Opportunities on EuroHPC JU systems for AI applications

Vangelis Floros, EuroHPC JU

26 September 2023

EuroHPC Supercomputers

- *Empowering European Scientific Research, Academia, Industry & SMEs*
- *Providing the necessary computing power to accelerate discovery and innovation in Europe*

EuroHPC Infrastructure activities

- Hosting Entity Selection
- System Procurements
- Operation & Monitoring
- Access Time allocation
- Hyperconnectivity
- Federation
- High-Level Application Support

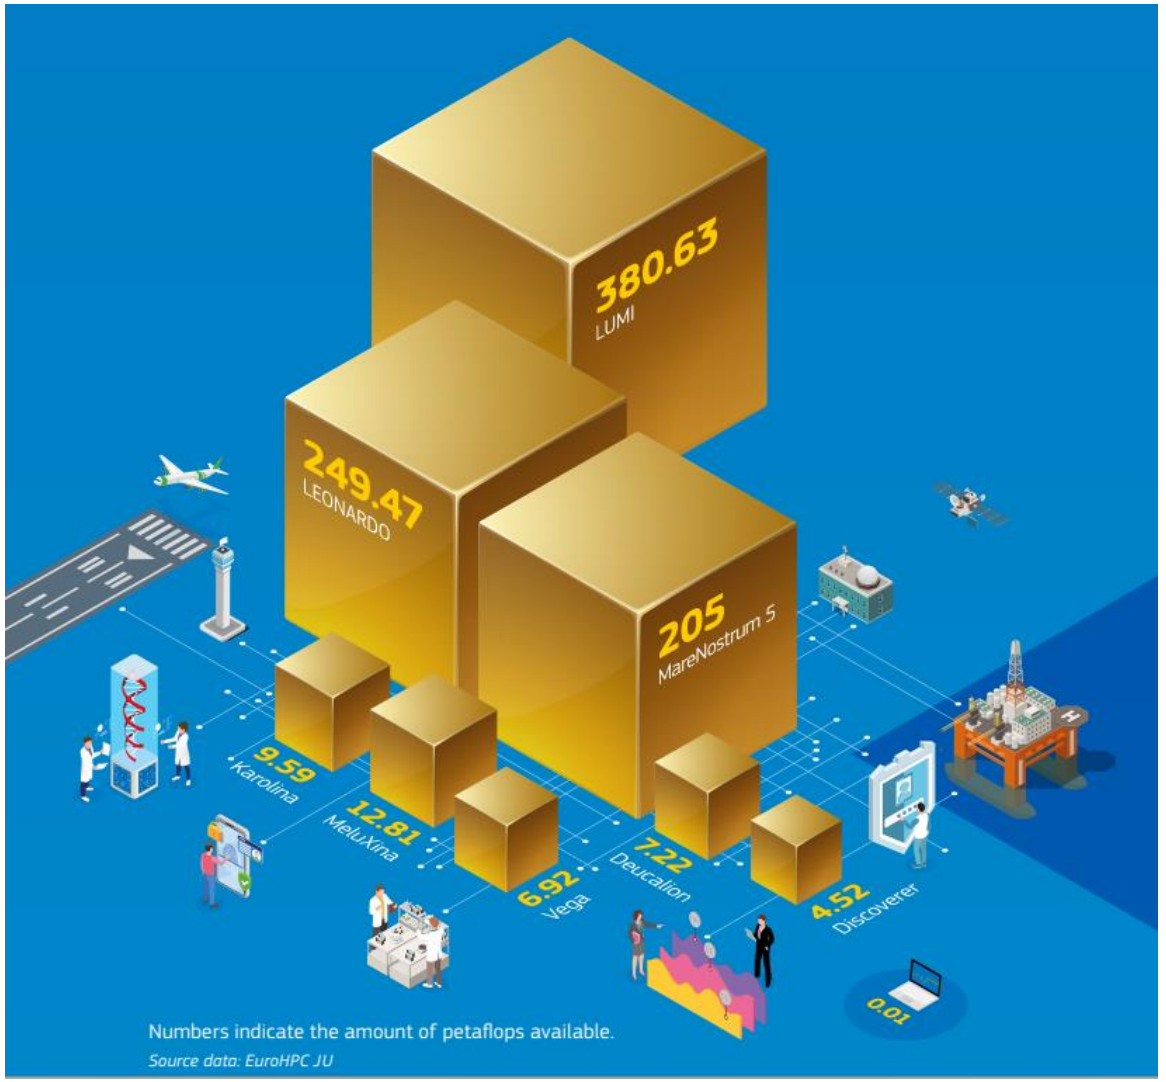

EuroHPC Infrastructure 2023

Pre-exascale
(owned by EuroHPC, operated by national Hosting Entity)

Petascale
(co-owned with national Hosting Entity)

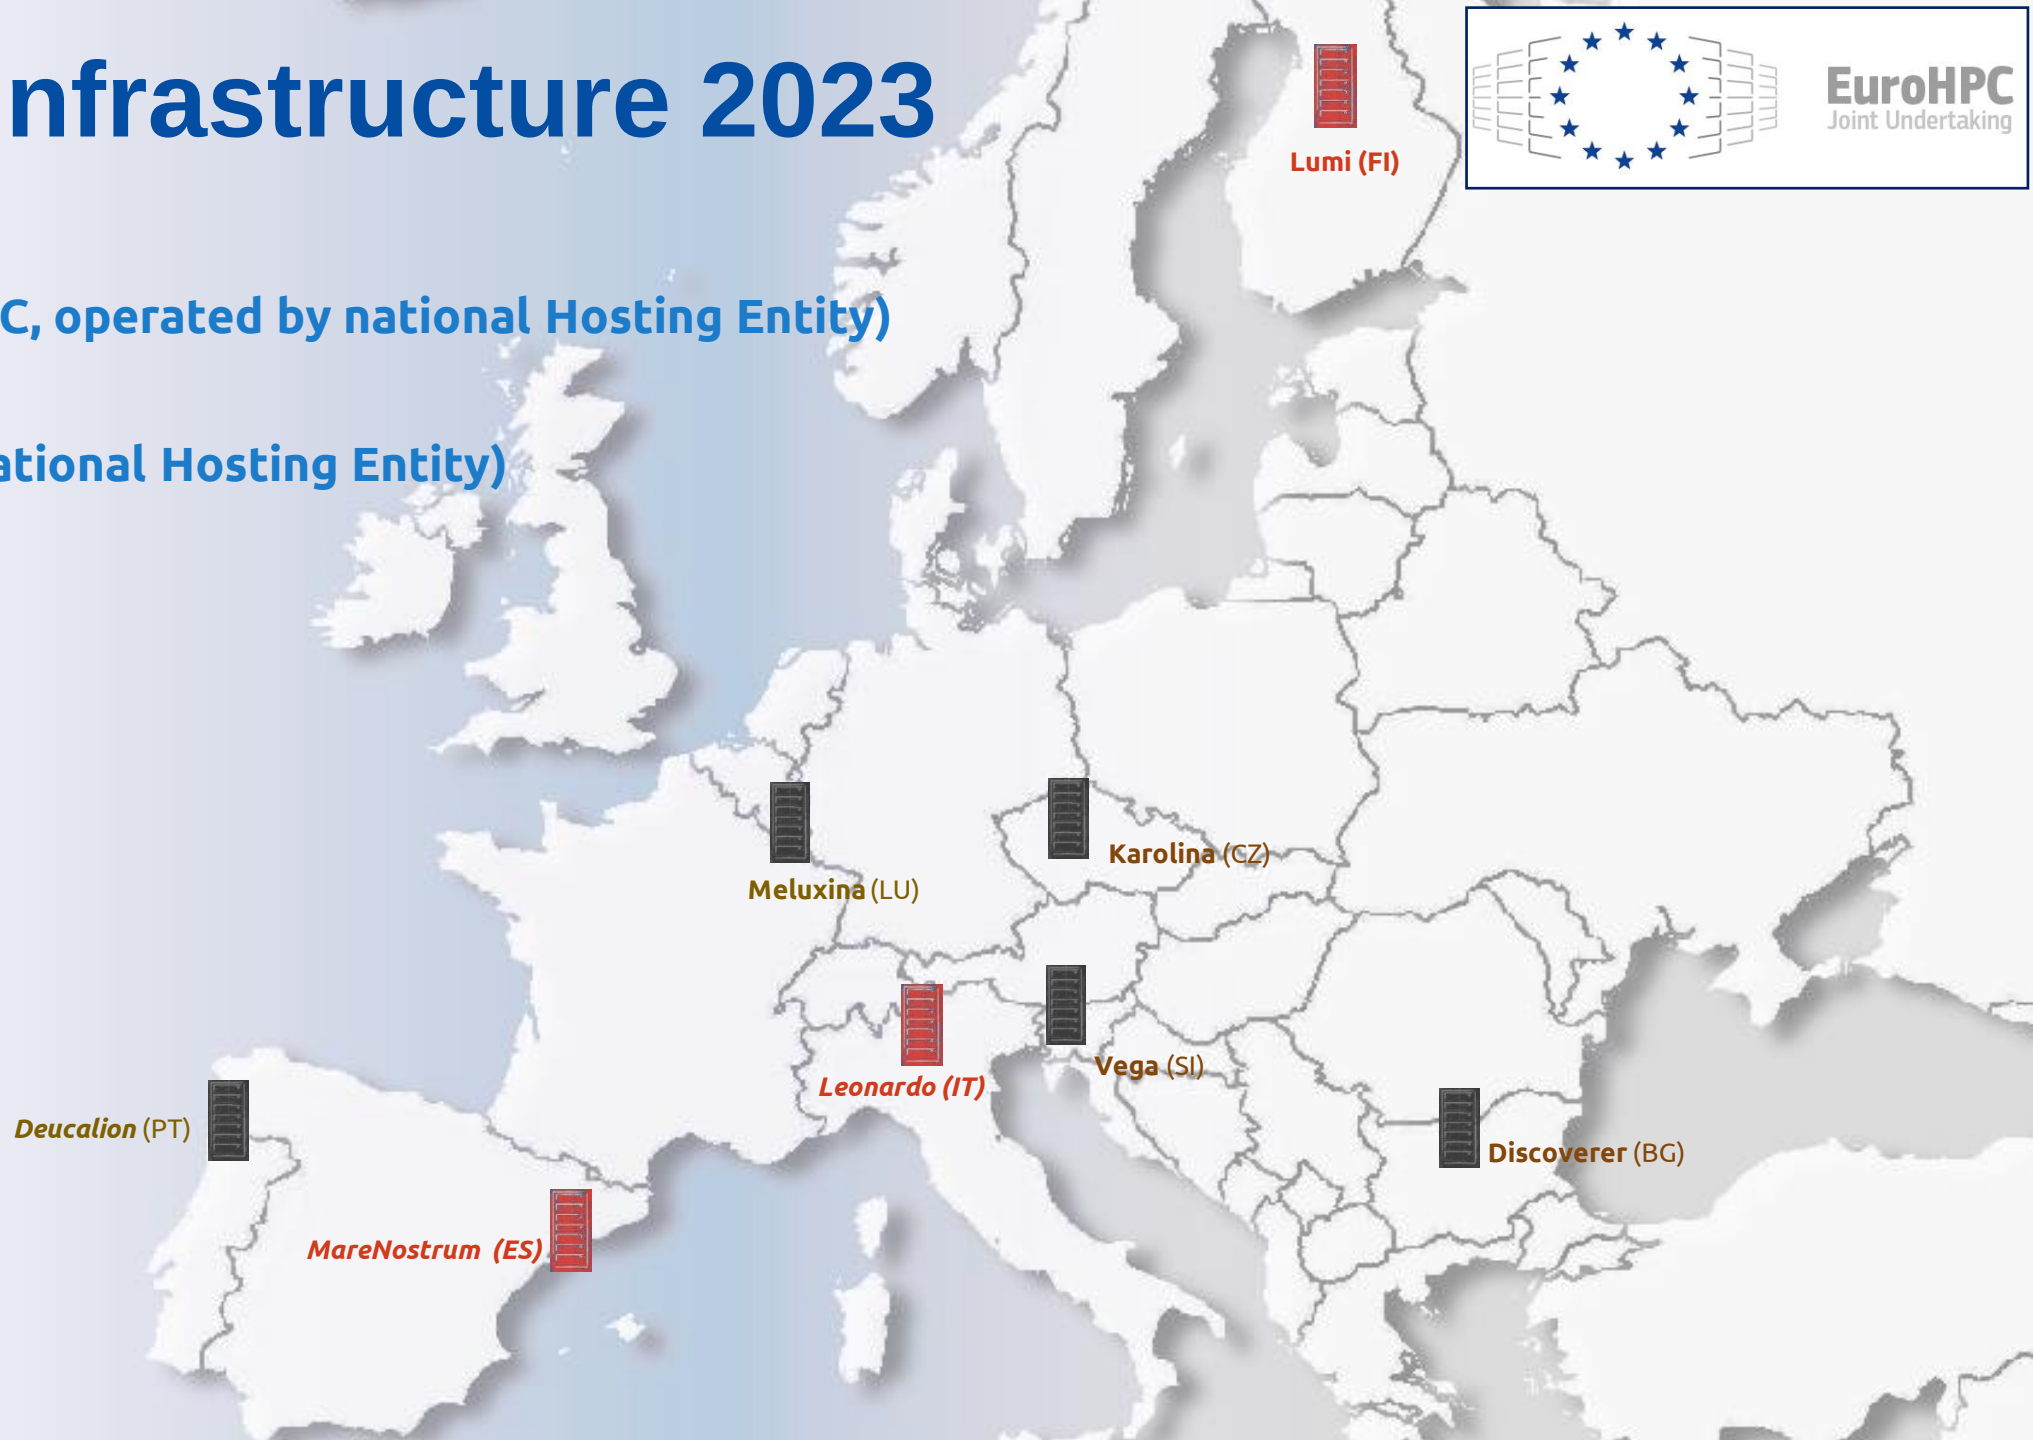

EuroHPC systems | Pre-exascale

LUMI (Lead by CSC)
Kayaani, Finland

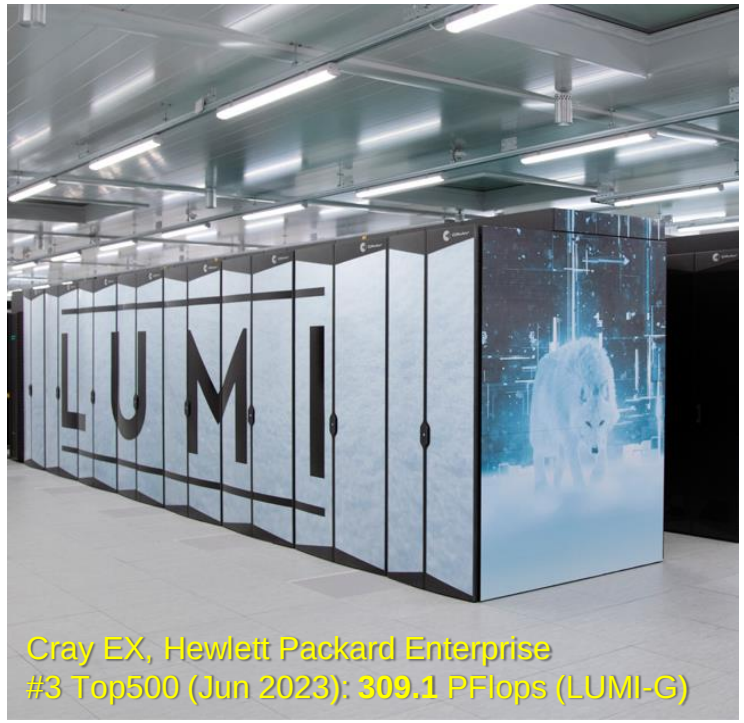

Cray EX, Hewlett Packard Enterprise
#3 Top500 (Jun 2023): **309.1** PFlops (LUMI-G)

- 4.976 Nodes (2,928 GPU + 2,048 CPU)
- 11,712 GPUs (AMD MI250X)
- Slingshot Interconnect (200 Gb/s)
- 117 PB Storage (Lustre + Ceph)

AMD platform
- CPU: 64-core AMD EPYC™
- GPU: AMD Instinct™ (MI250X)

Leonardo (Lead by CINECA)
Bologna, Italy

Atos BullSequana XH2000
#4 Top500 (June 2023): **238.7** PFlops (BOOSTER)

- 4,992 Nodes (3,456 GPU + 1,536 CPU)
- 13,842 GPUs (Nvidia A100)
- Quad-rail Infiniband HDR (200 Gb/s)
- 110 PB Storage (Lustre)

Intel/NVidia platform
- CPU: Intel Sapphire Rapids (56-core), Intel Ice Lake (32-core)
- GPU: Nvidia custom Ampere (A100)

MareNostrum 5 (Lead by BSC)
Barcelona, Spain

Atos BullSequana XH3000 / Lenovo Think Systems
205 PFlops (projected sustained - aggregate)

- 7,528 Nodes (1,120 GPU + 6,408 CPU)
- 4,480 GPUs (Nvidia H100)
- Quad-rail Infiniband NDR200
- 250 PB Storage (GPFS)

Intel/NVidia platform
- CPU: Intel Sapphire Rapids (56-core), Intel Sapphire Rapids (32-core),
- GPU: Nvidia Hopper (H100)

EuroHPC systems | Petascale

Sustained performance:	**6,9 petaflops**
CPU:	AMD Epyc Ro
GPU:	Nvidia A100
TOP500 ranking:	#32 in EU; #106 globally (June 2021)
Vendor/model	Atos BullSequana XH2000
Operated by	IZUM, Maribor, Slovenia

Sustained performance:	**12,8 petaflops**
CPU:	AMD Epyc Rome
GPU:	Nvidia A100
TOP500 ranking:	#10 in EU; #36 globally (June 2021)
Vendor/model	Atos BullSequana XH2000
Operated by	EuroCC National Competence Centre Luxembourg

Sustained performance:	**9,13 petaflops**
CPU:	AMD Epyc Rome
GPU:	Nvidia A100
TOP500 ranking:	#20 in EU; #69 globally (June 2021)
Vendor/model	HPE Apollo 2000Gen10 Plus
Operated by	IT4I, Ostrava, Czech Republic

Sustained performance:	**4,45 petaflops**
CPU:	AMD Epyc Rome
GPU:	Coming soon
TOP500 ranking:	#27 in EU; #91 globally (June 2021)
Vendor/model	Atos BullSequana XH2000
Operated by	PSB consortium, Sofia, Bulgaria

Petascale systems in numbers

43.83 Petaflops sustained (54,41 Petaflops R_{peak})

- **14** partitions
- **5533** CPU Nodes
- **365** GPU Nodes
- **24PB** Lustre Storage
- **6802** AMD EPYC Rome CPUs / **1632** Fujitsu ARM A64FX CPUs
- **1748** Nvidia A100 GPUs
- Other: FPGA, Visualisation and Cloud capabilities

Sustained performance:	**7,22 petaflops (projected)**
CPU:	Fujitsu ARM A64FX, AMD Epyc Rome
GPU:	Nvidia A100
TOP500 ranking:	(TBC)
Vendor/model	Fujitsu FX700, Atos BullSequana
Operated by	MACC, Univ. Minho, Portugal

JUPITER Exascale System

Hosted by Julich Supercomputing Center (Germany)

- ***1 Exaflops*** *sustained (HPL) performance*
- *Implementing a dynamic* ***Modular Supercomputing Architecture (MSA)***
- *Hosted in* ***containerised*** *data center*
- *Integration of European hardware*

Procurement status

- *Competitive dialogue (now in Tendering Phase).*
- *Total budget:* ***273 Million Euro*** *(including options)*
- *Contract signature target:* ***Q4 2023***
- *Start of installation:* ***Q1 2024***
- *Acceptance (Phase 1):* ***Q4 2024***

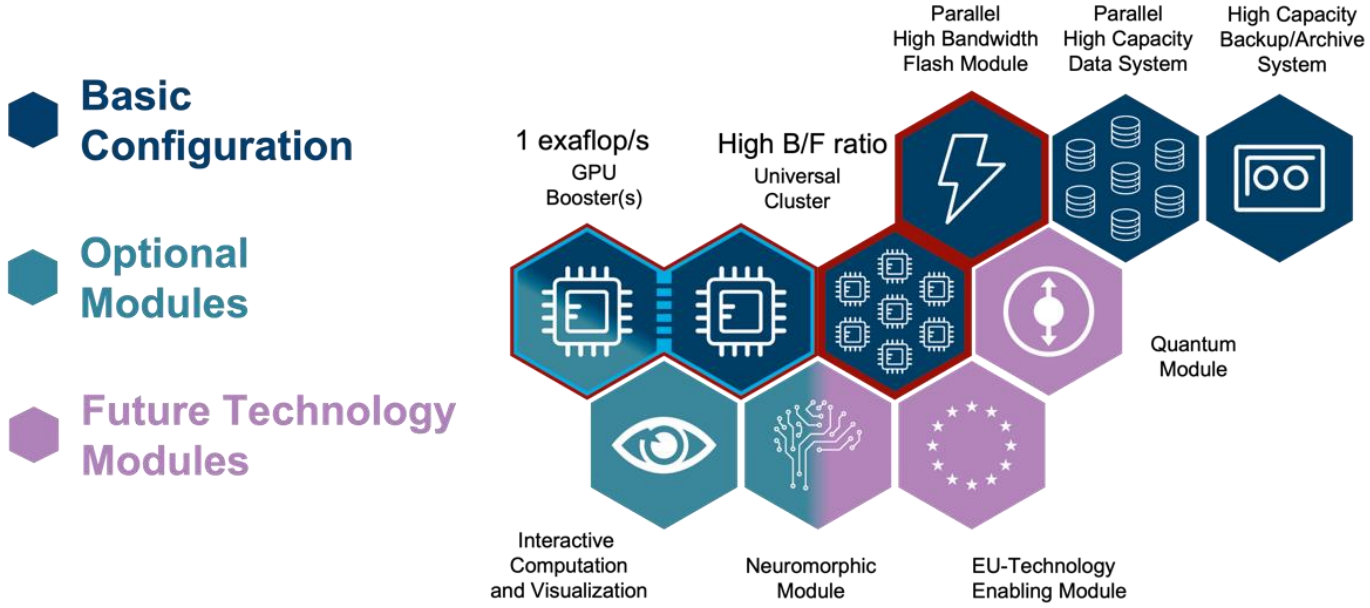

ACCESS TO EUROHPC SUPERCOMPUTERS

EuroHPC Joint Undertaking

WHO IS ELIGIBLE?

- Academic and research institutions (public and private)
- Public sector organisations
- Industrial enterprises and SMEs
- Established in the EU or H2020 affiliated country

→ Open to all fields of research

WHICH TYPES OF ACCESS EXIST?

- Regular access
- Extreme scale access
- Benchmark access
- Development access

…

WHAT ARE THE CONDITIONS FOR ACCESS?

Access is free of charge. Participation conditions depend on the specific access call that a research group has applied to.

In general users of EuroHPC systems commit to:

- use computing resources primarily for research and innovation
- acknowledge the use of the resources in their related publications
- contribute to dissemination events
- produce and submit a report after completion of a resource allocation

More information on EuroHPC access calls available at: https://eurohpc-ju.europa.eu/participate/calls_en

Access Policy

Access Policy v1.1 as adopted by the EuroHPC GB

- **4 Access Modes offering resources - continuously open with periodic cut-off dates.**
  - **Extreme scale:** Large applications, 2xYear. Peer-reviewed. Separate track for industry.
  - **Regular:** Medium to large applications, 3xYear. Peer-reviewed. Separate track for industry.
  - **Development.** All systems. Up to 1 year access. Monthly cut-offs
  - **Benchmark.** All systems. Up to 3 months access. Monthly cut-offs

Commercial Access (*)

- **Pay-per-use model**
- **No restrictions for open research – applications for civilian purpose**

New revision end of 2023

- **PRACE supports EuroHPC** in the implementation of the Access Policy!

Visit https://pracecalls.eu/ and https://prace-ri.eu/hpc-access/eurohpc-access/

AI Application Considerations

EuroHPC Supercomputers provide the perfect platform to AI applications

- **AI codes may need to be tested/ported on EuroHPC systems before allocations – Consider Benchmark/Development calls**

EuroHPC allocations are project based – fixed period of time – predefined usage schedule

- **Not for production usage**
- **Appropriate for research and for demanding model training but not for (production) inference runs**

EuroHPC Supercomputers are multitenancy environments

- **Applications run as jobs, submitted through a shared queuing system (SLURM) - Large allocations my take time to start running**
- **Jobs typically run for max 48hrs – Large runs require implementation of snapshotting functionality**

EuroHPC Supercomputers provide high-speed connectivity to the external world (x100 Gbit links to GEANT), **however**:

- **Large data transfers need to be coordinated with the hosting site**

EuroHPC Supercomputers provide large storage capabilities, **however**:

- **No archiving / long-term storage**
- **Extremely large storage requirements need to be agreed with the hosting site**

Application Support Teams (ASTs)

EuroHPC High-Level support teams to provide Level-2 and -3 application support:

- **Application Enabling**
- **Code Scaling and Optimisation**
- **Best practices and guidelines**
- **Training**
- **Specific consideration for AI applications**

EPICURE Project

- **4-year duration – starting February 2024**
- **15 beneficiaries (all EuroHPC Hosting Entities)**
- **10 Million Euro budget (50% EC co-funding – Digital Europe Program)**
- **Coordinated by CSC**

Thank you!

Keep up with EuroHPC news:

https://eurohpc-ju.europa.eu @EuroHPC_JU EuroHPC Joint Undertaking

EuroHPC systems 2023-2025

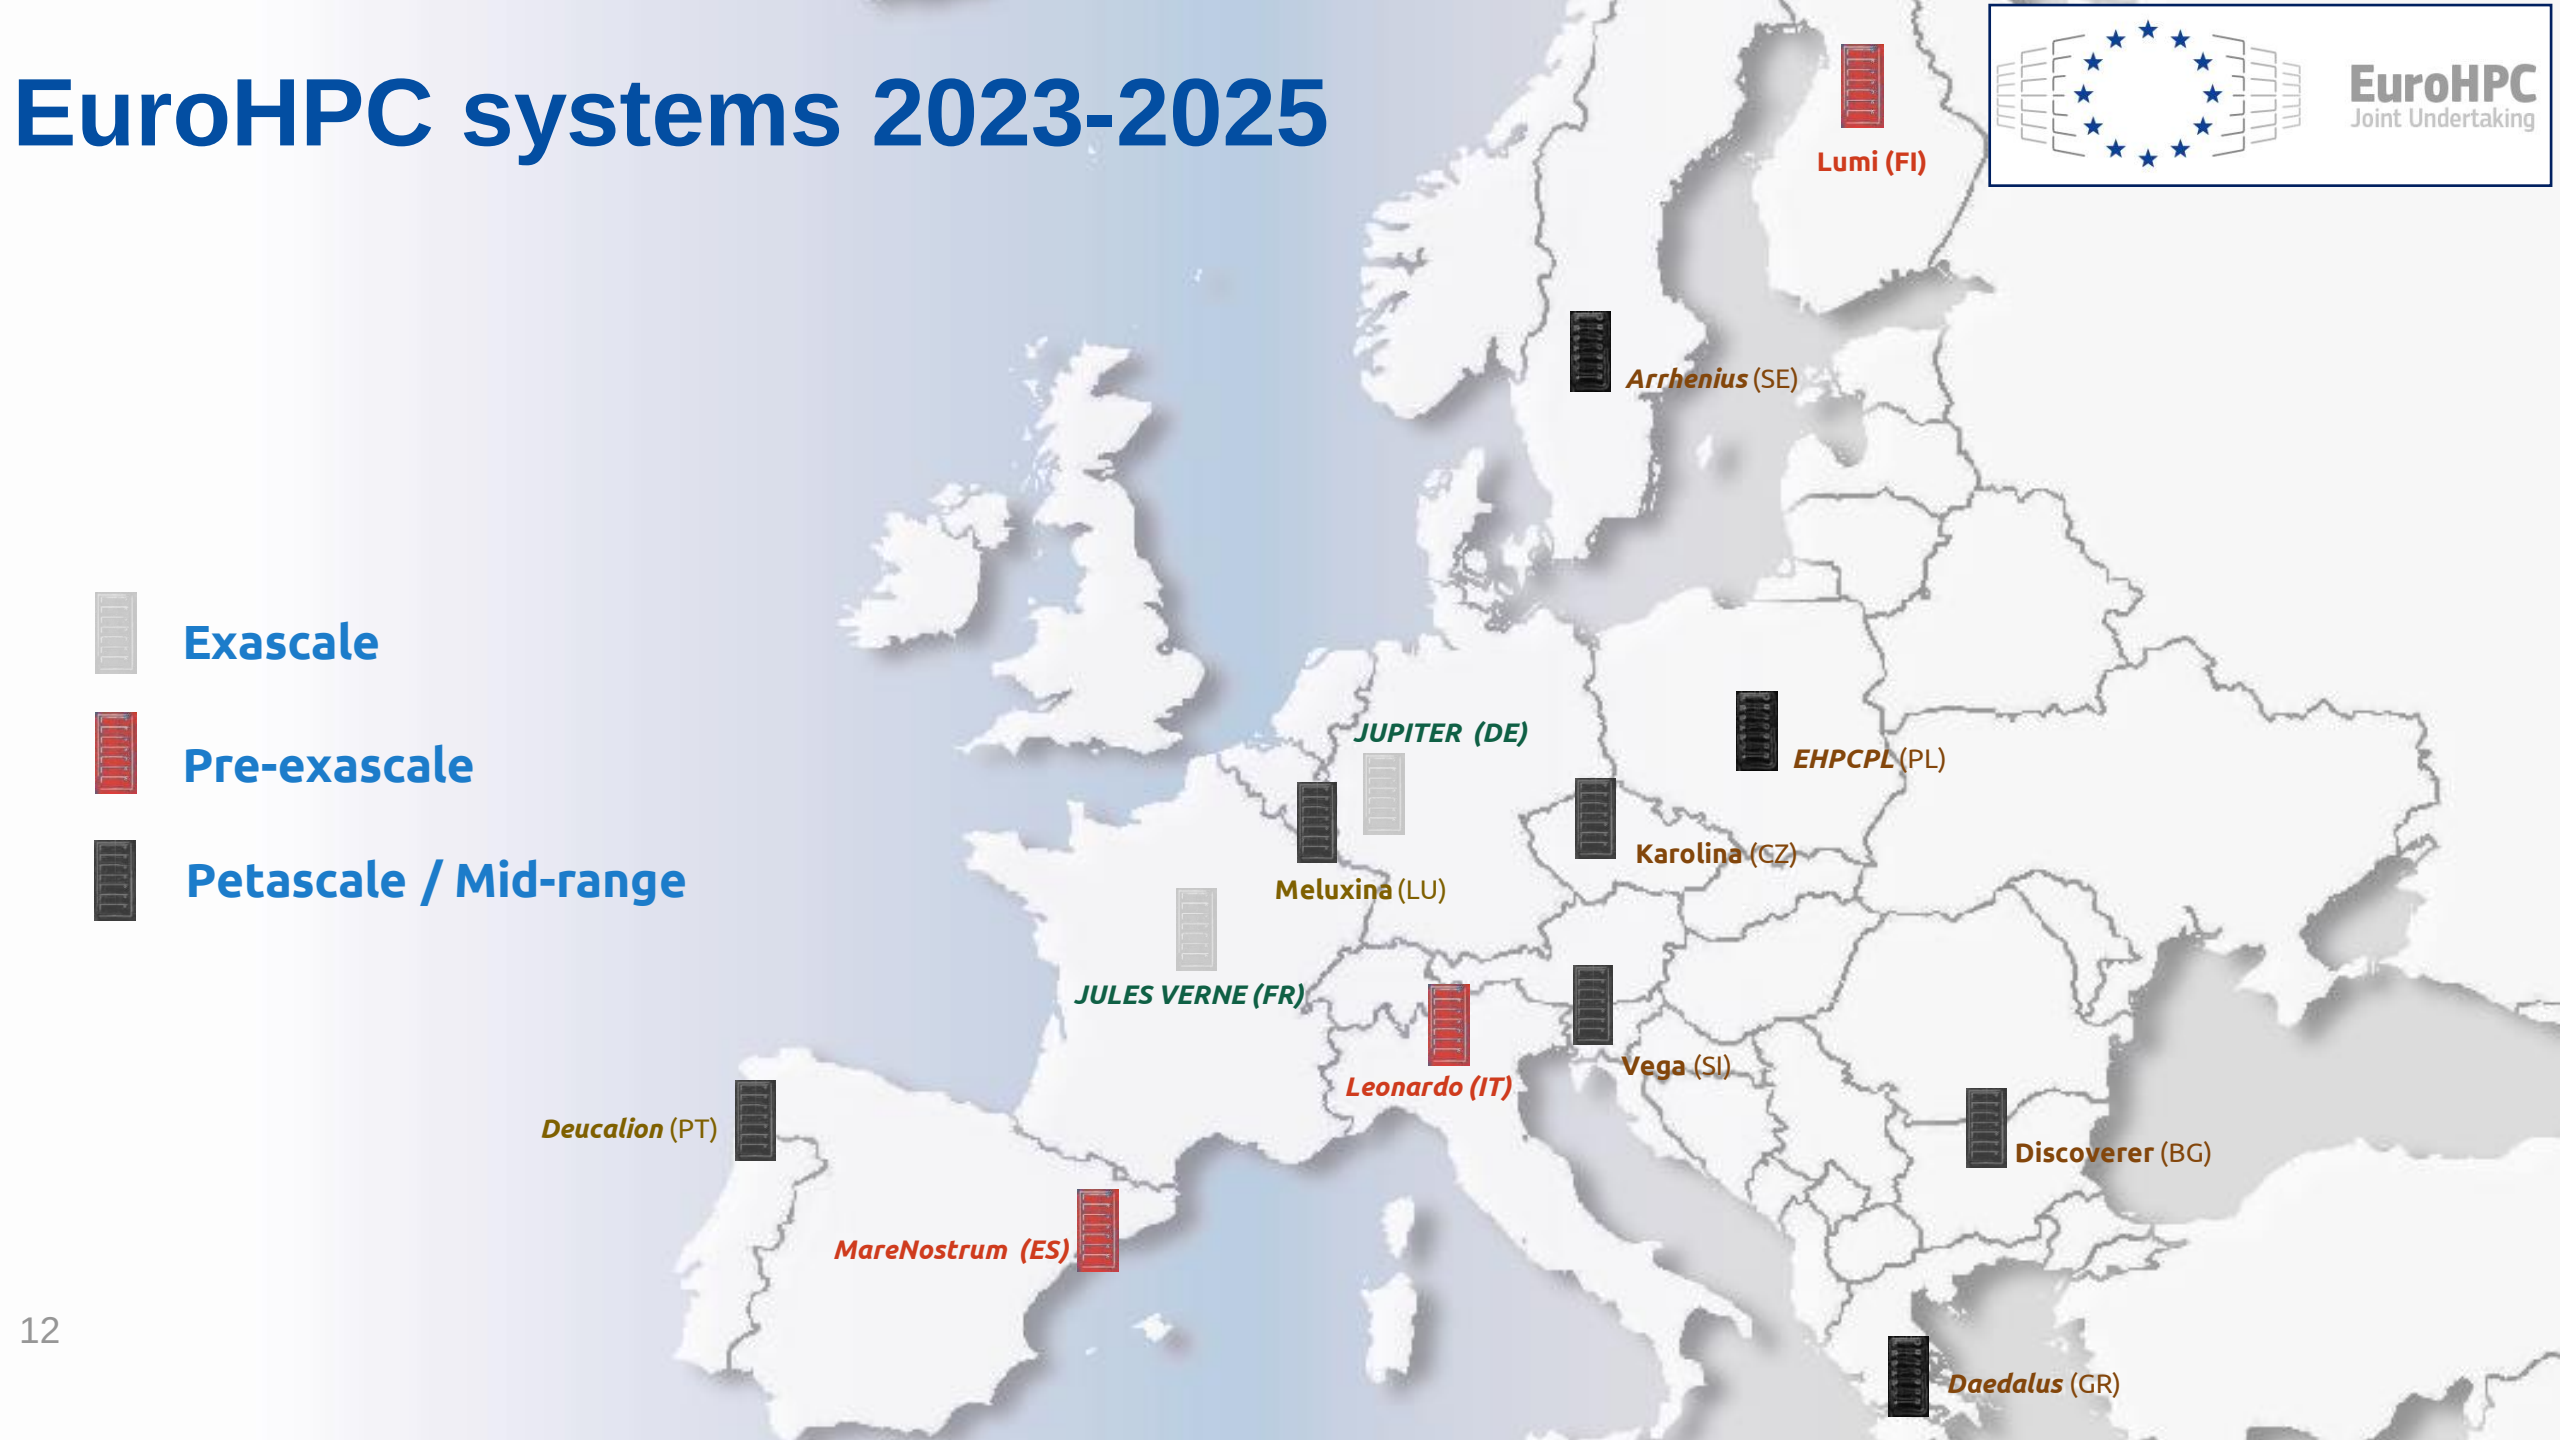

Federation & Hyperconnectivity

Federate HPC resources accross all EuroHPC systems

- **Authentication, Authorization and Identification services (AAI)**
- **Computing services**
  - Interactive Computing
  - Cloud access – Virtual Machines - Containers
- **Data services**
  - Archival Services and Data repositories
  - Data mover / transport services
- **User and Resource management**

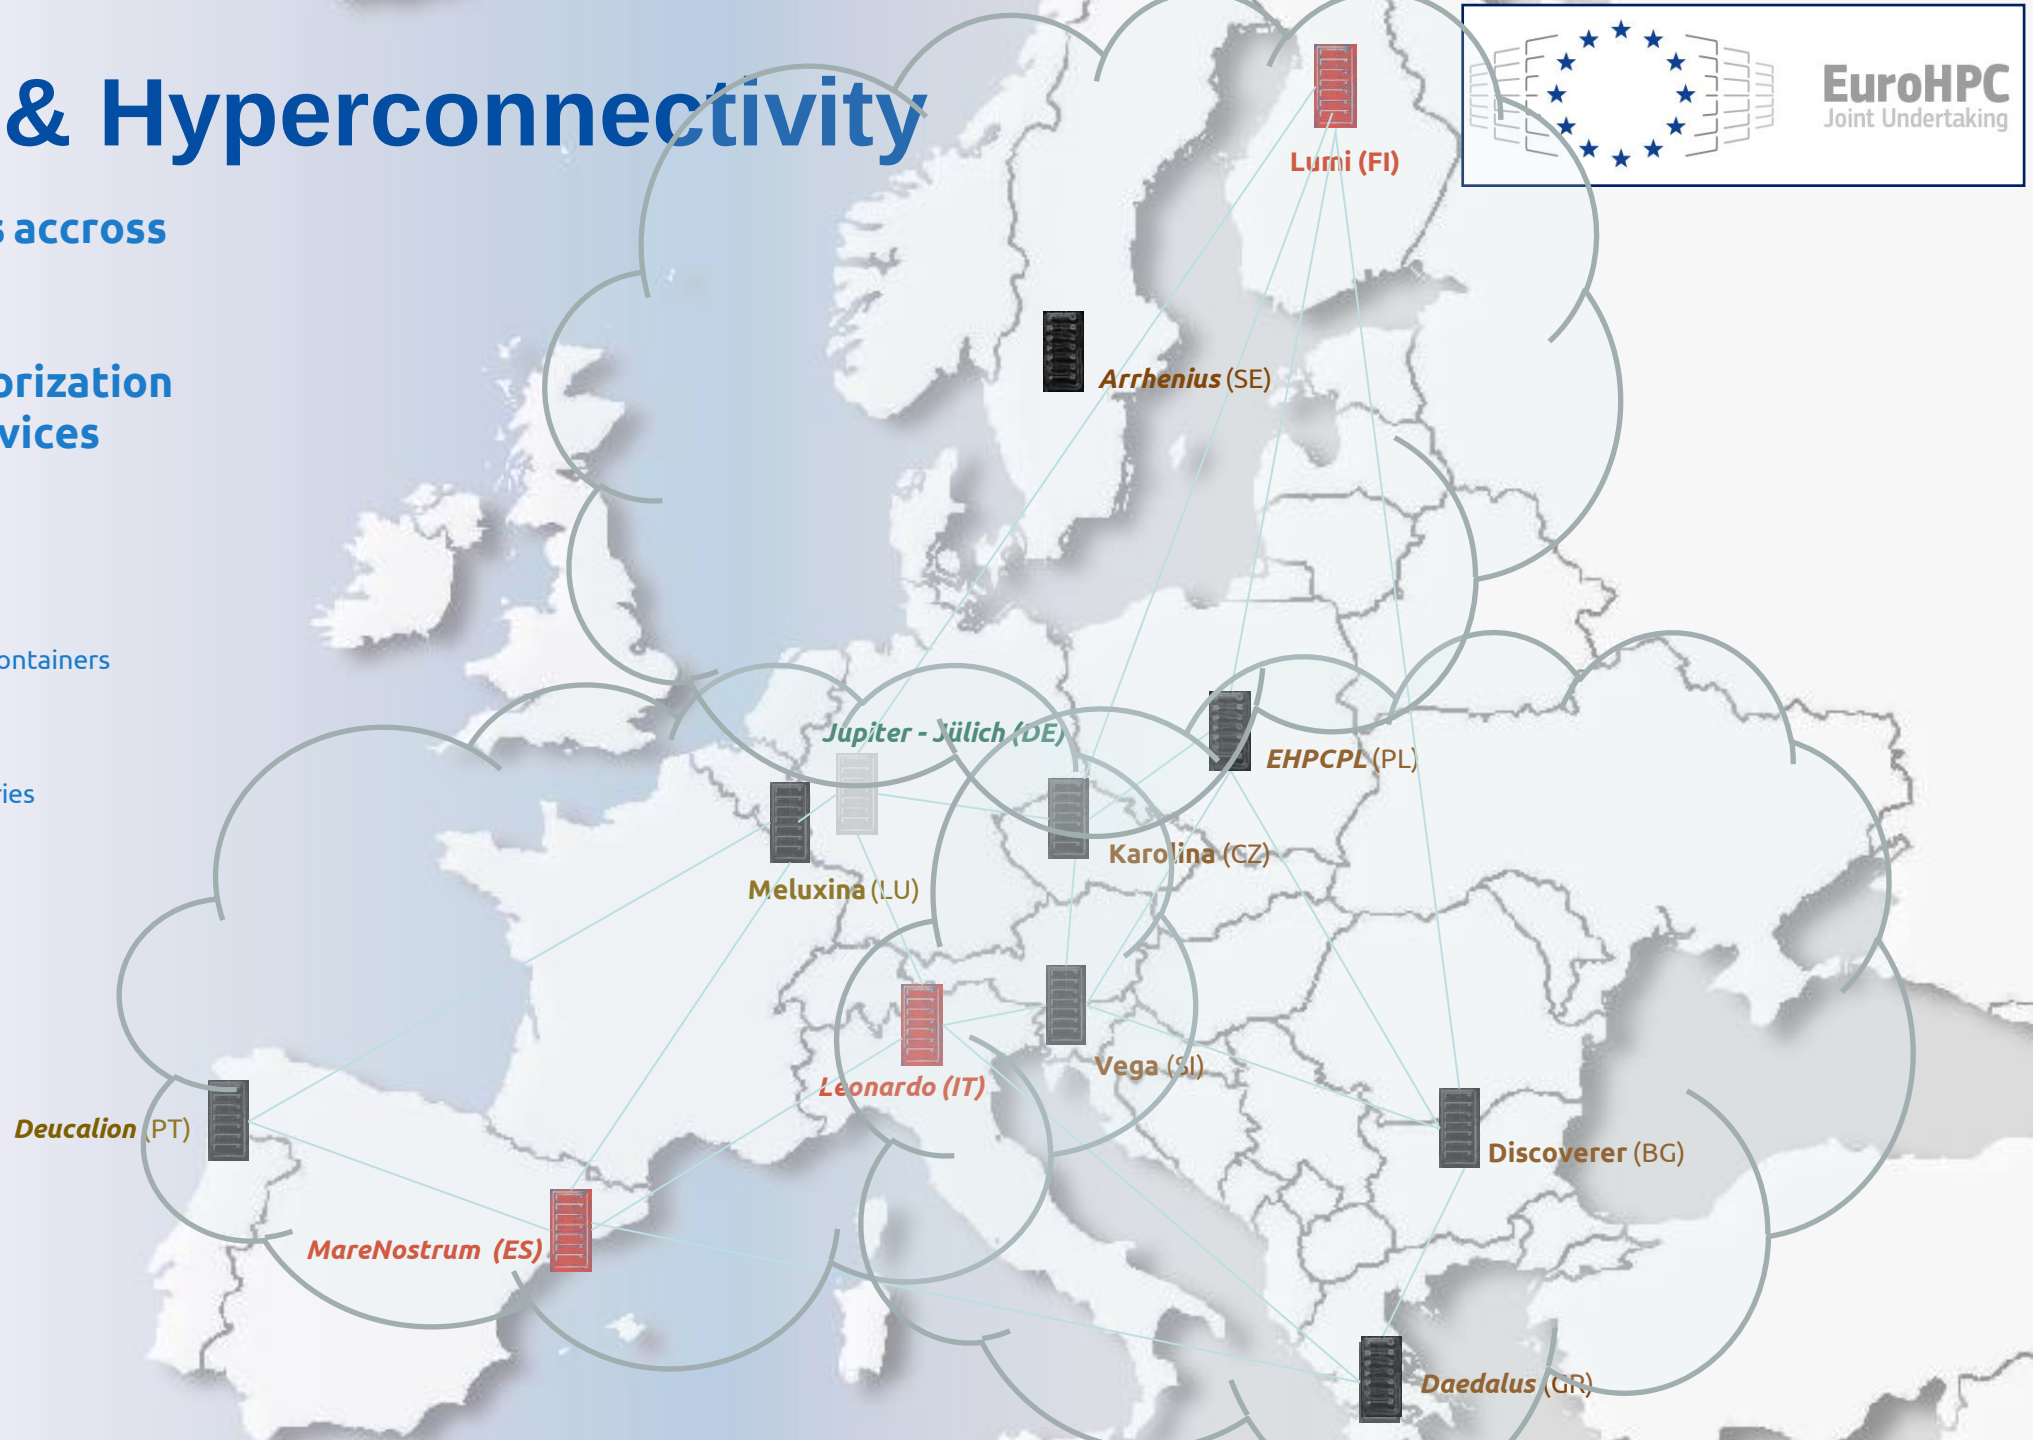

EuroHPC systems – expected timeline

(ranking top500 – not installation)

Deucalion

LUMI (#3)
CSC, Kajaani, FI
375/6.3 PF, HPE Cray EX
(AMD RI/AMD EPYC)

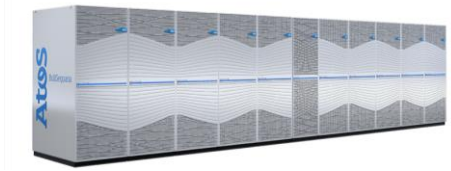

Leonardo (#4)
CINECA, Bologna, IT
240/9.5 PF, Atos
(NVIDIA A100/Intel SPR)

MareNostrum 5
BSC, Barcelona, ES
***(205 PF)**, Atos*
163/4 PF (Accelerated)
(NVIDIA Hopper/((?))
36/2 PF (GP)
(Intel SPR/NVIDIA Grace)

CASPIr
Daedalus
EHPCPL
Levente
Arrheniu

JUPITE

***JUPITER** (end of Sept)*
Jülich (DE)
(1 EF)
1 EF (Accelerated)
(-)
+8 PF (GP)
(-)

Jun Jul Aug Sep Oct Nov Dec Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan Feb Mar

2023 2024 2025