# Accelerating GPU-acceleration on Supercomputers in Japan and Application Development Support for Next Generation

**Taisuke Boku**
Director, HAIRDESC
(Advanced HPC-AI Research & Development Support Center)
/ Professor, CCS at University of Tsukuba



*https://www.hairdesc.jp*

# NFS (National Flagship System) in Japan so far

| System | Year | Processor | Architecture | Peak Performance |
|---|---|---|---|---|
| (Earth Simulator) | 2002 | NEC Vector (SX-5 base) | Vector | 40TF |
| K | 2012 | Fujitsu SPARC64 VIIIfx | Multicore CPU (8c) | 10PF |
| Fugaku | 2020 | Fujitsu A64FX | Manycore CPU (48c) | 488PF |
| Fugaku-NEXT | 2030? | Fujitsu MONAKA-X + NVIDIA GPU | CPU+GPU (multi-socket/node) | ??? (for HPC + AI) |

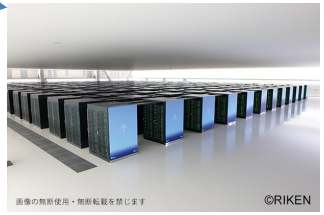


- We need a rapid code development ready for FugakuNEXT and currently running GPU systems in **NIS (National Infrastructure Systems)**
  - Each site has their own education course and materials
  - Users need a “standard” course for efficient learning of GPU computing, advanced performance tuning and introducing latest technology (ex. GH200 or MI300A)
- We need a central organization to coordinate these activities to catch up the advanced GPU computing technology in the world for HPC/AI/AI4S
- NIS (National Infrastructure Systems) in 9 national universities + α : several systems are fully equipped with GPU and leading GPU-computing in Japan (ex: TSUBAME4.0, Miyabi-G, Genkai, ...)

# Expected Timeline of FugakuNEXT R&D

(Slide courtesy by M. Kondo@RIKEN)

## FugakuNEXT development and deployment schedule

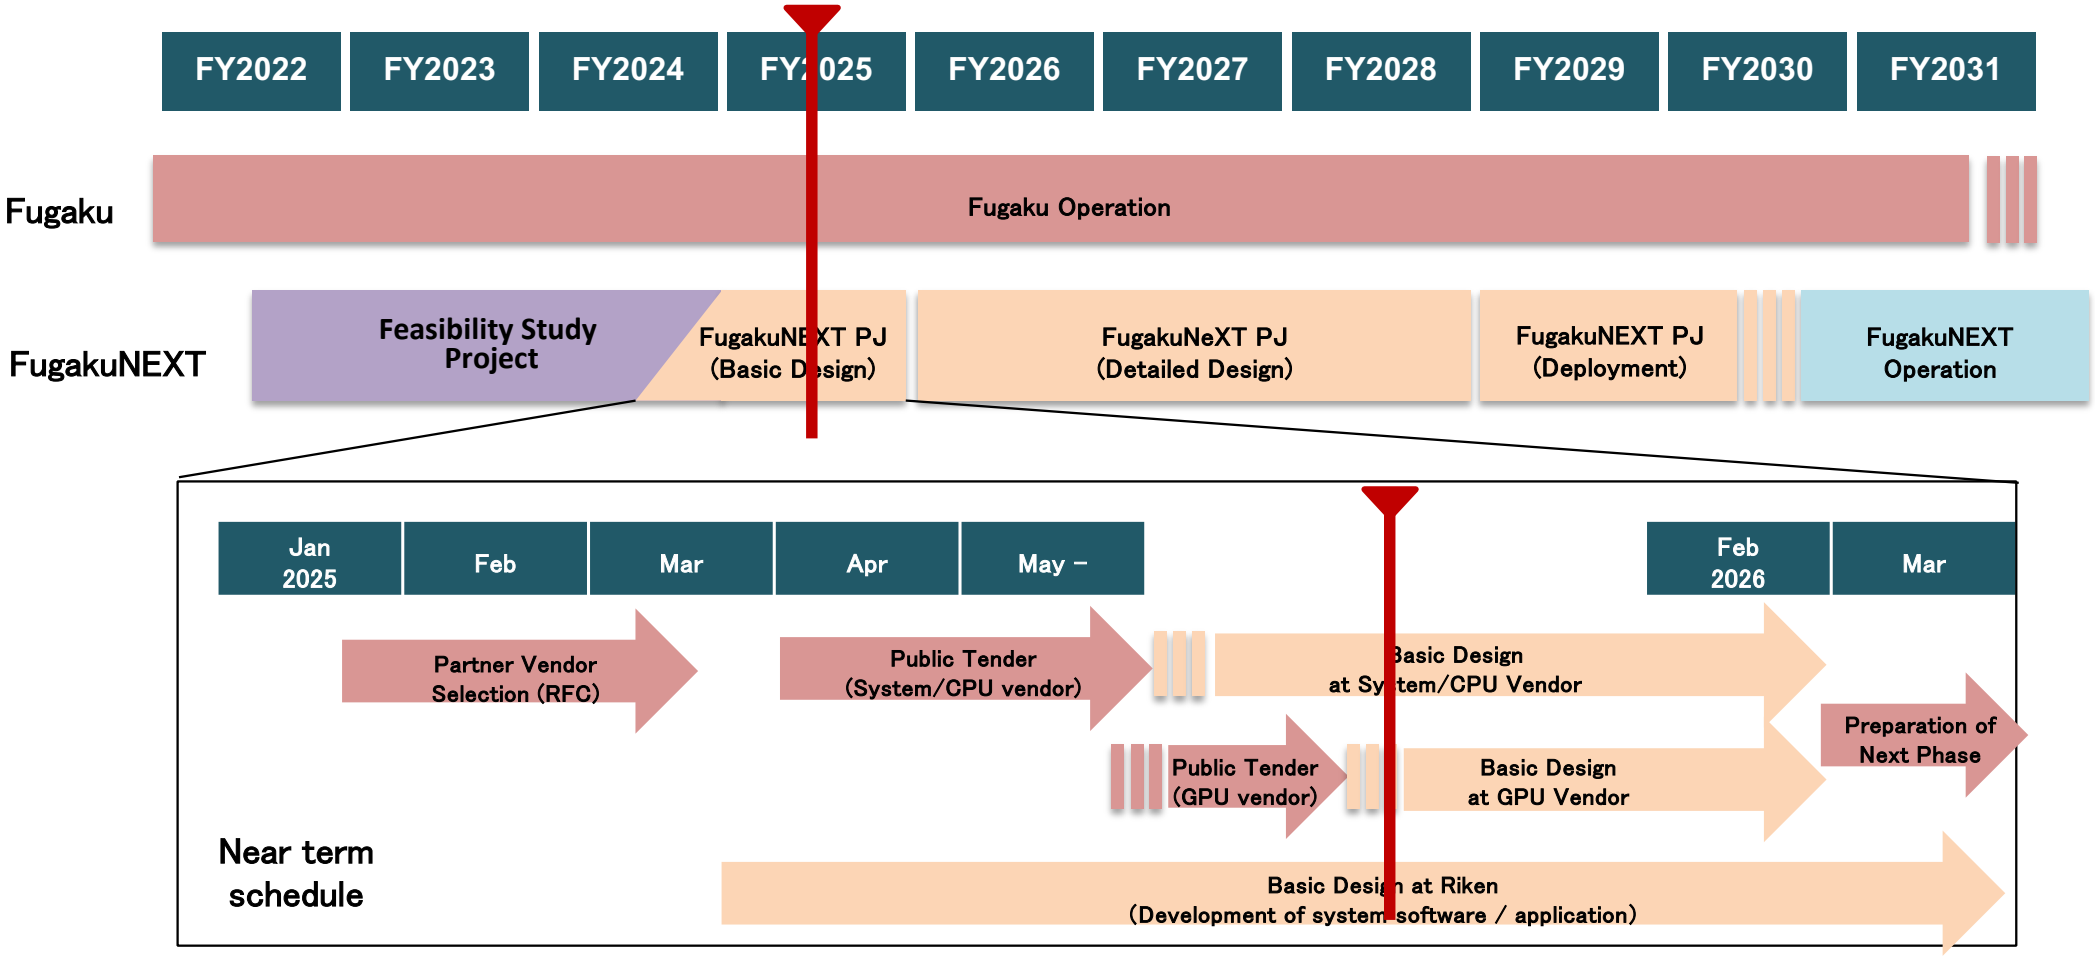

# Aim of the Next-Generation Computing Infrastructure, Built on the Legacy of “Fugaku”

(Slide courtesy by M. Kondo@RIKEN)

- **The value and national significance of the flagship system**
  - Ensuring global leadership in computational science through a Zetta-scale HPC environment
  - Leveraging economies of scale for deployment, operation, and community engagement
  - Japan’s most powerful research platform for AI development, where computation is decisive

**Zetta-scale computing resources to drive value creation by expanding computational frontiers**

- **Development of “FugakuNEXT” to enhance industrial competitiveness**
  - Built on the ARM architecture of “Fugaku” while pursuing further advancements in system design
  - Pioneering computation-driven problem solving, incl. “AI for Science” and quantum comp. integration
  - Introducing GPUs as accelerators under a Japan–U.S. joint development framework, modernizing Japan’s applications with active use of AI to accelerate social implementation of research outcome
  - Establishing a “Made with Japan” devel. framework in alignment with Japan’s semiconductor strategy

**Promoting the advancement and succession of competitive domestic technologies, securing their strategic indispensability in IT industry, and driving their expansion into global markets**


4

## Aim of the Next-Generation Computing Infrastructure, Built on the Legacy of "Fugaku"

(Slide courtesy by M. Kondo@RIKEN)

- **The value and national significance of the flagship system**

NFS so far: CPU-only solution for "K" and "Fugaku"

⇩

GPU + CPU Heterogeneous System for "FugakuNEXT"
toward successful HPC/AI/AI4S

No appropriate GPU device in Japan

⇩

Made in Japan ⇨ Made with Japan

**Promoting the advancement and succession of competitive domestic technologies, securing their strategic indispensability in IT industry, and driving their expansion into global markets**


5

# Overview of FugakuNEXT System Architecture

(Slide courtesy by M. Kondo@RIKEN)

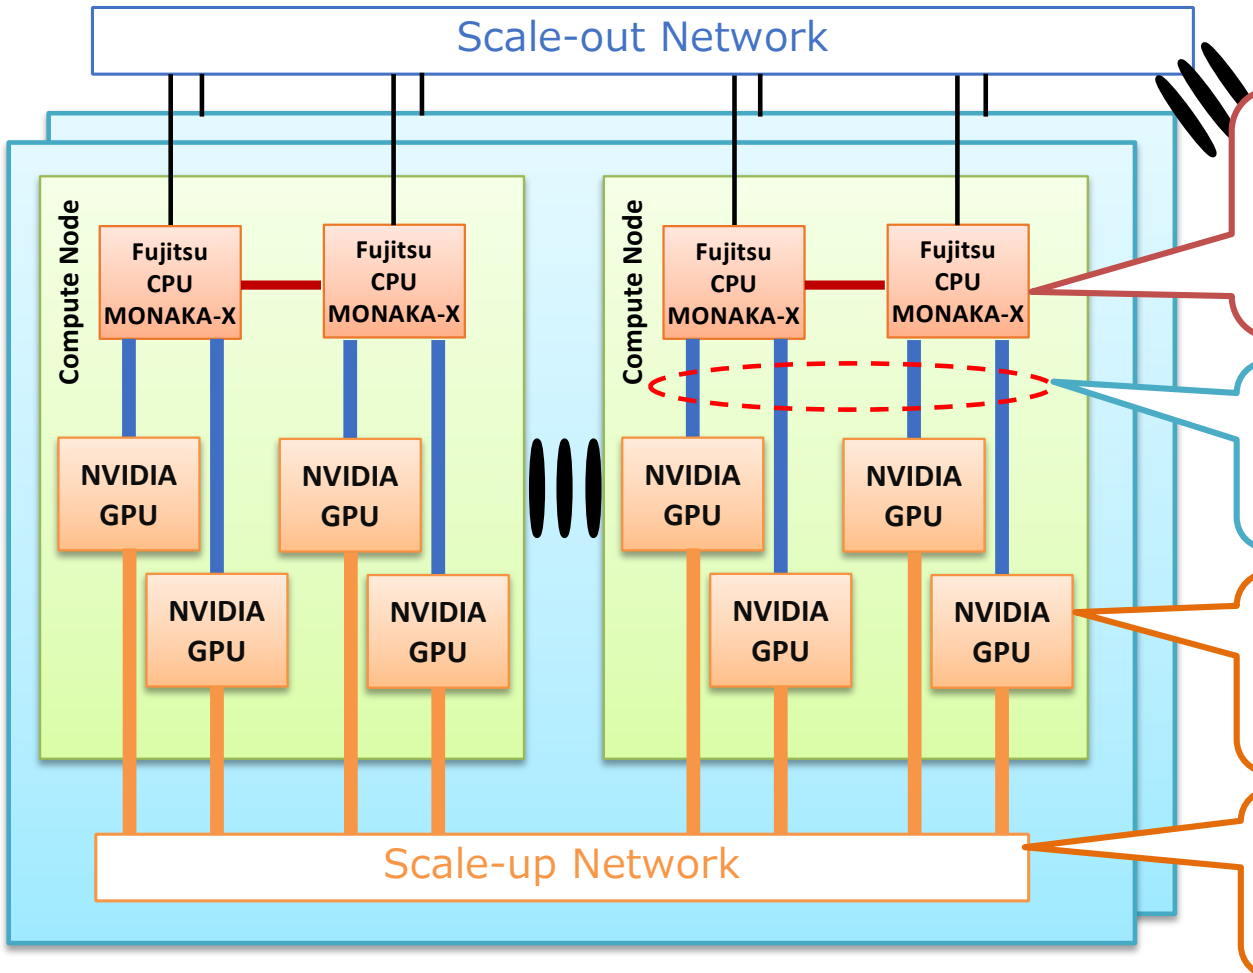


Responsibility

- **Fujitsu MONAKA-X CPU** — Fujitsu, Riken
  - ARM-based manycore CPU
  - Considering the integration of an NPU based on ARM SME
- **CPU-GPU connection** — Fujitsu, NVIDIA, Riken
  - Considering adoption of a high-bandwidth, coherent CPU–GPU interconnect
- **NVIDIA GPU** — NVIDIA, Riken
  - High-perf. vector and matrix units
  - Considering advanced 3D staking memory
- **Scale-up Network** — NVIDIA, Riken
  - multi-GPU interconnect via NVLink, either within a node or across nodes

**Detailed system architecture configuration will be determined in basic design phase**


6

# Architectural Specifications

(Slide courtesy by M. Kondo@RIKEN)

## High BW & heterogeneous node arch and whole system overview

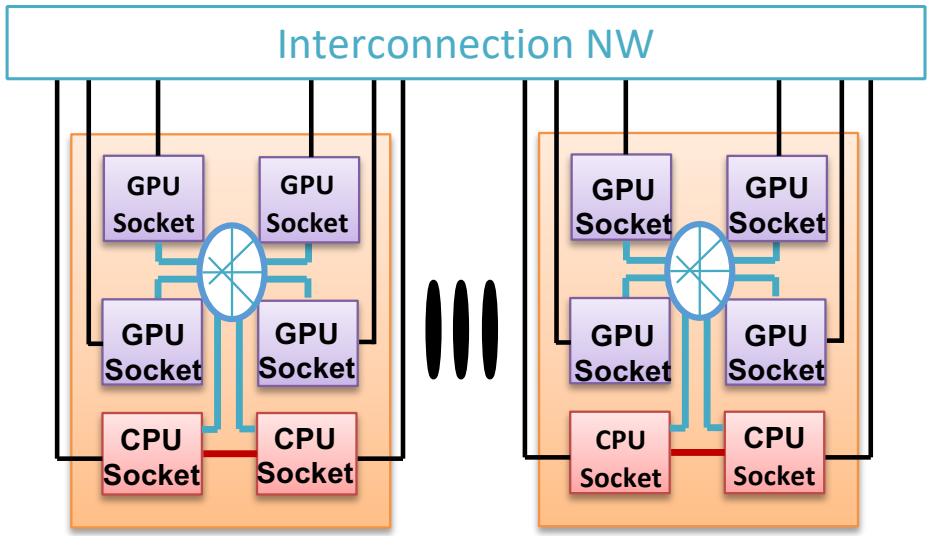


## Performance target of the entire system

| | CPU | GPU |
|---|---|---|
| Total Num. of Nodes | >= 3400 Nodes | |
| FP64 Vector FLOPS | >= 48PFLOPS | >= 2.6EFLOPS |
| FP16/BF16 AI FLOPS | >= 1.5EFLOPS | >= 150EFLOPS |
| FP8 AI FLOPS | >= 3.0ELOPS | >= 300EFLOPS |
| FP8 AI FLOPS (w/ sparsity) | – | >= 600EFLOPS |
| Memory Size | >= 10PiB | >= 10PiB |
| Memory Bandwidth | >= 7PB/s | >= 800PB/s |
| Total power consumption | < 40MW (compute and storage) | |

**(Assigned) Target: More than 5-10x effective performance gain in HPC apps, more than 50EFLOPS effective AI performance (needs Zetta-scale low-precision perf.)**

**Our Goal: ～100x apps performance improvement by combining simulation and AI**

# Governmental program for supporting GPU application development

- For urgent requirement of GPU-ready application on National Flagship System (NFS) and National Infrastructure Systems (NIS), MEXT decided to launch a new supporting organization to shift CPU apps to GPU
- Call for proposal on “Next Generations HPC/AI Development Support Organization” Project, Jul. 2025 by MEXT
- Sep. 2025: a proposal from Research Organization for Information on Science and Technology (RIST) was awarded

**Advanced HPC-AI Research and Development Support Center (HAIRDESC)**
**(Director: Taisuke Boku)**

- New center was launched on Oct. 1st 2025, at Kobe Japan (close to RIKEN R-CCS where Fugaku is located)

# HAIRDESC organization

Head quarter (RIST)
- HAIRDESC: Advanced HPC-AI Research & Development Support Center
(Director: Taisuke Boku)

Supporting with each other to share the technological information, human resources, GPU education materials, operation of hackathon/symposium

Core Supporting Organizations
- U. Tsukuba (Center for Computational Sciences)
  PI: Akira Nukada
- U. Tokyo (Information Technology Center)
  PI: Takashi Shimokawabe
- Inst. of Science Tokyo (Supercomputing Research Center)
  PI: Rio Yokota

Collaborating Organizations
- National Universities
  Supercomputer Center:
  Hokkaido, Tohoku, Nagoya, Kyoto, Osaka and Kyushu
- RIKEN R-CCS
- GPU vendors: AMD and NVIDIA

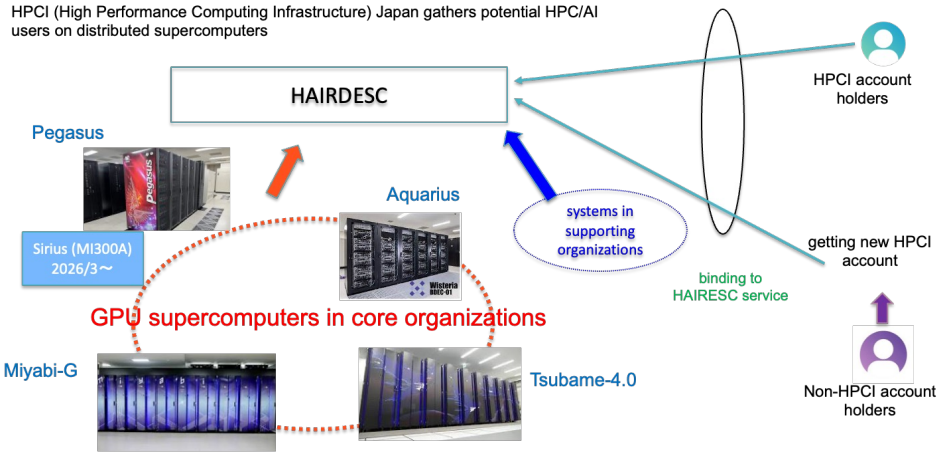

# HAIRDESC (≠ HAIR DESK)

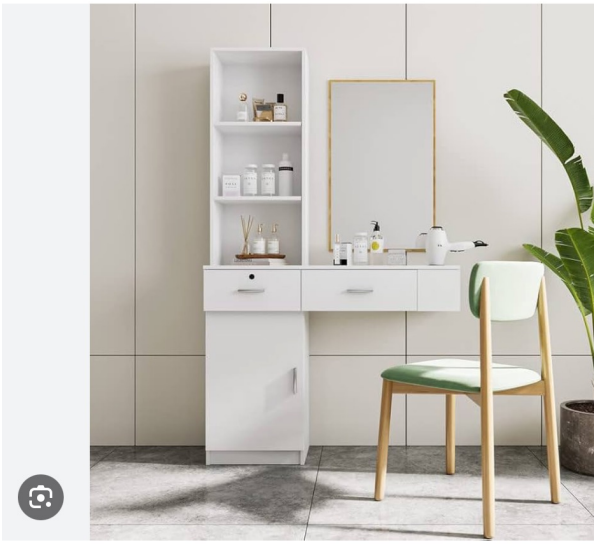

Amazon.com: Modern Hair Desk with Multi-Layer Storage ...

HAIRDESC
Advanced HPC-AI R&D Support Center

Let‘s make up yourself for new meeting with GPU-computing!

# Where we are ?



Port Island (artificial) out of Kobe downtown
HAIRDESC and R-CCS is at next station or
just in 10min by walk
Also 10min by train (PortLiner) from Kobe Airport

# What we do ?

- Construction of standard GPU education program for all HPCI users in Japan (HPCI: HPC Infrastructure Japan)
  - GPU coding education materials: slides, documents, video, etc.
  - Constant events such as tutorials, hackathons, symposium in on-site/on-line/on-demand style with video archiving
- Networking on GPU computing community in Japan, crossing over application domains
  - Networking researchers to share same GPU technologies even in different applications and fields through symposium and on-line tools
- Porting support on CPU-oriented applications to GPU systems for governmental projects including "Grand Reach" and "AI for Science"
- Advanced research on state-of-the-art GPU technologies for efficient, effective and smart GPU computing
  - Core organization (three universities) proceed advanced GPU computing research
  - AI-supported GPU coding, heterogeneous-vendor GPU coding system, multi-precision computation, GPU-CPU unified module utilization, etc.
- International collaboration with any GPU computing activities, events, supercomputer centers
  - Networking with all international HPC/AI computing community: DOE-MEXT, HANAMI, E4S
  - Information exchange on SC, ISC, etc.
  - Invitation of researchers in seminar, symposium, etc.

# GPU education materials for standard learning course

- GPU tutorial materials everywhere, but not for easy to understand for novice users
- Most of GPU education materials explain how to use GPU, how to describe directives, …, but in our course:
  - how GPU works fundamentally and what's different from CPU, with synthetic codes
  - multiple levels and vendors GPU – common issues and different issues (OpenACC/OpenMP target/CUDA/HIP, …)
  - good code examples with efficient use of directives, with actual codes
  - bad code examples with inefficient use of directives, with actual BAD codes
  - easiest manner of porting OpenMP code to OpenACC/OpenMP target
- Most of governmental support project just focuses on "how to improve science and how to train researchers"
  ➔ "how to improve science and society and how to train GPU-ready people"
  - first-step trainees are researchers but also expanding it to industry people
  - increasing much of GPU-ready people (researchers and industry workers) because GPU is a common language for HPC/AI/AI4S in the world (in Japan, Society5.0 and digital twin)

# Standard GPU education course building

- NIS in Japan has been installed several large scale GPU-ready systems, especially at JCAHPC (U. Tsukuba + U. Tokyo) and ScienceTokyo, with Miyabi-G and TSUBAME4.0, and RIKEN R-CCS follows to introduce largest scale of GPU system (AI4S System) soon.
  (MITI lab. such as AIST introduces another strategic GPU system for AI.)
- Universities for NIS are performing tutorials and hackathasons for GPU coding guidance individually.
  ➔ Gathering all GPU education materials under support by GPU vendors (AMD and NVIDIA) to establish a standard GPU education course in Japan, with materials, tutorials, hackathons and symposiums.
- Recording all the video recorded through all these events to support all users for 365/24 on-demand service.

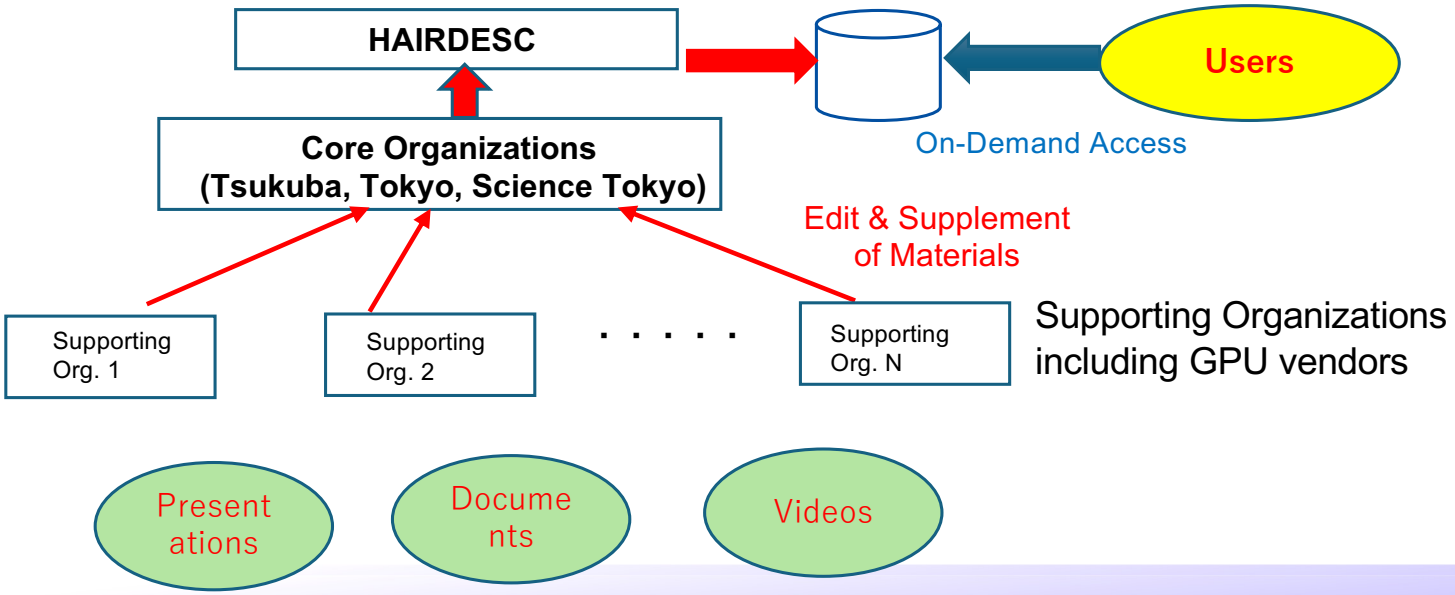


**Education Materials**
**- programming guide for various levels**
**- various languages and GPUs**
**- sample codes and explanation including good codes (and bad codes:-))**

# GPU resource sharing for education

HPCI (High Performance Computing Infrastructure) Japan gathers potential HPC/AI users on distributed supercomputers

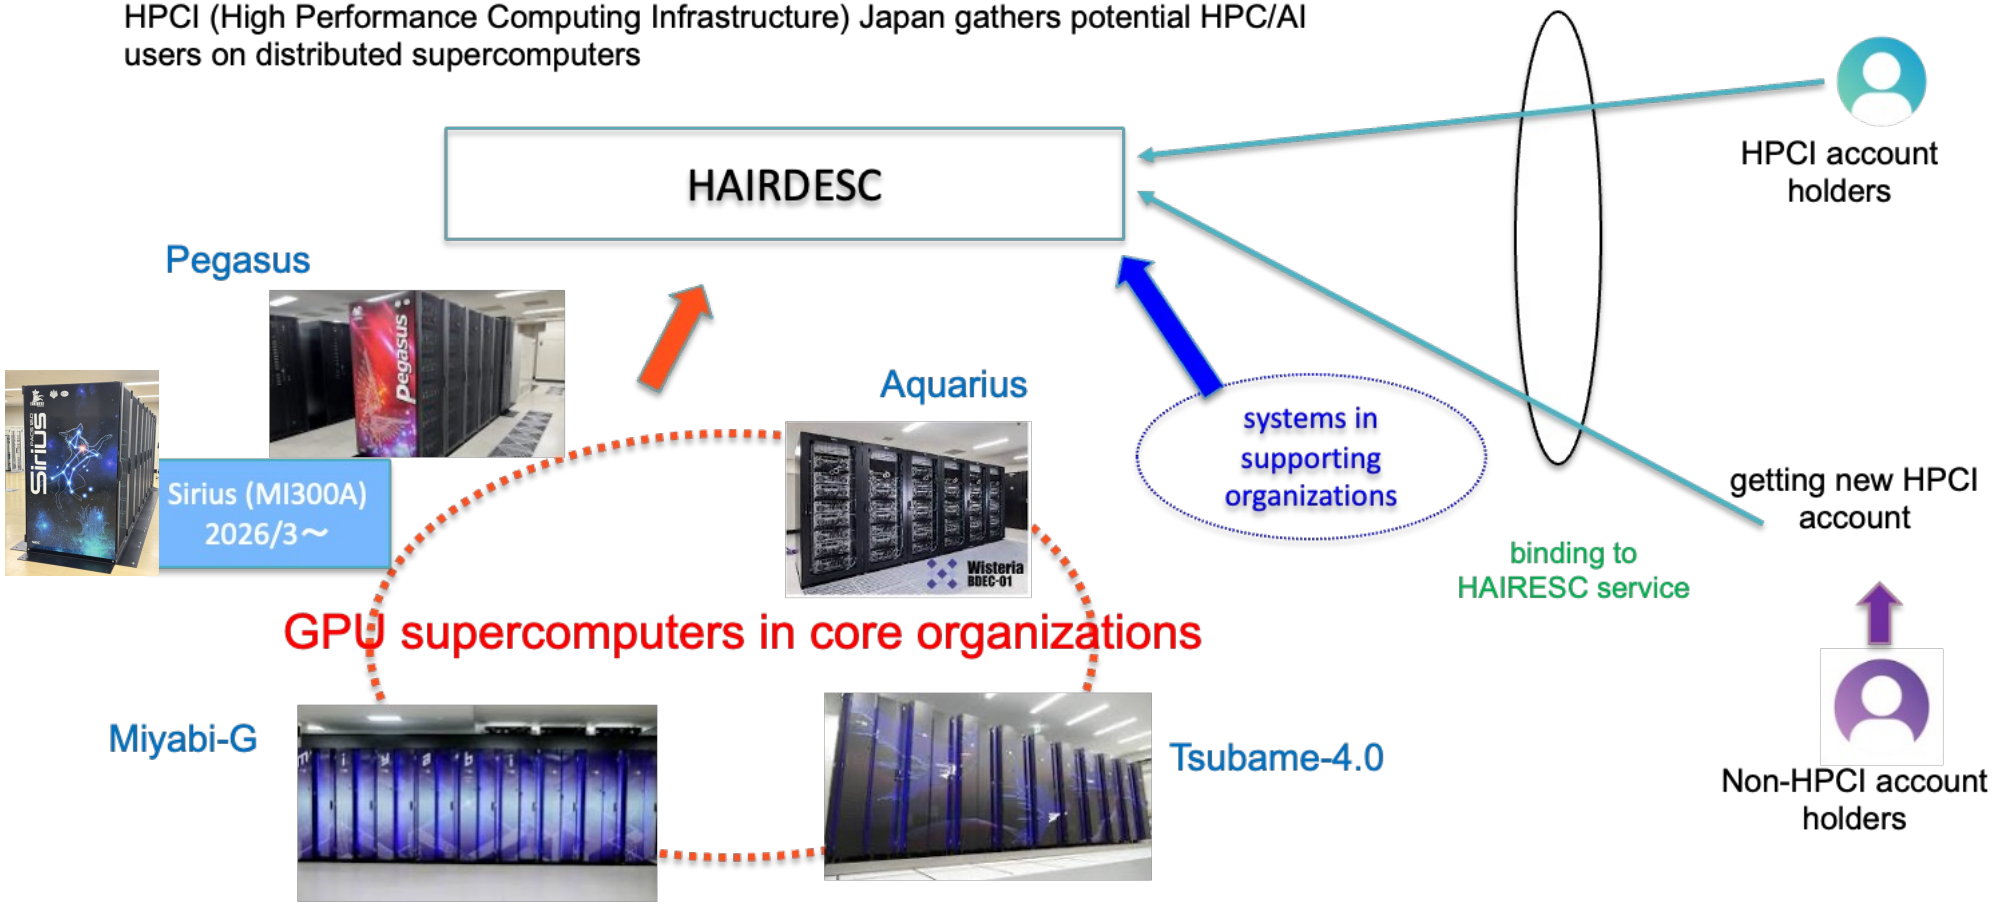

# Construction of networking community over various fields

- Various scientific fields such as physics, chemistry, material, bio, climate, etc. share the similar computing algorithms and tuning methods on GPU computing
- Gathering all the fields to share the knowledge and skills for GPU coding and performance optimization through constant tutorial and symposium
- Sharing methodologies: matrix calculation, stencil computing, multi-kernel sharing a GPU, GPU-GPU communication with MPI, etc.
- AI must be applied for every field, and AI researchers and HPC researchers must collaborate

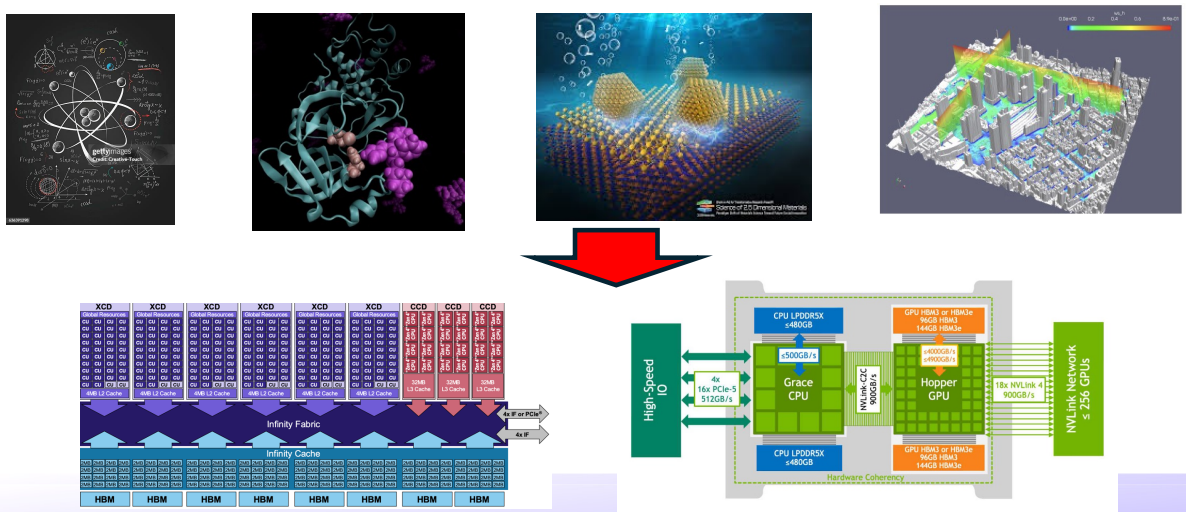

# Governmental program coming soon (1) “AI for Science”

## AI for Scienceによる科学研究革新プログラム

令和7年度補正予算額　370億円

文部科学省

**課題・取組の方向性**

- タンパク質の構造予測を行うAlphaFold（ノーベル賞）は研究にかかる時間とコストを劇的に削減するなど、**AIは、研究力の生産性の向上のみならず、科学研究の在り方そのものを変革**。国際的にAIの研究開発や利活用への投資が進む中、**自国でAI研究開発力を保持することは安全保障上極めて重要**。科学研究におけるAI利活用（AI for Science）において、米国・EU等は国家的な取組として、リソース（計算資源・研究資源・人材・データ等）を有効活用し、戦略的に推進。
- 我が国においては、世界最高水準の情報基盤を有するとともに、**ライフ・マテリアル等の重点分野において次のAI開発・利活用の要となる質の高い実験データを持つ等の強み**を有しており、これらのリソースを最大限活用し、**科学基盤モデル・AIエージェント開発、次世代AI駆動ラボシステム開発、これらの実装に向けた取組を進めることで、第７期科学技術・イノベーション基本計画で目指す研究力向上を牽引。**

**事業内容**　**事業実施期間**　～令和10年度

- 国のコミットメントの下で、我が国が有する**計算資源等のリソースを戦略的かつ機動的に配分しながら**、重点領域への集中投資により世界をリードすることを目指す**プロジェクト型（基金事業）**と、あらゆる分野における波及・振興及び先駆的な研究を目指す**チャレンジ型**を**両輪**とし、**AI for Science先進国**の地位を確立する。

① **プロジェクト型：**我が国の勝ち筋となる重点領域において、シミュレーションデータに加え、実験データの取得・活用による我が国発の最先端AI基盤モデル・AIエージェント開発、次世代AI駆動ラボシステム開発、これらの実装に向けた取組を一体的に推進。我が国の研究力を抜本的に強化するとともに、産学の協働により、研究開発投資を促進し、先駆的取組の早期実装・ビジネス化により科学研究を変革するイノベーションを創出。

② **チャレンジ型：**あらゆる分野の研究者がAIを活用して科学研究の高度化・加速化を図るため、計算資源の確保等の研究環境を整備し、アカデミア全体にAI for Scienceの波及・振興を促進し、意欲ある研究者による次の種や芽となる新たなアイデアへの挑戦への支援を行うとともに、我が国独自の競争優位を築く先駆的な研究を創出。

※上記の他、AI for Scienceに不可欠な計算資源の戦略的増強として、76億円を別途計上。

【事業スキーム】

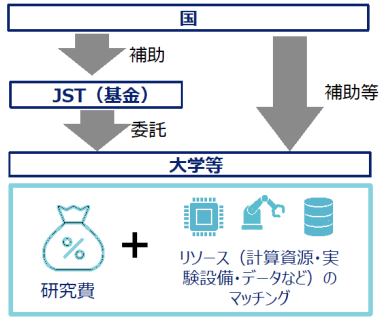


【取組のイメージ】

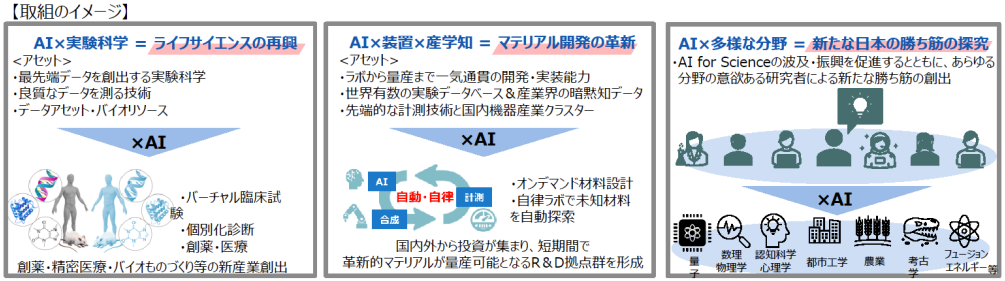


**«プロジェクト型»**　320億円

- 支援件数：5領域×3チーム程度（又は個人）
- 支援規模：20億円程度/件
- 支援期間：原則3年

**«チャレンジ型»**　50億円

- 支援件数：1,000件程度
- 支援規模：500万円程度/件
- 支援期間：～1年

（担当：研究振興局参事官（情報担当））

14

**Innovative scientific research by AI for Science**

- up to 15 teams with US$13M/3yrs for each
- up to 1,000 teams with US$35,000/1yr for each

➔ **US$240M in total**

**Developing new scientific methodology supported by AI for Science**

# Governmental program coming soon (2) “Grand Reach”

## 次世代計算科学グランドリーチプログラム

**演算部の多様化やシミュレーションとAIの融合といった新たな技術動向を踏まえつつ、近年の開発環境に適応した、ポスト富岳時代における「世界とつながり世界に普及する成果創出」に向けた戦略的なソフトウエア開発を推進**

### （A）エコシステム創出課題

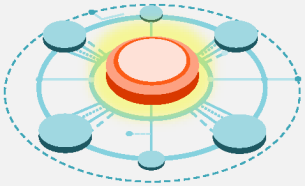

**ソフトウエアエコシステムの中核となりえる革新的なアプリケーション等の大規模な研究開発を行い、国際市場やコミュニティにおける高いシェアの獲得・維持を目指す。**

支援総額：最大1.2億円／件
支援期間：最長５年間、 支援件数： ３件程度

ポスト富岳時代（2030年代）のエコシステムの中核となりえる世界最先端の研究開発であることや、その成果物の具体需要の高さ及びコミュニティの拡大と維持に向けた計画等を評価

- ✓ ポスト富岳時代のヘテロジニアス・コンピューティングやAI for Science等の最新の潮流を踏まえた新規性のある研究課題であるか
- ✓ ユーザニーズや具体需要の高さ（ステークホルダー評価書を提出）
- ✓ ポスト富岳時代の汎用的な計算資源や演算システムの想定に立った戦略や設計（相互運用性やマルチスケーラビリティ等）
- ✓ 成果物の普及やソフトウエアの維持・管理に向けた体制や計画 等

### （B）エコシステム連動課題

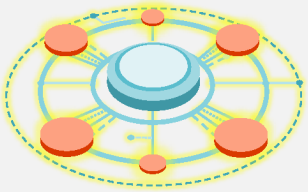

**支配的なソフトウエアエコシステムを基盤とした派生的な研究開発を行い、国際コミュニティにおけるプレゼンス向上や優位性確保等を通じた多面的便益創出を目指す。**

支援総額：最大5,000万円／件
支援期間：最長５年間、 支援件数：10件程度

波及効果の高いプラットフォーム上での先端的研究開発（派生的開発）であることや、その成果物の普及に向けた計画等を評価

- ✓ ポスト富岳時代のヘテロジニアス・コンピューティングやAI for Science等の最新の潮流を踏まえた新規性のある研究課題であるか
- ✓ 影響力の大きいソフトウエアエコシステム上の派生開発等であり、成果の公開によって我が国のプレゼンス向上等が期待できるか
- ✓ 成果物の普及に必要な体制や計画（海外プレイヤーとの連携等）
- ✓ 成果創出による便益の最大化に向けた計画 等

### （C）一般課題（計算資源の優先利用のみ）

A課題又はB課題の趣旨に基づき、計算科学の発展・利用加速に資する先端基盤的な研究開発を行う。
支援期間：最長５年間、 支援件数：最大15件

【A～C共通事項】 中間評価の実施： 事業期間が３年を超える場合、支援開始３年目に中間評価を行い、以降の課題継続可否を評価・判断する。
計算資源配分： 令和８年度における課題ごとの計算資源配分量の上限は、課題の種類を問わず、70百万ノード時間積とする。

17

**Grand Reach**

- up to 3 teams with US$0.7M/year x 5 years for each
- up to 10 teams with US$0.3M/1yr for each x 5 years for each
- up to 15 teams with computation resources only

→ **US$27M in total**

**Developing scientific applications to share internationally**

# Collaboration with RIKEN R-CCS for applications on FugakuNEXT

- Collaboration with R-CCS toward application development and performance tuning for FugakuNEXT
  - GPU technology
  - Programming method, including AI-support
  - Large scale parallel execution with inter-node communication
  - …
- Regionally close (HAIRDESC is at next stop of Port-Liner with R-CCS)
  - Deep discussion among researchers
  - In person networking
- Bridging application fields not just focused by FugakuNEXT with R-CCS
  - New governmental program “Grand Reach” to develop large scale GPU-accelerated codes for AI for Science
  - HAIRDESC supports these code development teams and bridging with FugakuNEXT

# International activities

- **Support on Export/Import of GPU codes for HPC/AI to enhance international collaboration over “GPU is a common language”**
- Collaboration with various international community on GPU/HPC/AI research
  - DOE-MEXT
  - HANAMI
  - E4S
  - …
- Promotion of activities at major international conferences: SC, ISC, SCA,…
- Collaboration with various HPC/AI centers (planned)
  - DOE
  - EuroHPC JU
  - JHPCN
  - …
- Sharing educational materials ?

# GPU vendor support and collaboration

- Support from GPU vendors is quite important to share advanced technical information, performance tuning skills, scientific code examples, etc.
- Tutorial/Hackathon on practical systems with GPU requires support of lecturers and education materials
- Tightly communicating with GPU vendors through any opportunity

- Let's move it toward advanced GPU code production and scientific result shared in the world

# Advanced GPU coding and utilization technologies

- Three core organizations perform their own advanced research on GPU coding, libraries, optimization, AI application, etc.
- Main target and role of three organizations (but not limited to them)
  - HPC: University of Tsukuba
  - AI: Institute of Science Tokyo
  - AI for Science: University of Tokyo
- They collaborate with each other to enhance the research and make synergy on HPC, AI and AI4S

# University of Tsukuba

**Only organization who operates both NVIDIA- and AMD- based supercomputers in HPCI**

- Easy GPU-ization by stdpar construct and unified memory

**standard Fortran loop to be parallized**

```
do concurrent (i=1:n)
   dx(i) = vx(i+1)-vx(i)
end do
```

Reducing data transfer between GPU and CPU by Unified Memory

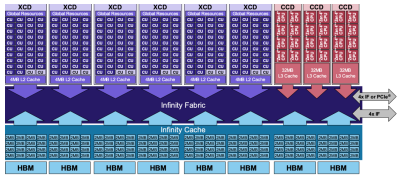


AMD MI300A
(post-Cygnus)

(Coming on March 2026)

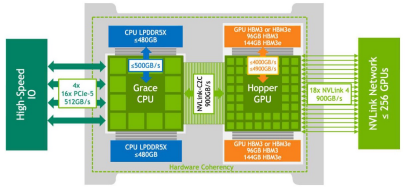


NVIDIA GH200
(Miyabi-G)

**Just using this construct to generate GPU code by compiler**
**Automatic data transfer go GPU memory (uncontrollable)**

- Performance comparison over AMD and NVIDIA GPUs, and common programming framework with performance portability → UniSYCL, collaboration with ORNL

**SYCL as a common language for all accelerators to apply both GPUs, and make possible to construct heterogeneous GPUs on a node or heterogeneous nodes on interconnect**

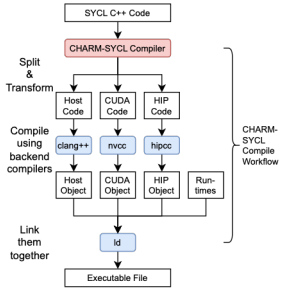


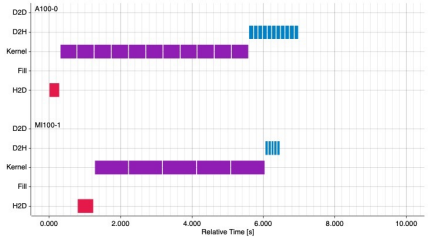


Automatic load balancing over NVIDIA and AMD GPUs by UniSYCL

D-to-H Copy
Kernel Exec.
H-to-D Copy

- Research on both kind of GPUs to be common or different on performance characteristics and coding know-how

# Three GPU systems in U. Tsukuba


Pegasus (2022-)
H100 GPU + Xeon +
Optane memory


Sirius (2026-)
MI300A APU

Miyabi-G (2025-)
GH200 GPU+CPU
jointly operated with U. Tokyo

## Main research topics at CCS, U. Tsukuba

- HPC application tuning especially for in-house code on various fields
- Code transform and tuning over multiple vendor GPUs
- Multi-device coordination system over different vendors and generations of accelerators
- etc.

# University of Tokyo

## HPC applications

✓ Mesh (structured/unstructured) and multi-body

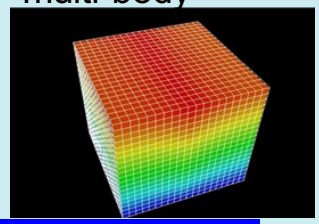

structured lattice (difference method)

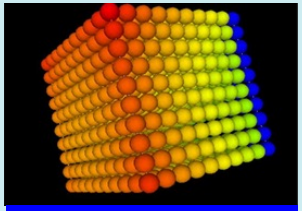

Multi-body

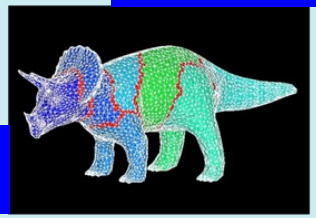

Unstructured mesh (FEM/???)

## Shift from Fortran

kokkos

✓ Code shift to Kokkos (performance portable framework over GPUs, C/C++), Solomon

✓ Applying OpenACC/OpenMP unified library, and incrementally escaping from Fortran

OpenACC とOpenMP の両方に対応した通常の実装

```
#if defined(OFFLOAD_BY_OPENACC)
#if defined(OFFLOAD_BY_OPENACC_KERNELS)
#pragma acc kernels
#pragma acc loop
#else
#pragma acc parallel
#pragma acc loop
#endif
#endif
#if defined(OFFLOAD_BY_OPENMP_TARGET)
#if defined(
   OFFLOAD_BY_OPENMP_TARGET_LOOP)
#pragma omp target teams loop
#else
#pragma omp target teams distribute
   parallel for
#endif
#endif
  for (int i = 0; i < N_i; i++) {
    // loop body A
  }
```

Solomon を用いた等価な実装

```
OFFLOAD()
for (int i = 0; i < N_i; i++) {
  // loop body A
}
```

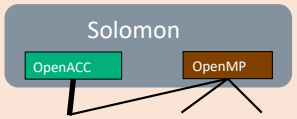


## AI for Science

✓ Simulation + Data + Learning (S+D+L) unified advanced software platform “h3-Open-BDEC”

✓ QC-HPC hybrid computing

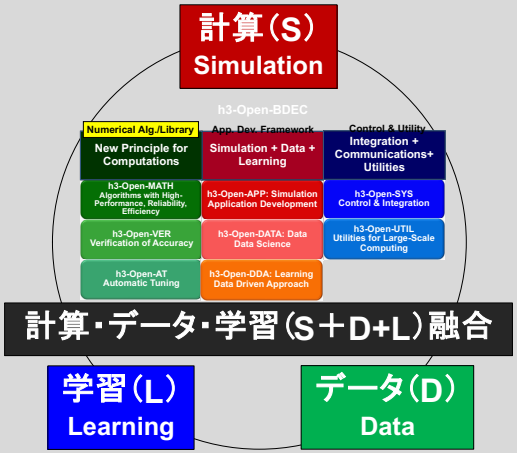


# Institute of Science Tokyo

## Deep Learning Software Stack

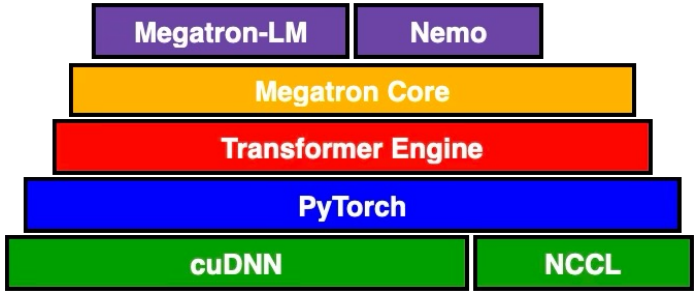


Megatron-LM: frontend for experts
NeMo: frontend for general users
Transformer Engine: backend for tensor core
Megatron Core: backend for parallel/distributed environment
cuDNN: matrix solver library
NCCL: inter-GPU communication library

・development/deployment of these software platforms to next-gen systems including NIS
・material distribution to users for easy understanding of software stack
・AI-oriented GPU utilization is specially focused

## GPU-shift support by Generative AI

・code generation accuracy improvement by practical condition, I/O example, explicit constraints

・references to the existing code base and related files

・breaking down the Complicated implementation in to simple steps

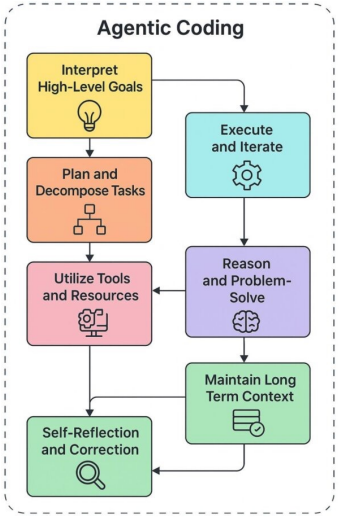


・unit test and combined test for verification on generated code is automatically performed
・feedback system where AI automatically proposes code optimization

# Summary

- Supercomputers in Japan including NFS and NIS are going to implement GPUs mainly, especially FugakuNEXT will be equipped with Fujitsu “MONAKA-X” CPU and NVIDIA GPU
  ➔ shifting from “Made in Japan” to “Made with Japan”
- It is critical to develop GPU-ready applications from traditional CPU-ready ones (OpenMP+MPI) to catchup all the development of next generation systems including NIS (National University’s Supercomputers)
- We launched a new governmental support organization named “HAIRDESC: Advanced HPC-AI Research and Development Support Center” in Kobe, and three core organizations (U. Tsukuba, U. Tokyo and Inst. Science Tokyo) collaborate with HAIRDESC
- Constructing a standard GPU education course over a collection and newly created “GPU education materials” for efficient GPU coding training for nation-wide computational and AI scientists
- Supporting GPU code development for top-level applications under governmental programs
- HPC, AI and AI4S with GPU technologies are going to research and develop by collaboration of HAIRDESC and three core organizations