FS3.0: Overview (Lead organization: RIKEN R-CCS)

Background

- The development of the next flagship system that will succeed the supercomputer Fugaku (hereinafter referred to as “FugakuNEXT”) and the enhancement of the HPCI (High Performance Computing Infrastructure for Innovation), which integrates diverse computational resources across Japan, have been progressing in parallel.
- In the FugakuNEXT era, the rapid advancement of generative AI has significantly increased the importance of AI for Science, and the computational resources required to support such workloads are becoming increasingly diverse.
- To meet these emerging requirements, the FS3.0 project conducts feasibility studies on computer architecture, system software, and next-generation application development, and proposes development plans for HPCI systems in the FugakuNEXT era and beyond. (FS1.0: Feasibility Study for Fugaku; FS2.0: Feasibility Study for FugakuNEXT; FS3.0: Feasibility Study for Systems Beyond FugakuNEXT.)

Feasibility Study Areas

- **Architecture Feasibility Study**
  - In collaboration with participating vendors, we investigate advanced technologies such as next-generation semiconductor technologies, 2.5D/3D/3.5D stacked packaging, and chip-to-chip optical interconnects, and explore a range of possible system configurations for HPCI systems in the FugakuNEXT era.
- **System Software Feasibility Study**
  - Building on the software ecosystem inherited from Fugaku and FugakuNEXT, we will define Japan’s system-software development positioning within the international community and examine system software that supports AI for Science in addition to traditional large-scale simulations.
- **Application Feasibility Study**
  - We investigate trends in innovative application use and development enabled by generative AI—such as automated computational science experiments and automatic code generation—across diverse scientific domains, and clarify the system requirements needed to enable such next-generation HPC applications.

Applications
enabled by AI and HPC computational resources

System software
bridging applications and next-generation architecture

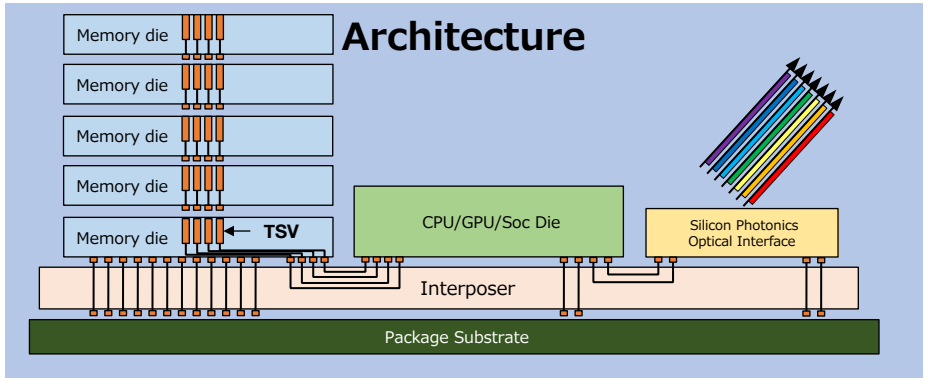

FS3.0: Project Organization

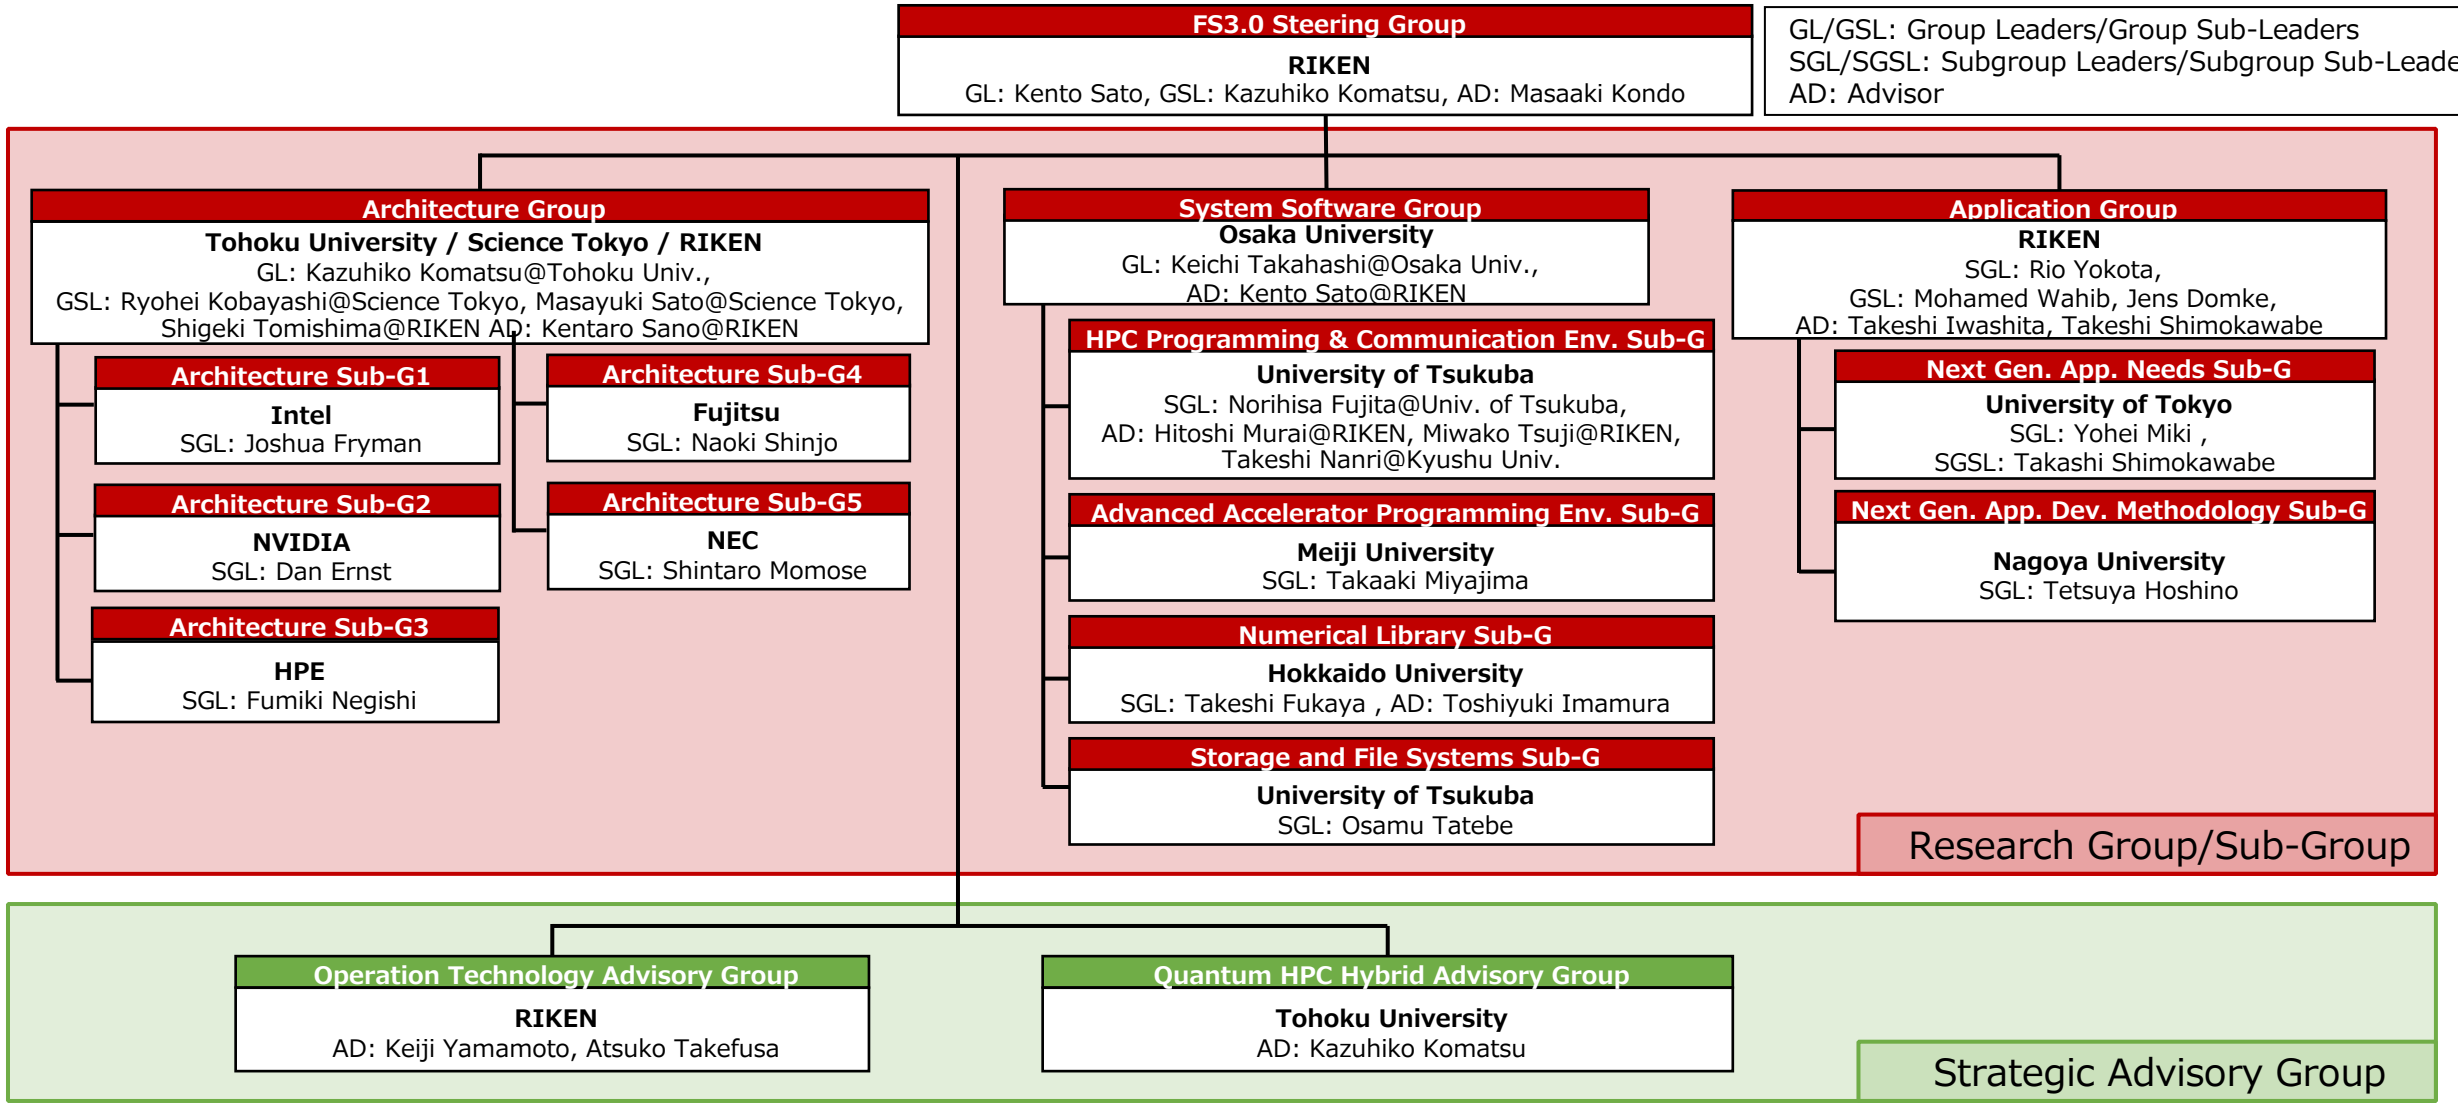

V1.2

Stratigic Project Organization

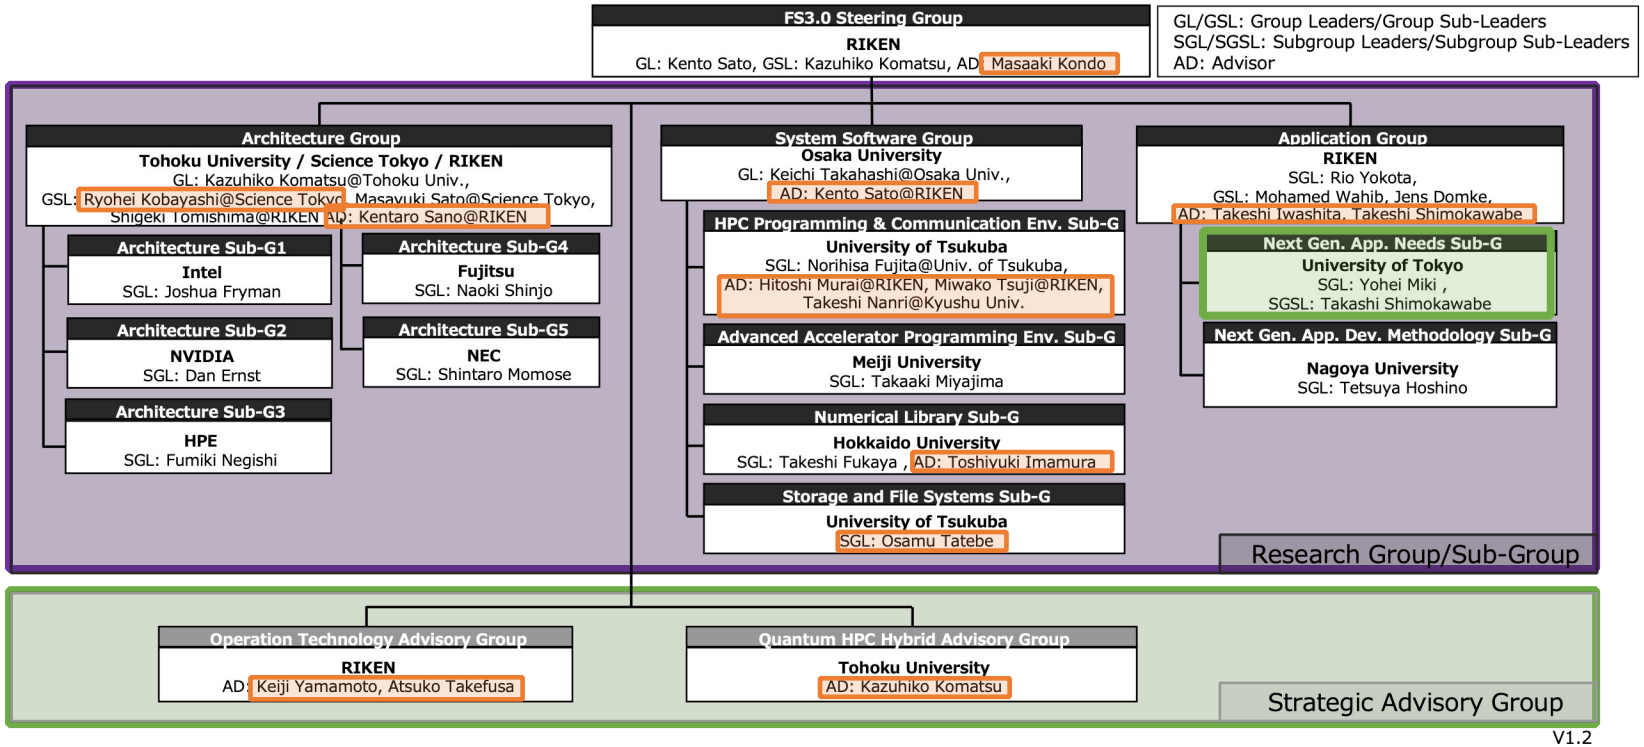

V1.2

- Point 1 : **Participation and collaboration of major vendors and a wide range of HPCI institutions (9 HPCI institutions)**
- Point 2 : **Continuity of expertise from FS2.0 to FS3.0 and collaboration with the FutakuNEXT development project**
  - **FS2.0→FS3.0 Continuity:** FS2.0 leaders participate as advisors in groups and subgroups for continuity from FS2.0 to FS3.0
  - **Collaboration w/ FugakuNEXT:** The participation of key members of the FugakuNEXT development project for consistent FS between FugakuNEXTGand FS3.0
- Point 3 : **Collaboration with teams and HAIRDESC (Advanced HPC-AI Research and Development Support Center)**
  - **Collaboration with other teams:** Establishment of a strategic advisory group will facilitate collaboration with other research teams and ensure consistency.
  - **HAIRDESC**: The next-gen. application needs subgroup is responsible for collaboration with HAIRDESC and for collecting needs and benchmark codes, and feeding them back to this FS3.0 team.

Architecture Research Group

Objectives

- Through a survey of technology and research trends related to architecture, the group will investigate and study the system requirements that the next-generation HPCI system suite should have for the development of Japan's computational infrastructure.

Feasibility Study Areas

- **Architecture Group (Tohoku University, Institute of Science Tokyo and RIKEN):** To conduct a comprehensive feasibility study on the next-gen. systems that meets the needs in HPC, the architecture group collaborates with venders and lead the investigation for elements of architectural technologies that are required for various applications such as AI for Science and large-scale scientific computation.
- **Architecture subgroup 1 (Intel Corporation):** Investigate and study various architectures based on current/future roadmaps for heterogeneous node configurations with Xeon CPUs and GPUs; evaluate chiplet configurations using advanced packaging technologies such as EMIB and Foveros; Study of CXL; Analyze cross-architecture programming environment based on oneAPI
- **Architecture subgroup 2 (NVIDIA Corporation):** Provide our technical information on future architecture, system configuration, related software, etc. to maximize application and power performance, investigate elemental technologies, and study the possibility of utilizing domestic technologies
- **Architecture subgroup 3 (HPE)**: Study on system architecture (processor, memory, storage, I/O, etc.), OS, compilers, file systems, libraries, frameworks, development languages, and other related technical fields.
- **Architecture subgroup 4 (Fujitsu Limited):** Study future available semiconductor manufacturing technologies (chiplet packaging technology, memory device technology, optical interconnect technology), study the roadmap of FUJITSU-MONAKA-X, and research on the target application domains and the system software for the target architectures and application domains
- **Architecture subgroup 5 (NEC Corporation):** Study various architectural options and roadmaps, including heterogeneous node configurations with vector, scalar, accelerator, etc with the applicability to next-generation HPC, AI, quantum, etc. based on the characteristics of vector architecture (Trade-off among compute density, power saving, etc.). Study the direction of system architecture and system software development for efficient operation and utilization. Also, study key applications and benchmark performance targeted by HPCI systems in the FugakuNEXT era.

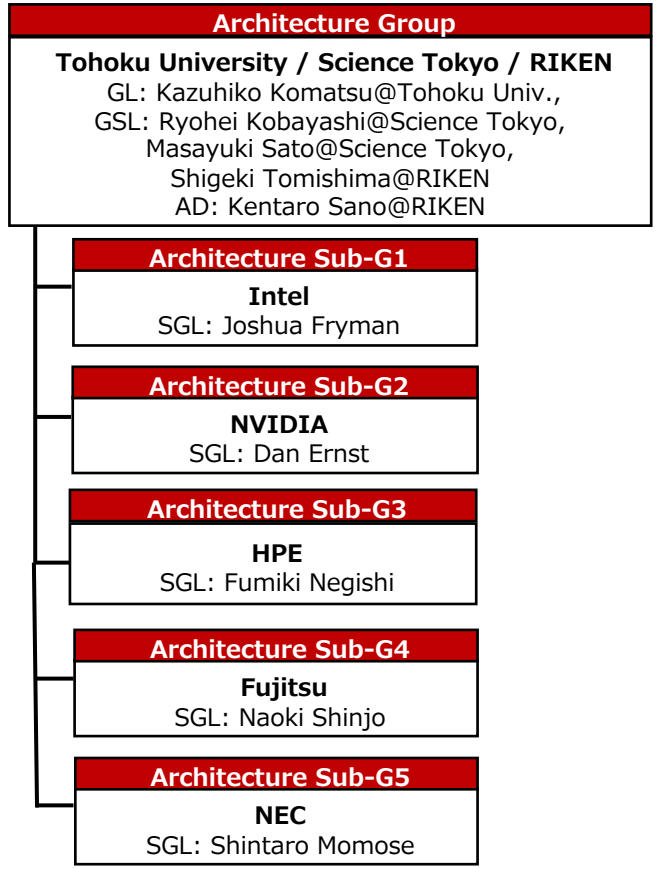

System Software Research Group

Objectives

- We will investigate the system-software ecosystem that next-generation HPCI systems should provide and examine succession strategies for sustaining the software stack of Fugaku, FugakuNEXT, and systems beyond FugakuNEXT (ensuring sustainable development). In addition, based on technology trends in architectures and applications, we will study the software ecosystem required for the integrated operation of future HPCI systems across sites (enabling cross-cutting deployment).

Feasibility Study Areas

- **System Software Research Group (Osaka University):** Survey latest system-software trends and, together with the Archi and App groups, define requirements for next-generation hardware, while considering reuse and modernization of existing software assets.
- **HPC Programming and Communication Env. Subgroup (University of Tsukuba):** Survey programming and communication environments for future HPCI systems, focusing on accelerator-oriented methods and needs from HAIRDESC, including emerging environments such as Kokkos and Julia.
- **Advanced Accelerator Programming Env. Subgroup (Meiji University):** Investigate applicability of advanced AI chip (e.g., Cerebras, SambaNova) to HPC, comparing programming environments and identifying performance characteristics using representative HPC applications.
- **Numerical Library Subgroup (Hokkaido University):** Provide guidelines for future numerical library development by reviewing next-generation architectures, application needs, and new mathematical techniques, including trends in mixed-precision and randomized algorithms.
- **Storage and File System Subgroup (University of Tsukuba):** Study storage and file-system technologies for HPCI systems in the FutureNEXT era, identifying trends, issues, and requirements, and exploring methods to meet the requirements

System Software Group
Osaka University
GL: Keichi Takahashi@Osaka Univ.,
AD: Kento Sato@RIKEN

- **HPC Programming & Communication Env. Sub-G**
  University of Tsukuba
  SGL: Norihisa Fujita@Univ. of Tsukuba,
  AD: Hitoshi Murai@RIKEN, Miwako Tsuji@RIKEN,
  Takeshi Nanri@Kyushu Univ.
- **Advanced Accelerator Programming Env. Sub-G**
  Meiji University
  SGL: Takaaki Miyajima
- **Numerical Library Sub-G**
  Hokkaido University
  SGL: Takeshi Fukaya , AD: Toshiyuki Imamura
- **Storage and File Systems Sub-G**
  University of Tsukuba
  SGL: Osamu Tatebe

Application Research Group

Objectives

- We investigate trends in innovative applications that will drive future HPC utilization, aiming to clarify—from the application perspective—the functional and performance requirements needed for next-generation computing infrastructure. Looking toward 2030, we work with academic and industrial communities to define the vision for next-generation HPC applications.

Feasibility Study Areas

- **Application Research Group (RIKEN):** The group focuses on four pillars: AI for Science, AI agents, automatic code generation, and automated computational experiments. We study how advanced AI technologies can be integrated with scientific computing, including autonomous agent-based AI, automatic GPU-porting and code generation, and AI-driven experiment-automation workflows.

- **Next-Generation Application Needs Subgroup (University of Tokyo):** The subgroup coordinates with HAIRDESC and the FugakuNEXT developers to identify user needs for next-generation HPC applications. These findings inform system requirement definitions for architecture and system software groups. To support broad adoption, the subgroup develops representative benchmark codes usable for performance evaluation and as training material for application developers and users, including NIS users. We analyze diverse programming approaches—from CUDA and OpenMP target to abstraction layers like SYCL and Kokkos—and explore generative-AI-assisted development in collaboration with the Next-gen. App. Dev. Methodology subgroup.

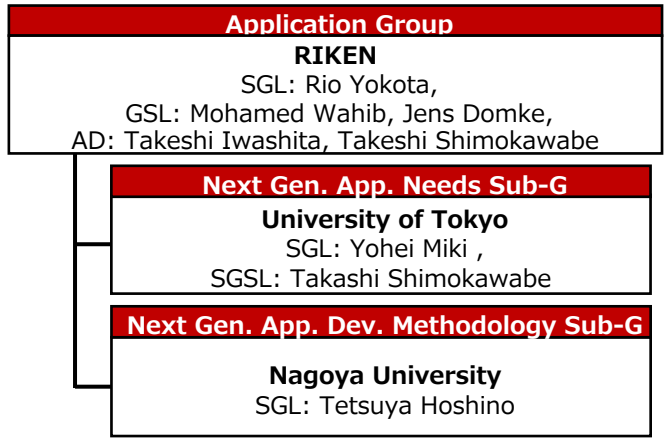

- **Next-Generation Application Development Methodology Subgroup (Nagoya University):** We examine LLM-based iterative development workflows, where large language models propose, revise, and test code. Through this, we extract system requirements needed to support efficient LLM-assisted development on future HPCI systems. The subgroup evaluates both commercial LLMs (e.g., Claude Code) and local deployment of domestic LLMs (e.g., Swallow) for confidential applications. Evaluations use benchmark suites from the Next Gen. App. Needs Subgroup and key applications from HAIRDESC. We also assess multi-agent LLM workflows, clarifying the functions required for next-generation HPCI systems to effectively support such advanced development environments.

Direction of architectural research

- **Study of Architectures for Advancing Next-Generation HPC and AI Research**
- **Node Architecture**
  - **Heterogeneous systems**: Explore heterogeneous node designs that combine domestic and international CPUs, GPUs, FPGAs, and AI accelerators, in addition to the FugakuNEXT architecture (domestic Arm CPU + GPU).
  - **Next-generation memory hierarchies**: Study configurations featuring large-capacity, high-bandwidth memory using HBM and CXL, along with hierarchical memory-access management.
  - **Adoption of 2.5D/3D/3.5D packaging and chiplet technologies**: Investigate the potential for high-density integration to alleviate data-movement bottlenecks and enable low-power operation (Collaborators: Tadahiro Kuroda, Atsutake Kosuge).
  - **New computing principles**: Assess the applicability of high-bandwidth, low-latency, and energy-efficient computing and communication technologies such as optical computing (e.g., silicon photonics) (Collaborator: Jun Shiomi).
- **System Architecture**
  - **Large-scale HPC/AI systems**: Supporting HPC simulations, AI processing, and their convergence at the scale of several thousand to ten thousand nodes, emphasizing high performance and efficiency through tightly coupled and homogeneous system architectures.
  - **Small- to medium-scale AI-utilization systems**: Consider flexible cluster or distributed environments capable of deploying large numbers of AI agents or distributed inference workloads, enabling diverse loosely coupled and heterogeneous system configurations, including cloud-like operational models.

Node Architecture

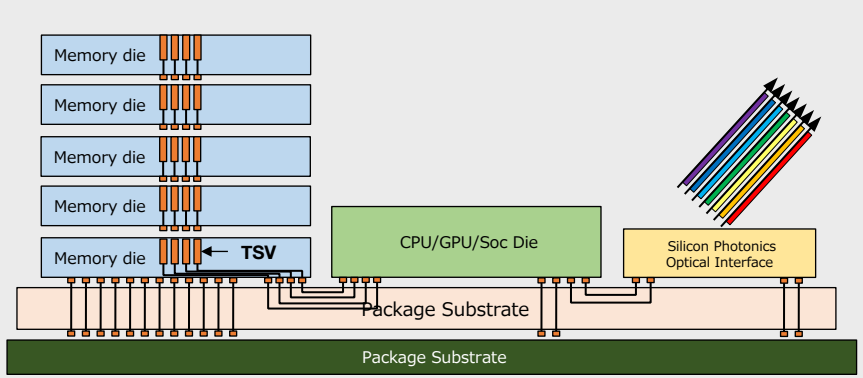

Example of Processing Element Configuration

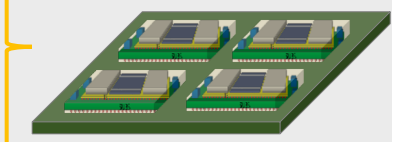

Example of Mounting
on Substrate

System Architecture

For large-scale HPC and AI experiments
Tightly coupled and homogeneous system architecture

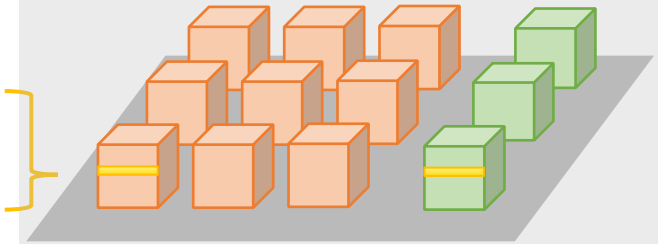

For AI utilization
Loosely coupled and heterogeneous system

Planned Deliverables: Continuous FS Platform for HPCI

- **(1) HPCI-RB: Specification Development Reference (HPCI Reference Blueprint)**
  → Reference describing system requirements for HPCI in the of FugakuNEXT era (multiple drafts as needed)
  - **Technical Specification Reference (TS ref.)**: Describes guidelines for architecture and system software requirements
  - **Performance Evaluation Criteria Reference (PEC ref.)**: Describes guidelines for effective performance targets to be achieved by the application.
    - Develops methods to extrapolate future performance based on the technical specification guidelines and performance measurements from current systems
  - Using this blueprint—while taking the circumstances of each HPCI institution into account—we will support the procurement and deployment of HPCI systems appropriate for the FugakuNEXT era.
- **(2) HPCI-CB: HPCI Continuous Benchmarking**
  → An environment that enables continuous performance evaluation of applications deemed important by each HPCI institution.
  - The benchmarking platform currently being developed under FugakuNEXT will be extended to other HPCI sites, enabling performance assessment not only of the key applications targeted in FugakuNEXT but also of applications considered important across different HPCI institutions.
- **(3) HPCI-CFSP: Establishment of a continuous research platform (HPCI Continuous FS Platform)**
  → A mechanism to ensure long-term preservation and utilization of FS3.0 outcomes, including results and know-how from items (1) and (2).
  - (1) Consider methods for managing and maintaining HPCI-RB, including technical information covered by NDAs.
  - (2) Develop long-term operational policies for HPCI-CB.