

# Fugaku**NEXT** Project: next Japanese flagship supercomputer

**Yasumichi Aoki**

Principal, Field Theory Research Team

Leader, Next-Gen HPC Application Development Unit (**FugakuNEXT**)

RIKEN Center for Computational Science (R-CCS), Kobe, Japan

# K → Fugaku → FugakuNEXT

- **K:** SPARC64 VIIIIfx CPUs (Fujitsu) x 88k 11 PFlops
- **Fugaku** A64FX ARM v8.20A (Fujitsu) x 159k 540 PFlops

- **FugakuNEXT : Successor of Fugaku**

- CPU Fujitsu Monaca (ARM)
- **GPU NVIDIA** ← Big change

- **Lattice Field Theory...**

K



2012-2019

Fugaku

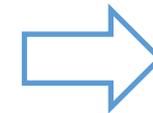
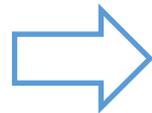


2021-

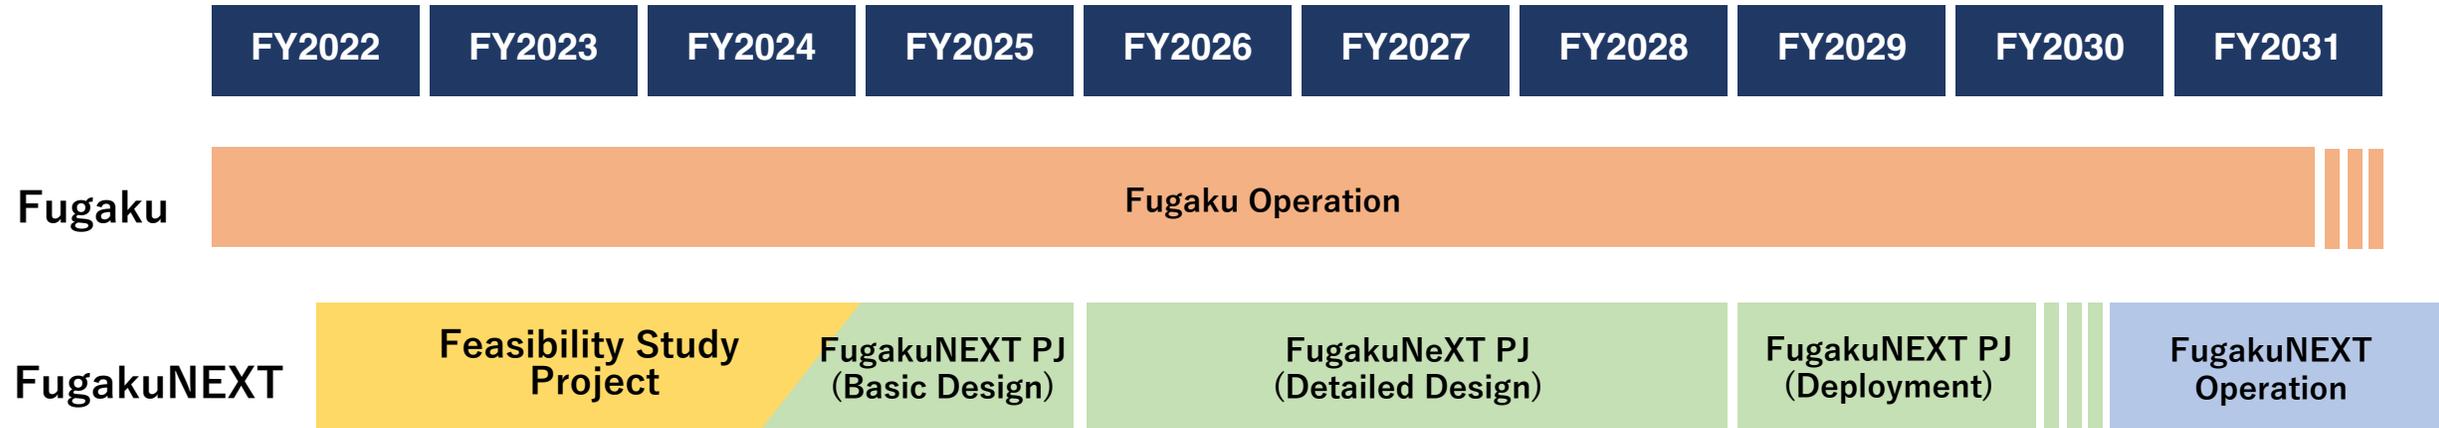
“FugakuNEXT”



~2030-

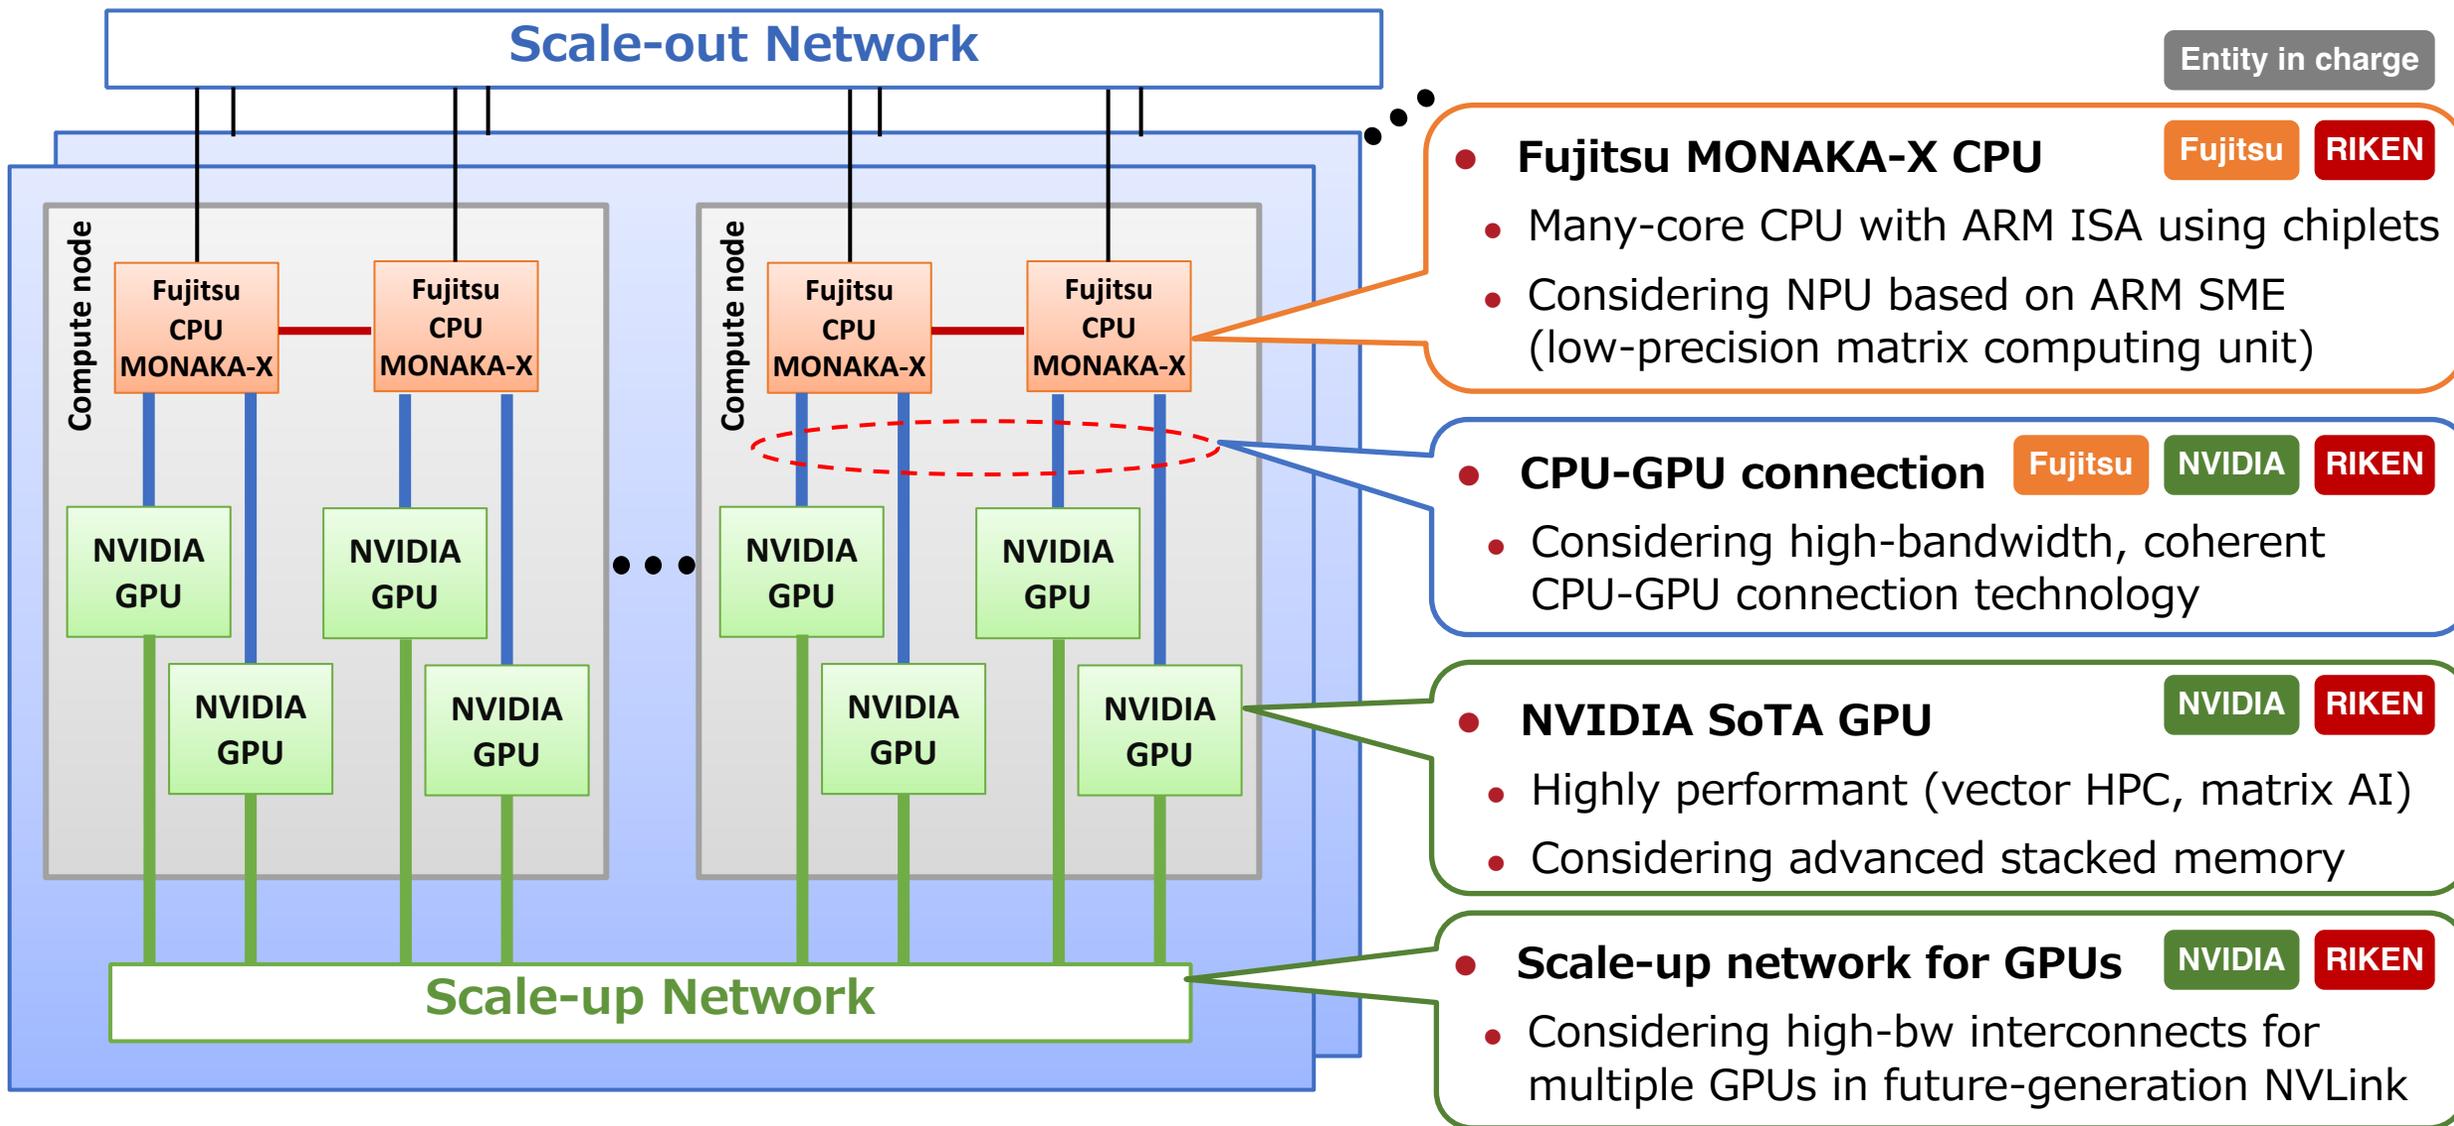


# Expected Timeline of Fugaku-NEXT R&D



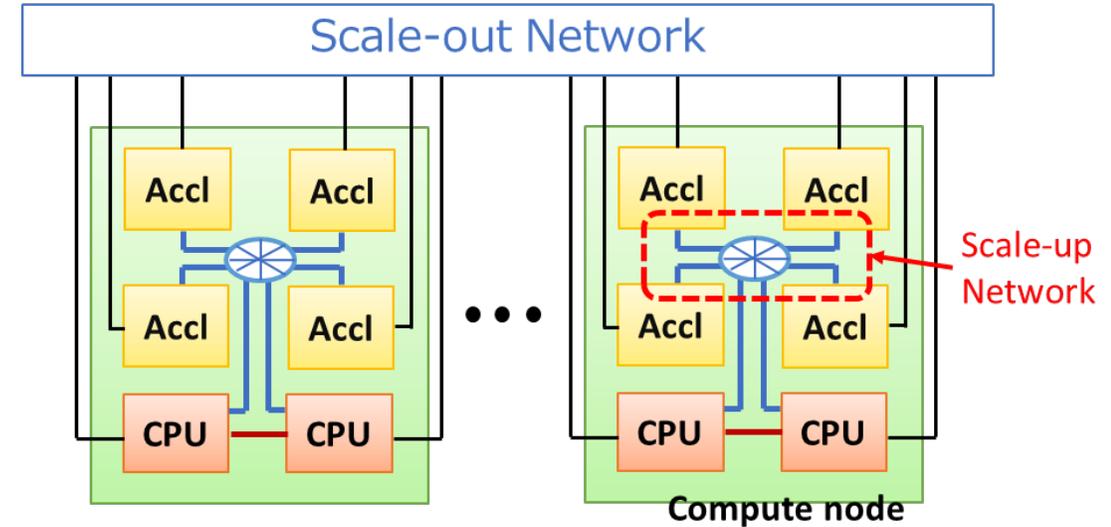
- **Launched development project of FugakuNEXT, for starting operation in 2030**
  - Establish a project organization built upon the existing collaboration structure
- **Initial targets made based on findings in the feasibility study**
  - System performance, system software development, application challenges, and others
  - In RIKEN, studying HPC-Quantum hybrid platform continuously to extend computable problems
- **Public tenders for basic design**
  - Overall system, computing nodes, and CPUs
  - Accelerators (GPUs)

# Overview of FugakuNEXT System Architecture



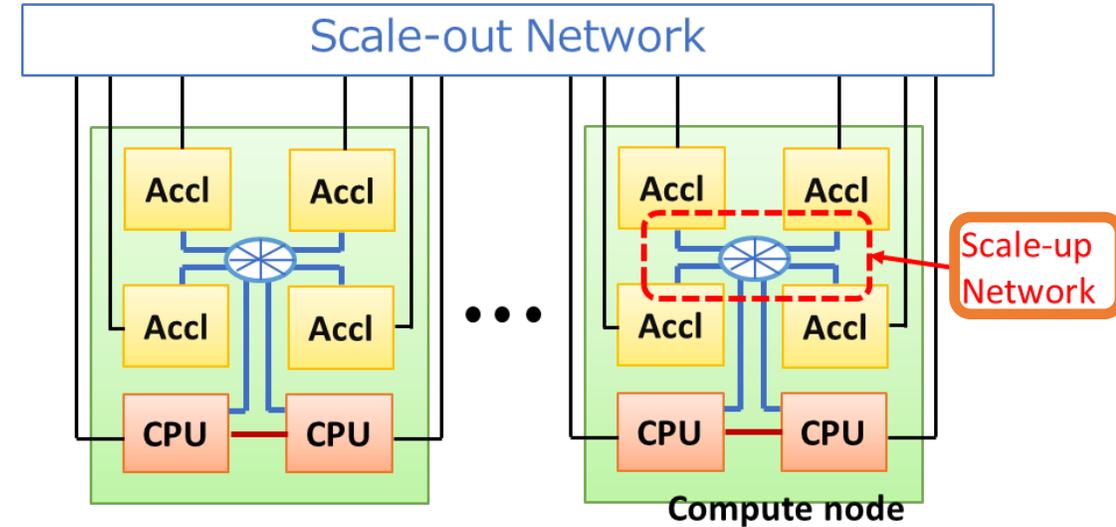
Detailed configurations and specs. will be determined in Basic design phase.

- **FugakuNEXT: Successor of “Fugaku” 2021~**
- **Launched development project**
  - Target performance based on the findings of FS, starting operation in 2030
  - “Application First” philosophy
  - Heterogeneous architecture of CPU & GPU with AI for science capability
  - Use of assets such as app software for “Fugaku”



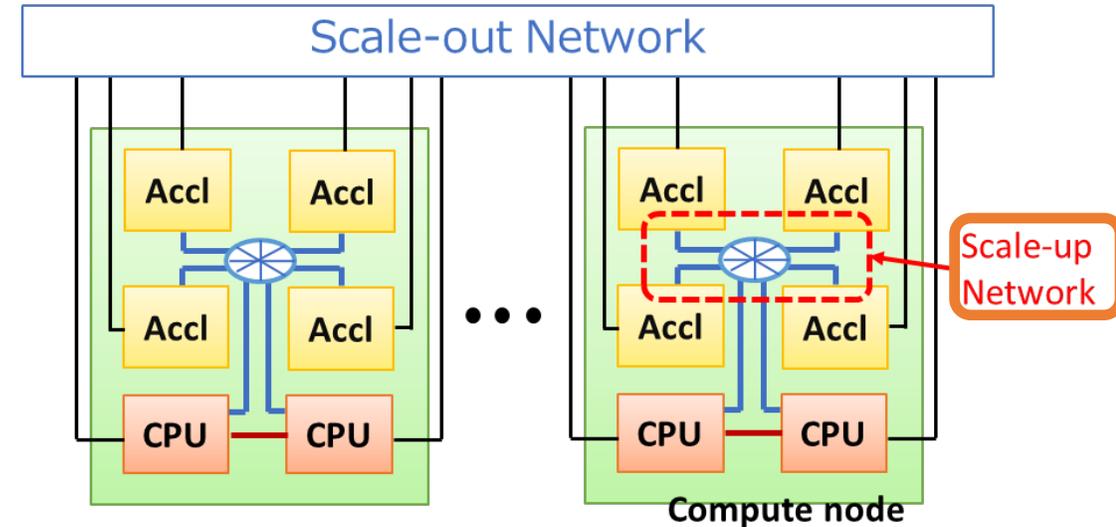
Items	CPUs	Accelerators	Fugaku	Improvement
Total node count	>= 3,400		158,976	
Theoretical peak perf. of FP64 vector	>= 48 PFLOPS	>= 3.0 EFLOPS	488-537 PFLOPS	x5.7~
Theoretical peak perf. of FP16/BF16 matrix	>= 1.5 EFLOPS	>= 150 EFLOPS	1.95-2.15 EFLOPS	x70.5~
Theoretical peak perf. of FP8 matrix	>= 3.0 ELOPS	>= 300 EFLOP	—	
(Sparsity considered FP8 matrix)	—	>= 600 EFLOPS	—	
Main memory capacity	>= 10 PiB	>= 10 PiB	4.85 PiB	x4.1~
Main memory bandwidth	>= 7 PB/s	>= 800 PB/s	163 PB/s	x4.9~
System power consumption	<= 40 MW (compute node and storage)		about 30 MW	

- **FugakuNEXT: Successor of “Fugaku” 2021~**
- **Launched development project**
  - Target performance based on the findings of FS, starting operation in 2030
  - “Application First” philosophy
  - Heterogeneous architecture of CPU & GPU with AI for science capability
  - Use of assets such as app software for “Fugaku”



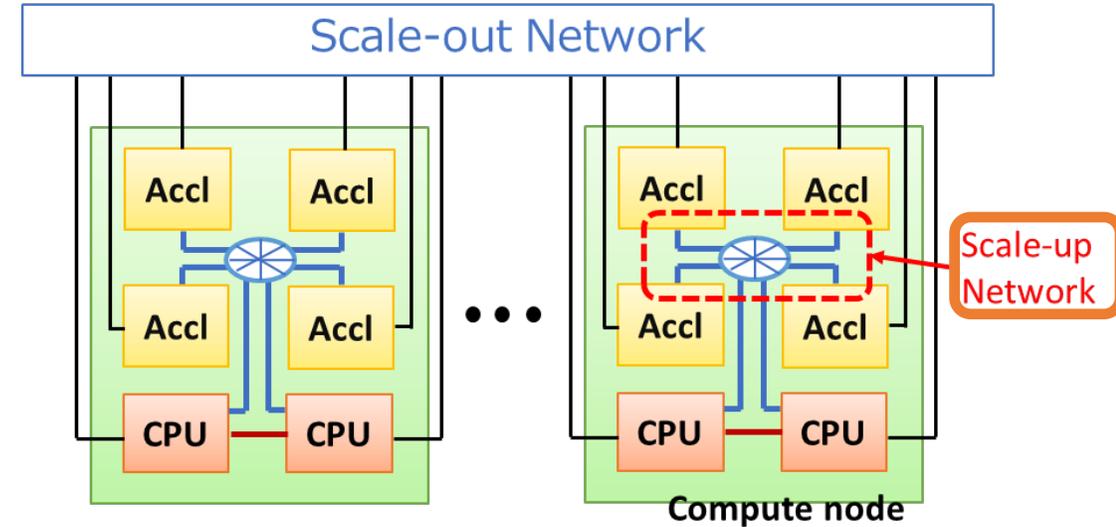
Items	CPUs	Accelerators	Fugaku	Improvement
Total node count	>= 3,400		158,976	
Theoretical peak perf. of FP64 vector	>= 48 PFLOPS	>= 3.0 EFLOPS	488-537 PFLOPS	x5.7~
Theoretical peak perf. of FP16/BF16 matrix	>= 1.5 EFLOPS	>= 150 EFLOPS	1.95-2.15 EFLOPS	x70.5~
Theoretical peak perf. of FP8 matrix	>= 3.0 ELOPS	>= 300 EFLOP	—	
(Sparsity considered FP8 matrix)	—	>= 600 EFLOPS	—	
Main memory capacity	>= 10 PiB	>= 10 PiB	4.85 PiB	x4.1~
Main memory bandwidth	>= 7 PB/s	>= 800 PB/s	163 PB/s	x4.9~
System power consumption	<= 40 MW (compute node and storage)		about 30 MW	

- **FugakuNEXT: Successor of “Fugaku” 2021~**
- **Launched development project**
  - Target performance based on the findings of FS, starting operation in 2030
  - “Application First” philosophy
  - Heterogeneous architecture of CPU & GPU with AI for science capability
  - Use of assets such as app software for “Fugaku”



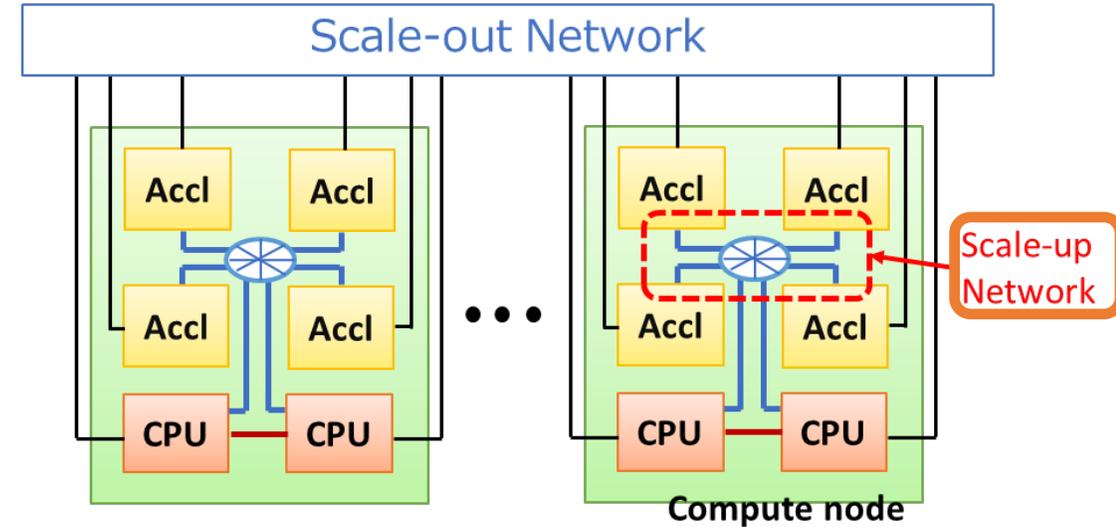
Items	CPU	Accelerators	Fugaku	Improvement
Total node count	>= 3,400		158,976	
Theoretical peak perf. of FP64 vector	>= 48 PFLOPS	>= 3.0 EFLOPS	488-537 PFLOPS	x5.7~
Theoretical peak perf. of FP16/BF16 matrix	>= 1.5 EFLOPS	>= 150 EFLOPS	1.95-2.15 EFLOPS	x70.5~
Theoretical peak perf. of FP8 matrix	>= 3.0 ELOPS	>= 300 EFLOP	—	
(Sparsity considered FP8 matrix)	—	>= 600 EFLOPS	—	
Main memory capacity	>= 10 PiB	>= 10 PiB	4.85 PiB	x4.1~
Main memory bandwidth	>= 7 PB/s	>= 800 PB/s	163 PB/s	x4.9~
System power consumption	<= 40 MW (compute node and storage)		about 30 MW	

- **FugakuNEXT: Successor of “Fugaku” 2021~**
- **Launched development project**
  - Target performance based on the findings of FS, starting operation in 2030
  - “Application First” philosophy
  - Heterogeneous architecture of CPU & GPU with AI for science capability
  - Use of assets such as app software for “Fugaku”



Items	CPUs	Accelerators	Fugaku	Improvement
Total node count	>= 3,400		158,976	
Theoretical peak perf. of FP64 vector	>= 48 PFLOPS	>= 3.0 EFLOPS	488-537 PFLOPS	x5.7~
Theoretical peak perf. of FP16/BF16 matrix	>= 1.5 EFLOPS	>= 150 EFLOPS	1.95-2.15 EFLOPS	x70.5~
Theoretical peak perf. of FP8 matrix	>= 3.0 ELOPS	>= 300 EFLOP	—	
(Sparsity considered FP8 matrix)	—	>= 600 EFLOPS	—	
Main memory capacity	>= 10 PiB	>= 10 PiB	4.85 PiB	x4.1~
Main memory bandwidth	>= 7 PB/s	>= 800 PB/s	163 PB/s	x4.9~
System power consumption	<= 40 MW (compute node and storage)		about 30 MW	

- **FugakuNEXT: Successor of “Fugaku” 2021~**
- **Launched development project**
  - Target performance based on the findings of FS, starting operation in 2030
  - “Application First” philosophy
  - Heterogeneous architecture of CPU & GPU with AI for science capability
  - Use of assets such as app software for “Fugaku”



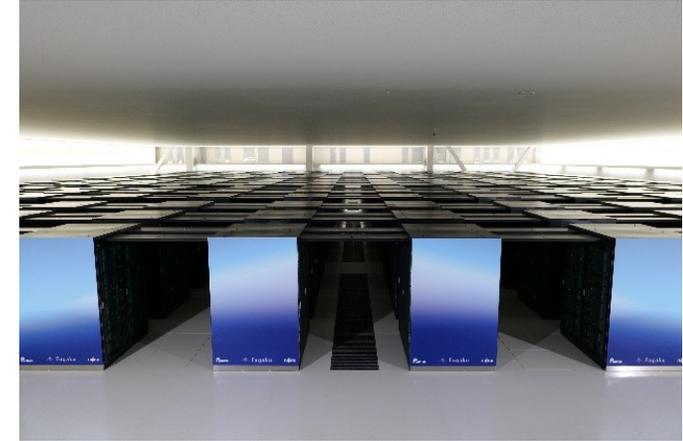
Items	CPUs	Accelerators	Fugaku	Improvement
Total node count	>= 3,400		158,976	
Theoretical peak perf. of FP64 vector	>= 48 PFLOPS	>= 3.0 EFLOPS	488-537 PFLOPS	x5.7~
Theoretical peak perf. of FP16/BF16 matrix	>= 1.5 EFLOPS	>= 150 EFLOPS	1.95-2.15 EFLOPS	x70.5~
Theoretical peak perf. of FP8 matrix	>= 3.0 ELOPS	>= 300 EFLOP	—	
(Sparsity considered FP8 matrix)	—	>= 600 EFLOPS	—	
Main memory capacity	>= 10 PiB	>= 10 PiB	4.85 PiB	x4.1~
Main memory bandwidth	>= 7 PB/s	>= 800 PB/s	163 PB/s	x4.9~
System power consumption	<= 40 MW (compute node and storage)		about 30 MW	

# Co-design and development

- **EEA – Early Evaluation Application – 14 applications**
- **Lattice QCD is one of them: Bridge++ (DWF-HMC)**
- **Plan to have an open environment for evaluation**
  - Utilizing CX framework : BenchKit, Benchmark
- **Challenge and opportunity for developers**
  - GPU porting and tuning
  - Utilize hierarchical memory and network
  - Use of highspeed, but low precision arithmetic engine
    - Dense matrix multiplication (use of Ozaki scheme etc)
    - AI

**back up**

# Quick Recap : Supercomputer "Fugaku", 2021~



**432 racks**

158,976 general-purpose many-core processors Fujitsu A64FX, 7 M Cores

**Half exaflops in DP**

1 Exaflops in SP!

**Very power-efficient system**  
with 20 MW power-consumption!

The Public Service has started  
in March, 2021.

# Fugaku, No.1 for Four Rankings, June 2020 to Nov 2021

- June 2020: partial system → Nov. 2020 : full system (#nodes, frequency 2.2GHz)

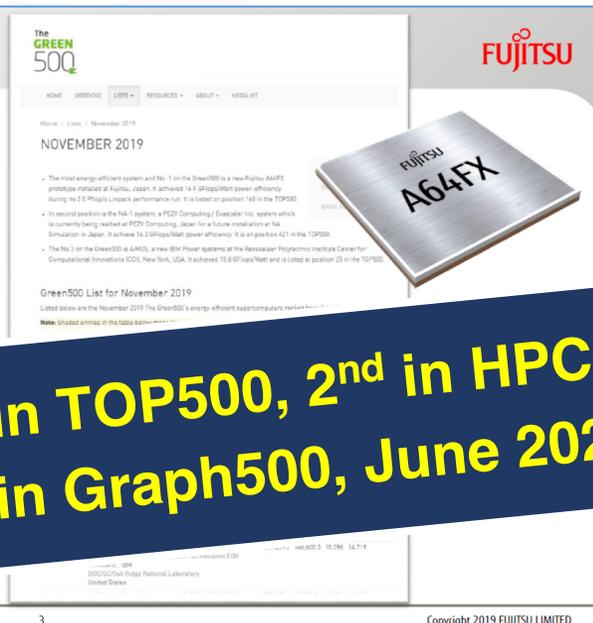
	Measured	Peak Perf	Efficiency	(June 2020)	Perf of 2 <sup>nd</sup> System	1 <sup>st</sup> / 2 <sup>nd</sup>
<b>LINPACK</b>	442.01 PF	537.21 PF	82.3%	(415.5 PF)	148.60 PF	<b>3.0</b>
<b>HPCG</b>	16.00 PF	537.21 PF	3.0%	(13.4PF)	2.92 PF	<b>5.5</b>
<b>HPL-AI</b>	2.00 EF	2.14 EF	93.2%	(1.42EF)	0.55 EF	<b>3.6</b>
<b>Graph500</b>	102.95 Tsteps			(70.98)	23.75 Tsteps	<b>4.3</b>

## Green500, Nov. 2019

A64FX prototype –  
Fujitsu A64FX 48C 2GHz  
ranked **#1** on the list

768x general purpose A64FX  
CPU w/o accelerators

- 1.9995 PFLOPS @ HPL
- 16.876 GF/W
- Power quality level 2

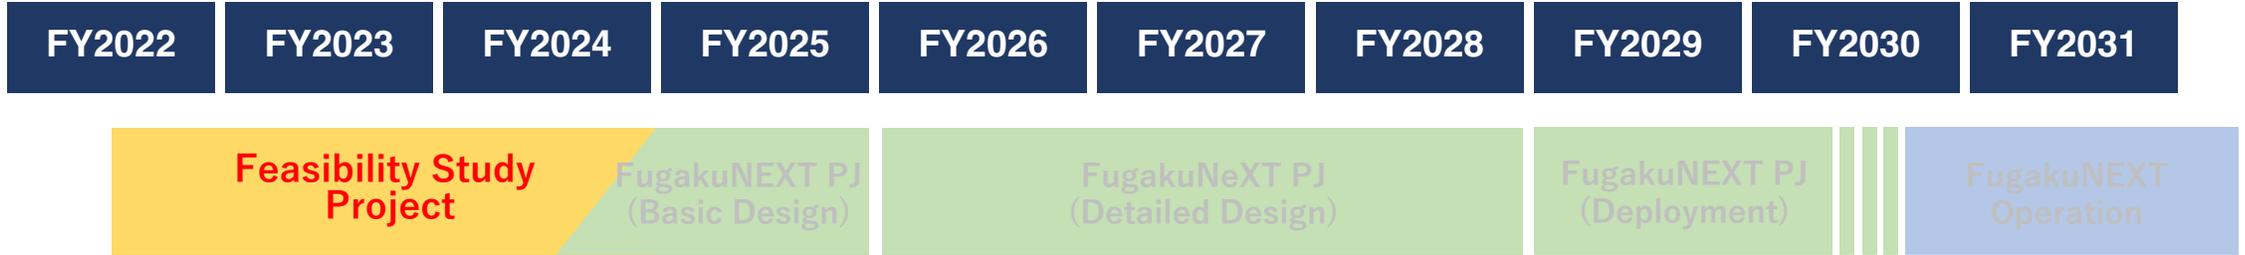


**7<sup>th</sup> in TOP500, 2<sup>nd</sup> in HPCG  
1<sup>st</sup> in Graph500, June 2025**

**3 to 4 times faster  
in every benchmark!**

A64FX prototype –  
No.1 in Green500 at SC19 (Nov/2019)  
16.87 GFLOPS/W

# FugakuNEXT: from Feasibility Study To R&D Project



- **Feasibility study project on next-gen HPC infrastructure**
- **Commissioned by MEXT\*** (2.5 years; Aug 2022 - Mar 2025)
- **Four teams selected for**
  - **System study** (two teams)
  - **New computing paradigm study** (actually, quantum computing)
  - **Operation technology study**
- **RIKEN team for system study**

**システムチーム** 次世代計算基盤として想定されるアーキテクチャ（プロセッサ、メモリ、ストレージ等）、システムソフトウェア、アプリケーションを提案

**代表機関：理化学研究所（近藤 正樹）**  
 オールジャパンかつ国外ベンダーも参画するデジタルツイン実現の基盤として、電力制約の厳しい環境下で、幅広いアプリ分野に適用可能な、高効率化し、低コストなシステムを提案

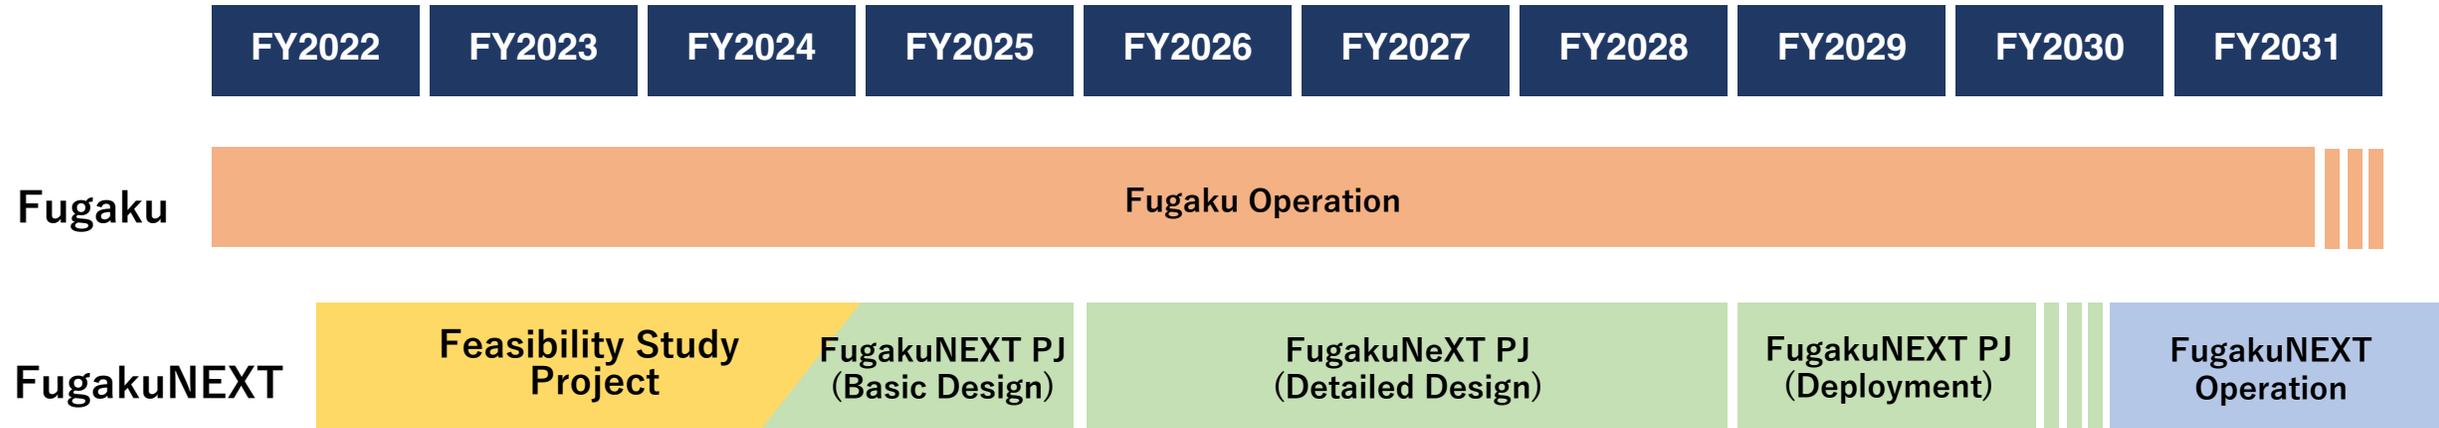
**代表機関：神戸大学（牧野 淳一郎）**  
 世界最高の電力当たり性能を実現するAI技術、AI応用技術を活用し、従来よりも高い実行効率を実現するシステムを提案

**代表機関：慶應義塾大学（R4～5：天野 英晴）**  
 量子コンピューティングとスーパーコンピューティングの連携可能なシステムの実現可能なシステムを提案

**代表機関：東京大学（堀 誠司）**  
 大学情報基盤センターが推進するHPCI第二階層システムと有機的に連携したシステムの実現可能なシステムを提案

\*) MEXT: Ministry of Education, Culture, Sports, Science and Technology

# Expected Timeline of Fugaku-NEXT R&D



- **Launched development project of FugakuNEXT, for starting operation in 2030**
  - Establish a project organization built upon the existing collaboration structure
- **Initial targets made based on findings in the feasibility study**
  - System performance, system software development, application challenges, and others
  - In RIKEN, studying **HPC-Quantum hybrid platform** continuously to extend computable problems
- **Public tenders for basic design**
  - Overall system, computing nodes, and CPUs
  - Accelerators (GPUs)

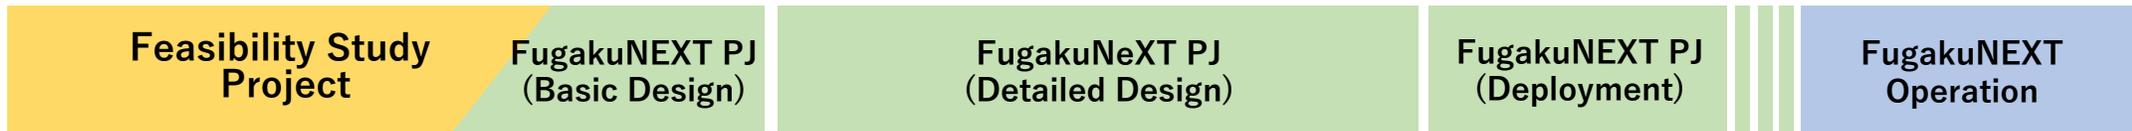
# Expected Timeline of Fugaku-NEXT R&D



Fugaku

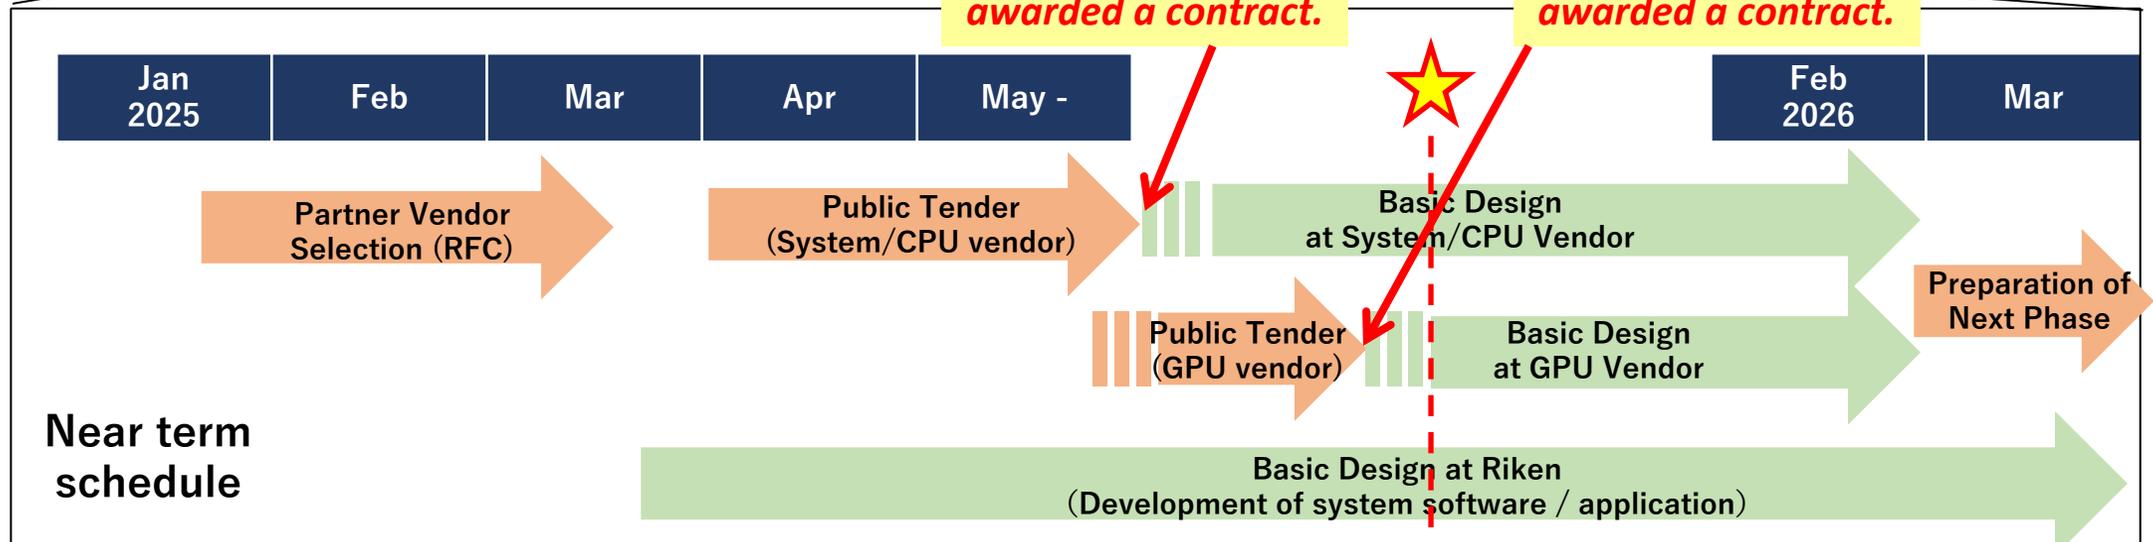


FugakuNEXT

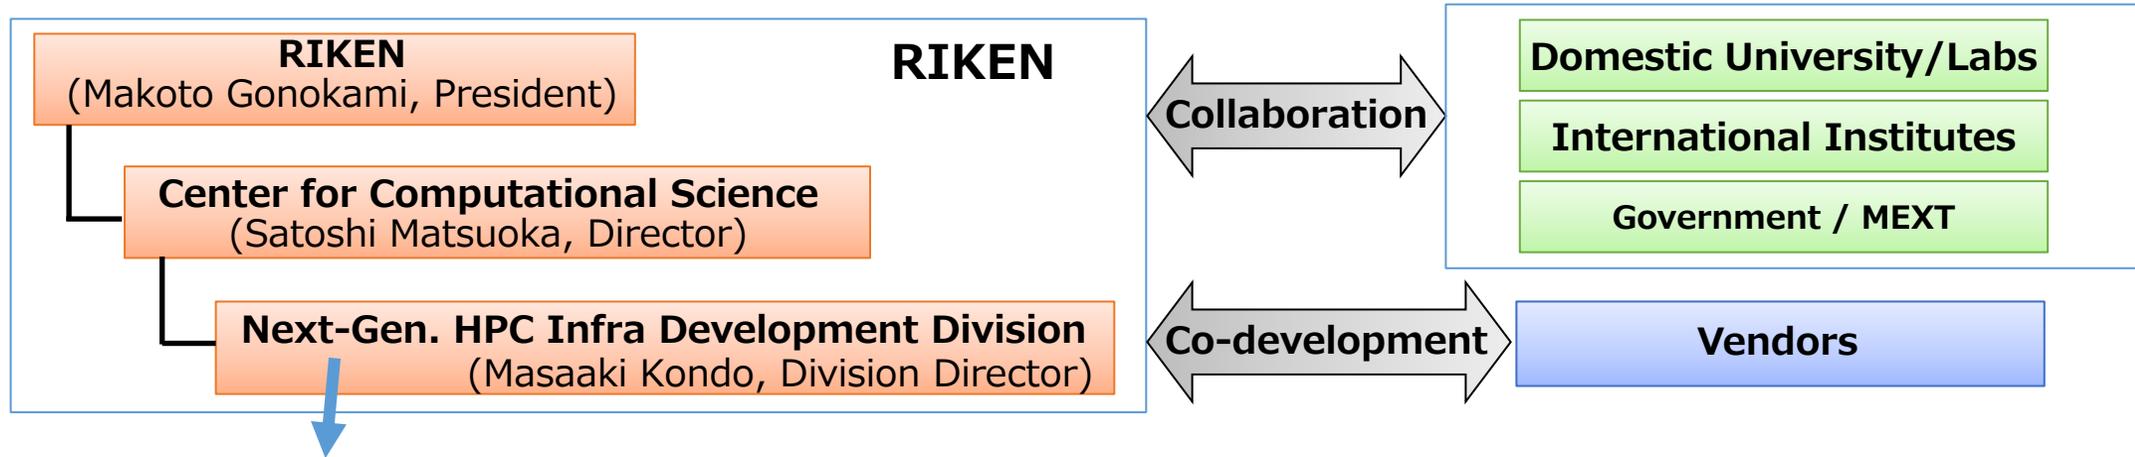


*Fujitsu Ltd. was awarded a contract.*

*NVIDIA was awarded a contract.*



# New Division for FugakuNEXT Project and Beyond



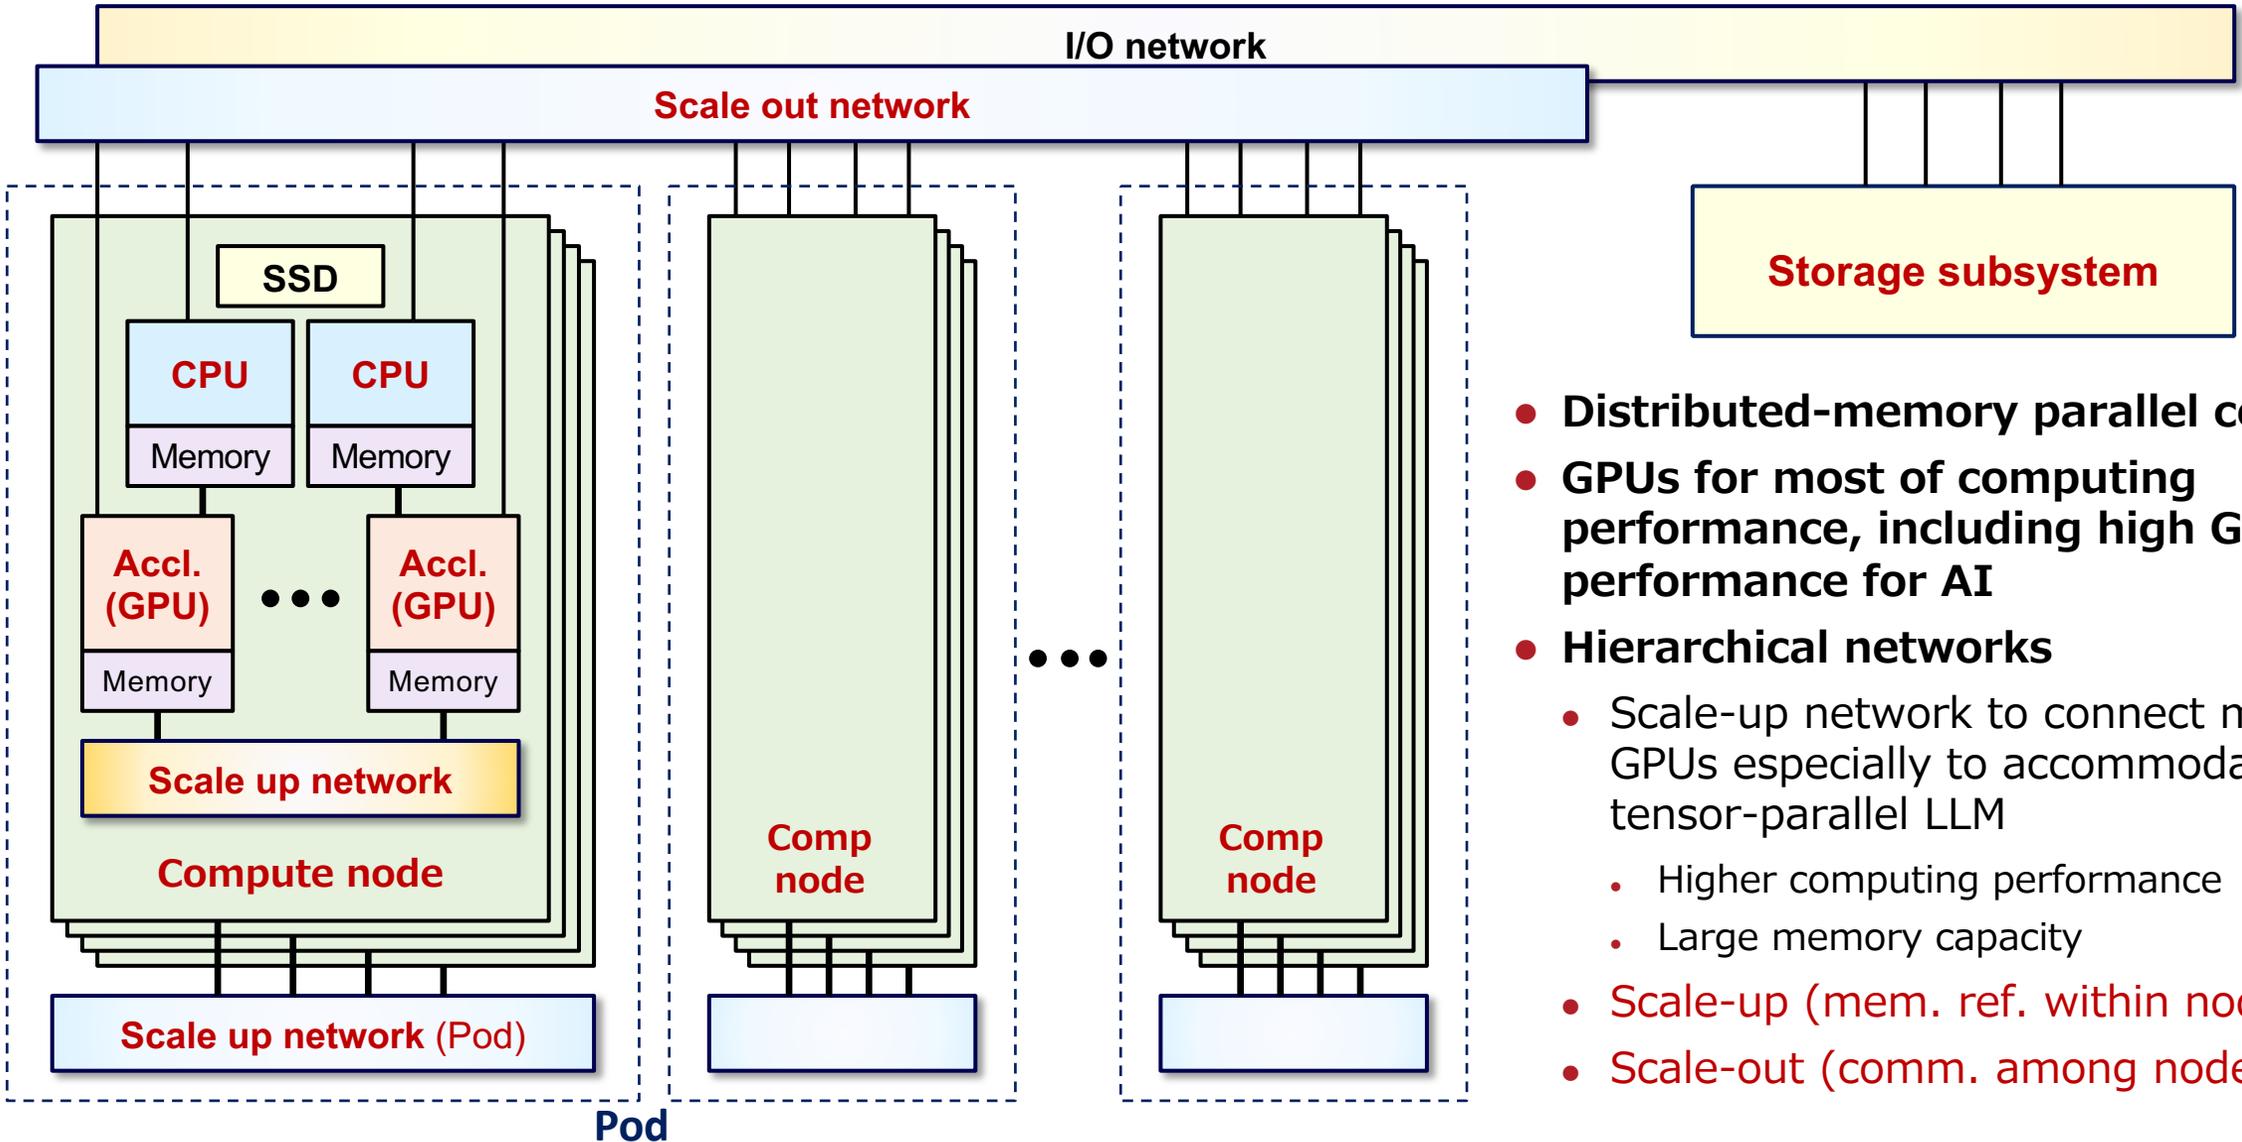
Division	Units		Tasks
<b>Next-Generation HPC Infrastructure Development Division</b>    <b>Masaaki Kondo</b> , Division Director <b>Fumiyoshi Shoji</b> , Deputy Division Director	<b>Next-Generation HPC Infrastructure System Development Unit</b>	<b>Kentaro Sano</b> , Unit Leader 	Development of FugakuNEXT system including architecture and system SW
	<b>Next-Generation HPC Application Development Unit</b>	<b>Yasumichi Aoki</b> , Unit Leader 	Development, support, and co-design of next generation apps for FugakuNEXT
	<b>Next-Generation HPC Operation Technologies Unit (Under Consideration)</b>	<b>Keiji Yamamoto</b> , Unit Leader 	Development of environment and operation technologies for FugakuNEXT
	<b>Advanced HPC Technologies Development Unit</b>	<b>Kento Sato</b> , Unit Leader 	Feasibility study and elementally technology research for advanced system
	<b>Next-Generation HPC Management Office</b>	<b>Yoji Shimada</b> , Office manager 	Responsible of the projet management of FugakuNEXT and beyond

# Overview of **FugakuNEXT**

## Architecture and System

### System software

### Application



- **Distributed-memory parallel computer**
- **GPUs for most of computing performance, including high GEMM performance for AI**
- **Hierarchical networks**
  - Scale-up network to connect multiple GPUs especially to accommodate tensor-parallel LLM
    - Higher computing performance
    - Large memory capacity
- **Scale-up (mem. ref. within node/pod)**
- **Scale-out (comm. among nodes/pods)**

## Goal / Concept of FugakuNEXT

- **5 to 10x higher effective performance** than Fugaku for existing HPC applications
- **Effective performance of 50 EFLOPS or higher for AI processing** with **Zetta-FLOPS-scale** peak performance in mind
- **Target 10-100x improvement for apps** by combining simulation and AI

- **"Application First" philosophy** that prioritizes application performance – **"co-design"**
- **Heterogeneous architecture** combining power-efficient CPUs and bandwidth-oriented GPUs to achieve the goal under power constraint
- CPUs that **can utilize the assets such as application software** developed for "Fugaku"

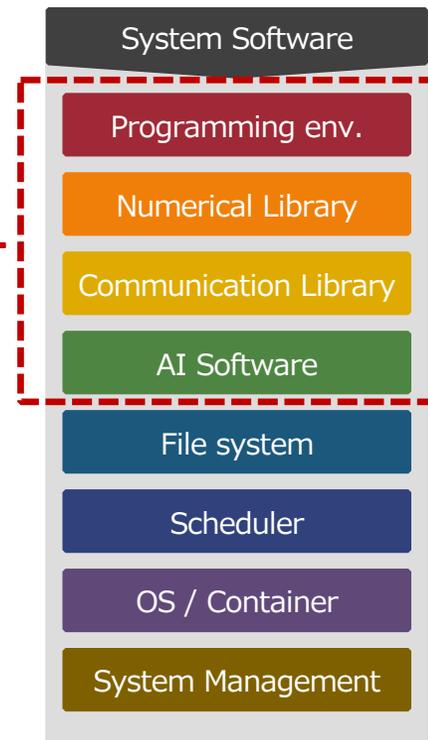
Items	CPUs	GPUs	Fugaku	Improvements
Total node count	>= 3,400		158,976	
Theoretical peak perf. of FP64 vector	>= 48 PFLOPS	>= 3.0 EFLOPS	488-537 PFLOPS	x5.7~
Theoretical peak perf. of FP16/BF16 matrix	>= 1.5 EFLOPS	>= 150 EFLOPS	1.95-2.15 EFLOPS	x70.5~
Theoretical peak perf. of FP8 matrix	>= 3.0 EFLOPS	>= 300 EFLOP	—	
(Sparsity considered FP8 matrix)	—	>= 600 EFLOPS	—	
Main memory capacity	>= 10 PiB	>= 10 PiB	4.85 PiB	x4.1~
Main memory bandwidth	>= 7 PB/s	>= 800 PB/s	163 PB/s	x4.9~
System power consumption	<= 40 MW (compute node and storage)		about 30 MW	

## Goals:

- **System Software Development Aiming to Achieve Target Performance**
  - Develop system software to **maximize hardware performance** in order to achieve the target performance of the FugakuNEXT project (HPC applications: 5~10x effective performance of Fugaku, AI processing: 50 EFLOPS in effective performance)
- **Continuous and Sustainable Development for FugakuNEXT and beyond**
  - Release the developed software as **Open-Source Software (OSS)** and develop the eco-system through international OSS communities
  - All the OSS software is supposed to be integrated into the upstream of OSS
  - The final goal is **a set of developed software is used** in domestic and international supercomputing centers and clouds
  - When commercial software is adopted and deployed, they must be carefully selected based on the sustainability

## System Software WGs

- **Programming Environment WG** (WGL: Hitoshi Murai, WGS: Miwako Tsuji)
  - Support the majority of existing HPC applications in Fugaku and programs implemented in the following languages and framework environments (Languages: C/C++, Fortran, Python and others /Frameworks: OpenACC, OpenMP)
  - Provide better development tool environment (e.g., debugger, profiler) than Fugaku
  - Support frameworks to efficiently work across different CPUs/GPUs with high compatibility (Kokkos, StdPar, SYCL, etc.)
- **Numerical Computing Library and Middleware WG** (WGL: Toshiyuki Imamura, WGS: Akiyoshi Kuroda)
  - Develop numerical libraries that efficiently use hardware and accelerate applications (incl. OZAKI-scheme)
  - Middleware supported by Fugaku should also be supported by FugakuNEXT depending on the needs/seeds
- **Communication Library WG** (WGL: Yoshifumi Nakamura, WGS: Shinji Sumimoto)
  - Support communication libraries deployed on Fugaku and also following libraries (MPI, UCX/UCC, etc.)
  - Develop system software to **efficiently use hierarchical networks with scale-up and scale-out** networks so that existing frameworks and applications can efficiently run in FugakuNEXT
- **AI Software WG** (WGL: Wahib Mohamed, WGS: Akiyoshi Jinguashi, Masaru Nagaso)
  - Support individual AI software which is widely used in 2029-30 and tune to efficiently run on FugakuNEXT
  - Develop **AI and data processing pipelines/workflows** that seamlessly integrate individual AI tools to automate a wide range of AI-for-Science research tasks.



# FugakuNEXT Application Development Area

- **Application Development and Development Support**

- GPU porting, tuning
- Use of low precision arithmetic
- Development of use case of AI in application algorithm
- Utilizing other advanced technologies
- (semi-)automatic porting / tuning

- **Benchmark**

- Benchmark on present machines
- Computing requirement → performance projection of FugakuNEXT

- **Co-design**

- Collaborating with Architecture and System Software
- Extracting information from applications for co-design

- **Flamework**

- CI/CD flamework build and operation

- **Relation with Application Community for Ecosystem**

- Community for current leading HPC applications
  - HPC performance building, tuning
  - AI use case for HPC
- Expanding reach of potential users of FugakuNEXT – HPC → AI, non-academia
- Close collaboration with HPC・AI開発支援拠点

