

InPEX 2026 *Brilliant* Concluding Remarks

Jack

A **Haiku** is a type of short-form poetry that originated in Japan.

Traditional Japanese Haiku consist of three phrases composed of 17 syllables in a 5, 7, 5 pattern;

line 1: 5 syllables

line 2: 7 syllables

line 3: 5 syllables

**AI models train,
Quantum and HPC align,
Roadmaps steer compute.**

The background is a dark, abstract digital illustration. It features a network of glowing lines in shades of teal, orange, and red, resembling a circuit board or data flow. Various sized circles and squares are scattered throughout, some appearing as nodes or data points. The overall aesthetic is futuristic and technological.


InPEX

The International Post-Exascale Project (InPEX)

Masaaki KONDO RIKEN

- presented Japan's comprehensive approach to AI-driven scientific research and next-generation supercomputing infrastructure
- New projects 23 Billion Yen (144M US Dollars)
- Research projects 5 B Yen (31M US Dollars)
- 2B/project (12.5 M US Dollars) for 3 years
- FugakuNext - Fujitsu and Nvidia parts operations 2030


Planned Schedule for FugakuNEXT and New Computer Building Development



Masaaki KONDO RIKEN

- presented Japan's comprehensive approach to AI-driven scientific research and next-generation supercomputing infrastructure
- New projects 23 Billion Yen (144M US Dollars)
- Research projects 5 B Yen (31M US Dollars)
- 2B/project (12.5 M US Dollars) for 3 years
- FugakuNext - Fujitsu and Nvidia parts operations 2030

Planned Schedule for FugakuNEXT and New Computer Building Development




Mitsuhisa SATO – RIKEN

JHPC Quantum Project

Combining quantum + supercomputer (IBM and Quantumuum (ion trap))

QC qubits	Characteristics	Targets
Superconducting Qubits (IBM)	Medium qubit count (100 qubits or more) Fast operating speed (a few ns). Medium Fidelity.	Development of utilization technology and system software for the utilization and practical use of large- and medium-scale NISQ machines.
Trapped Ion Qubits (Quantinuum)	High fidelity, the number of qubits is not large.(about 20 qubits). Slow operation speed (a few ms). Efficient all-to-all qubit operation.	Software development using small scale but high fidelity. Use of quantum computers with properties different from superconducting qubits.

Rick Stevens, ANL

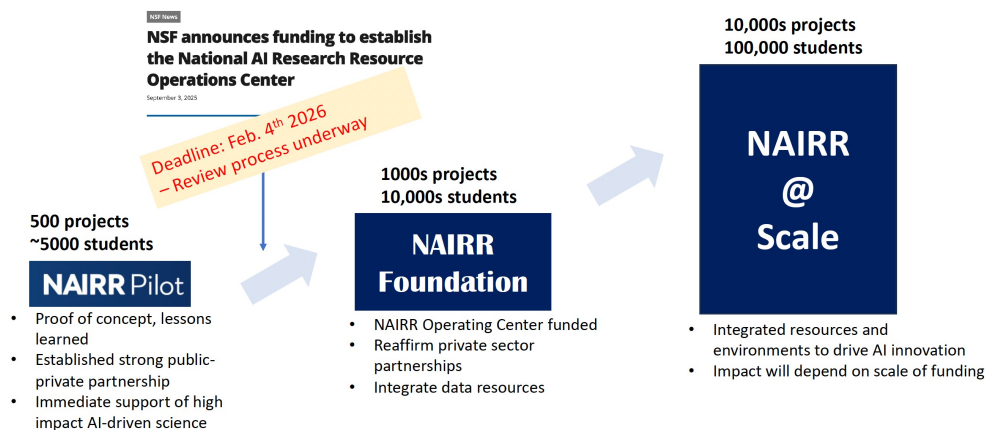


1. A November 2024 executive order launched the **Genesis Mission**, led by DOE and OSTP, aiming to **double U.S. R&D productivity within 10 years** using AI, HPC, and quantum computing.
2. A **consortium of 25+ major tech companies** (including NVIDIA, OpenAI, Google, Microsoft, Amazon, IBM) is being formed, with working groups structured similarly to the **IETF** model.
3. The mission focuses on **three application pillars**:
 1. Basic science discovery
 2. Energy technologies (fusion, fission, geothermal)
 3. National security
4. A key vision is **large, integrated research teams (~200 people)** using frontier AI models, exascale computing, and robotic labs to **accelerate decade-long problems by ~5×**.
5. A prototype “**replication engine**” reproduced 70,000 scientific papers, with runtimes ranging from **seconds (theory) to days (complex simulations)**.
6. Novel scientific discovery is estimated to require **~30× more effort than reproduction**, but with **500–1,000 GPUs and 75–100 agents**, 10 years of computational science could be reproduced in **~5 weeks**.
7. Tokens are the new metric for compute.
8. Major research gaps remain: AI systems currently lack **scientific intuition**, leading to inefficient brute-force approaches and high compute costs.

Bottom line: Genesis represents a **massive, coordinated national push for AI-driven science**, with strong demand (10,000 proposals) but clear open problems in reasoning, efficiency, and trust.

Katie Antypas NSF perspective

- Talked about NAIRR National AI Research Resource
- To broaden access to AI compute, data, models, software, training, and user support so that researchers and educators outside a few large institutions can use advanced AI infrastructure.
- NAIRR is an administrative priority
- Public – Private partnership
- Around 700 projects supported
- 14 agencies and 25 industries
- Project agency supported:503 industry:278



Laura Schulz, ANL

1. National coordination through the National Quantum Initiative (2018)

Aligns DOE, NSF, and DOD, with major DOE quantum centers (Q-NEXT, SQMS, QSC), emphasizing integration with HPC systems.

2. Strong regional ecosystem centered on the Chicago Quantum Exchange

Brings together leading universities, national labs, and industry, forming one of the most concentrated quantum hubs in the U.S.

3. Infrastructure investment via the Illinois Quantum and Microelectronics Park

Anchored by PsiQuantum, creating a dedicated site for fabrication, development, and access to quantum technologies.

4. Focus on usability: integrating quantum into HPC workflows

Goal is to make quantum a usable resource without requiring domain scientists to understand low-level quantum mechanics (e.g., Hamiltonians).

5. Open software stack development (OpenQSE)

Led by DOE labs to create standardized, interoperable interfaces—reducing fragmentation and enabling seamless workflow integration across platforms.

Gosia Nikovska, European Commission's

1. AI Continent Action Plan targets broad sector adoption

Designed to lower barriers to entry and build a pan-European AI ecosystem across key domains like health, climate, agriculture, finance, and public services.

2. Deployment of 19 “AI Factories” across Europe

Provide integrated services including compute access, data platforms, advisory support, trustworthy AI frameworks, domain expertise, and workforce training.

3. Flexible access model via EuroHPC Joint Undertaking

Tailored access pathways—fast-track for startups and extended allocations for large-scale projects—aim to broaden participation and optimize resource use.

4. Federated ecosystem with AI Antennas

Extends reach to smaller countries by linking them to major AI Factories, with a coordinated network (19 factories + 13 antennas) moving toward unified access.

5. Emerging “Gigafactories” for large-scale AI compute

Planned as massive infrastructure complements to AI Factories, requiring careful coordination to balance access across small and large users.

AI models train,
Quantum and HPC align,
Roadmaps steer compute.

International Exascale Software Project

Overall goal:

- To develop a plan for producing a software infrastructure capable of supporting exascale applications

Ken Kennedy – Petascale Software Project (2006)

- SC08 (Austin TX) meeting to generate interest **Nov 2008**
- Funding from DOE's Office of Science & NSF Office of Cyberinfrastructure and sponsorship by Europeans and Asians **Apr 2009**
- US meeting (Santa Fe, NM) April 6-8, 2009 **Jun 2009**
- European meeting (Paris, France) June 28-29, 2009
 - Outline Report **Oct 2009**
- Asian meeting (Tsukuba, Japan) October 18-20, 2009
 - Draft roadmap and refine report **Nov 2009**
- SC09 (Portland OR) BOF to inform others
 - Public Comment: Draft Report presented **Apr 2010**
- European meeting (Oxford, UK) April 13-14, 2010 **Oct 2010**
- Maui Meeting October 18-19, 2010 **Nov 2010**
- SC10 (New Orleans) BOF **Apr 2011**
- San Francisco Meeting – April 6-7, 2011 **Oct 2011**
- Cologne, Germany – October 6-7, 2011 **Apr 2012**
- Kobe, Japan April 2012

Frequency was 3 f-to-f meetings per year with Working Groups, working in between




THE INTERNATIONAL JOURNAL of
**HIGH
PERFORMANCE
COMPUTING
APPLICATIONS**

The International Exascale Software Project roadmap

**Jack Dongarra, Pete Beckman, Terry Moore, Patrick Aerts,
Giovanni Aloisio, Jean-Claude Andre, David Barkai,
Jean-Yves Berthou, Taisuke Boku, Bertrand Braunschweig,
Franck Cappello, Barbara Chapman, Xuebin Chi, Alok Choudhary, Sudip Dosanjh,
Thom Dunning, Sandro Fiore, Al Geist, Bill Gropp, Robert Harrison, Mark Hereld,
Michael Heroux, Adolfo Hoisie, Koh Hotta, Zhong Jin, Yutaka Ishikawa, Fred Johnson,
Sanjay Kale, Richard Kenway, David Keyes, Bill Kramer, Jesus Labarta, Alain Lichnewsky,
Thomas Lippert, Bob Lucas, Barney Maccabe, Satoshi Matsuoka, Paul Messina,
Peter Michielse, Bernd Mohr, Matthias S. Mueller, Wolfgang E. Nagel, Hiroshi Nakashima,
Michael E Papka, Dan Reed, Mitsuhsa Sato, Ed Seidel, John Shalf, David Skinner,
Marc Snir, Thomas Sterling, Rick Stevens, Fred Streitz, Bob Sugar, Shinji Sumimoto,
William Tang, John Taylor, Rajeev Thakur, Anne Trefethen, Mateo Valero,
Aad van der Steen, Jeffrey Vetter, Peg Williams, Robert Wisniewski and Kathy Yelick**

The International Journal of High
Performance Computing Applications
25(1) 3–60
© The Author(s) 2011
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1094342010391989
hpc.sagepub.com



Today the paper has
971 citations

Need To Address

- This is a pivotal moment in the development and convergence of HPC and AI.
- It calls for international coordination to define a shared roadmap for the future.
- What are the key areas that such roadmaps should address?

AI -> scientific computing workflow, if we make it reliable, efficient, explainable, reproducible, and grounded in physics, data, and numerical analysis.

Roadmaps to work on:

- 1. AI agents for scientific workflows**
- 2. Mixed-precision and energy-efficient AI**
- 3. Trust, verification, and uncertainty quantification**
- 4. AI-HPC system co-design**
- 5. Scientific data infrastructure**
- 6. Benchmarks for time-energy-fidelity-to-solution**
- 7. AI-assisted numerical algorithms**