



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



“Celebrating EuroHPC JU ExaStars Alice Recoque and JUPITER” | Plenary

Moderator: **Addison Snell, CEO, Intersect360 Research**

Sections of the Plenary

Introduction

Part 1 – Project Status & Site Adaptation

Part 2 – Architecture & Performance

Part 3 – Software

Part 4 – Users support

Panel Discussion

Introduction

Quick overview of the two EuroHPC exascale supercomputers

JUPITER | Introduction

Booster

- 125 Racks
 - BullSequana XH3000
- 5884 Nodes:
 - NVIDIA Grace-Hopper
 - 23.536 NVIDIA GH200




ExaFLASH: 29PB (raw) NVMe, IBM SSS6000

ExaSTORE: 308PB (raw) HDD*, IBM SSS6000

ExaTAPE: 370PB Tape*, LTO9

Cluster

- 10 Racks
 - BullSequana XH3000
- 582 Nodes AMD Turin
 - Dual-Socket
 - Q3-Q4/2026
- 162 Node SiPearl Rhea1
 - Dual Socket
 - EO2027



ALICE RECOQUE | Introduction

1. **Compute nodes:** Unified partition with memory coherent accelerated nodes allowing multi tenant execution of scalar and accelerated HPC/AI workloads
2. **GPU:** Accelerated nodes, each with 1 AMD Venice 256 cores host processor and 4 AMD Instinct™ MI430X GPU with 432 GB of HBM4 memory
3. **CPU:** SIPEARL next gen ARM-based processors + AMD EPYC™ Venice 256 cores with 1 TB of RDIMM for scalar workload
4. **Interconnect:** BULL BXI v3
5. **Storage:** Under acquisition

Introduction | Summary Slide

COMMON GROUND

- Both EuroHPC JU exascale systems, co-funded with their hosting entities
- Both delivered by Bull, with European technology across the stack
- Both combine accelerated and scalar compute in a single system
- Both designed for converged HPC and AI workloads
- Both reaching or exceeding 1 ExaFLOP/s

KEY DIFFERENCES


- Design: modular Booster and Cluster vs a unified compute partition plus a scalar Rhea2 partition
- Accelerators: NVIDIA GH200 superchips vs AMD MI430X GPUs with Venice CPUs
- Interconnect: NVIDIA Quantum-2 InfiniBand NDR vs Bull BXI v3
- Storage: JUPITER's three tiers installed, Alice Recoque's under acquisition
- Maturity: JUPITER in pre-access since May 2025, Alice Recoque from 2027

Project Status & Site Adaptation

Part 1

Project Status & Site Adaptation


1. Overall timeline
2. Choices
3. Difficulties
4. Status
5. Key contributors



JUPITER


Presenting: **Benedikt von St. Vieth – Jülich Supercomputing Centre**

JUPITER | 1. Overall Timeline



JUPITER | 2. Choices

- JUPITER: Procurement of multiple modules with independent timelines
- JUPITER: Strong focus on benchmarking, especially large scale applications that can use full system
- Datacenter: Switch from classical datacenter to MDC because of timeline and price hikes
- Early Access: As soon as possible and as large as possible (JUREAP, AI Competition)
 - JEDI deployed 6 months after contract





JUPITER | Site Adaption

- Upgrade of campus transformer station
 - 2x40 MVA to 2x60-80 MVA
- Upgrade of 110kV power line to new cables
- Preparing of concrete slab to host the MDC
 - 85x42x0,5m concrete
 - 1MW backup cooling (from campus)
 - River-water connection (up to 30 cbm/h)
- 3,2 Tbit/s network connectivity to campus




JUPITER | Installation

- Delivery of *Datahall*, 4 containers to host management plane and **ExaFLASH/ExaSTORE**
- Deliver of 7 *IT modules*, 7x2 containers with 10 racks each, hosting the 125 **Booster** racks
- IT modules mostly pre-comissioned
 - Racks shipped inside of containers
 - Rack-local copper cabling pre-comissioned in Angers (FR)
- Blade insertion on-site
- Global Dragonfly+ cabling on-site
- Currently: JUPITER Cluster installation



JUPITER | 3. Difficulties

- Aggressive timeline overall (EoI to 1EF in <4 years)
- Complexity of the system due to its size
 - Management Software (Scalability, Zero Trust in the era of AI *agents*)
 - Human Resources to achieve availability
- First retrofit of hardware
 - Power board replacement
 - Thanks to JUREAP findings
- Parallel acceptance of both MDC and JUPITER
 - And parallel installation, with rain, and dust... :)




JUPITER | 4. Status

- Must have 1ExaFLOP/s achieved
- Running Simulations and AI in pre-access since May 2025 (with disruptions, currently stable)
- ... the provider Consortium still has lots of homework to do for acceptance



JUPITER | 5. Key Contributors



Funding by

- EuroHPC (50%)
- BMFTR (Federal, 25%)
- MKW (State, 25%)
- 500M€ TCO

Delivery by

- ParTec/Bull Consortium for JUPITER
Booster/Cluster and ExaFLASH
- IBM for ExaSTORE
- Pro-Com for ExaTAPE




ALICE RECOQUE
BEYOND HPC

Presenting: **Corinne Béal - GENCI** and **Anaïs Gaertner - CEA**


ALICE RECOQUE | 1. Overall Timeline

- A multi-year project that began in 2022 !
- Anticipation of the adaptation work of TGCC (Very Large Computing centre of CEA)





ALICE RECOQUE | 2. Choices

- Adaptation of the existing TGCC building by CEA specific infrastructure division, while still operating other supercomputers
- 2022: Pre-study
 - Room layout : what goes where ?
 - Adaptation of supercomputer room
 - Strengthening of the floor capacity
 - Increasing electrical power
- Better energy performance
 - New warm water-cooling system
- Better address user needs
 - Evolution of IT infrastructure
 - More open environment





ALICE RECOQUE | 3. Challenges and solutions

- Increasingly dense cabinet => need to reinforce the concrete slab ($1 \text{ t/m}^2 \rightarrow 2.8 \text{ t/m}^2$)
 - In the basement: Creation of a metal frame to sustain the computer room floor (67 micropiles)
 - In the computer room: Add local reinforcement, application of a resin to waterproof the slab
- Done while having many HPC systems in production (including Joliot Curie) !



ALICE RECOQUE | 3. Challenges and solutions

- Upgrading the electrical capacity: 24MW (TGCC already had a 2x30MW 225kV line)
- Adding new cooling towers (20MW) for the warm water-cooling system but nearby residence => construction of a noise-damping wall
- Dismantling Joliot-Curie supercomputer (Nov. 25) to finalize the work adaptation (Q2/26)
- Bull will handle: Power and water distribution / Water leak detection system adaptation / Raised floor strengthening



ALICE RECOQUE | 4. Status

- Computer room ready to host Alice Recoque supercomputer
- Final water and power distribution realized by Bull (started Q3/26)
- Delivery of first racks planned in Q4/26
- Alice Recoque ready for production end 27

ALICE RECOQUE | 5. Key contributors

Jules Verne consortium is composed of

- France represented by GENCI (Hosting Entity) and CEA (Hosting site)
- The Netherlands represented by SURF
- Greece represented by GRNET



Alice Recoque will be delivered

- by Bull
 - with AMD unified partition
 - and SIPEARL CPUs



Funding by

- EuroHPC
- France 2030 and French Ministry of Research for France
- Ministry of Education for the Netherlands
- Ministry of Digital Governance for Greece



Project & Site | Summary Slide

JUPITER

- Hosting Entity decision in June 2022, with JEDI early access six months after contract
- Plan shifted from a concrete Data Centre to a Modular Data Centre for faster deployment
- Power supply connected in November 2024
- Modular Data Centre declared ready in January 2025, with JUREAP moving applications onto early XH3000 racks
- First HPL runs in May and October 2025, with hardware and software work continuing

ALICE RECOQUE


- Hosting Entity decision in June 2023
- Pre-study launched in 2022 at CEA's TGCC Computing Centre
- Existing building adapted, with floor capacity, power and warm-water cooling upgraded while other systems kept running
- Computer room ready, with final power and water distribution by Bull from Q3 2026 and first racks in Q4 2026
- Jules Verne consortium brings GENCI, CEA, SURF and GRNET together

Architecture & Performance

Part 2

Architecture & Performance


1. System architecture
2. Compute components (GPUs + CPUs) for HPC and AI
3. Memory and bandwidth
4. Performance at scale
5. Interconnect and EU value



JUPITER

Presenting: **Estela Suarez – Jülich Supercomputing Centre**

JUPITER | 1. System architecture



JUPITER | 2. Compute components

1 ExaFLOP/s



Picture: FZJ-JSC

JUPITER Booster

>1 EFLOP/s HPL (GPU)

~**24000 NVIDIA Grace-Hopper**

CPU: 72× ARM Neoverse V2 cores,
4×128b SVE2, 120 GB LPDDR5

GPU: NVIDIA H100 GPUs, 96 GB HBM3

(4×) InfiniBand NDR DragonFly+

Bull Sequana XH3000



Picture: nvidia

JUPITER Cluster

>5 PFLOP/s HPL (CPU)


~1100 AMD Turin CPU

~**300 SiPearl Rhea1 CPU**

80× ARM Neoverse V1 cores, 2×256b SVE,
256 GB DDR5 + 64 GB **HBM2e** (1TB/s)

InfiniBand NDR DragonFly+

Bull Sequana XH3000





Picture: heisse.de

JUPITER | 3. Memory and bandwidth

JUPITER Booster: node architecture


- 4x GH200 (CPU + GPU)
- Tight CPU-GPU connection: 450 GB/s
- **Unified memory:** 120GB LPPDR + 96 GB HMB3

Picture: FZJ-JSC




Picture: nvidia

GH200




JUPITER | 3. Memory and bandwidth

Traditional PCIe attached GPUs



**values for JURECA-DC:
AMD7742 + NVIDIA A100*

JUPITER (GH200) design



**values for JUPITER: GH200*

JUPITER | 4. Performance at scale

JUPITER performance: 1 Eflop/s

- Nr. 5 Top500
- Nr.17 Green500

The most energy-efficient Exascale system



Gordon Bell Prize for Climate Modelling 2025

<https://dl.acm.org/doi/10.1145/3712285.3771789>


JUPITER | 5. Interconnect

JUPITER interconnect

Intra-node: NVlink 4 (100-150 GB/s)

Inter-node: InfiniBand NDR (200 Gbit/s)

- 4 NICs per node (1 per GH200)
- Dragonfly+
 - 25 cells: **Booster**
 - 1 cell: **Cluster**



JUPITER | 5. EU added value




Modular Data Centre	Racks and Blades	Software	JUPITER Cluster Processor
Planned and constructed in Europe	Planned and constructed in Europe	Developed in Europe	Designed and developed in Europe
Bull	Bull	ParTec, Bull, JSC	SiPEARL



ALICE RECOQUE
BEYOND HPC

Presenting: **Maïke Gilliot - CEA** and **Eric Boyer - GENCI**

ALICE RECOQUE | 1. System architecture



Modular architecture addressing :

- Needs of HPC/AI/QC communities
- New services
 - Resources combination for complex user scenarios
 - Relying on CEA cybersecurity expertise

ALICE RECOQUE | 2. GPUs & CPUs for HPC and AI

AMD EPYC “Venice” CPU

- x86 Zen 6 arch, 256 cores
- 4 GB/core => 1TB/node
- CPU/GPU unified memory
- Interconnect BXlv3 400Gb/s per CPU
- CPU socket for scalar and accelerated workflows => **Multi tenant mode**

AMD Instinct MI430X GPU

- 288 Tflops FP64 - 432 GB HBM4 memory
- High level performance of native FP64
- Advanced AI capabilities
- Previous generations deployed at CINES (FR), LUMI (FI), LLNL, .. (MI250X / MI300A)
- Interconnect 2x BXlv3 = 800Gb/s per GPU

SiPearl Rhea2 CPU partition : ARM architecture, >100 000 cores

BULL's newest BullSequana XH3500 platform

ALICE RECOQUE | 3. Memory and bandwidth

Unified Node

- 4x432Gb HBM4 @ 23,3 To/s (GPUs) + 1 To RDIMM@8000 (CPU)
- Memory coherent access GPUs/CPU
- Enabling large AI models & HPC

Interconnect

- High injection BXIv3 4 x (2x400) Gb/s (GPUs) & 400 Gb/s (CPU Venice)
- Dragonfly topology with low blocking factor

I/O

- LUSTRE & Object Storage (on going procedure)

ALICE RECOQUE | 4. Performance at scale

What we target :

- 1 ExaFLOP/s
 - Balanced for HPC and AI usage
 - Efficiency, based on
 - time-to-solution
 - energy-to-solution
- } → cf. benchmarking

ALICE RECOQUE | 5. Interconnect...

- **Bull's BXI v3 solution**
 - GPU: 2x400 Gbps
 - CPU: 400 Gbps
 - Inter routers connection (800 Gbps/cell)
- **Dragonfly+ topology**
 - 405 Switch (DragonFly+ Topology)
 - 8,467 BXI-Connected Compute Nodes
 - Intra-cell: fat tree topology



ALICE RECOQUE | 5. ...and EU added value

For the hardware

- European integrator : Bull
- European components:
 - SiPearl Rhea2 CPU
 - BXI v3 interconnect



For the software

- Leveraging the strong open European software ecosystem
 - Ocean, pcocc, SSHProxy, Phobos, wi4mpi,...
 - EuroFP software stack, SEANERGYS
 - EESSI, Scalasca, Score-P, Meric, ...



Architecture & Performance | Summary Slide

JUPITER

- Modular architecture combining GPU-accelerated and CPU-based partitions
- Around 24 000 GPUs in the Booster, plus 324 SiPEARL CPUs and 1 164 AMD CPUs in the Cluster
- Tight CPU-to-GPU links (450 GB/s per direction) and high-bandwidth unified memory
- 1 exaflop/s performance – No. 5 on Top500 and the most energy-efficient exascale system
- European value across the data centre, racks, blades, software and CPU technology

ALICE RECOQUE


- Modular environment for computing, artificial intelligence, quantum computing and new services, built on CEA cybersecurity expertise
- Accelerated nodes combine 1 AMD Venice CPU with 4 AMD MI430X GPUs, each with 432 GB of HBM4
- Coherent CPU and GPU memory in the accelerated nodes, with Bull BXI v3 federating the partitions
- Target above 1 exaflop/s, benchmarked from an application perspective – time-to-solution and energy-to-solution efficiency
- European value through Bull, SiPEARL Rhea2 CPUs and BXI v3, plus an open European software ecosystem

Software

Part 3

Software

1. Software stack
2. Benchmarking
3. Workflows
4. Integration with EFP
5. Tooling and developer experience



JUPITER

Presenting: **Andreas Herten – Jülich Supercomputing Centre**

JUPITER | 1. Software stack | Overview

	Technology	Challenges	Provider
Resource Manager	Slurm	Scalable API Known	  The Slurm Company SLURM
Parallel Storage	Storage Scale System (GPFS)	Performance Scalable Data security	
MPI Runtime	Message Passing Interface	Stable Performance GPU-support Bug-free	 
GPU Support	CUDA HPC SDK	Memory management Performance Integration	 
Monitoring & Logging	Prometheus + Thanos Syslog + Fluentd	Usable Scalable Handle data storm	  
Reference Database	Data Center Information Management	Automation API Coherent	

JUPITER | 1. Software stack | Management

SMC xScale

- By Bull
- Core of stack
- Based on OSS and cloud technology



ParaStation

- By ParTec
- Enhancements of core
- Integration of ParTec tools into xScale



xOPS

- By JSC
- Enhancements of core
- Ansible roles, JSC needs/requirements



JUPITER | 1. Software stack

- Operating System: RedHat Enterprise Linux 9
- Scientific stack: EasyBuild / EESSI
 - Rolled out in Stages (current: Stages/2026)
 - GCC, LLVM, NVIDIA Compilers → OpenMPI/ParaStationMPI
 - Module browser: <https://apps.fz-juelich.de/jsc/software/jupiter/>
 - Also: Apptainer
 - EESSI: Not yet active, but soon™



JUPITER | 2. Benchmarking | Procurement



JUPITER Benchmark Suite

- >20 Benchmarks (Applications, Synthetic)
- Full-system benchmarks, TCO, Synthetic
- All public / open-source; reference results
- SC24 paper

Application-Driven Exascale:

The JUPITER Benchmark Suite

DOI [10.1109/SC41406.2024.00038](https://doi.org/10.1109/SC41406.2024.00038)



JUPITER | 2. Benchmarking | Build-Up

JUREAP: Research & Early Access Program


- >100 applications participating
- Benchmarking-first onboarding approach
- From first rack until production
- Many successes already; example: Gordon Bell Prize for Climate Modelling with ICON
- Developed *exaCB* tool for Cont. Benchmarking
→ every JUREAP app is reproducible, benchmark



JUPITER | 3. Workflows

- Enablers
 - Slurm: Main access point for most jobs
 - Also UNICORE: Drives for example Jupyter Hub
- AI workflows via Slurm/SSH, but prototypes also via UNICORE
- CI: Jacamar

JUPITER | 4. Integration with EFP



JUPITER | 5. Tooling and developer experience





ALICE RECOQUE
BEYOND HPC

Presenting: **Maïke Gilliot - CEA** and **Eric Boyer - GENCI**

ALICE RECOQUE | 1. Software stack

- **Essentials**

- OCEAN: CEA administration SW stack (Kubernetes-based)
- Lustre for parallel storage
- AMD Support: e.g., ROCm stack, libraries, ...
- Bull Support for BXIv3 related elements (OpenMPI-BXI, IntelMPI-OFI-BXI)



- **“Classic” HPC stack**

- Slurm/flux, OpenMP/ MPI,...

- **AI stack**

- Pytorch, JAX, Ollama, vLLM, ...




ALICE RECOQUE | 2. Benchmarking

- Balanced HPC and IA weight benchmark portfolio
- Efficiency as main criteria
 - Applications : ETS and TTS 50/50 for scoring
 - Representative workload as object metric
- Objectives :
 - Incentive for bid with the best balance between components
 - Incentive for energy optimization environments
 - Include all energy & costs contributors
- Best value for energy and money spent

Full TCO (total cost of ownership) method :

Based on a representative set of applications, run in production conditions



Then :

$$E_T = E_C + E_D + E_G$$

$$E_T = \left[(E_{CI} \times \frac{c}{c_{wt}} \times PUE_{compute}) + (E_{DI} \times PUE_{disk}) + (E_{GI} \times PUE_{cool}) \right] \times \frac{T}{t}$$

$$TCO_g = E_T \times \text{EnergyCost}$$

$$TCO = \text{Purchase} + \text{Maintenance} + \text{Infrastructure} + \text{Staff} + TCO_g$$

Where:

- Purchase is the acquisition cost of the machine
- Maintenance is the maintenance cost over 5 years
- Infrastructure is the specific cost of infrastructure for the machine
- Staff is the cost for complementary staff for machine exploitation over 5 years

Cost of the workload (TCO_g) takes into account the number of time (N) it will be possible to run it over 5 years

$$N = \frac{c}{c_{wt}} \times \frac{T}{t}$$

$$TCO_b = TCO / N$$

Variable definition :

t : wallclock time of the workload

E_{CI} : Energy consumption by computational nodes

E_{DI} : Energy consumption from disks and service nodes

E_{GI} : Energy consumption from internal cooling devices

c : total core count of the proposed configuration

c_{wt} : core count used for the selected workload

T : Number of second in 5 years

N : Number of runs of the workload during 5 years

E_T : Total energy consumption during 5 years

E_C : Energy of computational part during 5 years

E_D : Energy of storage and service nodes during 5 years

E_U : Energy of internal cooling device during 5 years

TCO_g : TCO due to energy consumption during 5 ans

PUE_{air} : Power Usage Effectiveness for air cooled devices

PUE_{coldw} : Power Usage Effectiveness for "Cold water" cooled devices

PUE_{hotw} : Power Usage Effectiveness for "Hot water" cooled devices

Energy calculations :

$$E_{CI} = E_{CI} \times \frac{c}{c_{wt}} \times PUE_{compute} \times \frac{T}{t}$$

$$E_{DI} = E_{DI} \times PUE_{disk} \times \frac{T}{t}$$

$$E_{GI} = E_{GI} \times PUE_{cool} \times \frac{T}{t}$$

With

$$PUE_{compute} = 1 + \left[(PUE_{air} - 1) \times \eta_{air-for-compute} \right] + \left[(PUE_{coldw} - 1) \times \eta_{coldw-for-compute} \right] + \left[(PUE_{hotw} - 1) \times \eta_{hotw-for-compute} \right]$$

$$PUE_{disk} = 1 + \left[(PUE_{air} - 1) \times \eta_{air-for-disk} \right] + \left[(PUE_{coldw} - 1) \times \eta_{coldw-for-disk} \right] + \left[(PUE_{hotw} - 1) \times \eta_{hotw-for-disk} \right]$$

$$PUE_{cool} = 1 + \left[(PUE_{air} - 1) \times \eta_{air-cool} \right] + \left[(PUE_{coldw} - 1) \times \eta_{coldw-cool} \right] + \left[(PUE_{hotw} - 1) \times \eta_{hotw-cool} \right]$$


ALICE RECOQUE | 3. Workflows

Alice Recoque will serve multiple kind of workflows

- Multi-scale and multi-physics numerical simulations including HPC simulations coupled with AI models
- HPC simulations with in-situ / in-transit post processing and computational steering capacities
- Inference of AI models with continuous learning, AI pipelines, ...
- Ingestion and inference on the fly (stream mode) of data coming from large scale instruments (satellites, telescopes, sequencers, accelerators, IoT...) toward virtual twins
- Challenges : support of end-to-end workflows, security, co-scheduling,

ALICE RECOQUE | 4. Integration with EFP

- Involved in collaboration with the EFP consortium and the operational HE since the start of the project
- EFP services will be available
 - for Alice
 - for Lucy (and Ruby)



Source:
EuroFP consortium

ALICE RECOQUE | 5. Tooling & developer experience

- **How to connect**
 - SSH
 - Remote desktop (with accelerated 3D features)
- **SW deployment**
 - Spack for community-specific installations
 - OCI containers deployment (pcocc)
 - System-wide deployment mechanism based on EESSI
- **GPU-specific support**
 - HIP, Hipify, OpenACC, OpenMP Target, AMD ROCm, Kokkos, ...



Software | Summary Slide

JUPITER

- Bull, ParTec and Jülich stack around SMC xScale, ParaStation and xOPS
- Familiar user environment with Red Hat Enterprise Linux 9, Slurm and message passing tools
- EasyBuild and Apptainer support the scientific software stack, with EESSI prepared
- More than 20 public procurement benchmarks and more than 100 JUREAP applications
- European Federation Platform integration is staged, with Jupyter-JSC and UNICORE remaining access routes

ALICE RECOQUE


- Essentials: OCEAN admin stack for HPC and AI, Lustre storage, AMD ROCm and Bull BXI v3 support
- Workflows cover computing with artificial intelligence, in-situ processing and streaming data
- EESSI, Spack, containers and AMD/Bull tools for developers
- Benchmarking uses real workloads, full cost of ownership and energy-efficiency criteria
- European Federation Platform integration planned

Users Support

Part 4

User Support


1. Users support and training
2. AI user enablement on accelerators
3. Alternative access mechanisms beyond HPC
4. Industry and scientific use cases
5. Success stories and early adopters



JUPITER

Presenting: **Thomas Breuer – Jülich Supercomputing Centre**


JUPITER | 1. Users support and training



+ **Training:** 'HPC Introduction' (2x per year), 'AI on HPC systems' (4x per year) + many more


JUPITER | 2. AI user enablement on accelerators

- Dedicated **AI support** experts
 - e.g. performance analysis & optimisation
 - Dedicated **AI & accelerators training**:
 - 'AI on HPC systems' course (4x per year)
 - topic specific AI courses
 - GPU courses
 - Dedicated **AI documentation**:
 - e.g. recipes & guides
 - **Tools** development:
 - e.g. *LLview* job reporting
 - JUPITER AI Factory (**JAIF**) (www.jaif.ai)
 - e.g. services & knowledgebase
- 



JUPITER | 3. Alternative access mechanisms beyond HPC

- Hardware:
 - Several QC systems
 - JADE: available in 2027
 - JSC cloud -> JARVIS
- Software:
 - Interactive SC with *JupyterLab*
 - AI Inference Platform: Blablador
 - ...



*Depending on procurements and available functionality

**This is a subset of the available software

JUPITER | 4. Industry and scientific use cases

Project overview

Currently active or soon to be active projects from various research domains and industry:

- *EXT, REG, STA & AIF-SC: 48*
- *AIF PG & FL: 24*
- *BEN & DEV: 43*
- *further projects:*
 - *have already ended*
 - *are currently in review*
 - *applied via national calls*

DragonLLM: Industry use case example

- 7B Hybrid MoE LLM
- trained with 28 languages on JUPITER
- joining forces: talented startup & JSC's large-scale HPC/AI experiences
- measurable economic impact:
 - **DragonLLM was acquired by OVHcloud**



JUPITER | 5. Success stories and early adopters

- JUREAP: >100 projects; preparing on JEDI, JUWELS Booster (large-scale and full system runs during BigDays), early-access to JUPITER Booster

<https://www.fz-juelich.de/en/js/jupiter/jupiter-publications>


- Gordon Bell Prize for Climate Modeling (ICON)

<https://dl.acm.org/doi/10.1145/3712285.3771789>

- Quantum Simulation of 50 Qubits (JUQCS)

<https://arxiv.org/abs/2511.03359>

- World record in 340s walltime
- Earlier record: 48 Qubits
 - K computer (3102 s)
 - TaihuLight (8548 s)
 - JUPITER (286 s)



Phytoplankton Near-Surface Wind Air-Sea CO₂ Flux



ALICE RECOQUE
BEYOND HPC

Presenting: **Anaïs Gaertner - CEA** and **Corinne Béal - GENCI**

ALICE RECOQUE | 1. Users support and training

- A single point of contact to access user support: hotline.tgcc@cea.fr
 - Different teams of specialist : system, application
 - A new “Application Support Team (AST)” for specific needs
- A dedicated team of 20 experts for high level support starting Q2/2026
 - 12 FTE in French supercomputing national centres (TGCC/IDRIS/CINES)
 - 7 FTE from AMD and BULL, 0.5 FTE from SURF
 - Objective: long term support, training in AI/HPC/QC/HPDA
 - In complement to L1/L2 already in place @TGCC
- Companion projects:
 - NumPEX: porting software at exascale (72M€, 2021-2028).
 - Epicure, Minerva, AI2F, ...
- A wish to engage with all HPC / AI CoE in Europe



ALICE RECOQUE | 2. AI user enablement on accelerators



- Alice Recoque will be instrumental in AI Factory France alongside with national existing AI supercomputers (Jean Zay, Adastra)
- In 2025 GENCI **supported more than 1700 AI projects from academia & industry** (including leading French startups) for training and low scale inference (in AI For Science)
- Alice Recoque will foster the development at scale of European Frontier models in AI for Science, agentic AI and world models
- AI communities will access starting H2 2026 to early systems / prototypes to prepare the rise of Alice Recoque with support from AST members, AMD and Bull as well as AI Factory France







ALICE RECOQUE | 3. Alternative access mechanisms beyond HPC

Alice Recoque has the vision to address the ongoing convergence between HPC, AI and QC and develop innovative services

- TGCC Cloud using OpenStack instances : easy for making data/results available for (re)processing, strong interest from climate, SKA and CERN communities
- New agile access mechanisms for AI communities already proposed by GENCI since 6 years and proposed by EuroHPC
- TGCC Cloud for also for AI inference
- Stream based ingestion/inference of data from large scale instruments (satellites, telescopes, sequencers, network of sensors, IoT) toward digital twins
- And finally access to quantum computers/simulators acting as accelerators of HPC/AI workloads

ALICE RECOQUE | 4. Industry & scientific use cases



ALICE RECOQUE | 5. Success stories & early adopters

Download the French Scientific Case toward Exascale here ->
more than 70 scientific and industrial applications



Alice Recoque will arrive in 2027 but already

- Alice Recoque AST will start in Q2 2026
- An inception programme will rely on the AST & AI Factory France teams and access to production / prototype systems for preparing users communities
- Alice Recoque will benefit from the experience of other EuroHPC HostingEntities
- Alice Recoque will also provide new services toward federated learning (with Jupiter), ingestion/inference of instrumental data, link with quantum computers...

User Support | Summary Slide

JUPITER

- User support by a large number of experts in HPC & AI and through the JUPITER AI Factory (JAIF)
- Regular HPC and AI training courses
- Wide variety of access routes: EFP, secure shell, JupyterLab, Cloud, AI inference, QC, ...
- Active project portfolio across EuroHPC & GCS national calls for science & industry
- Early-access success stories, including a Gordon Bell Prize for climate modelling and a 50-qubit quantum simulation

ALICE RECOQUE

- Inception programme with 20 experts from Q2 2026, drawing on the AST and AI Factory France ahead of 2027 production
- More than 70 multi-industrial and scientific use cases in the French Scientific Case toward Exascale
- New agile access routes: TGCC Cloud, streaming ingestion toward digital twins, and quantum computers as HPC/AI accelerators
- Instrumental in AI Factory France alongside Jean Zay and Adastra, building on more than 1700 AI projects supported in 2025

Panel discussion with Q&A



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



THANK YOU!