

Forum TERATEC

20 - 21 OCTOBRE 2026

PARIS · PALAIS DES CONGRÈS

Beyond Exascale, toward the convergence of HPC, AI and quantum

Michaël KRAJECKI - GENCI

GENCI

Le calcul intensif au service de la connaissance

Grand Equipement national de calcul intensif

- >2,700 HPC, AI and Quantum projects (both academic and industrial) in 2025 – 15% provided by industry
- 1,700 AI projects in 2025 thanks to GPU, user support and tailored access modes

TGCC/CEA
- Hosting site of the **2nd EuroHPC exascale system**
- Hosting site for the **1st HPC-QC infrastructure** (HQI, HPCQS, EuroQCS-France)



EuroHPC Joint Undertaking

ALICE RECOQUE BEYOND HPC

IDRIS/CNRS
- **1st converged HPC-AI system > 3,800 GPUs in 2025**



FRANCE 2030

CINES/FU
- **> 90 PF** with next gen AMD GPUs (>2,000) and CPUs (>100k)
- **Used for HPC and AI**



Quantum Computing

FRANCE 2030 HQI France Hybrid HPC Quantum Initiative
Forum TERATEC 20 - 21 OCTOBRE 2026 PARIS · PALAIS DES CONGRÈS

Already available at TGCC :

Quantum emulator Qaptiva (40 qubits)

- **4 major programming environments:**
  - MyQLM-QLMaaS, Pulser, Perceval et QisKit

3 quantum systems:

1. **Pasqal system (Ruby) 100+-neutral-atom system**
   Available since Nov. 5, 2025
2. **Quandela system (Lucy), 12-qubit-photonic system**
   Available since May 2026
3. **Alice & Bob system (Kity) 18-cat-qubit system**

5 regional quantum ecosystem :

- QUORUM (IdF), Maison du Quantique Occitanie (MQO), Hybquant (Nouvelle Aquitaine), Maison du Quantique Alpes (Grenoble), MAQUEST (Grand Est)

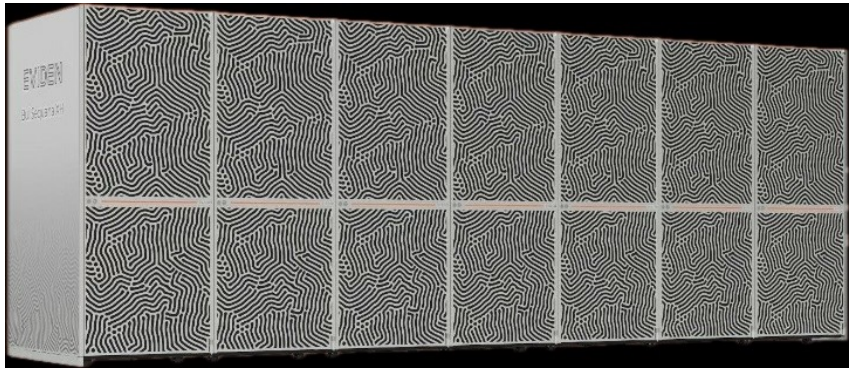


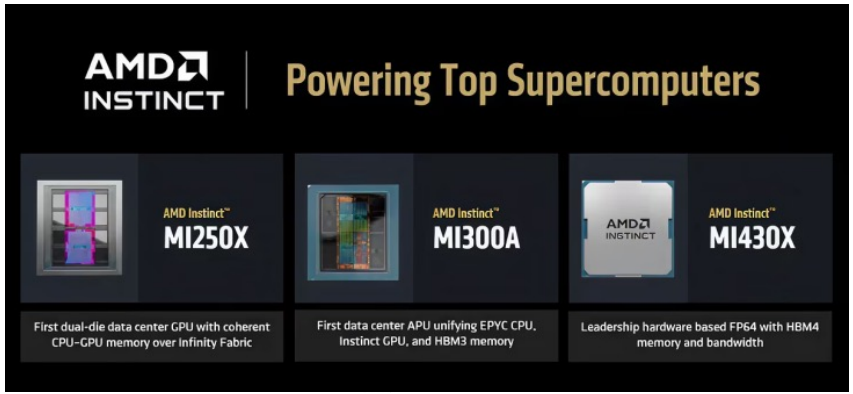
EuroHPC Joint Undertaking
GENCI

Alice recoque in a nutshell

- **Bull XH3500 integration**
  - 94 compute racks
  - 100% DLC
  - Production power expected between 12 and 15MW
- **Unified compute node**
  - Multi-tenant accelerated and scalar workloads
  - 1 CPU AMD Venice 256c tightly coupled to 4 GPUs AMD MI430x (432 GB HBM4 @ 19.6 TB/s)
  - 1 TB MRDIMM memory
  - 2 x 400 Gbs BULL BXIv3 per GPU + 1 x 400 Gbs per CPU
- **Scalar nodes**
  - ARM based SiPEARL Rhea2 128c
  - > 100k cores

ALICE RECOQUE
BEYOND HPC





L'Intelligence Artificielle générative

LE Grand Continent

Études · Puissances de l'IA

La vassalisation cognitive est une menace existentielle

« En abandonnant l'intelligence artificielle aux géants américains, nous sommes en train de perdre le contrôle de notre cerveau. »
https://legrandcontinent.eu/fr/2026/10/14/la-vassalisation-cognitive-est-menace-existentielle/







NORFOLK, VA – Cognitive warfare is not the means by which we fight; it is the fight itself. The brain is both the target and the weapon in the fight for cognitive superiority. Contesting in this environment comprises deliberate, synchronized military and non-military activities throughout the continuum of competition designed to gain, maintain and protect cognitive advantage.

Cognitive Warfare focuses on attacking and degrading rationality, which can lead to exploitation of vulnerabilities and systemic weakening. However, this becomes increasingly complex as non-military targets are involved. An example: Russian social media and public information operations targeted much of the international community in an attempt to label Ukraine as being at fault. Through a combination of communication technologies, fake news stories, and perceptions manipulation, Russia aims to influence public opinion, as well as decay public trust towards open information sources. These narratives have extensive reach, and often involve both offensive and defensive posturing.

https://www.act.nato.int/activities/cognitive-warfare/

GENCI

HPC + AI: The Ecosystem Challenge

- **Energy & Sustainability Limits**
  - The Efficiency Gap: The disconnect between model capability (AI) and system efficiency (HPC).
  - We are optimizing for accuracy, not compute-per-watt, leading to unsustainable scaling.
- **The "AI-HPC" Talent Gap**
  - Risk: Shortage of experts who understand both parallel algorithms and AI, leading to bottlenecks in scientific discovery.
- **Algorithmic Obsolescence & Hardware Fragmentation**
  - The Trap: Rapid AI model iteration + Proliferation of architectures (CPU, GPU, APU, Quantum) = Code Fragmentation
  - The Cost: High maintenance overhead, loss of portability, and "vendor lock-in" that stifles scientific agility.
- **Supply Chain Fragility**
  - Risk: Geopolitical tensions and export controls disrupt access to critical compute resources for global research
- **Reproducibility & Verification Crisis**
  - Risk: Proprietary black-box optimizations make it difficult to guarantee numerical reproducibility across different hardware/software generations.

Erosion of Scientific Autonomy

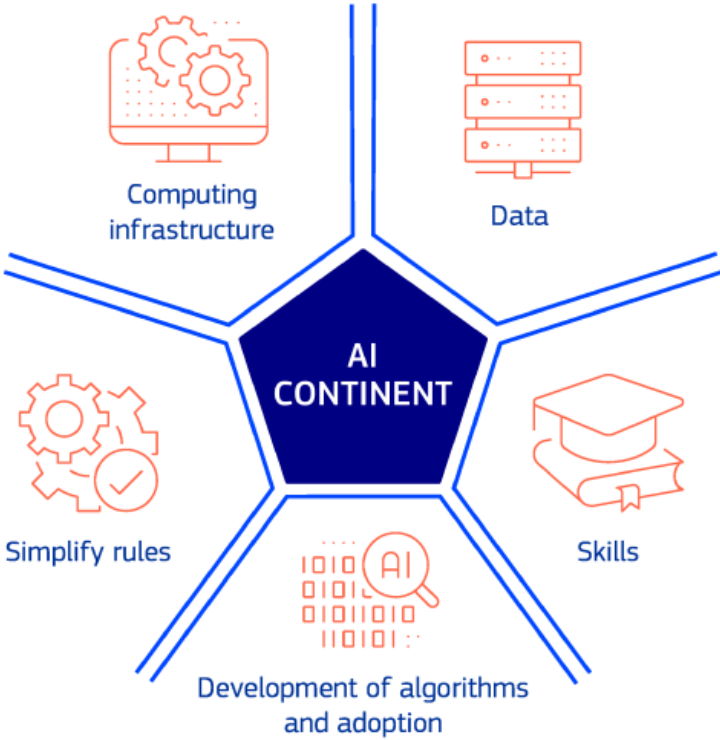
The gradual loss of Scientific Autonomy due to large reliance on a proprietary, hyperscaler-dominated HPC/AI stack.

1. **Convenience vs. Control**
   - The current stack is incredibly convenient and performant. Researchers prefer it because it works now. The risk is long-term: losing the ability to choose, modify, or replace the tools that drive science.
2. **The "Black Box" Problem**
   - Proprietary software stacks often hide optimizations. For science, transparency is key. If we cannot understand or verify the underlying algorithms because they are locked in proprietary binaries, we risk compromising scientific rigor.
3. **Innovation Stagnation**
   - A monoculture discourages alternative approaches. If everyone uses the same tools, we miss out on potentially better, more efficient, or more specialized solutions developed by other vendors or open-source communities.

AI Continent Action Plan

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS

1. **AI Factories / Gigafactories**
   - Train and develop complex AI models
   - Boost research in highly sustainable infrastructure
   - Encourage investments
2. **Data Union Strategy**
   - Improve data access for businesses and administrations and simplify data rules
   - Data Labs in AI Factories will gather and curate high-quality data from different sources
3. **AI talent**
   - Form partnerships to recruit internationally
   - Offer AI fellowships to top students, researchers, and AI professionals
4. **Development of algorithms and adoption**
   - Accelerate AI adoption in strategic sectors, such as healthcare, automotive and advanced manufacturing
   - Support businesses and public administrations in developing and deploying promising AI solutions
5. **Facilitation of the AI Act Implementation**



European Commission

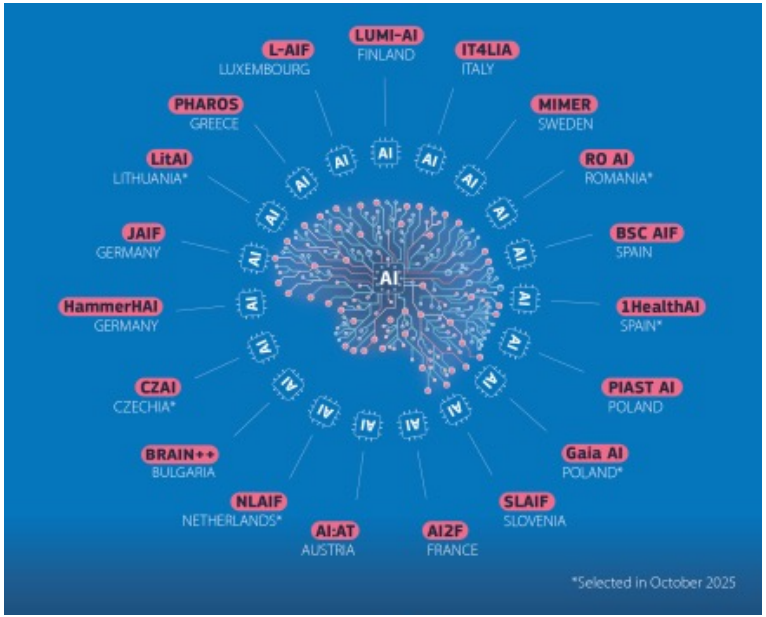
The 19 AI Factories (and their 13 AI Antennas)

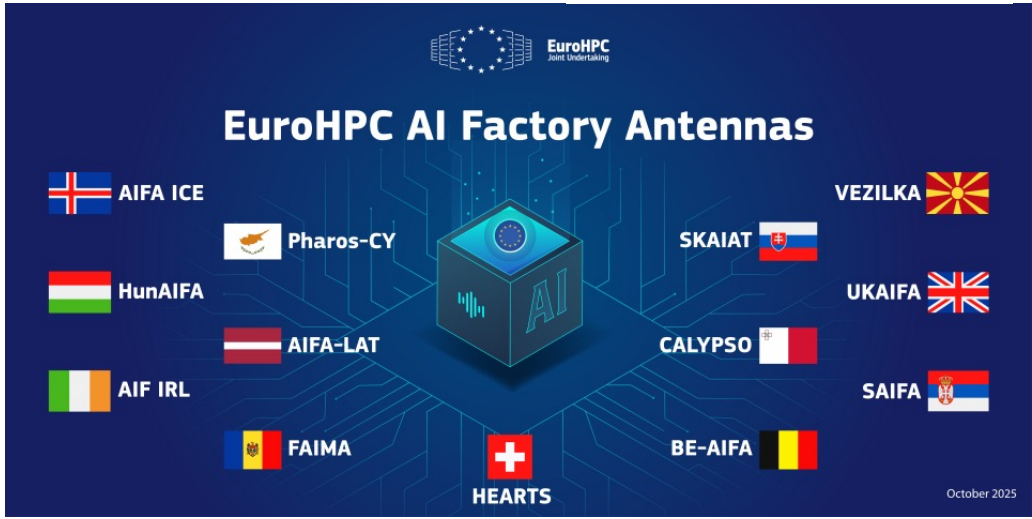
Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS

AI Factory France

Budget : €2.6 billion

EuroHPC Joint Undertaking





GENCI

What is an « AI Factory »?

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS

AI-optimized supercomputers + Associated resources & services





AI Factory France

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS

30M€, 3 years, 20 partners

Objectives

- Provide the AI community with a one-stop-shop of tailored AI services (compute, data, skills)
- Intensify efforts in domain-specific specialisation for widening the use of AI in industry, science and public services
- Foster a close collaboration between research and industry. Develop public / private continuum of facilities and services.
- Foster startup creation. Attract and retain skills with new career paths
- Invest massively in talent detection, education and permanent (re)training.

With more than 40 letters of support including

Serving 12 application verticals / AI + X

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS


Finance


Cyber/Defense


Materials science


Robotics


Health


Earth science


Aerospace


Energy


Edtech


Agriculture


Finance


Humanities

While waiting for Alice Recoque…

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS

Dalia (hosted at IDRIS)

- GB200 NVL72: 36 NVIDIA Grace CPUs and 72 Blackwell GPUs
  - User training via MesoNet
  - Job execution for open research (outside the ZRR)
- Co-investment with AMIAD
- In production since November 2025



HELIOS – The first AMD rack-integrated system

- 18 compute nodes, each equipped with:
  - 1× AMD Venice 256-core host CPU
  - 4× AMD MI455x GPUs – 432 GB HBM4, 23.3 TB/s bandwidth

Provided to GENCI by AMD and BULL

- First in Europe – deployment at CINES: Q4 2026 / Q1 2027
- Strategic partnership / AMD ROCm ecosystem
- Strengthening national & EU capabilities / HPC–AI convergence
- Readiness for Alice Recoque and Gigafactory workloads

GENCI