TPC WORKSHOP · ISC HIGH PERFORMANCE 2026 · HAMBURG

Same Mission, New Owner: Is the AI Factory HPC's Successor or Its Eviction Notice?

How the rise of AI Factories is reshaping HPC: a natural evolution of its mission and impact — or a fundamental shift that risks displacing scientific computing?

Moderated by Gokcen Kestor

Barcelona Supercomputing Center



TPC
Trillion Parameter Consortium
AI FOR SCIENCE



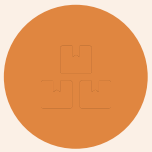
euro TPC
AI FOR SCIENCE

BRIDGING THE AI–HPC GAP

Closing the gap — but on whose terms?

TPC’s purpose is to put AI to work for science. The machine that does it — **GPUs, RDMA fabrics, the parallel-software stack —** were built by HPC, and are now the AI factory’s too.

The hardware gap is closing fast. **What’s unsettled is the mission it serves — and who sets its priorities.**



Same mission?

AI to accelerate science — TPC’s founding aim



New owner?

tuned for token output and revenue, not scientific discoveries



One machine

GPU, RDMA, collectives — AI and HPC share it

THREE READINGS

Three ways to read the same shift

The same convergence reads three ways from HPC’s natural successor to its eviction notice.



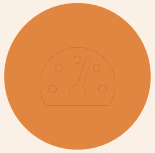
Evolution

The AI factory is HPC at greater scale — same mission, better tools. Its natural successor.



Convergence

AI and HPC as roommates — one shared, co-designed machine serving both.



Eviction?

Commercial AI priorities crowd out open science — scientific computing as a tenant.

HPC’s successor *its eviction notice*

THE PANEL

Three perspectives at the HPC–AI boundary

We have three panelists from different institutions/domains, each with a different perspective on ownership and mission drivers.

Each opens briefly; then we debate and take your questions. Listen for where they agree, and where they don’t.

Panelists seem to believe evolution/convergence is the way, did we rule out eviction?



Yanfei Guo

Evolution · NVIDIA — GPU comms (NCCL, NVSHMEM)



Shreeya Badhe

Convergence · C-DAC, India — Trinetra interconnect



Pedro Valero-Lara

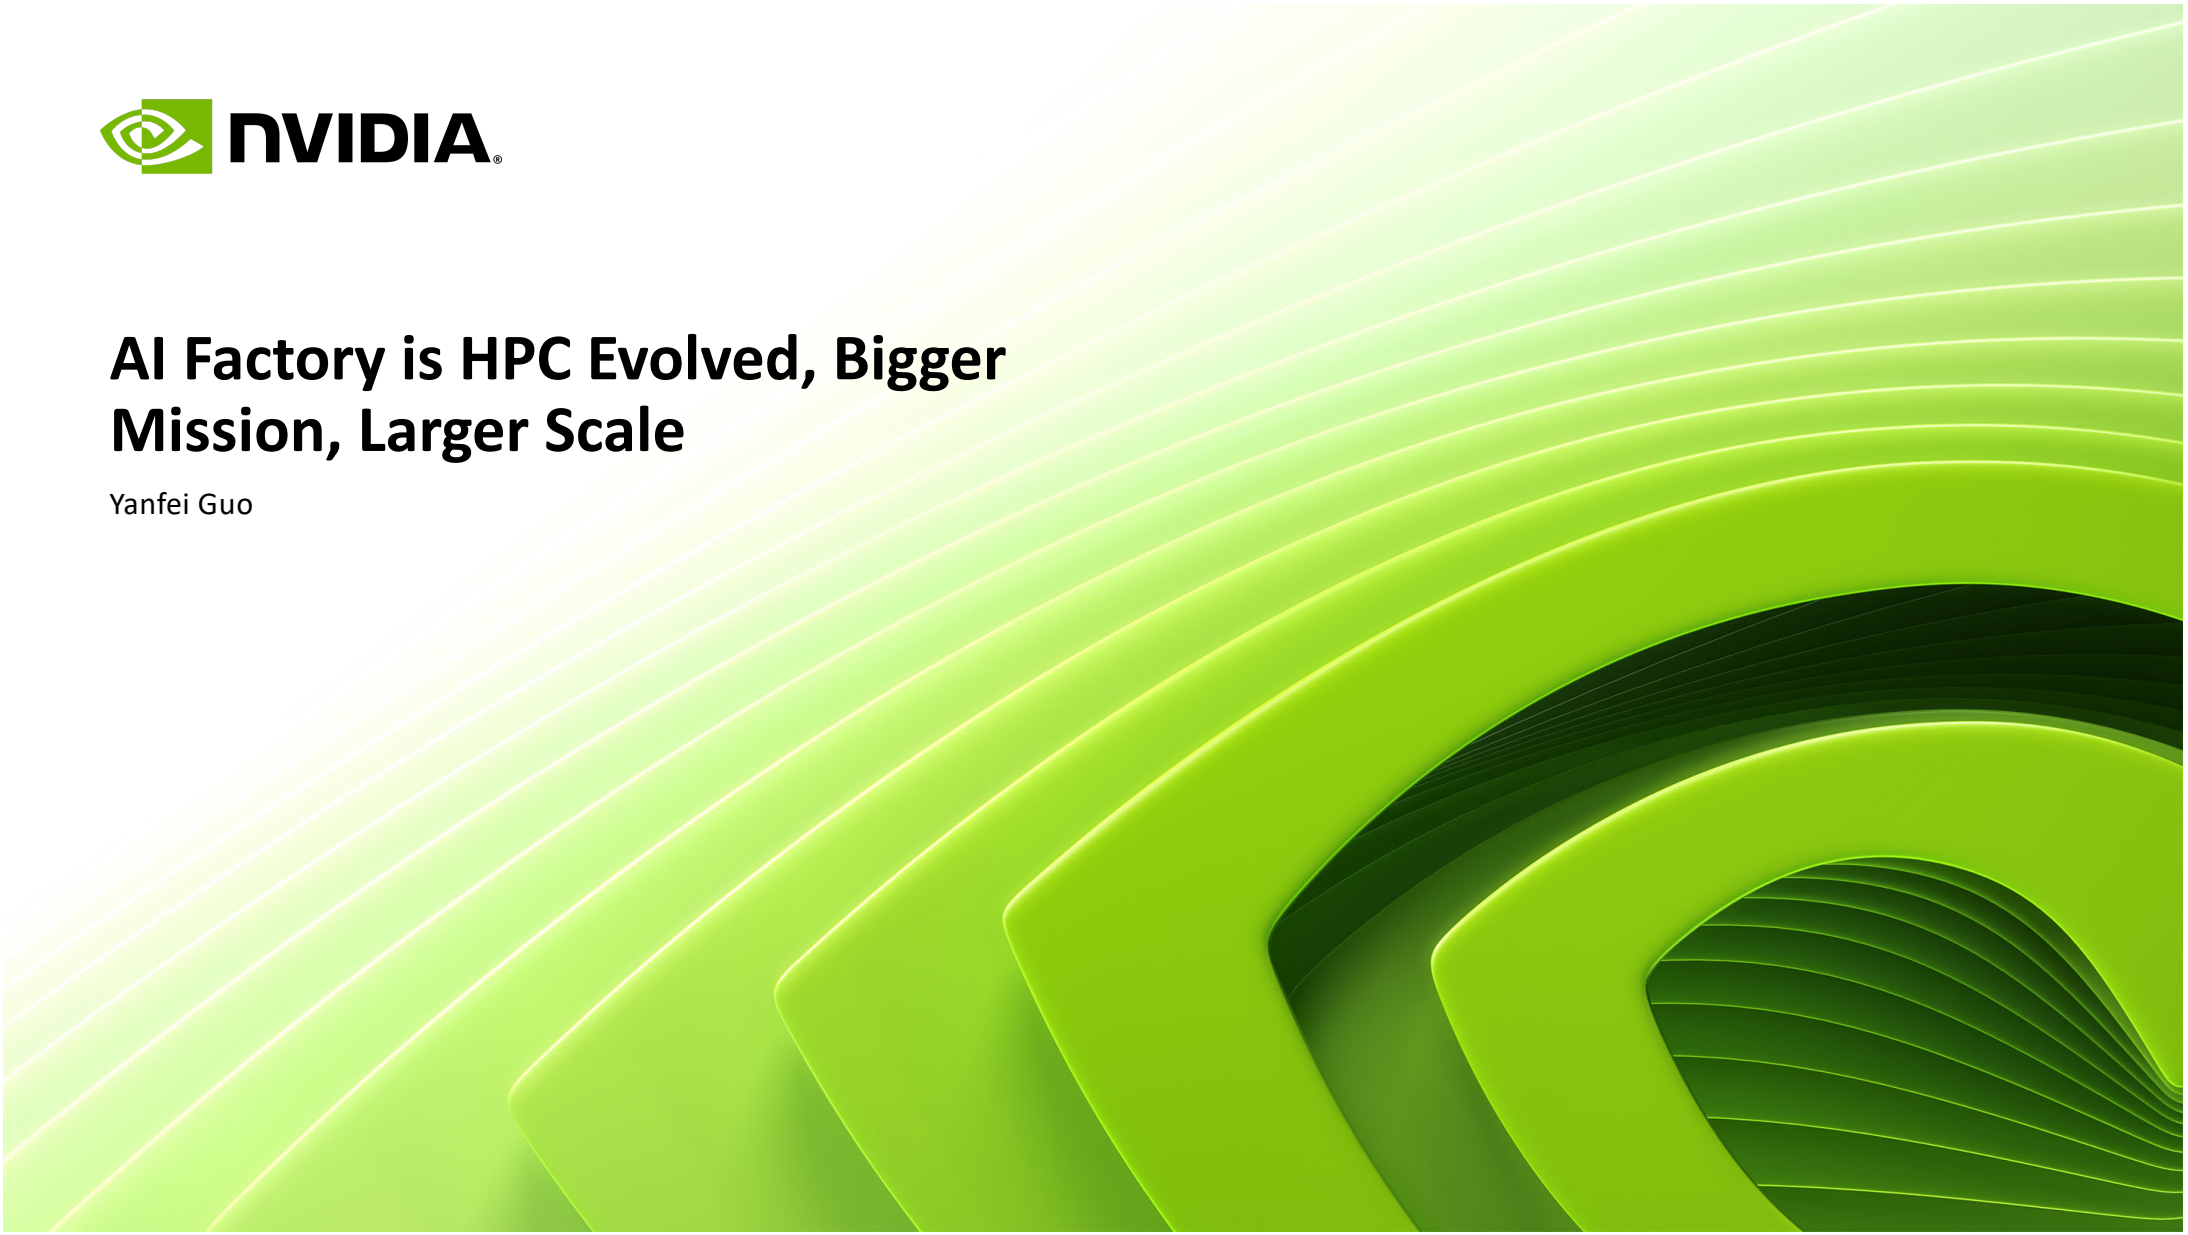
Continuity · ORNL — no new owner; science leads

OPENING REMARKS

Over to our panelists



Three short openings — then we put the question to the room.

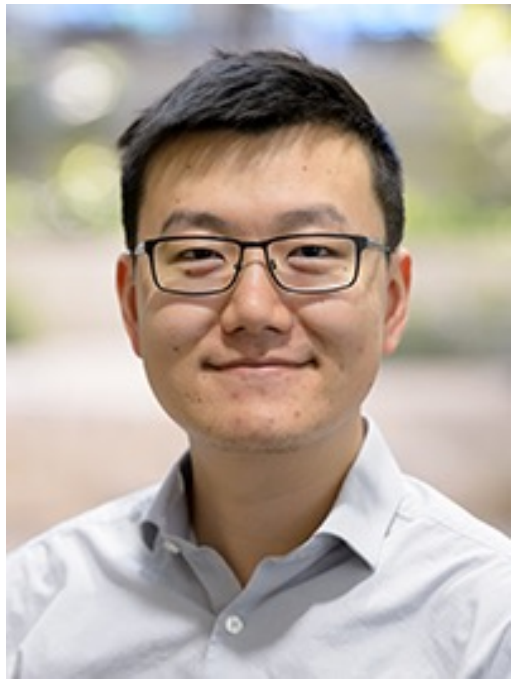


AI Factory is HPC Evolved, Bigger Mission, Larger Scale

Yanfei Guo

Yanfei Guo – NVIDIA

- Senior Software Architect, GPU Communication Team, NVIDIA Corporation
  - We are responsible for GPU communication libraries including NCCL and NVSHMEM
- My research focus is HPC Communication Libraries, Programming Models and Runtime Systems
  - I was a Computer Scientist at Argonne National Laboratory
  - 10 yrs in the MPICH project, one of the most popular MPI library



AI Factory is HPC Evolved, Same Mission, Larger Scale

- AI Factory is an evolution of HPC.
  - From the software perspective, many core infrastructure were evolved from traditional HPC. Parallel algorithm, collective communication, shared memory communication that drives AI today all came from HPC.
  - From hardware perspective, first GPU supercomputer was built in 2008. By 2012, the fastest supercomputer is built with GPU. RDMA networks like Infiniband was also HPC technology.
- AI applications fosters new technologies that benefits all
  - Double Precision FP -> Mixed Precision FP
  - Latency Optimized Communication -> Bandwidth Optimized Communication
- AI and HPC applications are going to be roommates
  - Powerful tools to help scientific research
  - Agentic AI is going to close the loop, HPC applications will be a crucial part to achieve AI for Science
- Opportunities for new algorithms, new programming models and new runtime systems
  - Increased system scale is always going to a challenge
  - Maybe it is a time to rethink and rebuild HPC applications differently

Shreeya Badhe

Associate Director
Centre for Development of Advanced Computing (C-DAC), India



Professional Background

25+ years at C-DAC, contributing to India's HPC ecosystem
Technical Lead for **Trinetra** – C-DAC's indigenous high-performance interconnect
Leading hardware development of scalable communication fabrics for next-generation HPC and AI systems

Expertise

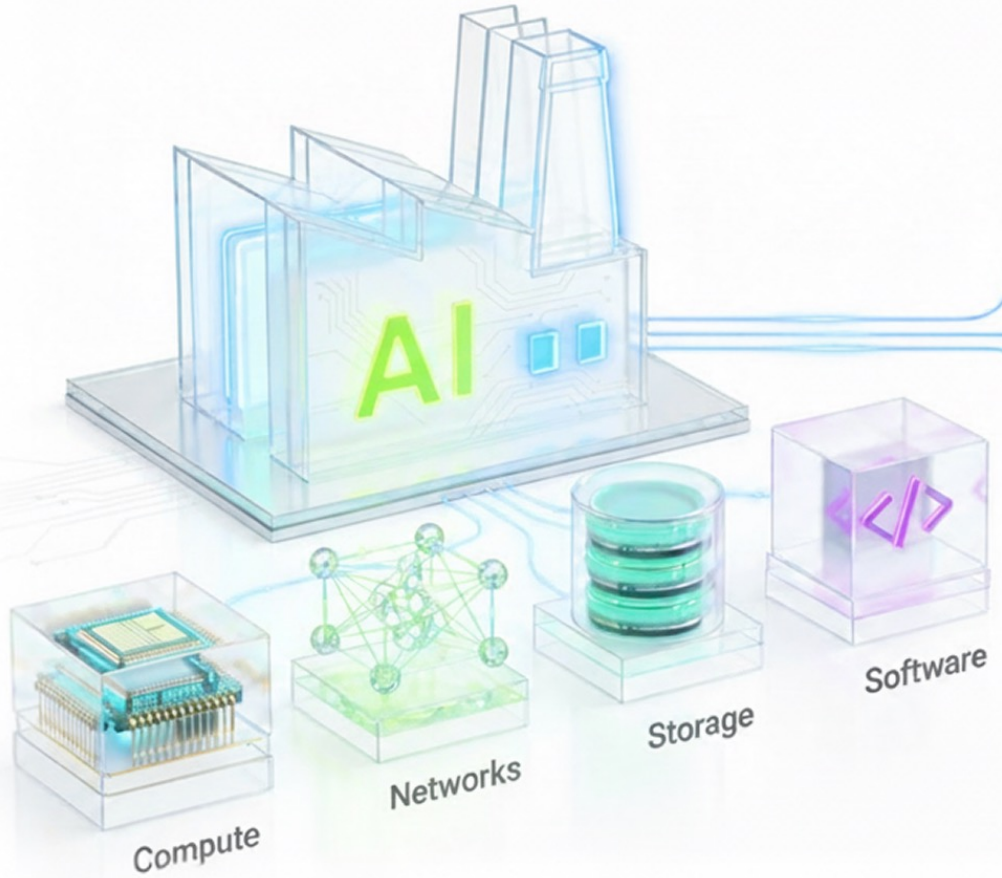
- 📌 High-speed Interconnect Architecture and Design
- 📌 Spec to SoC Design
- 📌 High-Speed Interface & Protocol Design
- 📌 FPGA Prototyping, Emulation & Hardware-Software Co-design
- 📌 End-to-End Performance Optimization of Communication fabrics for large-scale computing systems

Current Focus

Building indigenous technologies for sovereign AI/HPC infrastructure
AI-HPC convergence and scalable system architectures
Next-generation Communication fabrics for HPC and distributed AI

AI Factories: Evolution Of HPC, Not Its Replacement

Co-existence through Convergence



Same Foundation, New Workloads

AI Factories build on decades of HPC innovation –architecture, scale, reliability–optimized for unique demands.

Broader Mission, Greater Impact

From understanding the universe to enhancing human capability–AI amplifies the reach of HPC across science, industry, and society.

Complement, Don't Compete

Simulation and AI are synergistic. Together, they solve problems faster.

Responsible Scale

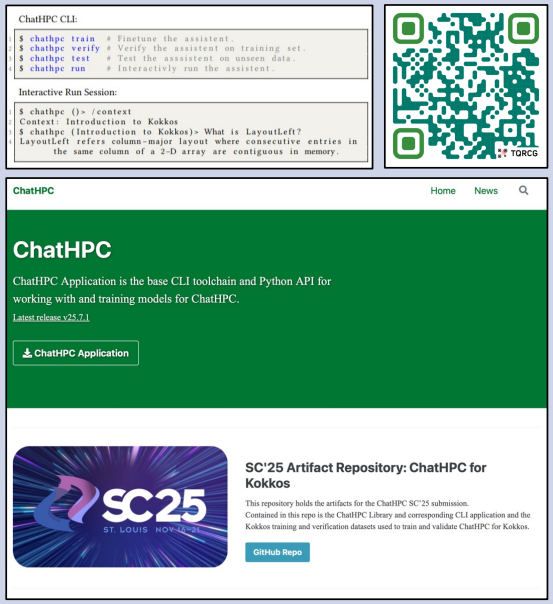
We must design AI Factories with sustainability and energy efficiency as core principles for the long term.

Valero-Lara, Pedro, PhD

Senior Computer Scientist at Oak Ridge National Laboratory, Programming Systems Group

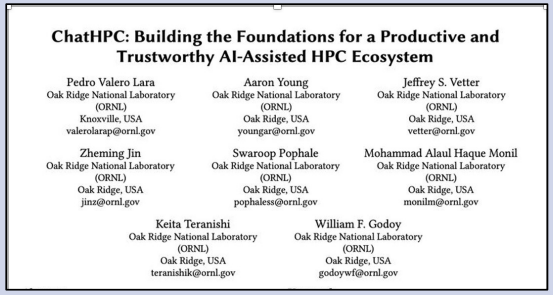
Transforming HPC with AI

- ChatHPC, an LLM-based framework for HPC capabilities
  - A CLI Python library on top of Code Llama, PyTorch, LoRA, …
  - Fine-tuning, Testing, Refinement, Merge, and Inference
  - AI-assistants for Programming Models, Math Libraries, I/O, Tooling, etc.
  - ChatBLAS, ChatHPC for Kokkos, ChatMPI, ChatTau, ChatADIOS, etc.
  - https://ornl.github.io/ChatHPC/
- Current Work
  - Expanding ChatHPC capabilities across other HPC domains and AI capabilities; Agentic, Multi-modality, ….
- Community
  - LLM4HPC Workshop (co-located with ISC-HPC'26):
  - https://ornl.github.io/events/llm4hpc2026/











OAK RIDGE National Laboratory

Valero-Lara, Pedro, PhD

Senior Computer Scientist at Oak Ridge National Laboratory, Programming Systems Group

Same Mission, New Owner: Is the AI Factory HPC's Successor or Its Eviction Notice?

Honestly, I think that there is no a new “owner”, the owners, domain scientists, computational scientists, and computer scientists are still the same.

I do not think that the big companies are very interested in the HPC and Scientific Computing mission.

Our mission is being influenced by AI. AI is changing everything, not only HPC. AI is influencing how we make HPC, but this is made, mainly by the HPC and Scientific Computing community, using and adapting AI capabilities to our mission.



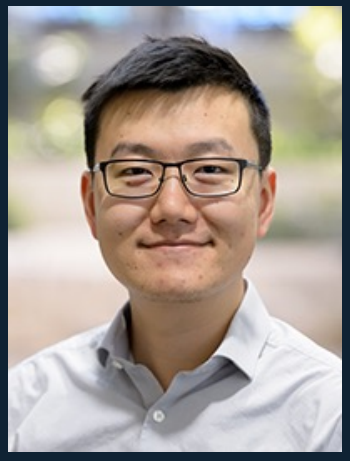
TPC WORKSHOP · ISC HIGH PERFORMANCE 2026 · HAMBURG

Trillion Parameter Consortium

AI FOR SCIENCE

Same Mission, New Owner: Is the AI Factory HPC's Successor or Its Eviction Notice?

How the rise of AI Factories is reshaping HPC: a natural evolution of its mission and impact — or a fundamental shift that risks displacing scientific computing?







euro TPC AI FOR SCIENCE