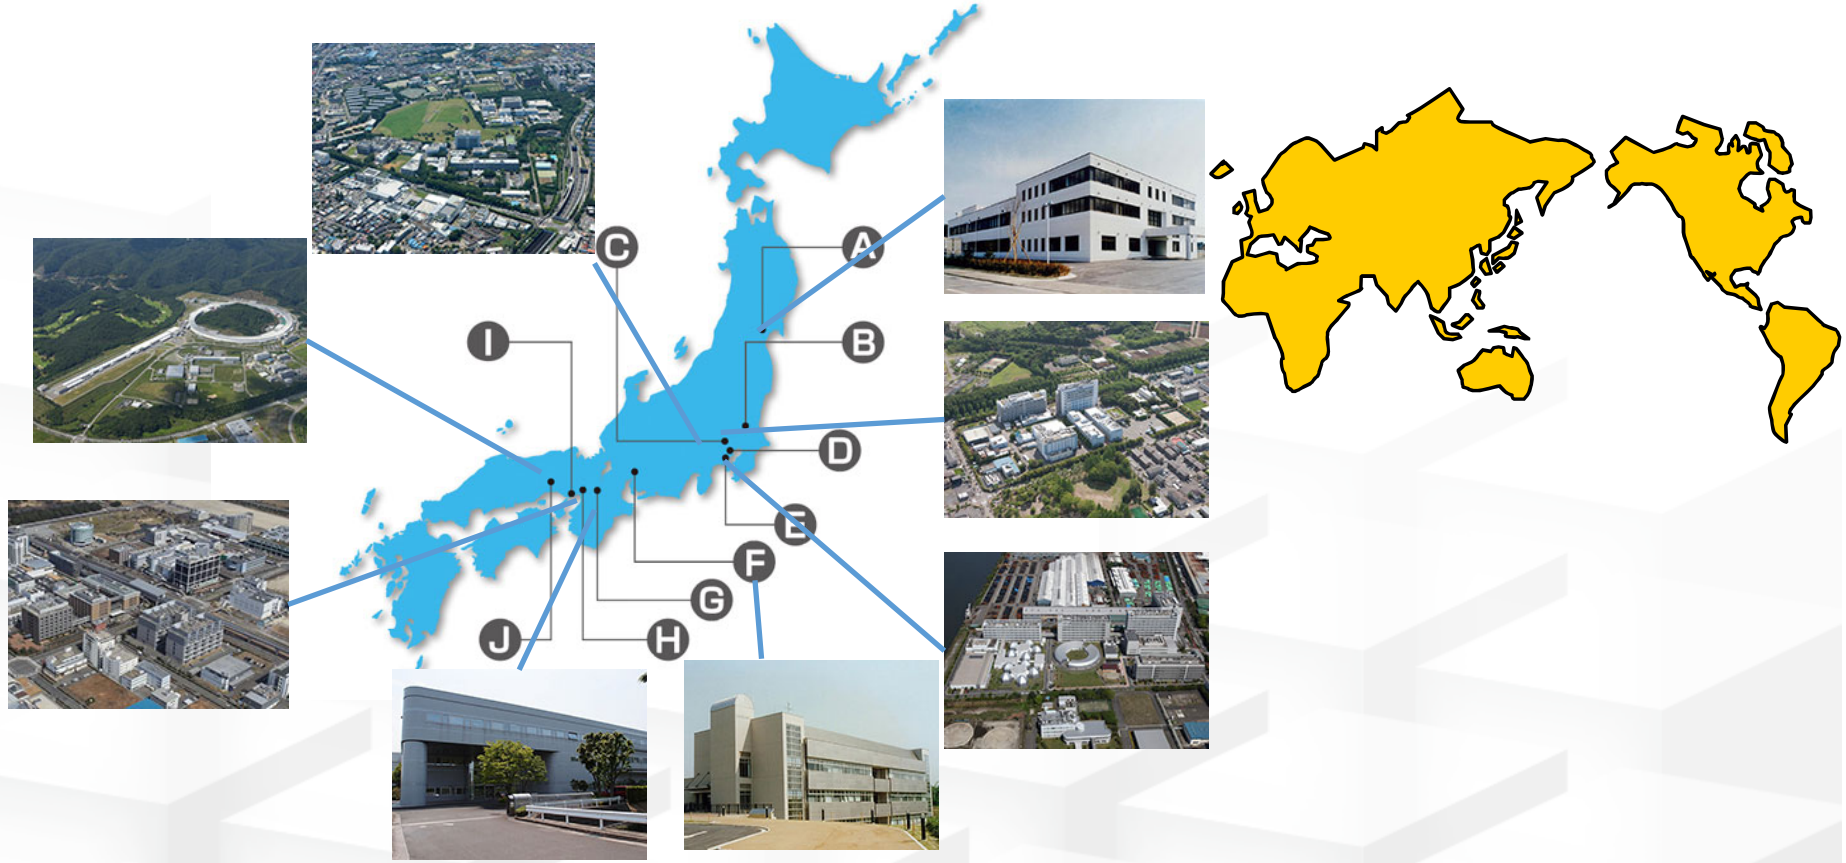

About RIKEN and R-CCS

- National institute for basic research in Japan covering physical to biological sciences
- Foundation in 1917
- 14 centers in 10 sites across Japan

- July 2010: RIKEN AICS (Advanced Institute of Computational Science) was established
- Sep. 2010: Installation of K computer began
- Jun 2012: #1 on TOP500
- Nov 2012: #1 (10PF+) on TOP500
- Apr. 2018: AICS is renamed to R-CCS (Riken Center for Computational Science)
- Aug. 2019: K computer was shutdown

Background: Flagship2020

- Missions
  - Building the Japanese national flagship supercomputer, post K, and
  - Developing wide range of HPC applications, running on post K, in order to solve social and science issues in Japan

- Project organization
  - Post K Computer development
    - R-CCS is in charge of development
    - Fujitsu is vendor partner.
    - International collaborations: DOE, CEA, JLESC (NCSA, ANL, UTK, JSC, BSC, INRIA, RIKEN)
  - Applications
    - The government selected
      - 9 social & scientific priority issues
      - 4 exploratory issues

      and their R&D organizations.

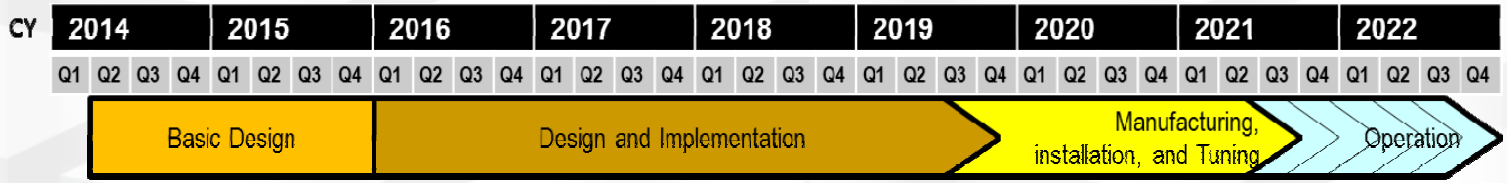

Background: Flagship2020

- Missions
  - Building the Japane
    and
  - Developing wide ra
    in order to solve sc
- Project organization
  - Post K Computer d
    - R-CCS is in charge
    - Fujitsu is vendor
    - International coll
      BSC, INRIA, RIKEN)
  - Applications
    - The government selected
      - 9 social & scientific priority issues
      - 4 exploratory issues

      and their R&D organizations.

	Target Applications	
	Program	**Brief description**
①	**GENESIS**	**MD for proteins**
②	**Genomon**	**Genome processing (Genome alignment)**
③	**GAMERA**	**Earthquake simulator (FEM in unstructured & structured grid)**
④	**NICAM+LETK**	**Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)**
⑤	**NTChem**	**molecular electronic (structure calculation)**
⑥	**FFB**	**Large Eddy Simulation (unstructured grid)**
⑦	**RSDFT**	**an ab-initio program (density functional theory)**
⑧	**Adventure**	**Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)**
⑨	**CCS-QCD**	**Lattice QCD simulation (structured grid Monte Carlo)**

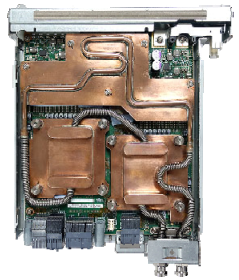

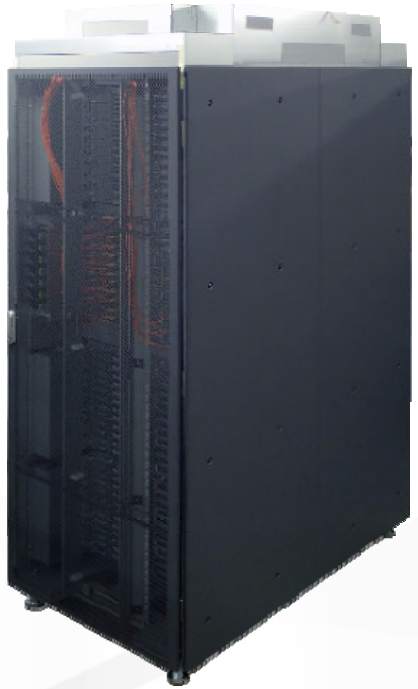

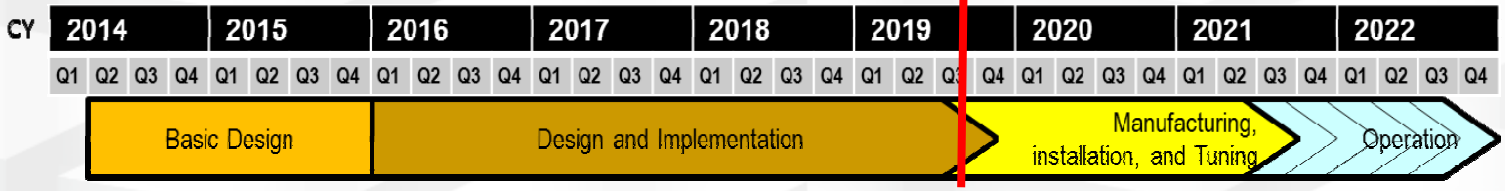

Fugaku

- A Fugaku prototype machine was built in Summer 2018. Since then, Fujitsu has been testing and evaluating the machine.
- Ten racks of Fugaku achieve almost the same performance of K computer (864 racks)

X 10 =

		Fugaku	K
CPU Architecture		A64FX (Armv8.2-A SVE +Fujitsu Extension)	SPARC64 VIIIfx
Node	Cores	48	8
	Peak DP performance	2.7+ TF	0.128 TF
	Main Memory	32 GiB	16 GiB
	Peak Memory Bandwidth	1024 GB/s	64 GB/s
	Peak Network Performance	40.8 GB/s	20 GB/s
Rack	Nodes	384	102
	Peak DP performance	1+ PF	< 0.013PF
Process Technology		7 nm FinFET	45 nm

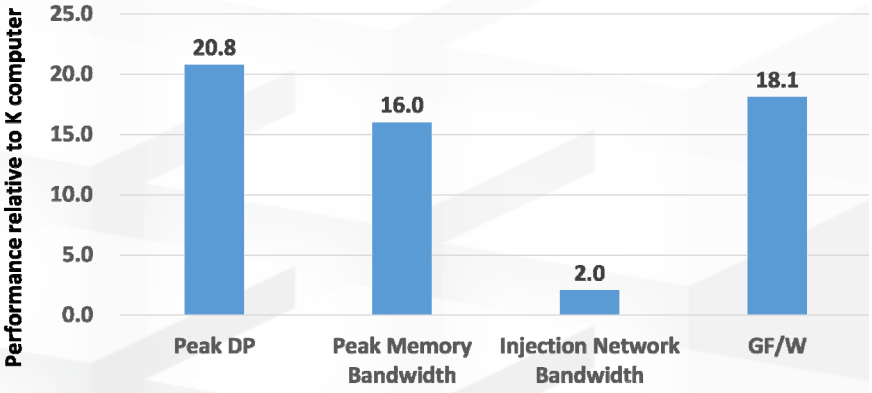

CPU A64FX

1.76+ GHz

urtesy of FUJITSU LIMITED

Architecture	Armv8.2-A SVE (512 bit SIMD)
Core	48 cores for compute and 2/4 for OS activities
	DP: 2.7+ TF, SP: 5.4+ TF, HP: 10.8+ TF
Cache L1	64 KiB, 4 way, 230+ GB/s(load), 115+ GB/s (store)
Cache L2	CMG(NUMA): 8 MiB, 16way Node: 3.6+ TB/s Core: 115+ GB/s (load), 57+ GB/s (store)
Memory	HBM2 32 GiB, 1024 GB/s
Interconnect	TofuD (28 Gbps x 2 lane x 10 port)
I/O	PCIe Gen3 x 16 lane
Technology	7nm FinFET

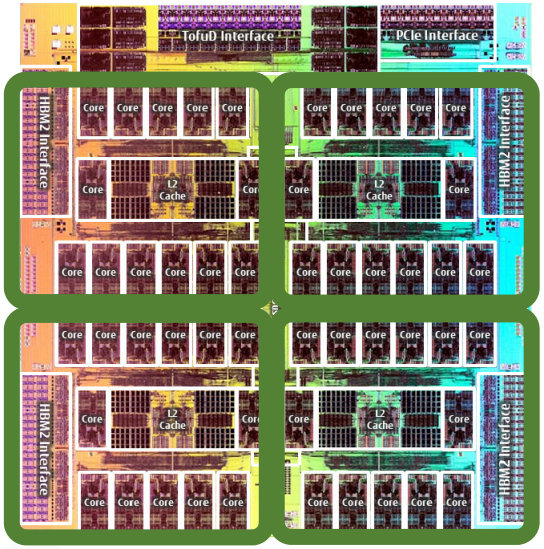

Performance

Stream triad: 830+ GB/s

Dgemm: 2.5+ TF (90+% efficiency)

ref. Toshio Yoshida, “Fujitsu High Performance CPU for the Post-K Computer,”
IEEE Hot Chips: A Symposium on High Performance Chips, San Jose, August 21, 2018.

TofuD Interconnect

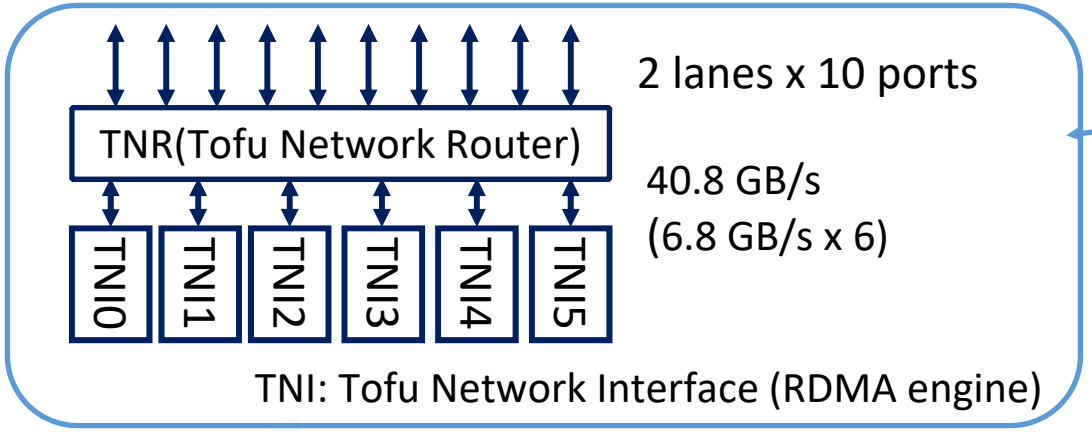

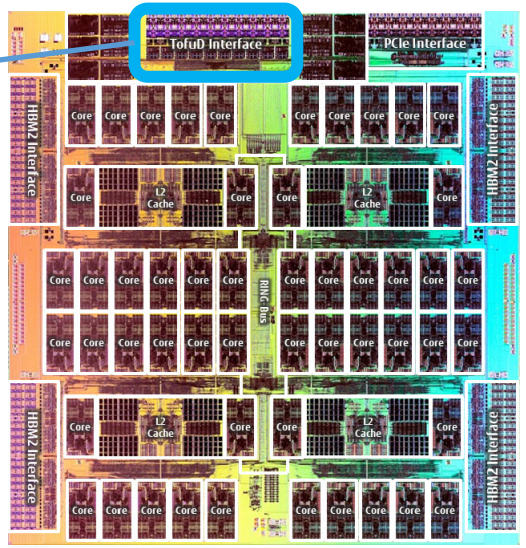

- 6 RDMA Engines
- Hardware barrier support
- Network operation offloading capability

8B Put latency	**0.49 – 0.54 usec**
1MiB Put throughput	**6.35 GB/s**

rf. Yuichiro Ajima, et al. , “The Tofu Interconnect D,” IEEE Cluster 2018, 2018.

An Overview of Fugaku Hardware

- **150k+ node**

2.7 TF x 150k+ = 405+ PF

- **Two types of nodes**
  - Compute Node and Compute & I/O Node connected by Fujitsu TofuD, 6D mesh/torus Interconnect
- **3-level hierarchical storage system**
  - 1st Layer
    - One of 16 compute nodes, called Compute & Storage I/O Node, has SSD about 1.6 TB
    - Services
      - ~ Cache for global file system
      - ~ Temporary file systems
        - Local file system for compute node
        - Shared file system for a job
  - 2nd Layer
    - Fujitsu FEFS: Lustre-based global file system
  - 3rd Layer
    - Cloud storage services

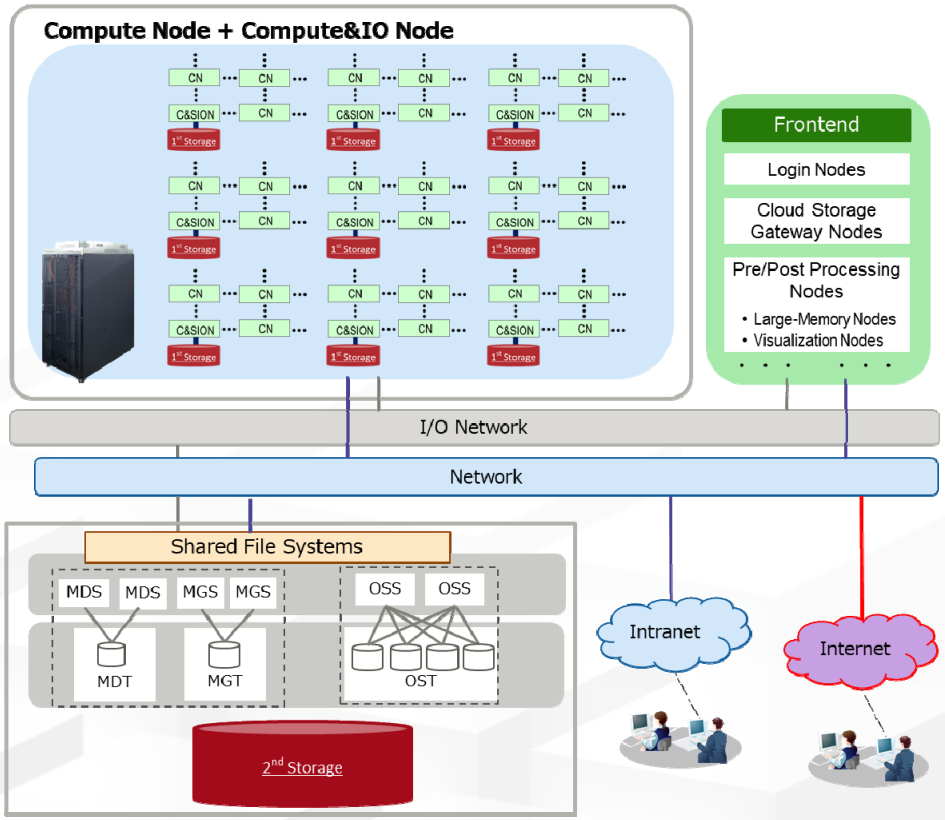

Recent Result by Fujitsu

ref. Toshiyuki Shimizu, “POST-K SUPERCOMPUTER DEVELOPMENT,” SCAsia 2019.

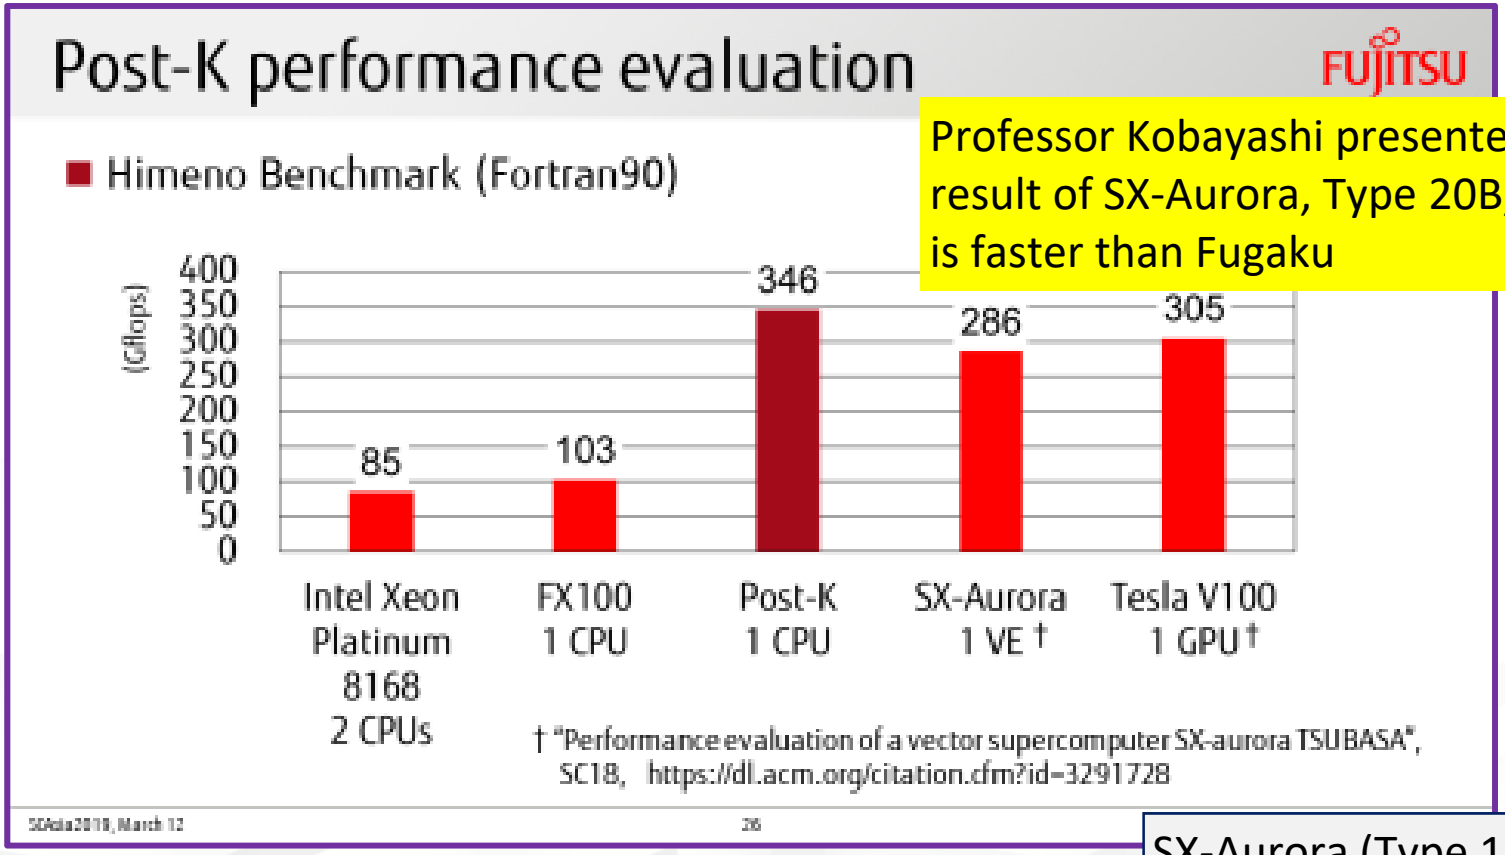

Professor Kobayashi presented the updated result of SX-Aurora, Type 20B, 355GF, which is faster than Fugaku

- Himeno benchmak is a solver for the Poisson's equation solution using the Jacobi iteration method

SX-Aurora (Type 10B)
1.4 GHz, 2.15 TF, 1228.8 GB/s
Tesla V100
1.245 GHz, 7 TF, 900 GB/s
Fugaku (Post-K)
1.76+ GHz, 2.7+ TF, 1024 GB/s

Estimated Performance Achievement

Performance Targets

- ✓ 100 times faster than K for some applications (tuning included)
- ✓ 30 to 40 MW power consumption

Peak Performance

	PostK	K
Peak DP (double precision)	400+ Pflops (34x +)	11.3 Pflops*
Peak SP (single precision)	800+ Pflops (70x +)	11.3 Pflops
Peak HP (half precision)	1600+ Pflops (141x +)	--
Total memory bandwidth	150+ PB/sec (29x +)	5,184TB/sec

* Reported in TOP500 (including I/O nodes)

Geometric Mean of Performance Speedup of the 9 Target Applications over the K-Computer

37x +

Predicted Performance of 9 Target Applications

As of 2019/05/14

Area	Priority Issue	Performance Speedup over K	Application	Brief description
Health and longevity	1. Innovative computing infrastructure for drug discovery	125x +	GENESIS	MD for proteins
	2. Personalized and preventive medicine using big data	8x +	Genomon	Genome processing (Genome alignment)
Disaster prevention and Environment	3. Integrated simulation systems induced by earthquake and tsunami	45x +	GAMERA	Earthquake simulator (FEM in unstructured & structured grid)
	4. Meteorological and global environmental prediction using big data	120x +	NICAM+ LETKF	Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)
Energy issue	5. New technologies for energy creation, conversion / storage, and use	40x +	NTChem	Molecular electronic simulation (structure calculation)
	6. Accelerated development of innovative clean energy systems	35x +	Adventure	Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)
Industrial competitiveness enhancement	7. Creation of new functional devices and high-performance materials	30x +	RSDFT	Ab-initio simulation (density functional theory)
	8. Development of innovative design and production processes	25x +	FFB	Large Eddy Simulation (unstructured grid)
Basic science	9. Elucidation of the fundamental laws and evolution of the universe	25x +	LQCD	Lattice QCD simulation (structured grid Monte Carlo)

An Overview of System Software Stack

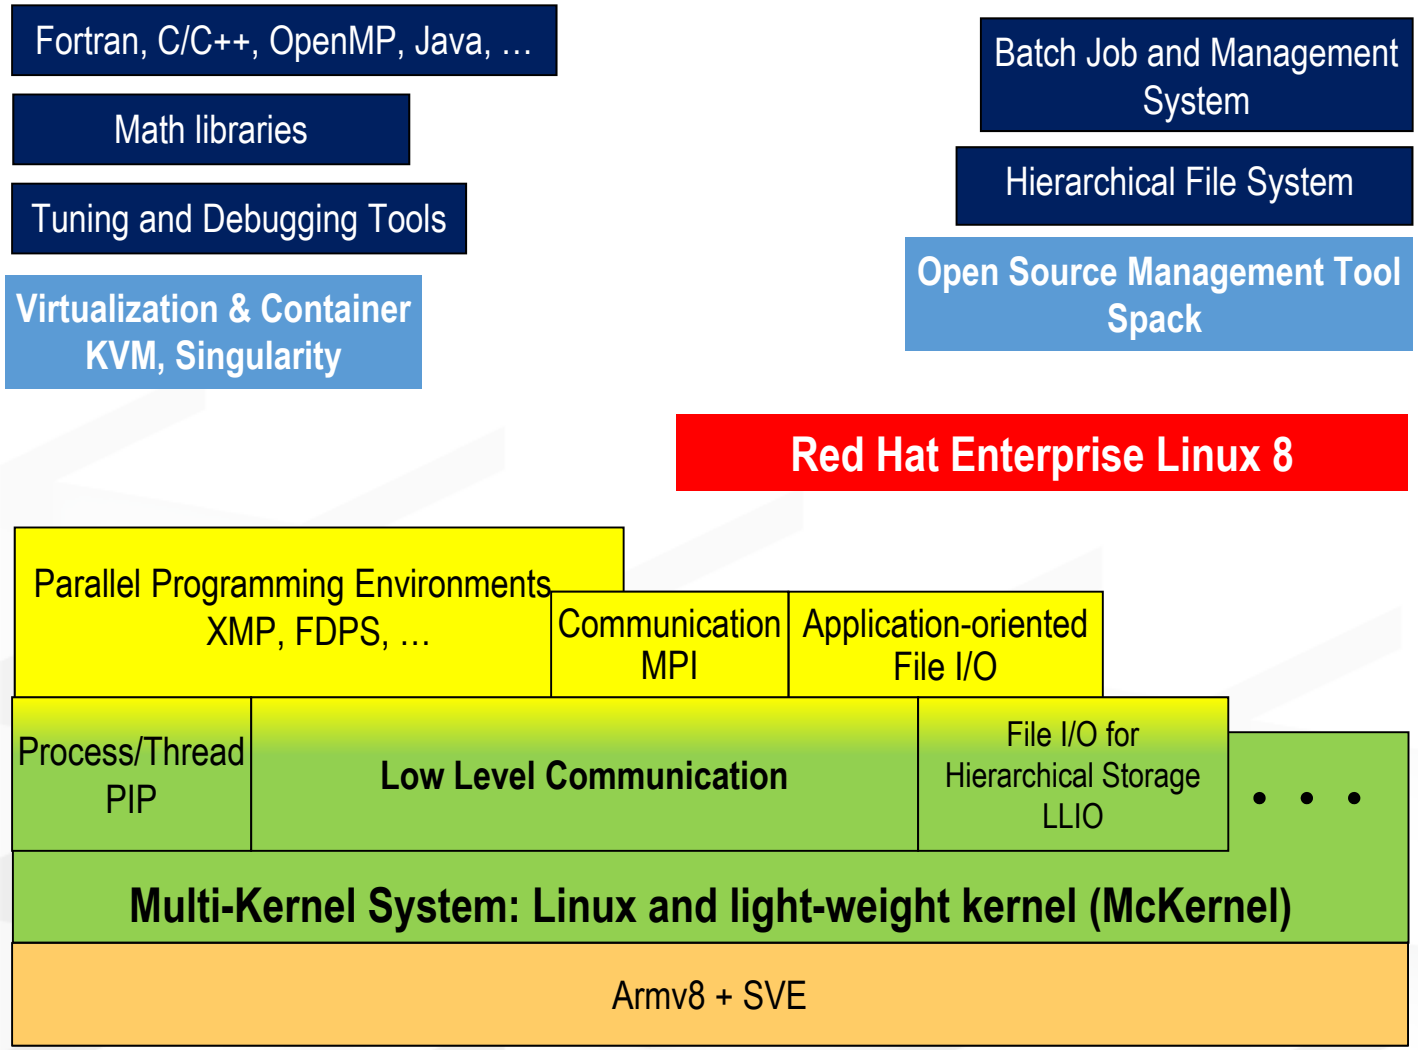

Concluding Remarks

- The K supercomputer is being removed. After reconstruction of electrical and water-cooling system, installation of Fugaku will begin in the end of 2019
- Early access program will start in 2020
- General operation will start around 2021