

HLRS

Thirty Years, One Vision

High-Performance
Computing Center
Stuttgart

European AI Factories: Germany


Dennis Hoppe
Head of Converged Computing, HLRS

Germany's AI Factory Strategy: A Dual-Pillar Approach

- Germany hosts **2 AI Factories**
 - **JAIF** (Exascale Pillar)
 - Coordinated by **JSC** with partners including RWTH Aachen, Fraunhofer, and Hessian.AI
 - Integrated with Europe's first exascale system JUPITER and the inference extension JARVIS is planned to go into operation in 2026
 - Focusing on massive-scale training and inference, and cross-disciplinary scientific AI
 - 2nd cut-off; services started on 31 October 2025
 - **HammerHAI** (Industrial Pillar)
 - Coordinated by **HLRS** with LRZ, GWDG, KIT, and SICOS
 - EuroHPC JU's first AI-optimized supercomputer
 - Tailored for industry, engineering, and global challenges
 - 1st cut-off; services started on 1 April 2025



The HammerHAI System Going Live Late 2026



NVIDIA GB200 NVL by HPE liquid-cooling system.
Image source: HPE

Compute Nodes

- 215x NVIDIA GB200 NVL-4 nodes (860 GPUs)
- 9x Axelera AI nodes (9 AIPUs in total)

Custom Nodes

- 13x NVIDIA L4 Tensor Core GPUs (24 GB memory) (suitable for inference)
- 5x NVIDIA RTX A1000 large memory nodes (4 TB memory)

Performance

- 15 Exaflops of peak AI inference performance (FP4)

Storage

- 10 PB VAST Data DASE

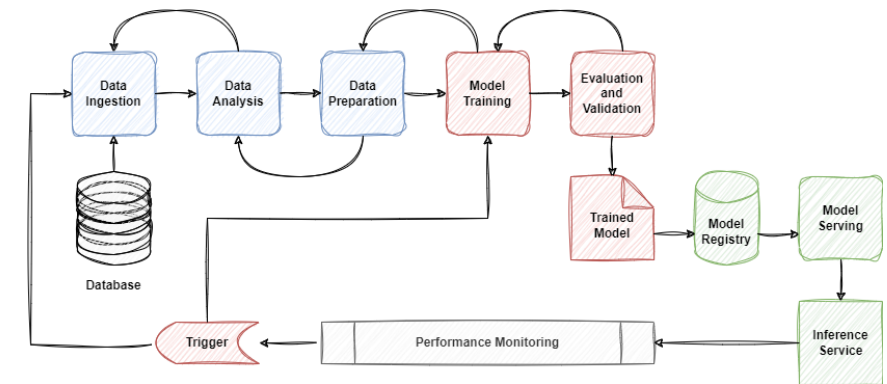
Network

- NVIDIA Quantum-X800 InfiniBand (800 Gbit/s)
- NVIDIA Spectrum X (800 Gigabit Ethernet)

Rethinking HPC for the AI Era

From Infrastructure to Innovation

- The Strategic Shift
 - From a minimal set of libraries to the AI software zoo
 - From bare-metal to containers and VM
 - From exclusive node access to multi-tenancy
 - From queues to on-demand resources
- The AI User Experience Journey
 - From complex governance to frictionless compute access
 - From providing a cluster to enabling a “plug-and-play” AI ecosystem
 - From raw compute power to real-world impact in 3-6 months
 - From hard-coded scripts to seamless automation



Thank you!



Dennis Hoppe

Head of Converged Computing, HLRS
Project Manager, HammerHAI

High Performance Computing Center (HLRS)
University of Stuttgart
Nobelstr. 19
D-70569 Stuttgart, Germany

email: hoppe@hls.de

phone: +49-711-685-60300

fax: +49-711-685-65832

web: www.hls.de