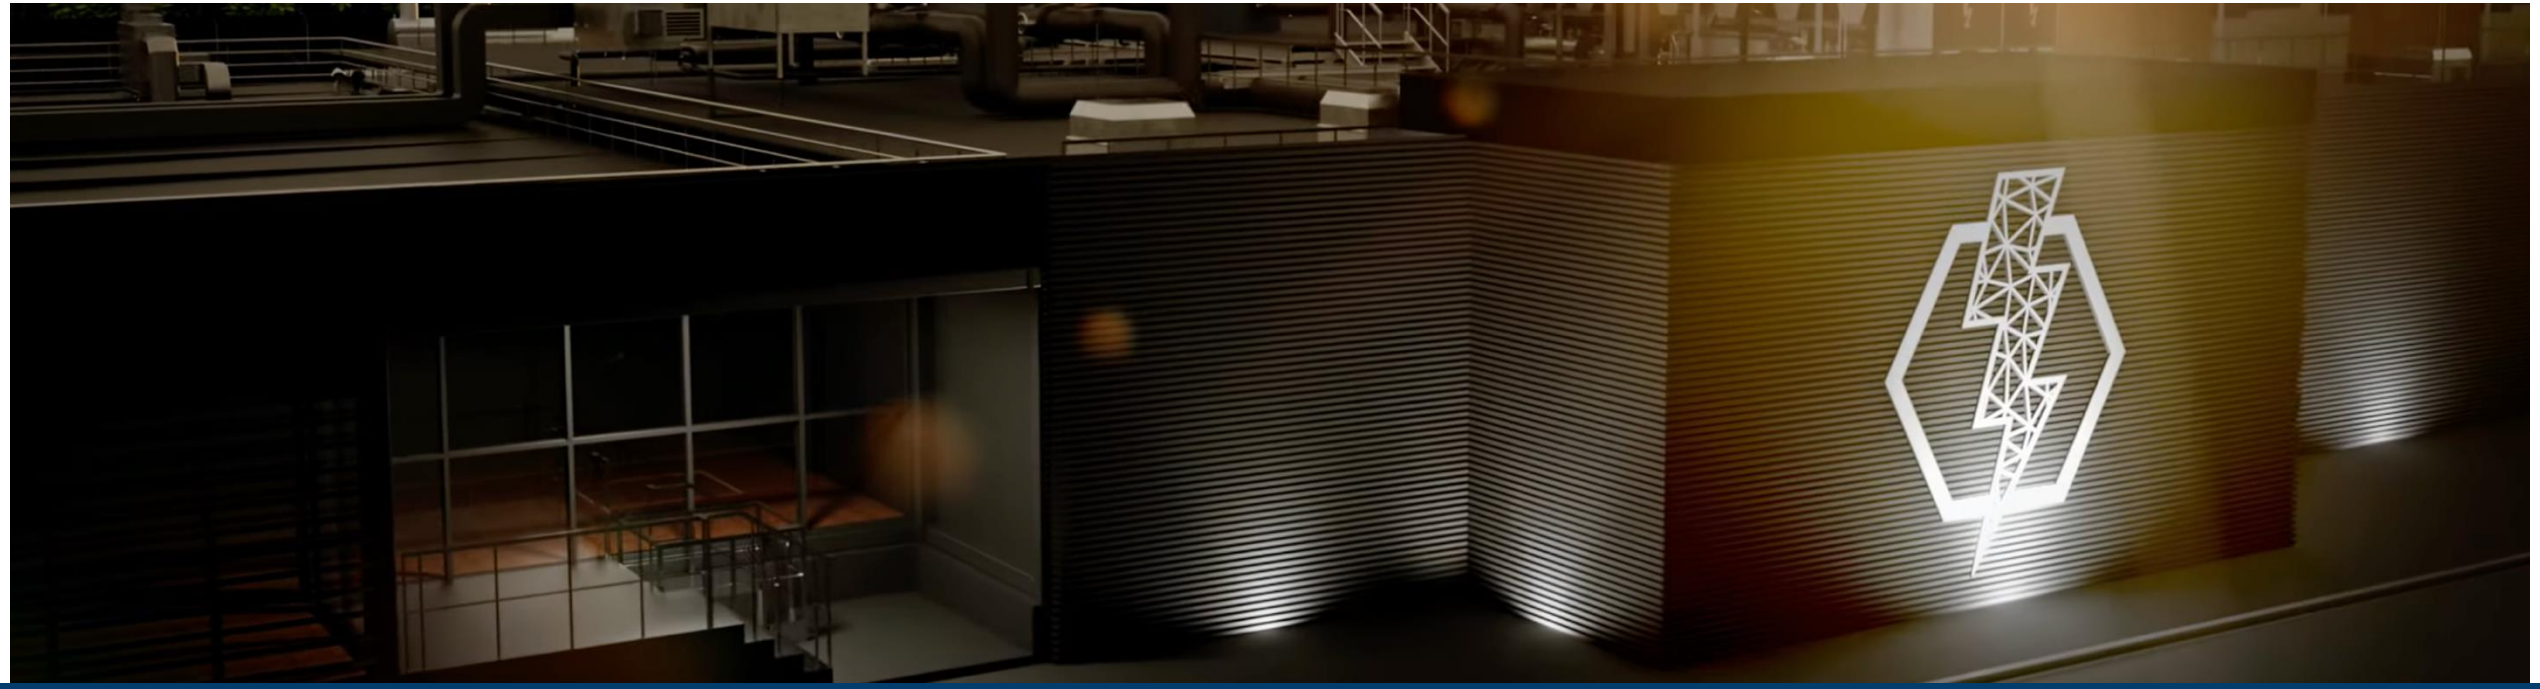

JUPITER - FROM LAUNCH TO LANDING

Accelerated Data Analytics and Computing Symposium (ADAC17)

2025-09-03 I JÜLICH SUPERCOMPUTING CENTRE | BENEDIKT VON ST. VIETH

Member of the Helmholtz Association

EuroHPC Joint Undertaking

Federal Ministry of Research, Technology and Space

Ministry of Culture and Science of the State of North Rhine-Westphalia

GCS Gauss Centre for Supercomputing

JÜLICH Forschungszentrum Shaping Change

A LONG TIME AGO

-

BECOMING A HOSTING ENTITY AND PROCURING JUPITER

EUROHPC AND SMART SCALING@GCS

Discussions and plannings for HPC systems funded directly by European Commission

EuroHPC Established
17.12.2021

Gauss Centre for Supercomputing Smart Scaling Strategy Developed
Exascale for JSC, HLRS and LRZ

JÜLICH Forschungszentrum
Shaping Change

Datacenter Studies and Campus Preparation planning

- New Datacenter planning, especially cooling/electricity
- Cost estimates
- New Grid connection and transformator station plannings

JUPITER – HOSTING ENTITY DECISION 15.06.2022

tagesschau

Startseite ▸ Wirtschaft ▸ Technologie ▸ Hochleistungs-Rechner: Supercomputer "Jupiter" kommt nach Jülich

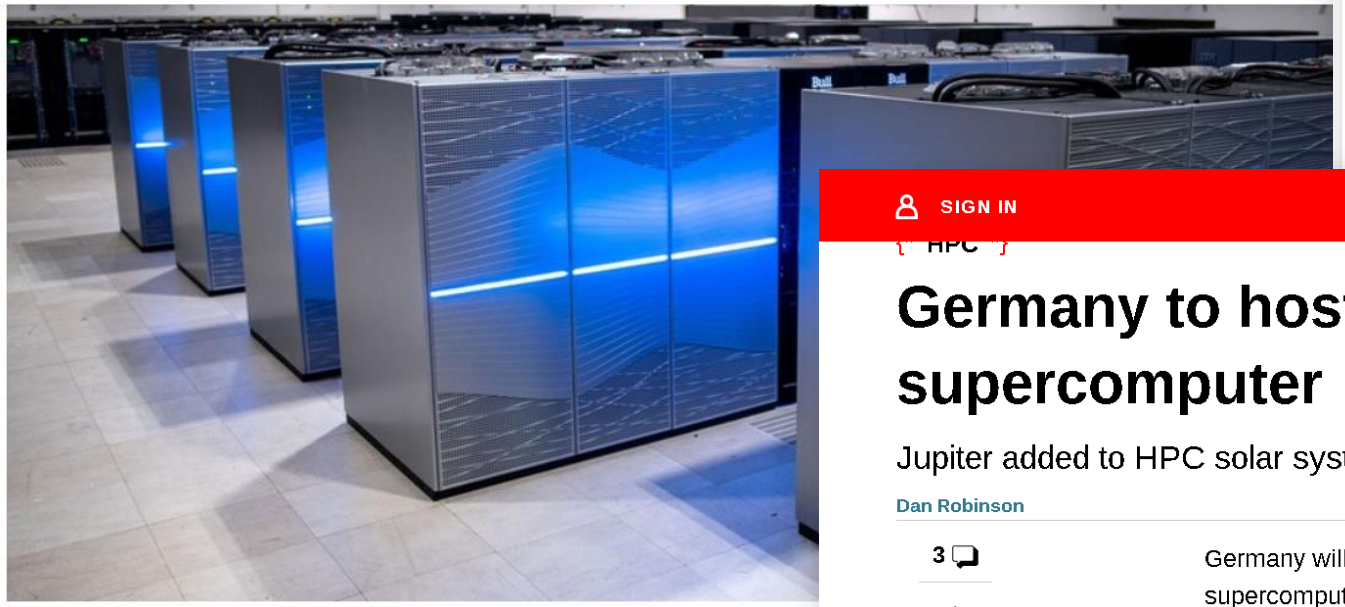

Hochleistungs-Rechner

Supercomputer "Jupiter" k

Stand: 15.06.2022 16:43 Uhr

**Das Forschungszentrum Jülich wird Standor
Exascale-Computers. "Jupiter" soll die Schal
Rechenoperationen in der Sekunde durchbr**

SPIEGEL Netzwelt

»Jupiter«

Jülich bekommt Europas ersten Exascale-Supercomputer

Das Forschungszentrum Jülich bekommt für eine halbe Milliarde Euro einen neuen Vorzeigerechner. Er soll helfen, Fragen zum Klimawandel zu beantworten – mit mehr als einer Trillion Rechenoperationen pro Sekunde.

15.06.2022, 16.52 Uhr

SIGN IN

The Register

{ HPC }

Germany to host Europe's first exascale supercomputer

Jupiter added to HPC solar system

Dan Robinson

Thu 16 Jun 2022 // 07:33 UTC

3

Germany will be the host of the first publicly known European exascale supercomputer, along with four other EU sites getting smaller but still powerful systems, the European High Performance Computing Joint Undertaking (EuroHPC JU) announced this week.

Germany will be the home of Jupiter, the "Joint Undertaking Pioneer for Innovative and Transformative Exascale Research." It should be switched on next year in a specially designed building on the campus of the Forschungszentrum Jülich research centre and operated by the Jülich Supercomputing Centre (JSC), alongside the existing Juwels and Jureca supercomputers.

LAYING THE FOUNDATION FOR TAKEOFF (2022)

- **Numerous calls/meetings/discussions**
- **Datacenter to Modular HPC Datacenter decision**
- **Preparation of the Descriptive Document to start procurement**

June | July | August | September | October | November | December

Hosting Entity Decision
15.06.2022

Hosting Agreement
14.12.2022

PREPARING FOR LAUNCH

- Mission planning
  - Preparing descriptions, conditions, requirements, evaluation
  - Regular meetings
  - Started already early in year
  - Location: **Earth**
- Target: **JUPITER**
  - Booster
  - Cluster
  - Storage
  - (Machine Hall)

Ready for take off

16. Jan 2023: Publish Call (Descriptive Document)

17. Feb 2023: Deadline for Request for Participation

22. Feb - 17. Mar 2023: Evaluation, Notification

4. Apr 2023: First Dialogue

3.-5. May 2023: Second Dialogue

1. June 2023: Invitation to Tender

3. July 2023: Deadline for final Tender

3.-7. July 2023: Evaluation by Technical Experts

until 20. Aug 2023: Governing Board Decision

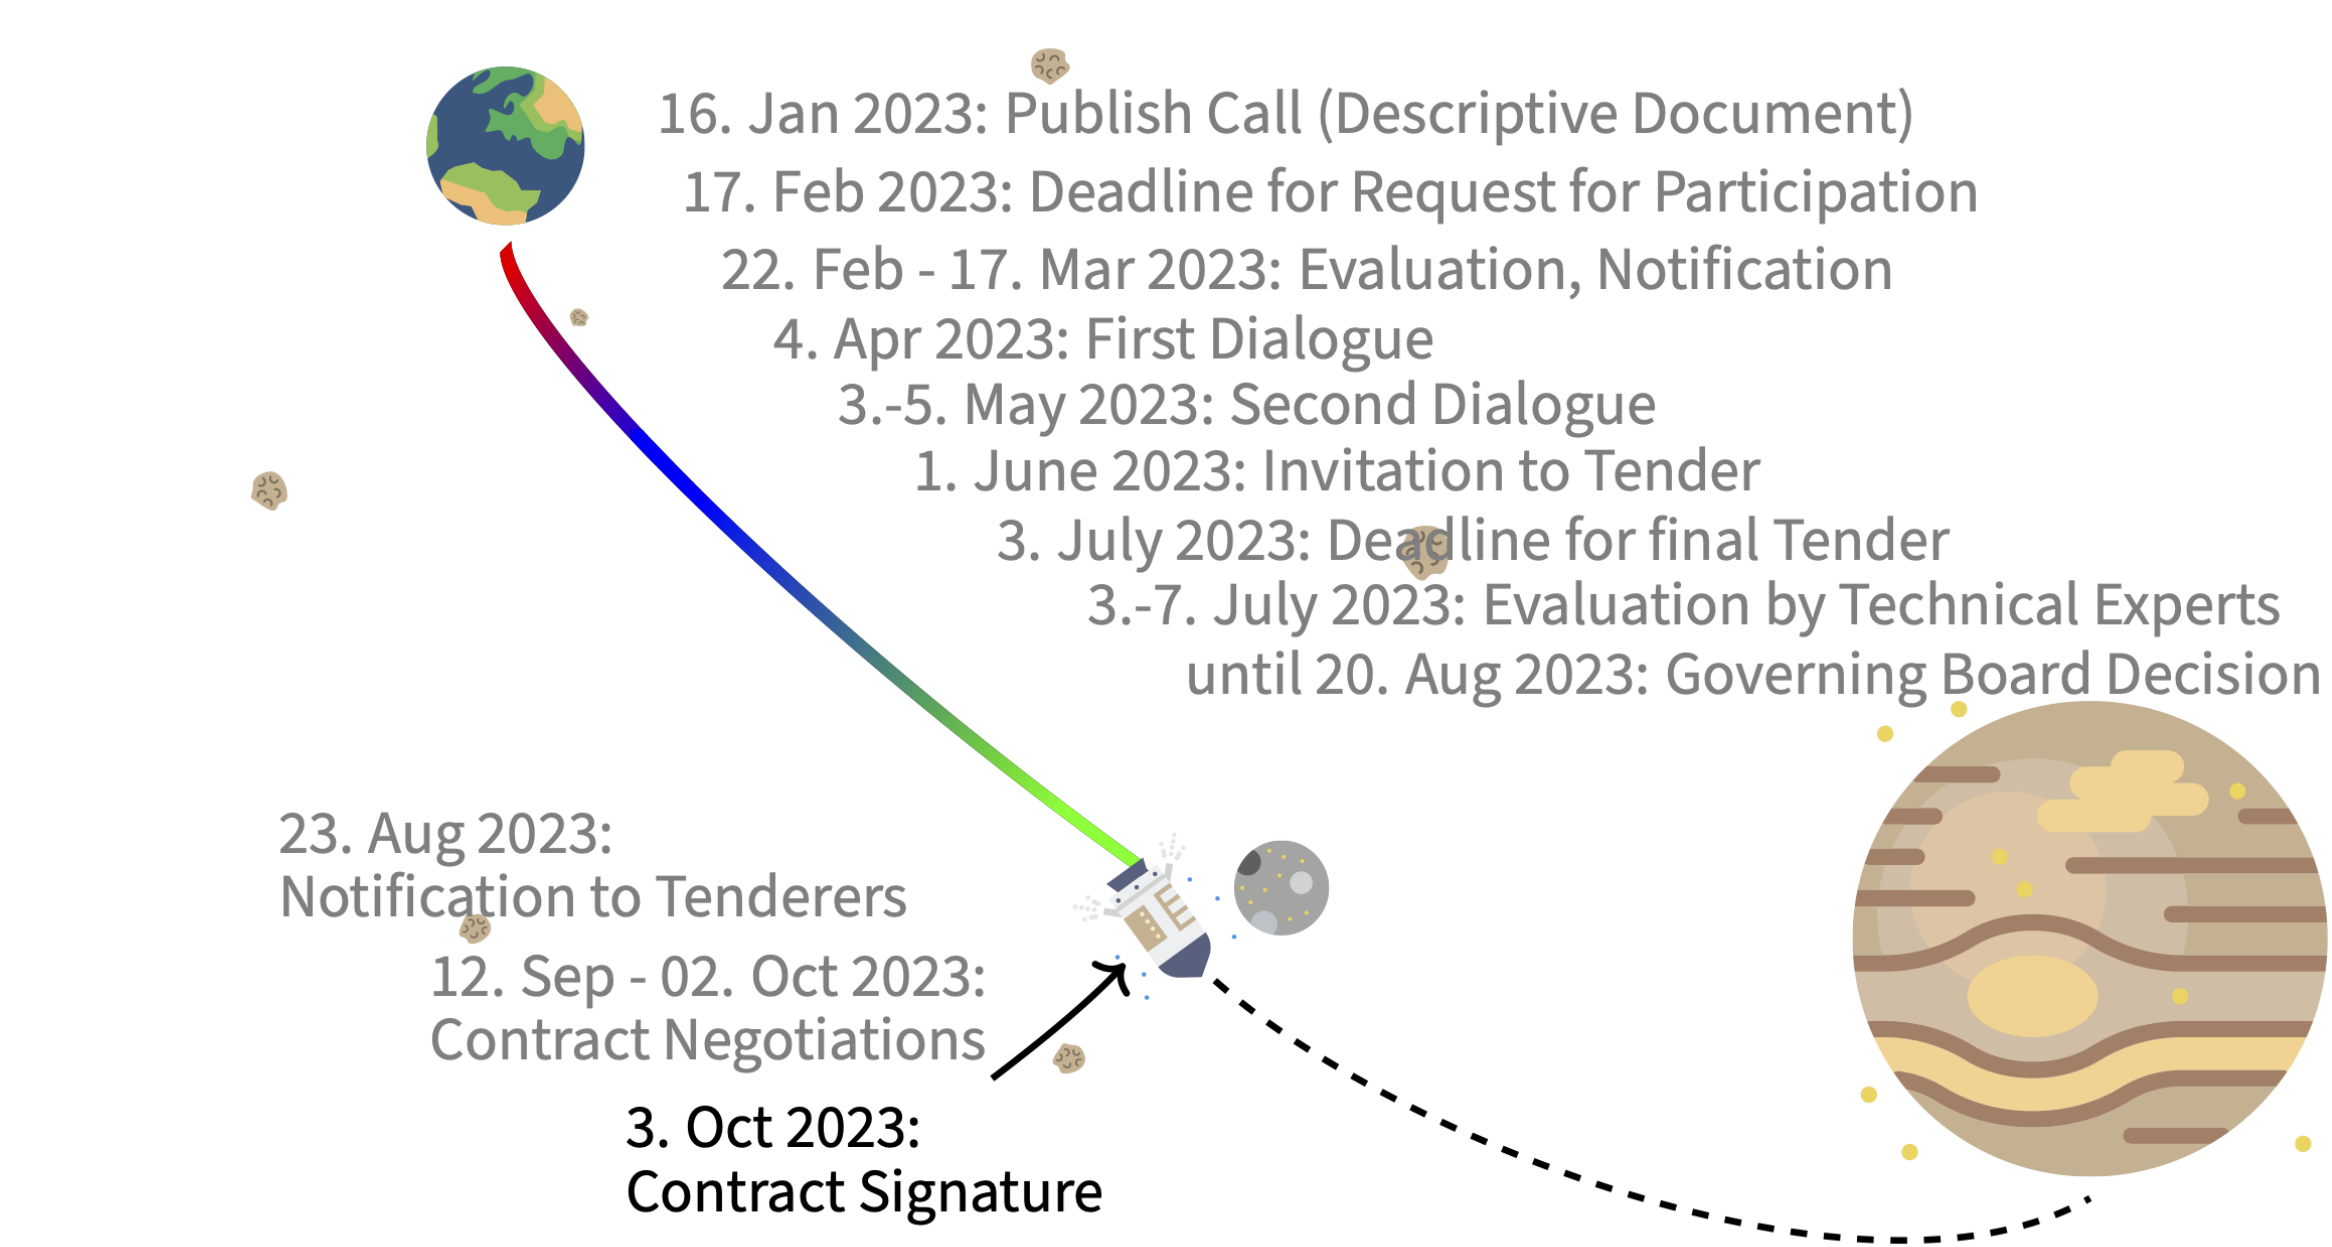

23. Aug 2023:
Notification to Tenderers

12. Sep - 02. Oct 2023:
Contract Negotiations

3. Oct 2023:
Contract Signature

JUPITER CONTRACT ANNOUNCEMENT 3.10.2023

HPC wire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

- Home
- Topics
- Sectors
- Exascale
- Specials
- Resource Library
- Podcast
- Events
- Solution Channels
- Job Bank
- About
- Subscribe

EU Grabs ARM for First ExaFLOP Supercomputer, x86 Misses Out

By Agam Shah

October 4, 2023

The configuration of Europe's first exascale supercomputer, Jupiter, has been finalized, and it is a win for Nvidia and a disappointment for x86 chip vendors Intel and AMD. The Jupiter supercomputer, which will cost €273 million to build, will pair SiPearl's Rhea processor, which is based on ARM architecture, with accelerator technology from Nvidia.

The supercomputer is being built by the European High-Performance Computing Joint Undertaking (EuroHPC JU) and a consortium including Eviden and ParTec. Eviden is an Atos business focusing on advanced computing initiatives tha

The Jülich Supercompu

(Correction: Jülich Supe

about 600 km or 375 m

Specifically, the superc

CPUs, and the initial co

supercomputers in the

on ARM.

That is a big disappoint

invest €33 billion to buil

development initiatives

leaders in a bid to get m

Jülich's fastest system,

November 2021 and is

Europe are the third-ran

performance of 309 pet

Off The Wire Industry Headlines

October 13, 2023

- Coherent File Format Accelerates Time-to-Solution with OpenFOAM
- HealthyCloud Project Unveils Roadmap to Maximize Impact of Health Data and Research Across Europe
- NCSA Welcomes 2023-24 Fellows
- Berkeley Lab CS Area to Share Computing Expertise at SC23

October 12, 2023

- Samsung Electronics to Host AI Forum 2023 Highlighting AI and Computer Engineering Innovation
- PacBio Announces Complete Computational Workflow for Human Whole Genome Sequencing Data Analysis
- SiFive Announces Differentiated Solutions for Generative AI and ML Applications
- EQTC 2023: Europe's Quantum Sector to Showcase Successes and Its Roadmap for Global Leadership
- EuroHPC JU Announces Procurement Call for Upgrading Discoverer Supercomputer
- Los Alamos Partners with AirMettle for Efficient In-Storage Data Analysis
- Caltech Researchers Demonstrate Quantum Eraser to Combat Erasure Errors in Quantum Systems
- Research Results, Computational Expertise Flourish at

SIGN IN / UP

The Register

HPC

Atos subsidiary Eviden scores contract win in Europe's first exascale system

$526M Jupiter set to rule EU's tech orbit by 2024

Dan Robinson

Wed 4 Oct 2023 | 16:45 UTC

The EU's supercomputing initiative, the European High Performance Computing Joint Undertaking (EuroHPC JU), has awarded a procurement contract for Europe's first exascale system, with installation due to start in early 2024.

Known as Jupiter (Joint Undertaking Pioneer for Innovative and Transformative Exascale Research), the system was announced last year followed by a call for tender in January of this year.

EuroHPC JU said the procurement contract for Jupiter was awarded to a consortium comprising of Eviden, the professional services side of French IT giant Atos, and ParTec, a German supercomputing hardware company.

The project is expected to have a total cost of €273 million ($287 million) covering the build, delivery, installation, and maintenance of Jupiter, according to the EuroHPC.

However, Eviden put the overall project cost at €500 million ($526 million), saying that this is the figure for the entire project, including the system manufacturing and its

THENEXTPLATFORM

HOME | COMPUTE | STORE | CONNECT | CONTROL | CODE | AI | HPC | ENTERPRISE | HYPERSCALE | CLOUD

LATEST > Intel To Set Its FPGA Unit Free To Pursue Its Own Path > COMPUTE

Search ...

HOME > HPC > Details Emerge On Europe's First Exascale Supercomputer

DETAILS EMERGE ON EUROPE'S FIRST EXASCALE SUPERCOMPUTER

October 5, 2023 Timothy Prickett Morgan

me details are emerging on Europe's first exascale system, codenamed "Jupiter" and to be installed at the ich Supercomputing Center in Germany in 2024. There has been a lot of speculation about what Jupiter ll include for its compute engines and networking and who will build and maintain the system. We now ow some of this and can infer some more from the statements that were made by the organizations rticipating in the Jupiter effort.

June 2022, the Forschungszentrum Jülich in Germany, which has played host to many supercomputers ce it was founded in 1987, was chosen to host the first of three European exascale-class percomputers to be funded through the EuroHPC Joint Undertaking and through the European national d state governments countries who are essentially paying to make sure these HPC and AI clusters are ere they want them. With Germany having the largest economy in Europe and being a heavy user of C thanks to its manufacturing focus, Jülich was the obvious place to park the first machine in Europe to eak the exaflops barrier.

at barrier is as much an economic one as it is a technical one. The six-year budget for Jupiter weighs in at 00 million, which is around $526.1 million at current exchange rates between the US dollar and the ropean euro. That is in the same ballpark price as what the "Frontier" exascale machine at Oak Ridge tional Laboratory and the "El Capitan" machine that is being installed right now at Lawrence Livermore tional Laboratory – both of which are based on a combination of AMD CPUs and GPUs and Hewlett ckard Enterprise's Slingshot variant of Ethernet with HPE as the prime contractor.

erybody knows that Jupiter was going to use SiPearl's first generation Arm processor based on the verse "Zeus" V1 core from Arm Ltd, which is codenamed "Rhea" by SiPearl and which is appropriate

JUPITER ARCHITECTURE

JUPITER MODULES

JUPITER Booster (~1ExaFLOP/s)

- ~125 Racks BullSequana XH3000
- Node design
  - ~6000 nodes, 4× NVIDIA CG1 per node
- CG1: NVIDIA Grace-Hopper
  - 72 Arm Neoverse V2 cores (4×128b SVE2); 120 GB LPDDR5
  - H100 (132 SMs); 96 GB HBM3
  - NVLink C2C (900 GB/s)

JUPITER Cluster

- ~15 Racks BullSequana XH3000
- Node design
  - ~1300 nodes
  - 2× SiPearl Rhea1 per node
- Rhea1
  - 80 Arm Neoverse V1 cores (2×256b SVE)
  - 256 GB DDR5, 64 GB HBM2e

ExaFLASH: 29PB (raw) NVMe, IBM SSS6000

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

ExaTAPE: 370PB Tape*, LTO9

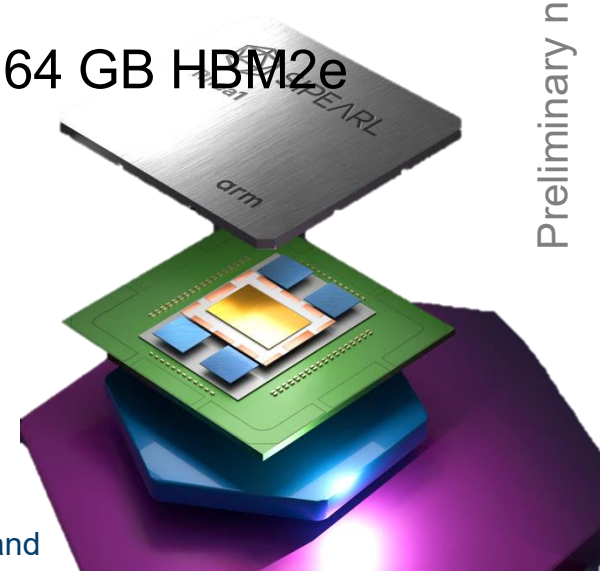

Preliminary numbers, might change during installation

*Upgrade budget available, depending on demand

JUPITER BOOSTER BLADE OVERVIEW

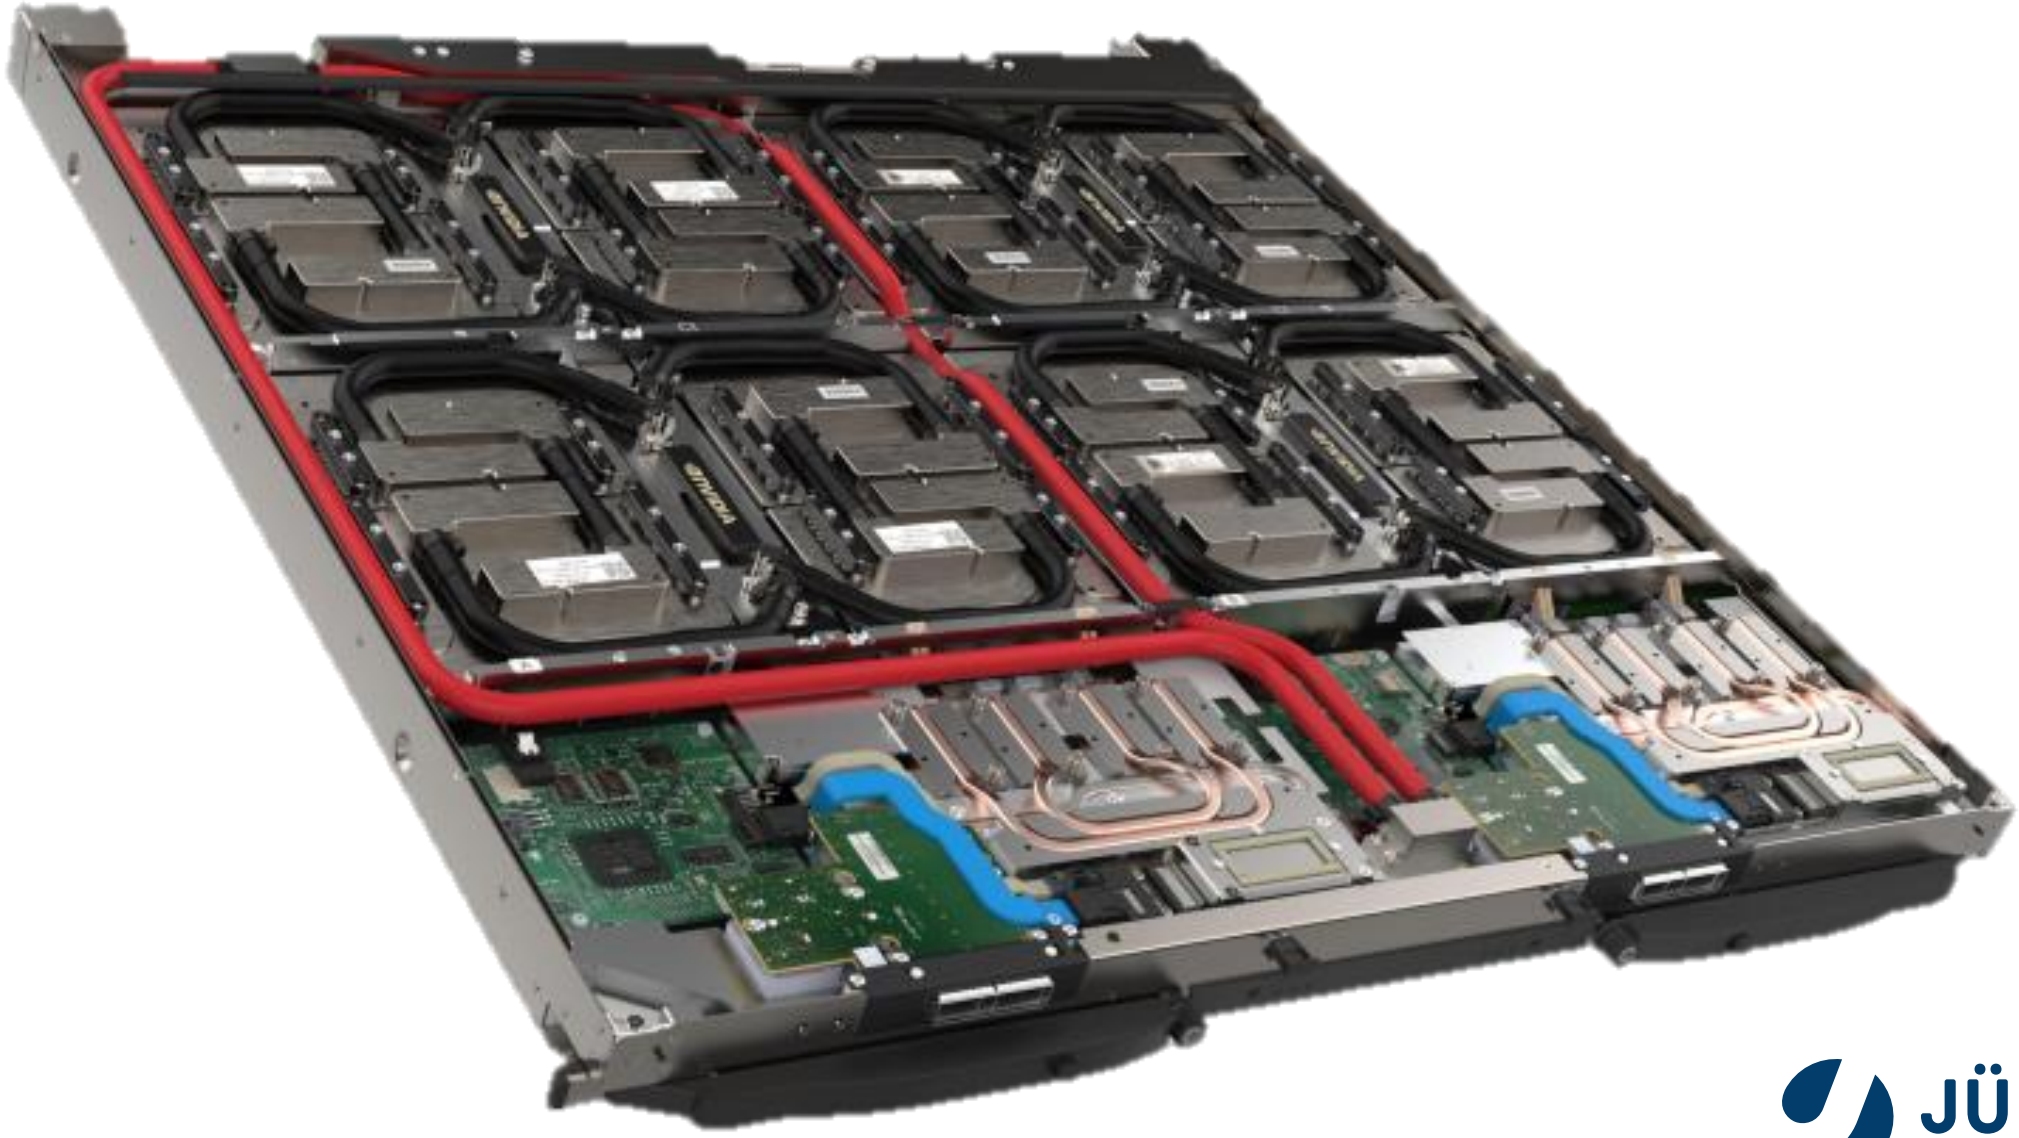

JUPITER – BOOSTER COMPUTE NODE ARCHITECTURE

Node Specs

- 4× NVIDIA Grace-Hopper in SXM5 Board (4× 680W)
- 4× NVIDIA InfiniBand NDR200
- 480 GB LPDDR5X / 360 GB HBM3 (usable)
- NVLink 4
  - GPU-GPU 150 GB/s per dir, CPU-GPU 450 GB/s per dir, CPU-CPU 100 GB/s per dir
- CG4 Motherboard (4× CG1 GH module + 4× CX7 HCA assembly)
  - All NVIDIA, except the BMC

CPU Specs

- ARM Neoverse V2
  - SVE2/NEON (4x 128 bit vector op)
- 72 cores @ ~2.4GHz (~3.2 GHz turbo)
- 120 GB LPDDR5X (8 channels)
  - *≥450 GB/s*
  - *~150 ns latency*

GPU Specs

- H100
  - 47.5 TFLOP/s (HPL Rmax single GPU)
- 90 GB HBM3
  - ≥3600 GB/s
  - ~450 ns latency

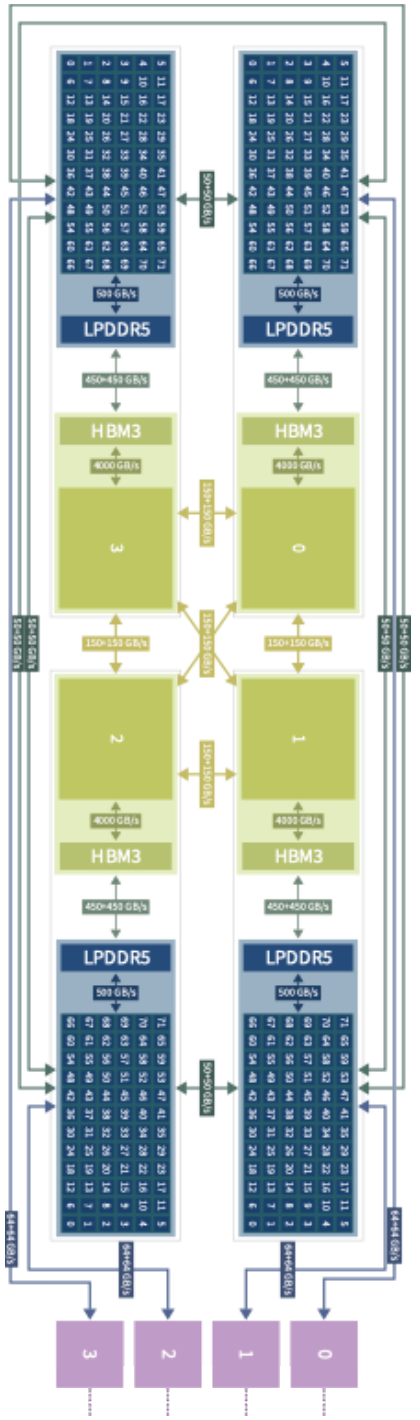

JUPITER – INTERCONNECT

One Network to Rule Them All

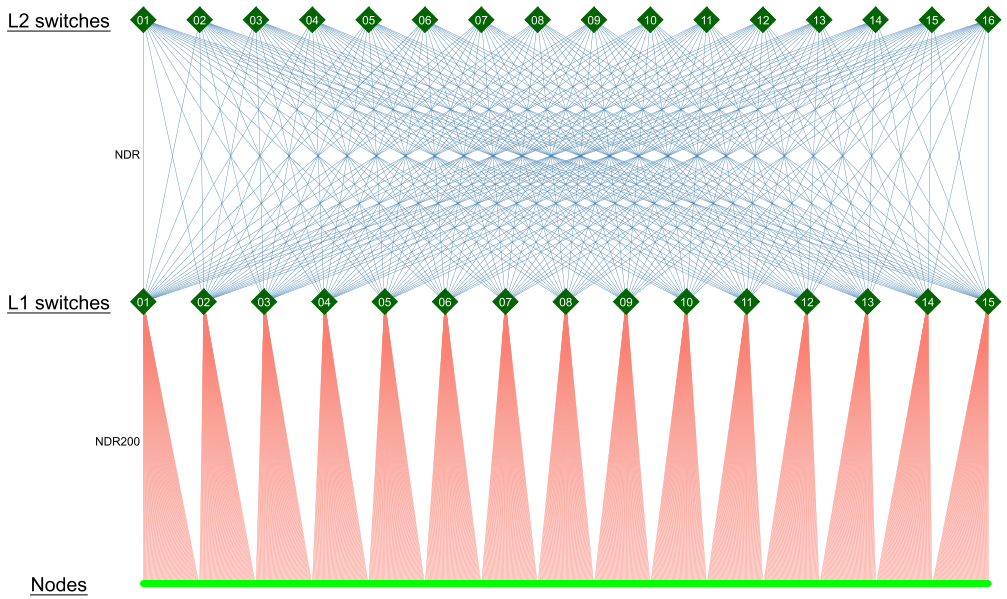

- NVIDIA Mellanox InfiniBand NDR/NDR200
  - NVIDIA Quantum-2 switches
  - NVIDIA Connect-X7 HCAs
- Dragonfly+ topology
  - 27 Dragonfly groups
  - Within each group: full fat tree
- 51000 links, 102000 logical ports, 25400 endpoints, 867 switches
- Adaptive Routing
- In-network processing on switch level (SHARPv3), tentatively

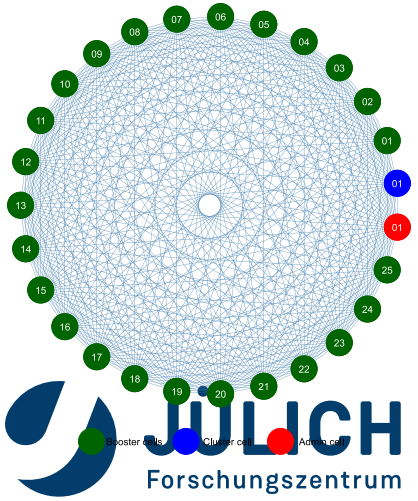

SYSTEM MANAGEMENT

JUPITER MANAGEMENT STACK

3 main pillars/actors

SMC xScale
Core part of the stack. Vast majority of components come from here.
Developed by Eviden
Heavily based on open source and cloud technologies

JUPITER MANAGEMENT STACK

3 main pillars/actors

SMC xScale	ParaStation
Core part of the stack. Vast majority of components come from here.	Enhancement of the core
Developed by Eviden	Developed by ParTec
Heavily based on open source and cloud technologies	Integrates ParTec tools in SMCx to streamline their support workflows

JUPITER MANAGEMENT STACK

3 main pillars/actors

SMC xScale	ParaStation	xOPS
Core part of the stack. Vast majority of components come from here.	Enhancement of the core	Enhancement of the core
Developed by Eviden	Developed by ParTec	Developed by JSC
Heavily based on open source and cloud technologies	Integrates ParTec tools in SMCx to streamline their support workflows	Extensive set of Ansible roles for HPC, targeting JSC's requirements and needs

JUPITER MANAGEMENT STACK – KEY AREAS

	Technology	Challenges	Provider
Operating System	Linux	Security, Performance, Stability, HW support	
Management Storage	Ceph	Multi-use, Performance, Scalable	IBM
Management Plane	Kubernetes	Scalable, 0 downtime, Flexible, Open	kubernetes
Configuration Management	Ansible	Standard, Easy to extend, Open	
Boot Image(s) Management	ImageBuilder	ARM / x86 support, Tracking, Integration	
Container(s)	UBI Universal Binary Images	Standard, Security, Consistency	

JUPITER MANAGEMENT STACK – KEY AREAS

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

THE PREPARATION PHASE

2023/2024 - PREPARING CAMPUS, SLAB, STORAGE, JEDI

… it is not only about waiting for JUPITER

ExaSTORE Procurement
300PB+ Spinning Disk Storage for JUPITER

MDC Procurement

MODULAR DATA CENTER FOR JUPITER

EVIDEN

- Vendor: Eviden
- Area: ~2300m^2
- 1x Datahall (Storage, Management)
- 7x IT Modules (20 Racks per module)
- UPS, Generator
- Entrance area
- Workshop, Warehouse
- 15x 2,5 Megawatt Power Stations

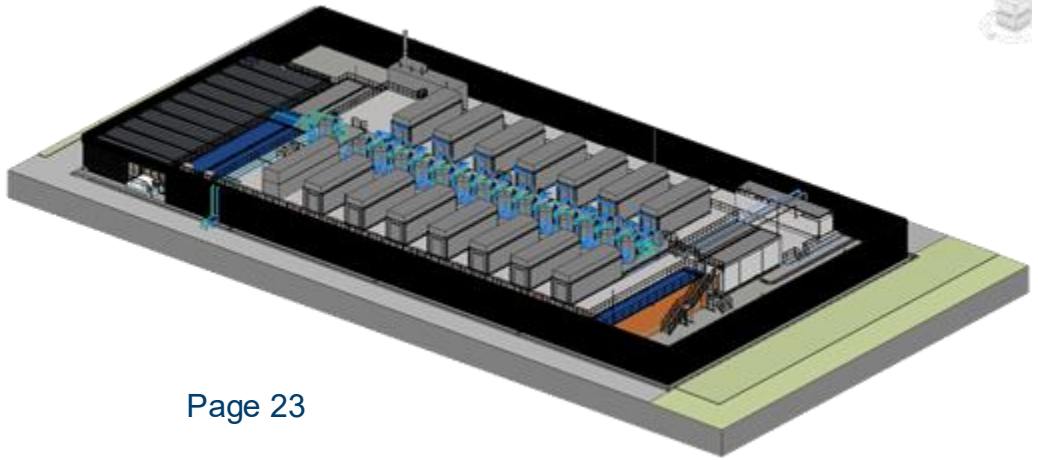

2023/2024 - PREPARING CAMPUS, SLAB, STORAGE, JEDI

… it is not only about waiting for JUPITER

ExaSTORE Procurement
300PB+ Spinning Disk Storage for JUPITER

MDC Procurement

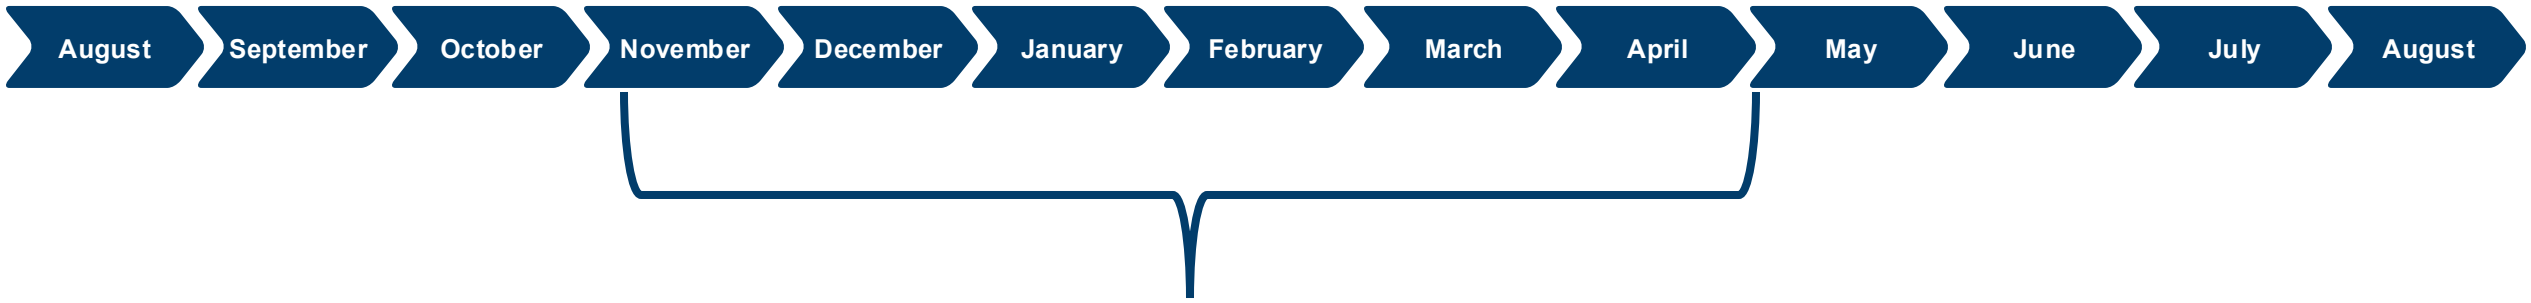

Concrete Foundation (Slab)

CONCRETE FOUNDATION

CONCRETE FOUNDATION

Construction of concrete slab 85 m x 42 m x 0.5 m


2024-08-15 17:29:56
August 2024
Backup Cold Water Cooling: 1MW
Power (Campus): 2*60-80 MVA
Water (River Rur): up to 30 cbm/h
Network:
3,2 Tbit/s


JUWELS Booster:
• 3744 NVIDIA A100 GPUs
• 44 PetaFLOP/s HPL
• 2,4 MW Peak, 1,1MW Average
• Direct-Liquid-Cooling
• 36°in, 42-44°out
• free cooling
• heat-reuse
Power (Datacenter): 15*2,5MVA
14 Free Cooling Towers, 1 Chiller
JUPITER Booster:
• ~24000 NVIDIA H100 GPUs
• 1 ExaFLOP/s HPL
• 17 MW Peak, 9-11MW Average
• Direct-Liquid-Cooling
• 36°in, 42-44°out
• free cooling
• heat-reuse
• Up to 30 cbm/h river water
• !!!Worst-case!!!
• Only on very-hot summer days
April 2025

2023/2024 - PREPARING CAMPUS, SLAB, STORAGE, JEDI

… it is not only about waiting for JUPITER

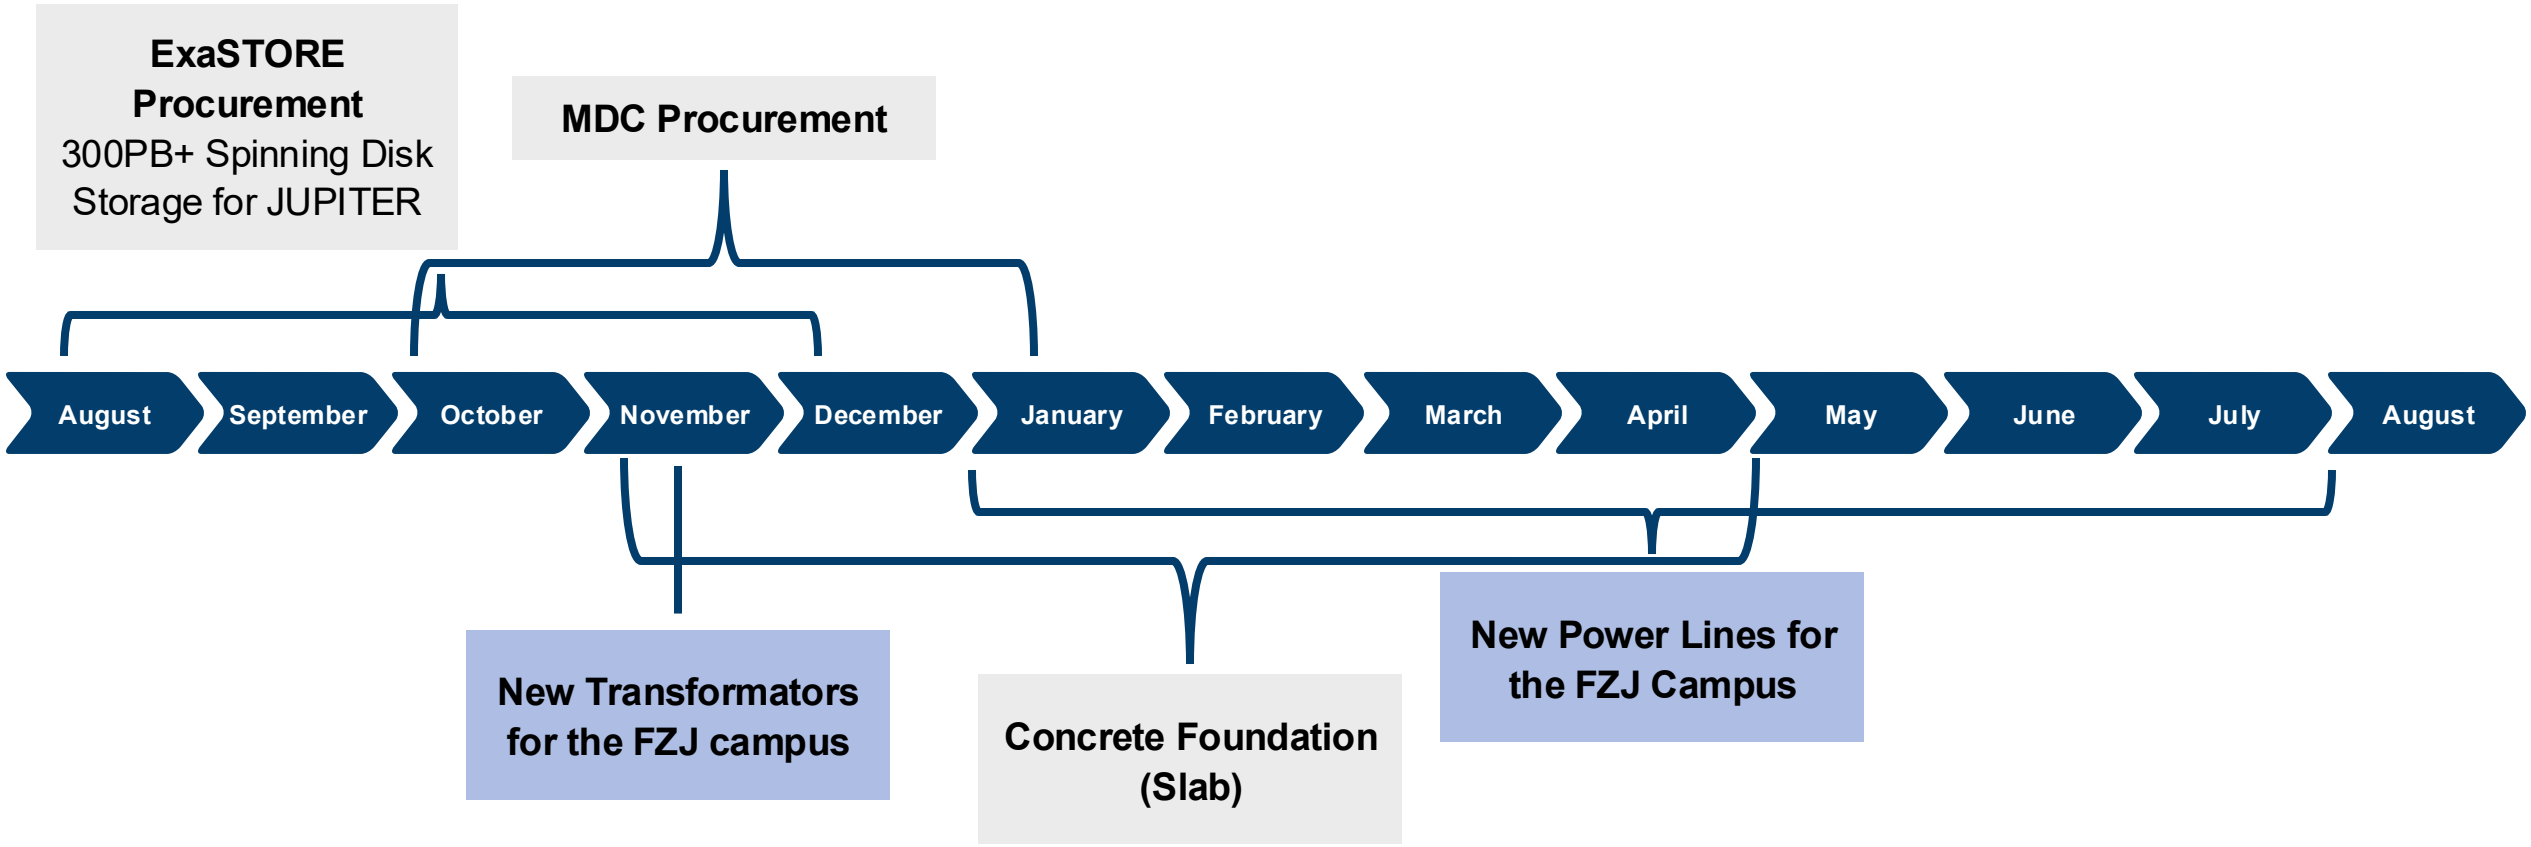

POWER TRANSFORMER SUBSTATION AND LINES

Upgrade of transformers 110 kV / 35 kV from 2 x 40 MVA to 2 x 60-80 MVA and upgrade 110kV power line

RECENT DEVELOPMENTS

2024/2025 - MDC AND JUPITER INSTALLATION

Always in motion is the future…

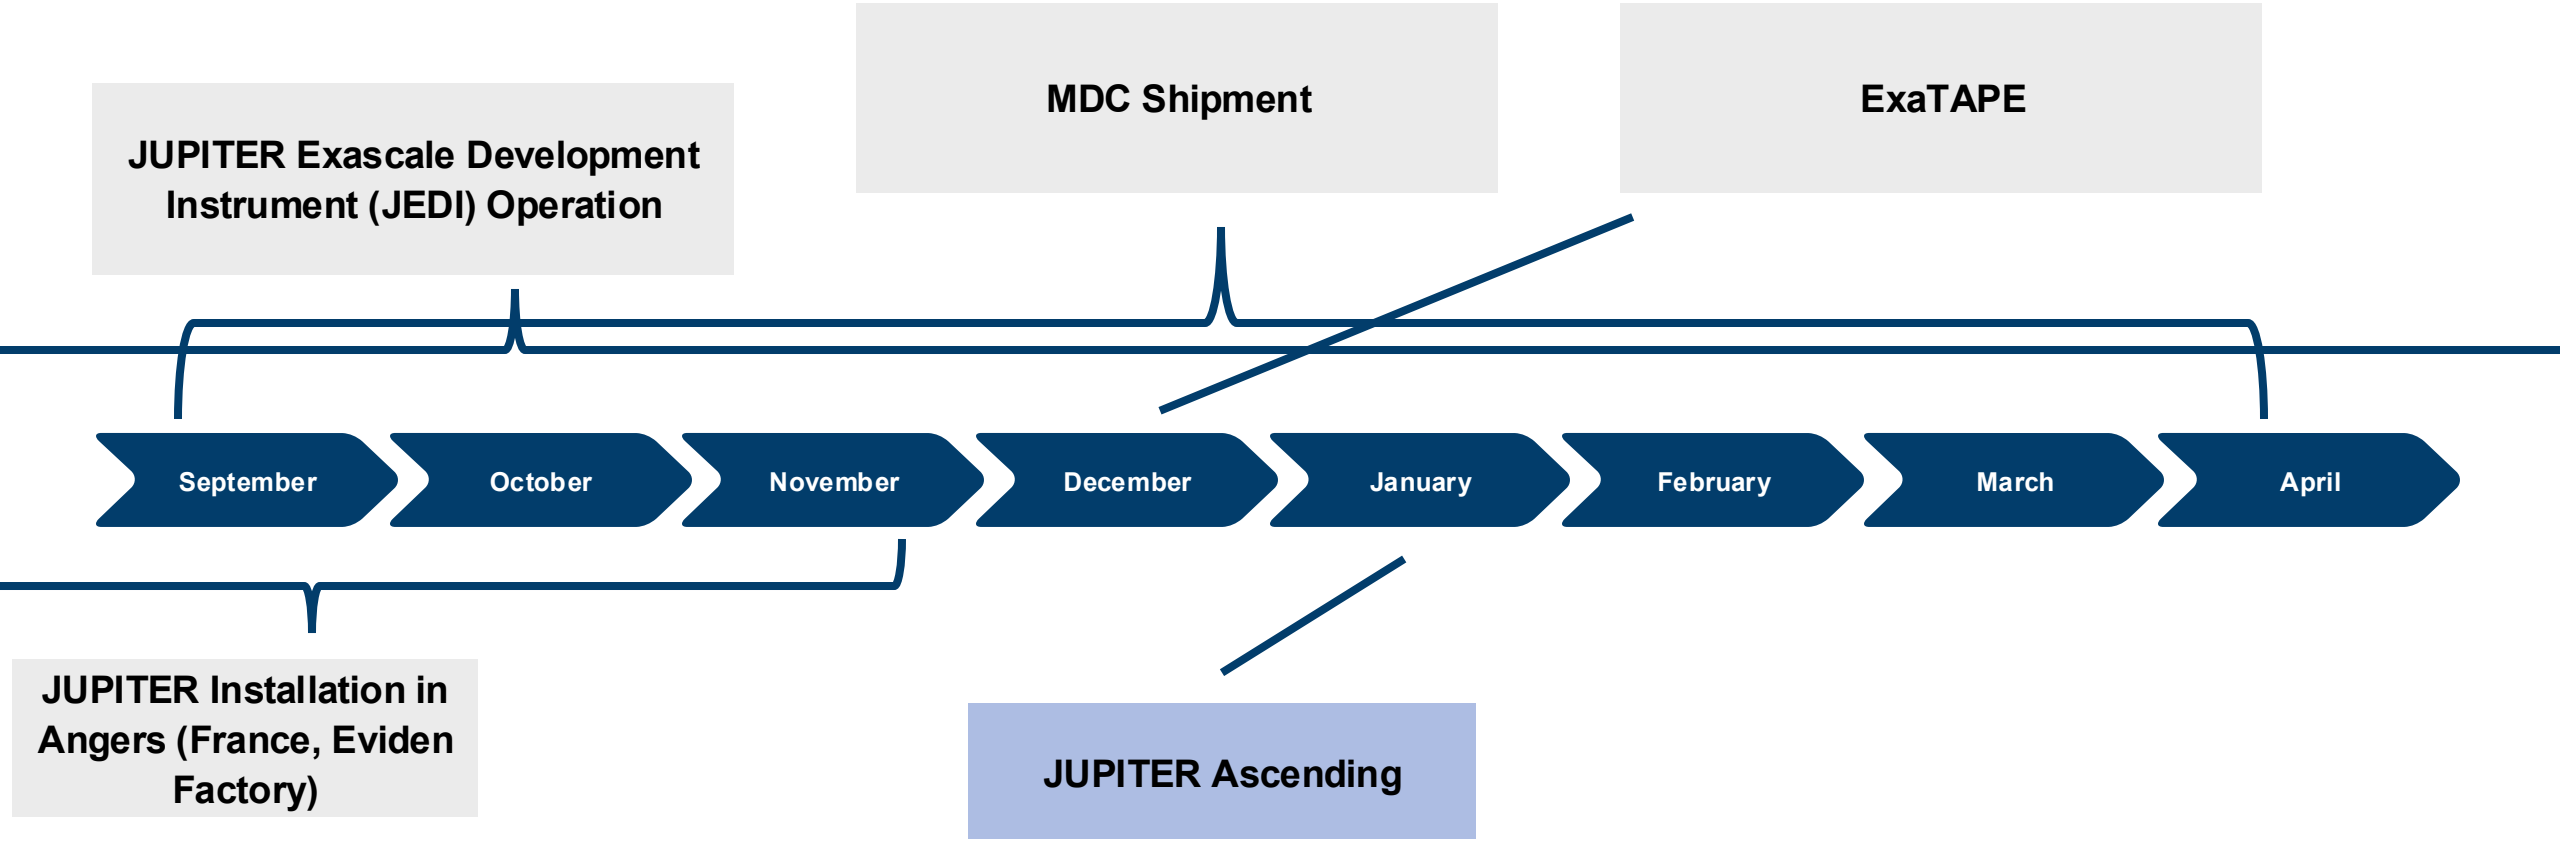

JUPITER ASCENDING

Since January 2025

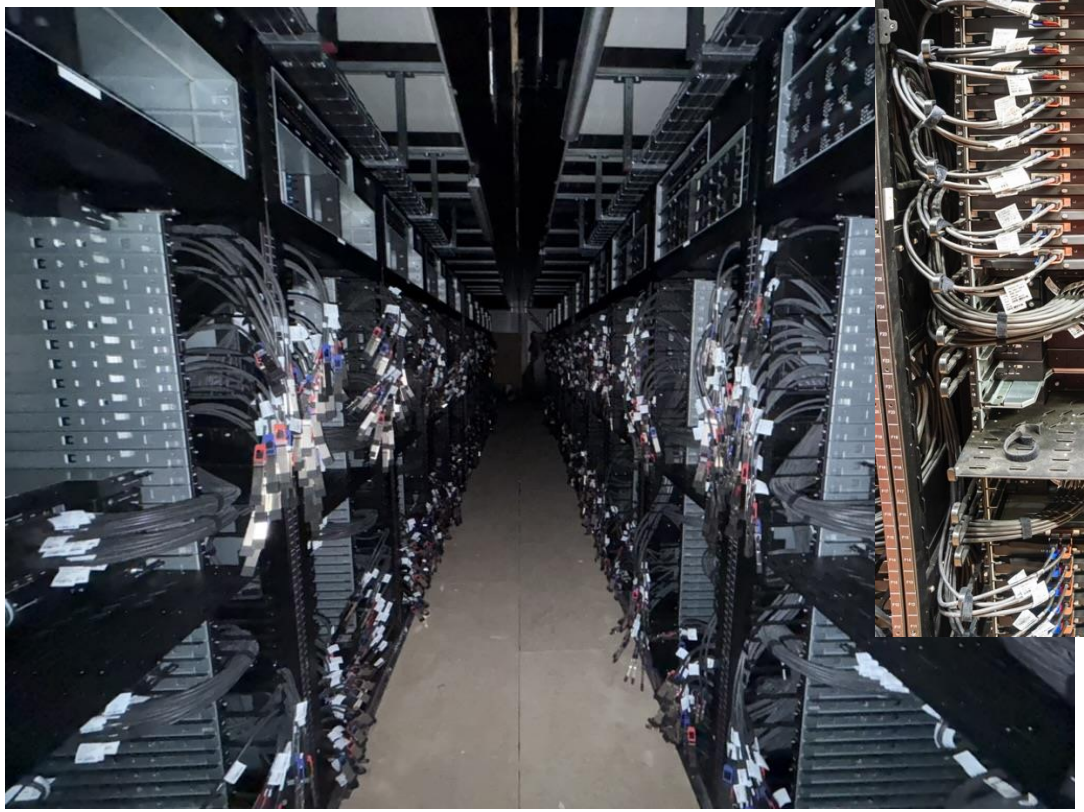

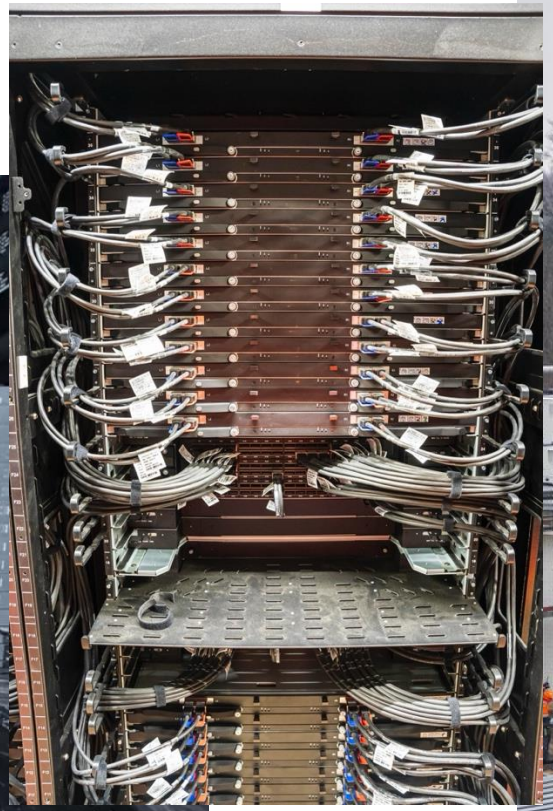


0221 / 956 86-0
www.COLONIA.info

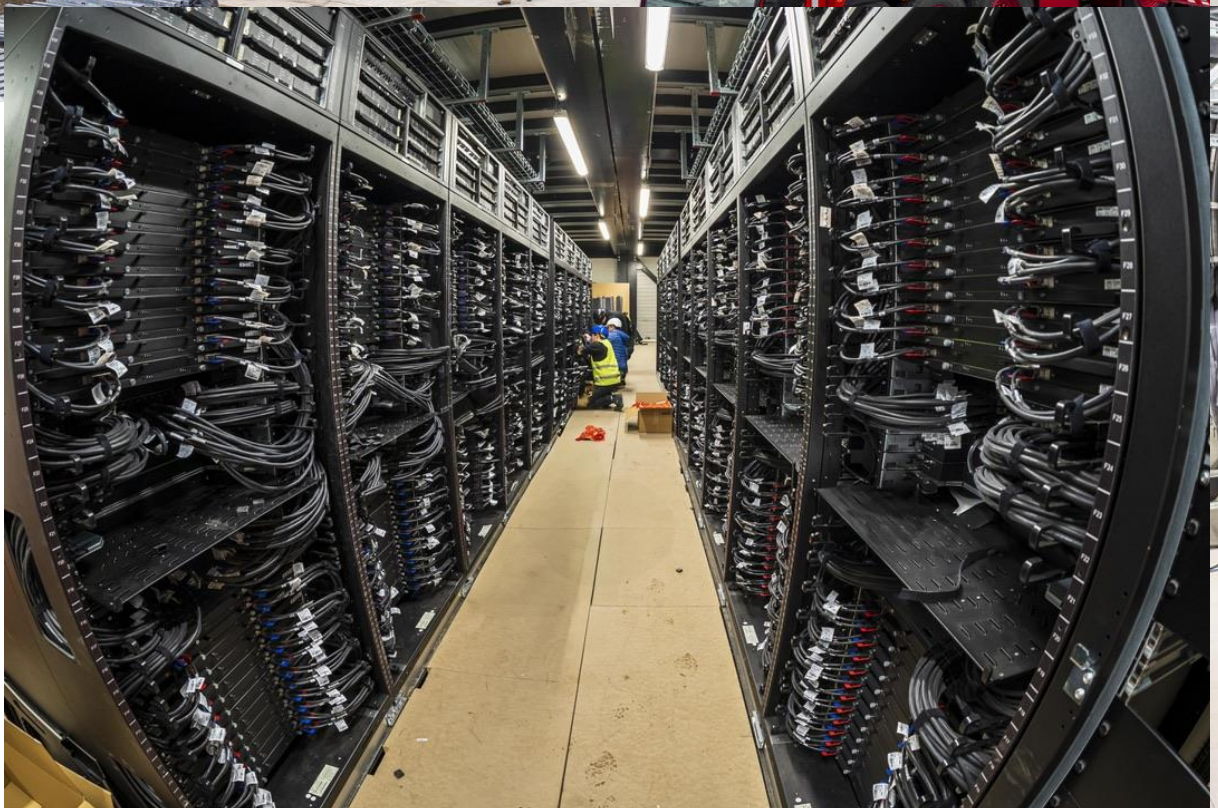

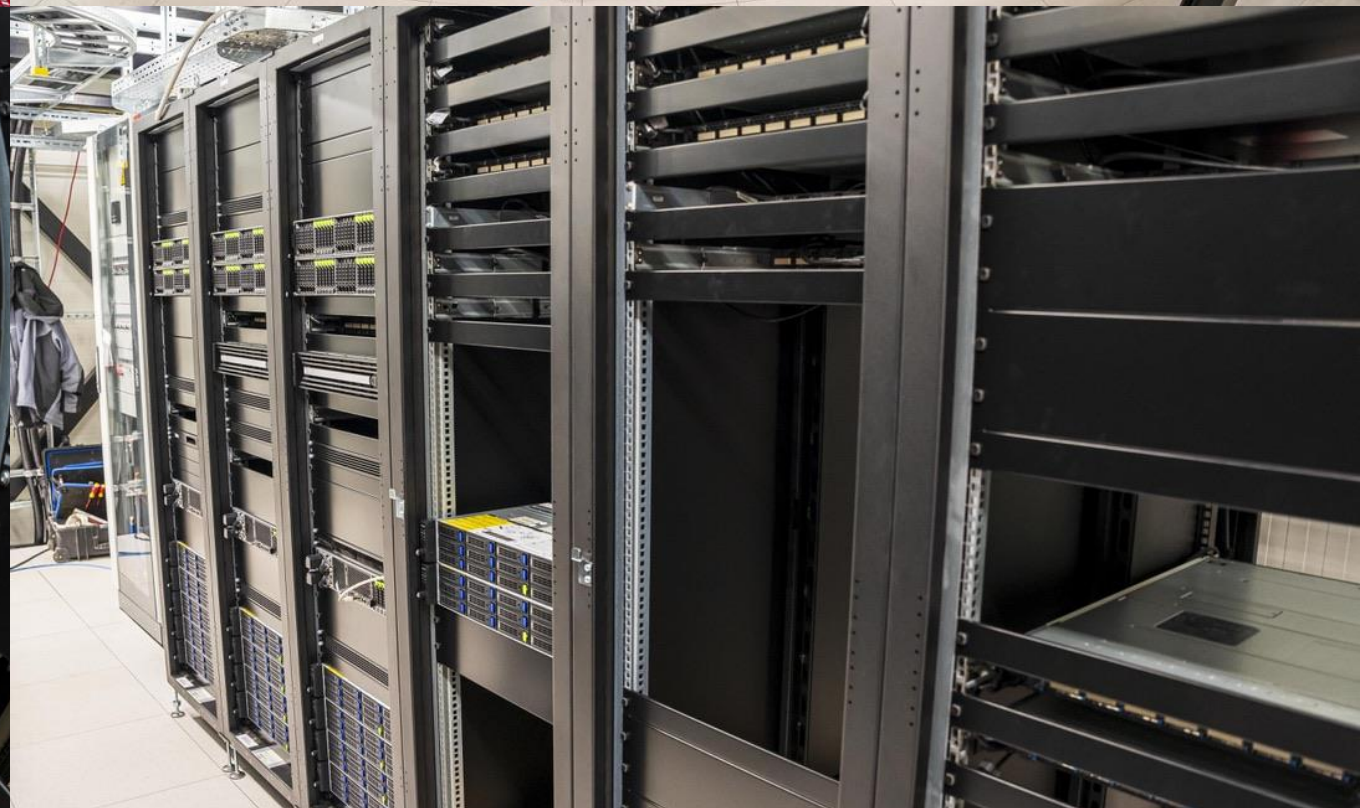

2024/2025 - MDC AND JUPITER INSTALLATION

Always in motion is the future.

JUPITER Exascale Development Instrument (JEDI) Operation

ExaTAPE

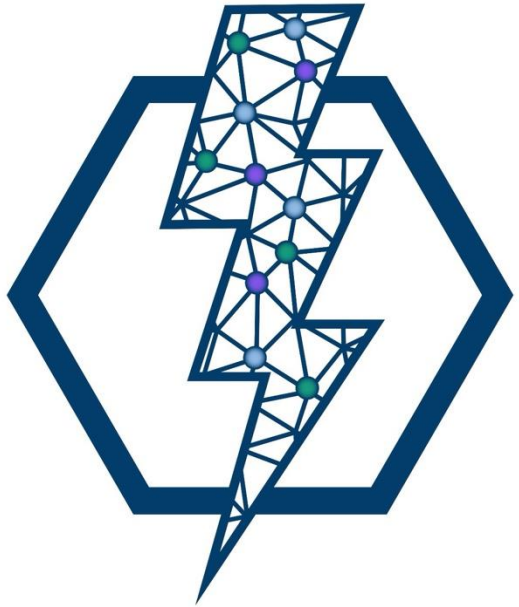

JUPITER AI FACTORY

September | October | ... | ebruary | March | April

JUPITER Installation in Angers (France, Eviden Factory)

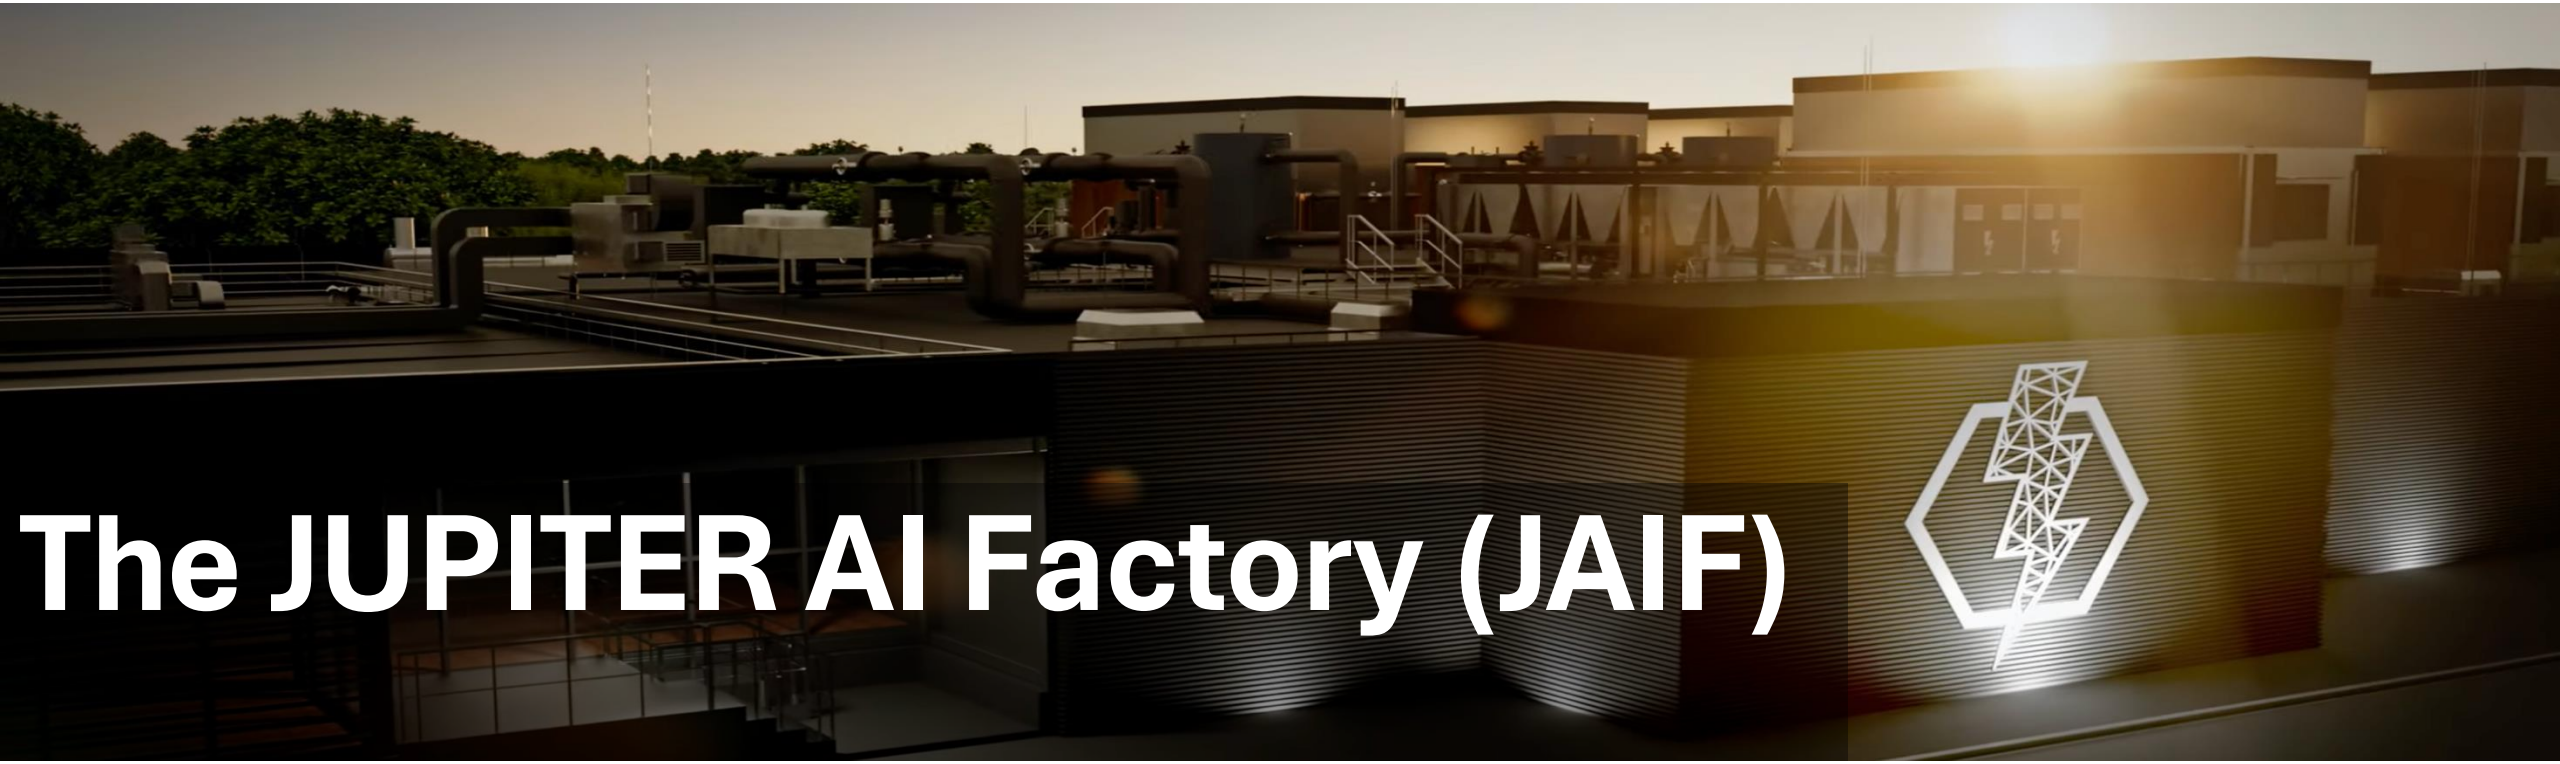

The JUPITER AI Factory (JAIF)

JÜLICH Forschungszentrum

RWTH AACHEN UNIVERSITY

Fraunhofer

hessian.AI

Associated partners:

LAMARR INSTITUTE FOR MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE

WEST AI KI-Servicezentrum

hessian.AISC

KI BUNDESVERBAND

THE JAIF ONE STOP SHOP

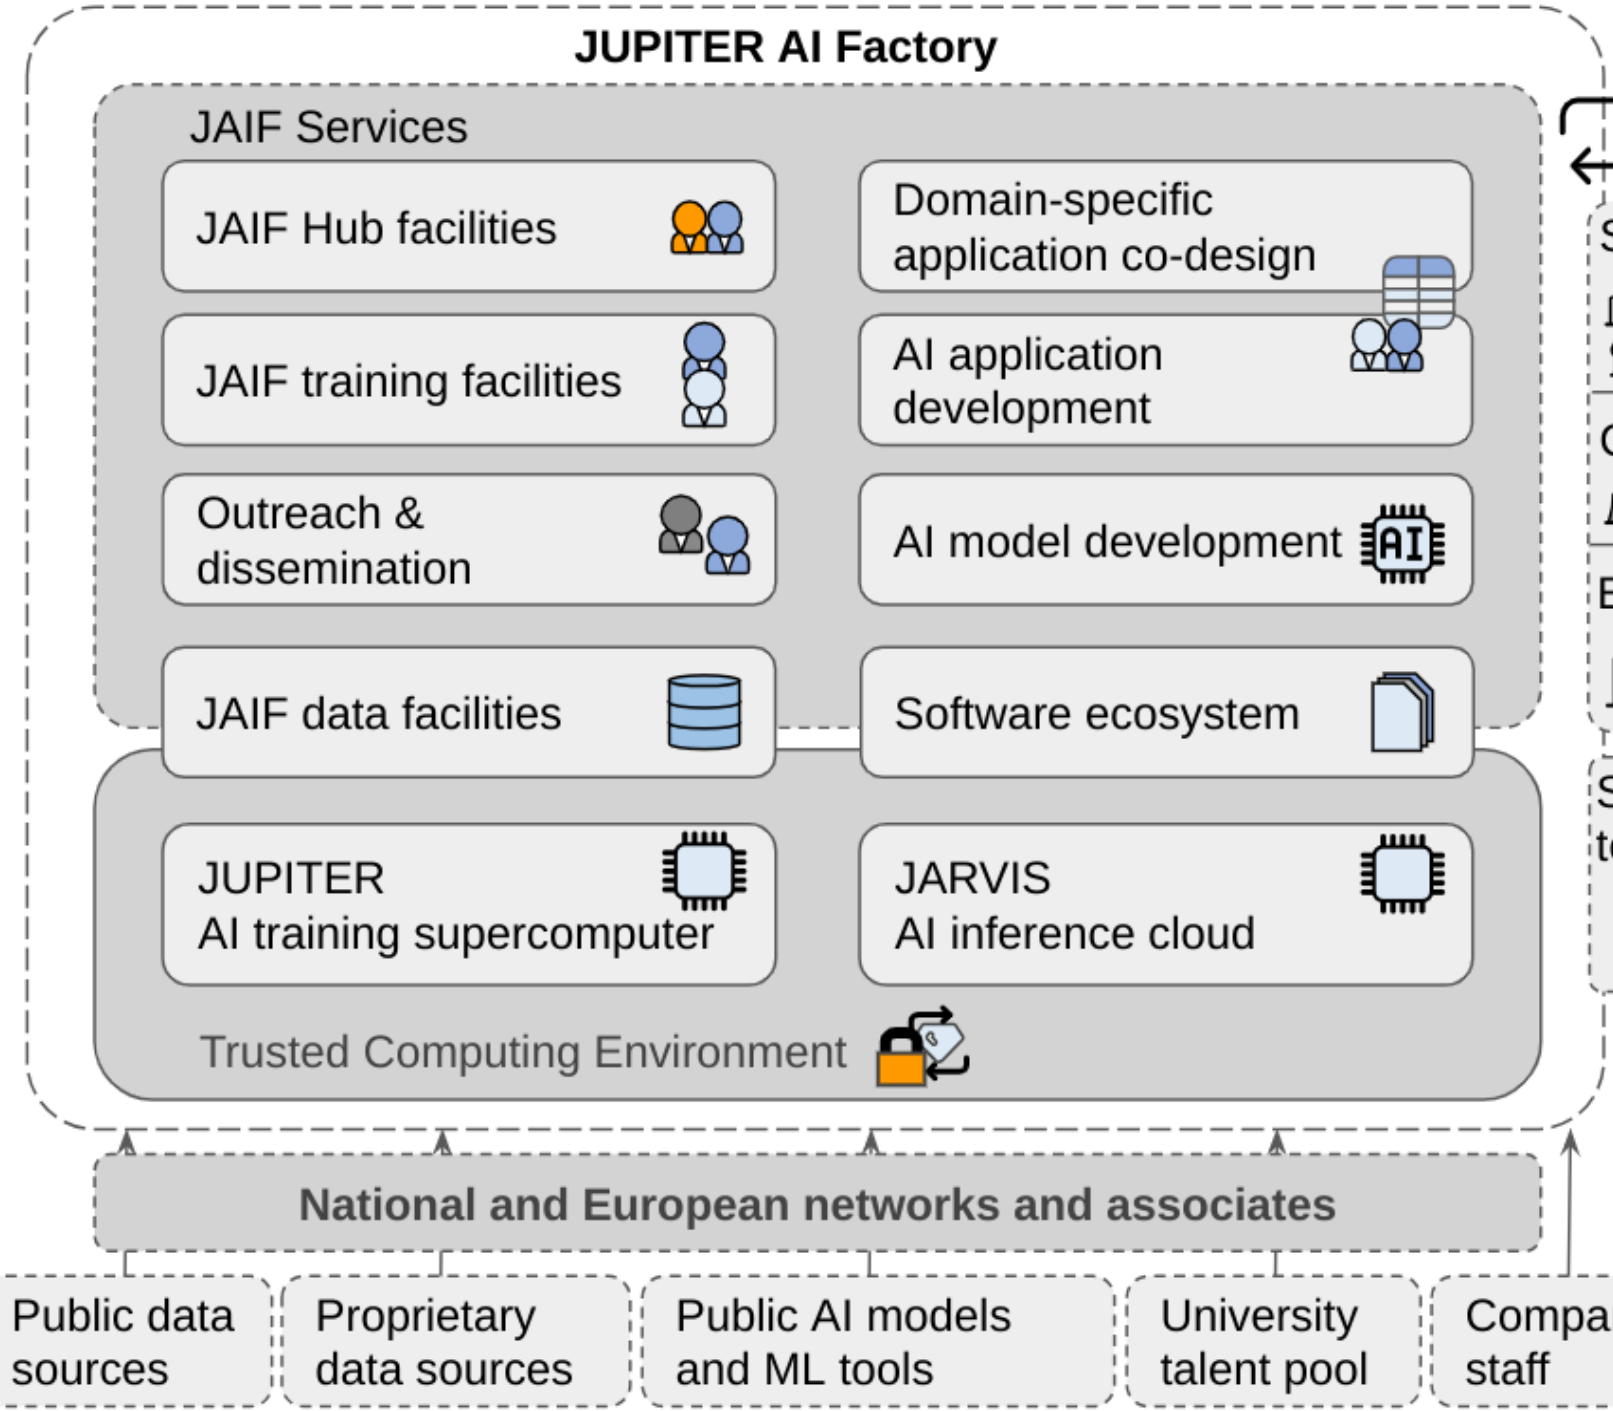

Modular JUPITER
Hybrid Training/Inference AI System

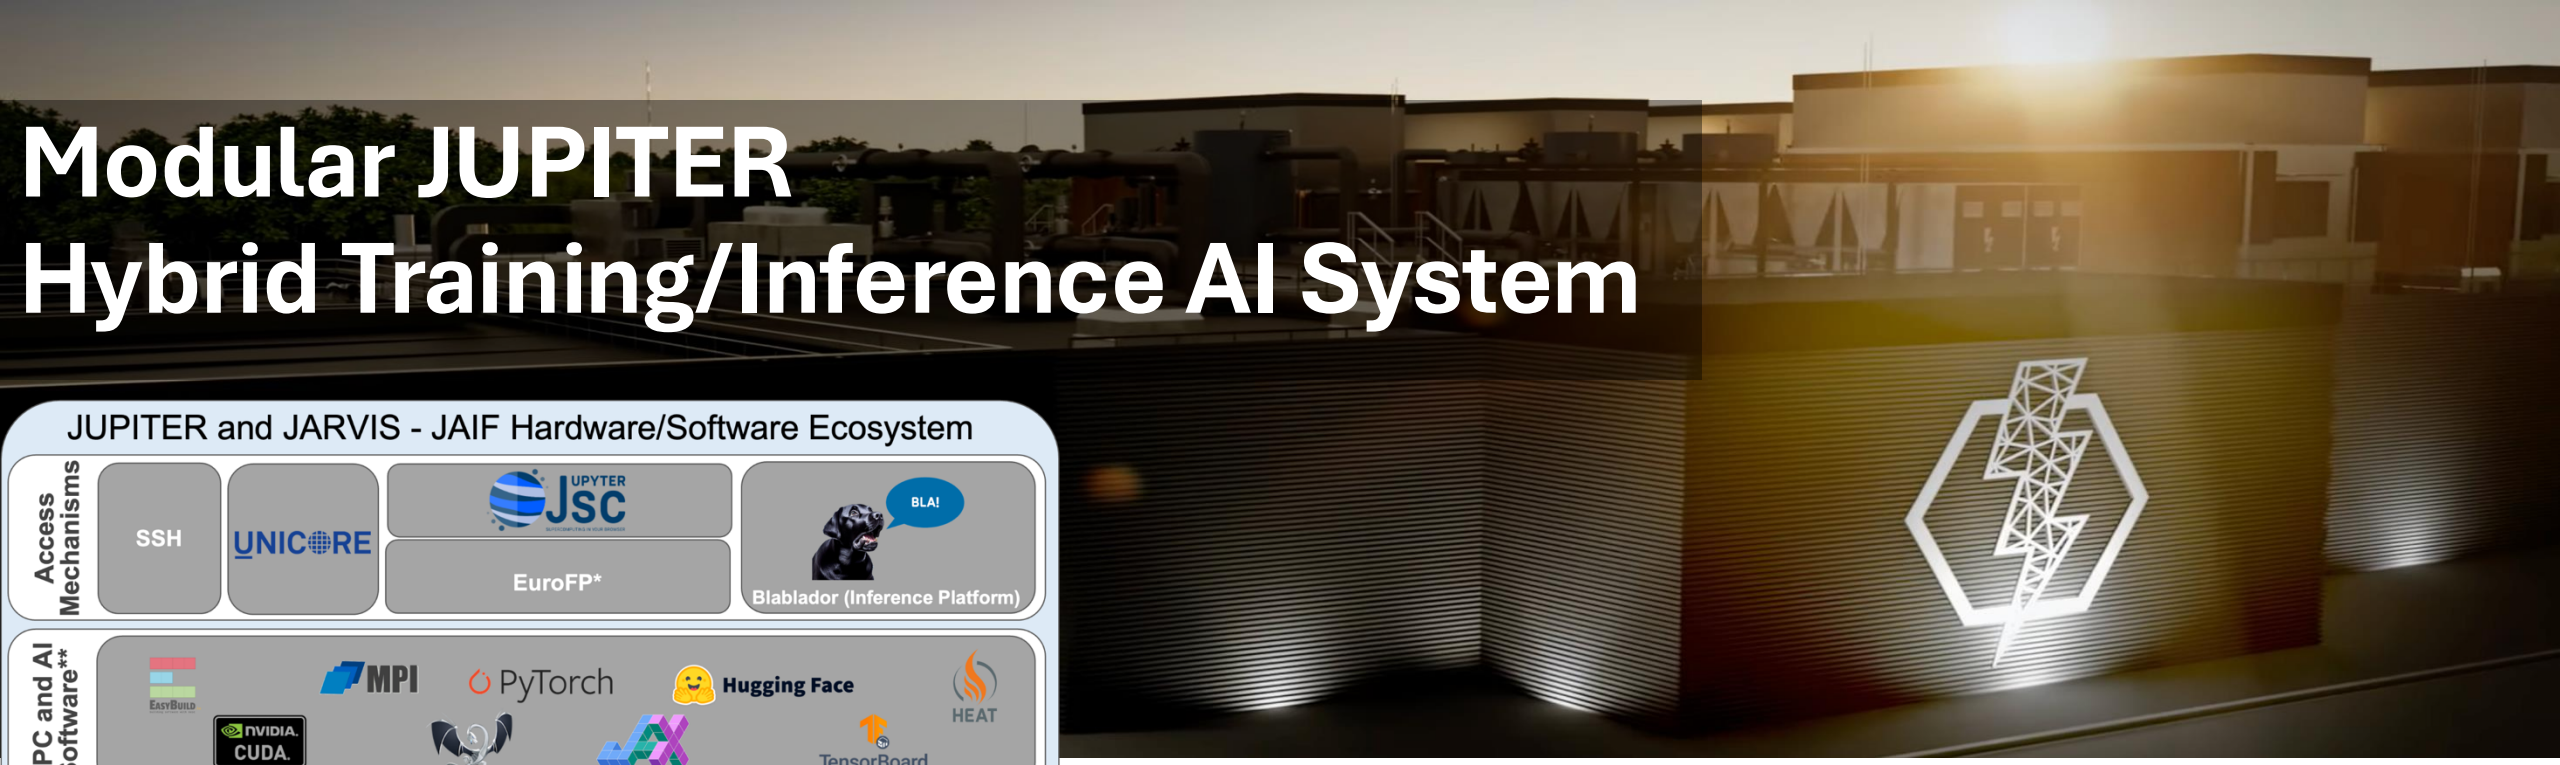

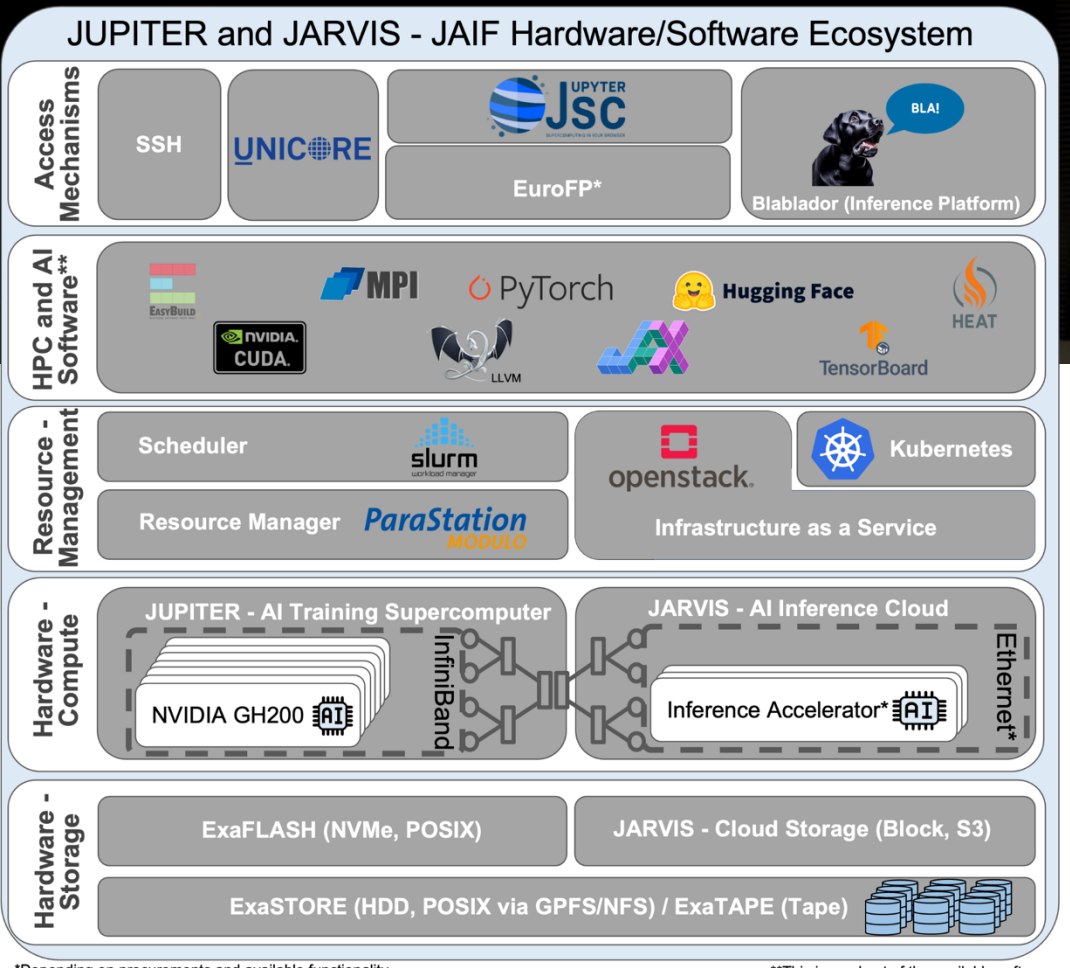

*Depending on procurements and available functionality

**This is a subset of the available software

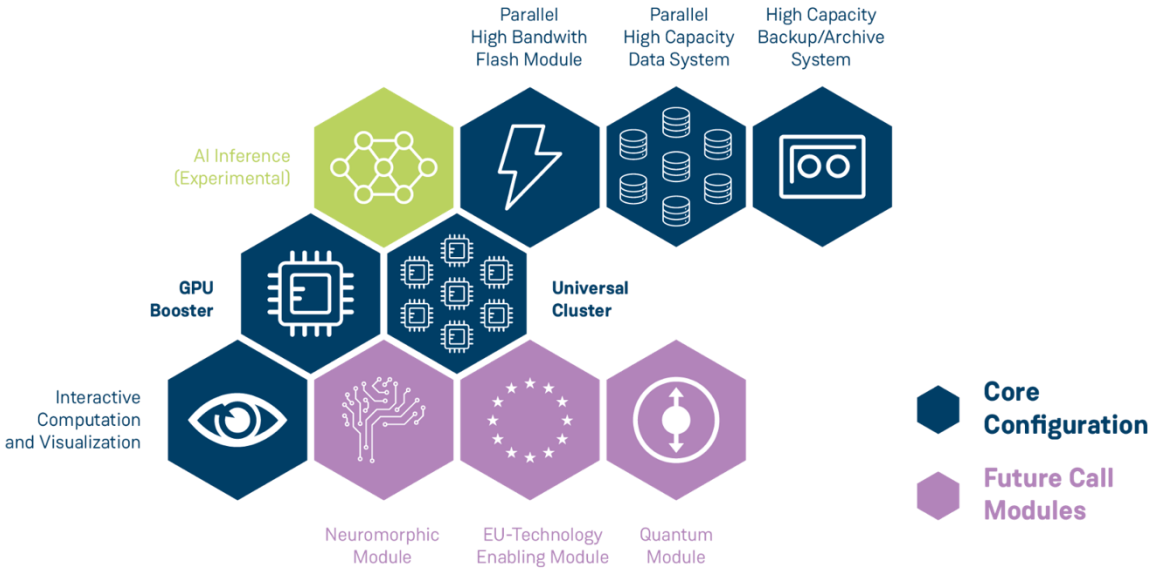

JUPITER AI FACTORY

2024/2025 - MDC AND JUPITER INSTALLATION

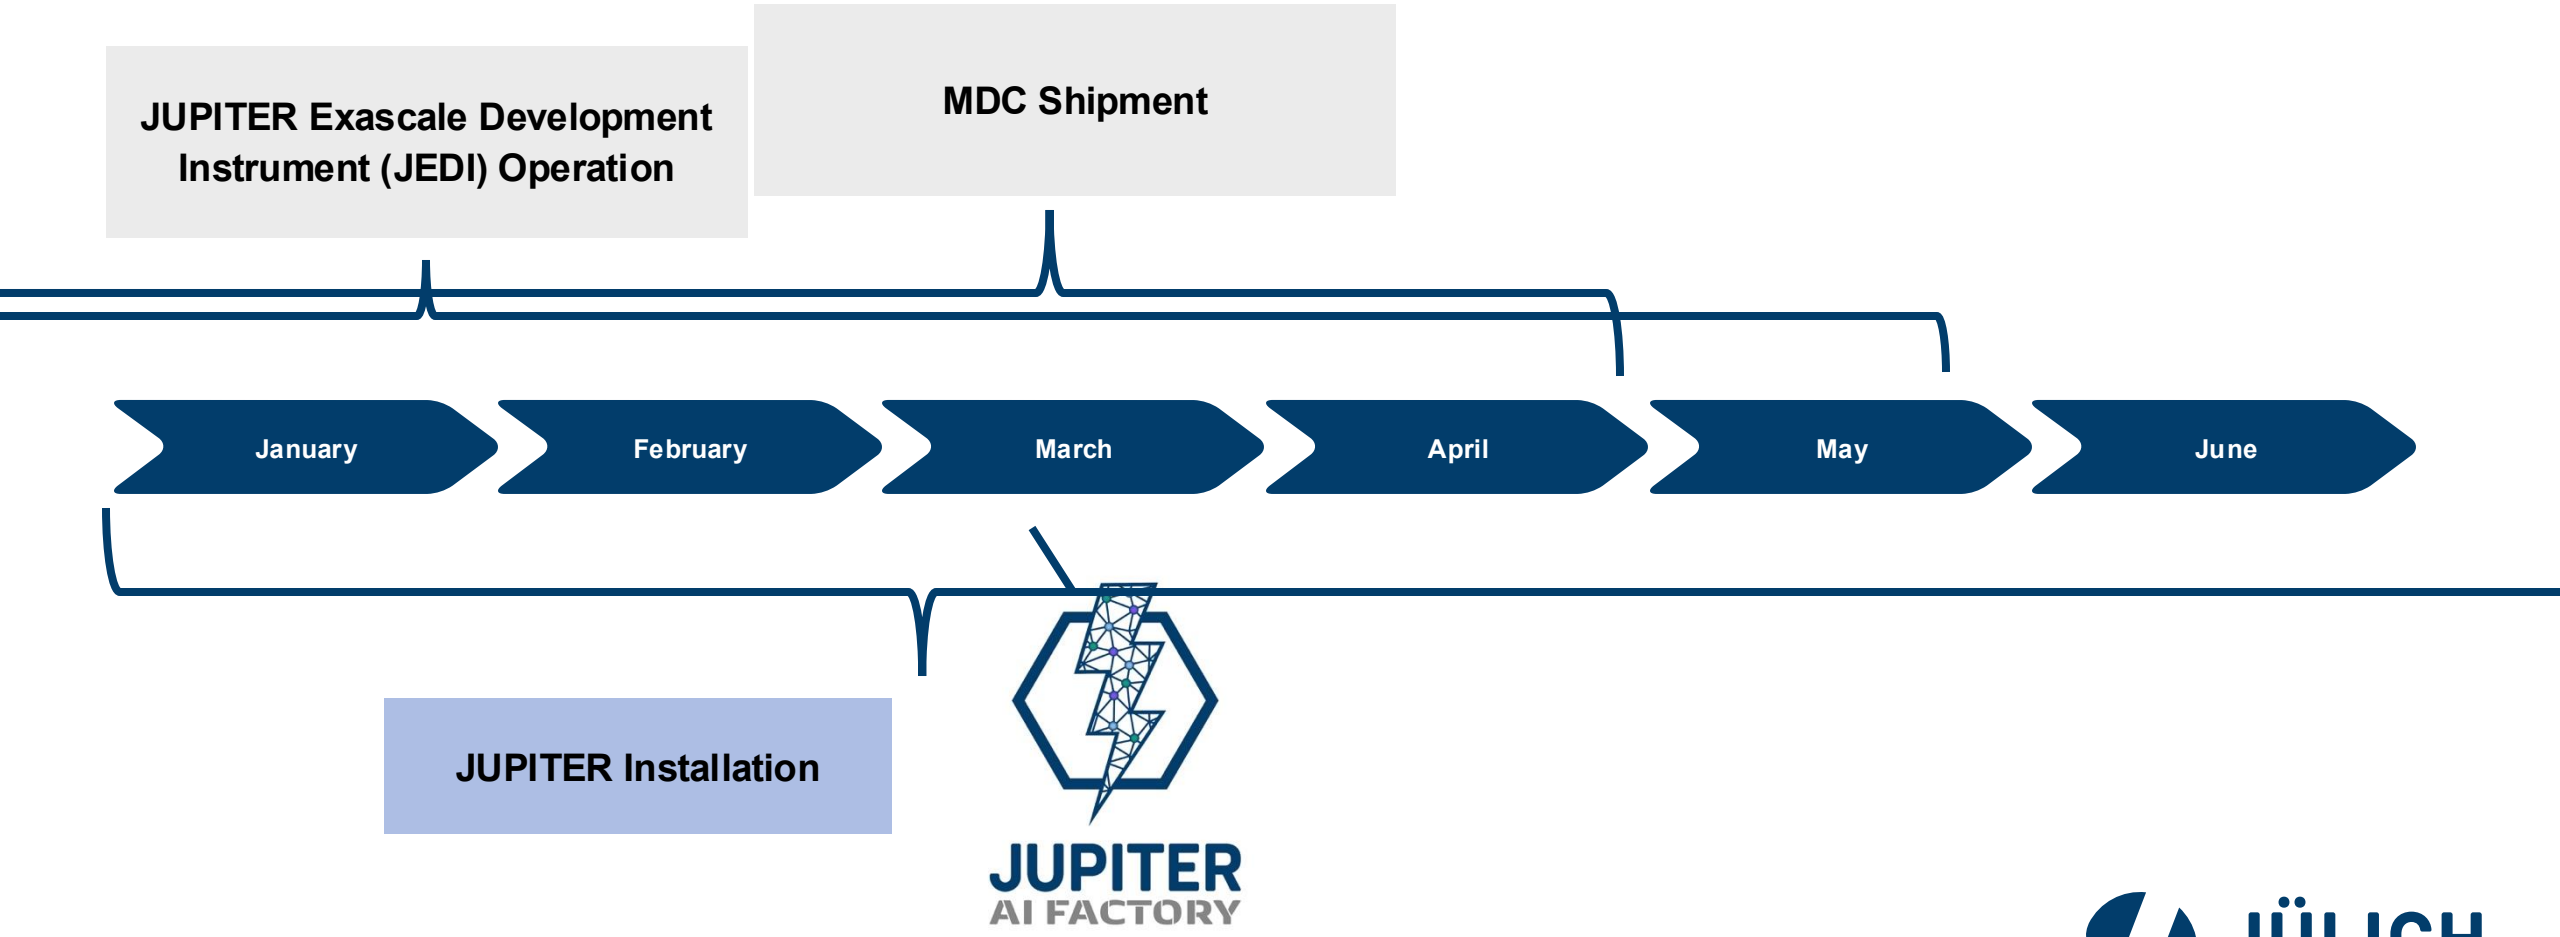


2025-08-29 10:59:57
COFINAIR
DHM

TOP500

- Entry to TOP500 June 2025 as #4 world, #1 EU
- 793 PFLOP/s HPL of 930 PFLOP/s th. peak
- Achieved with 4650 nodes

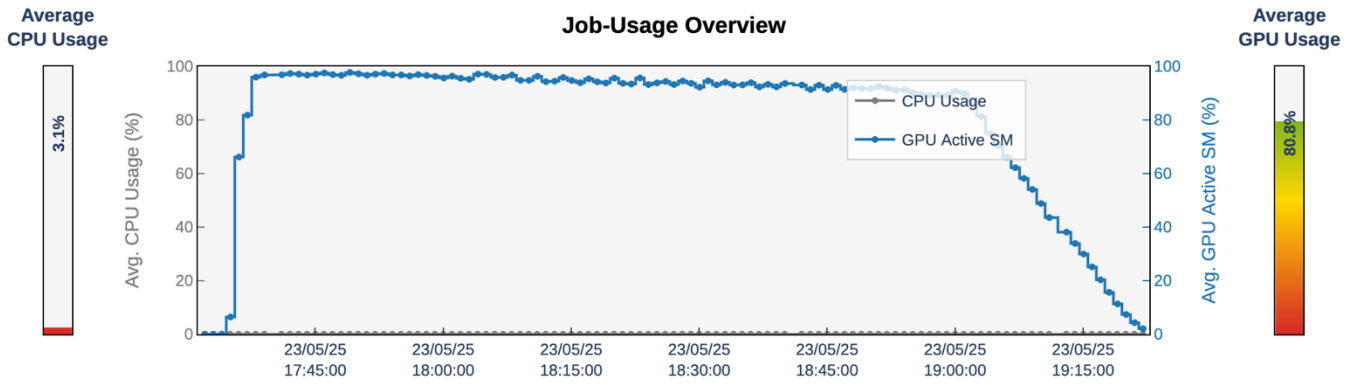

LLview job report of (a) HPL run

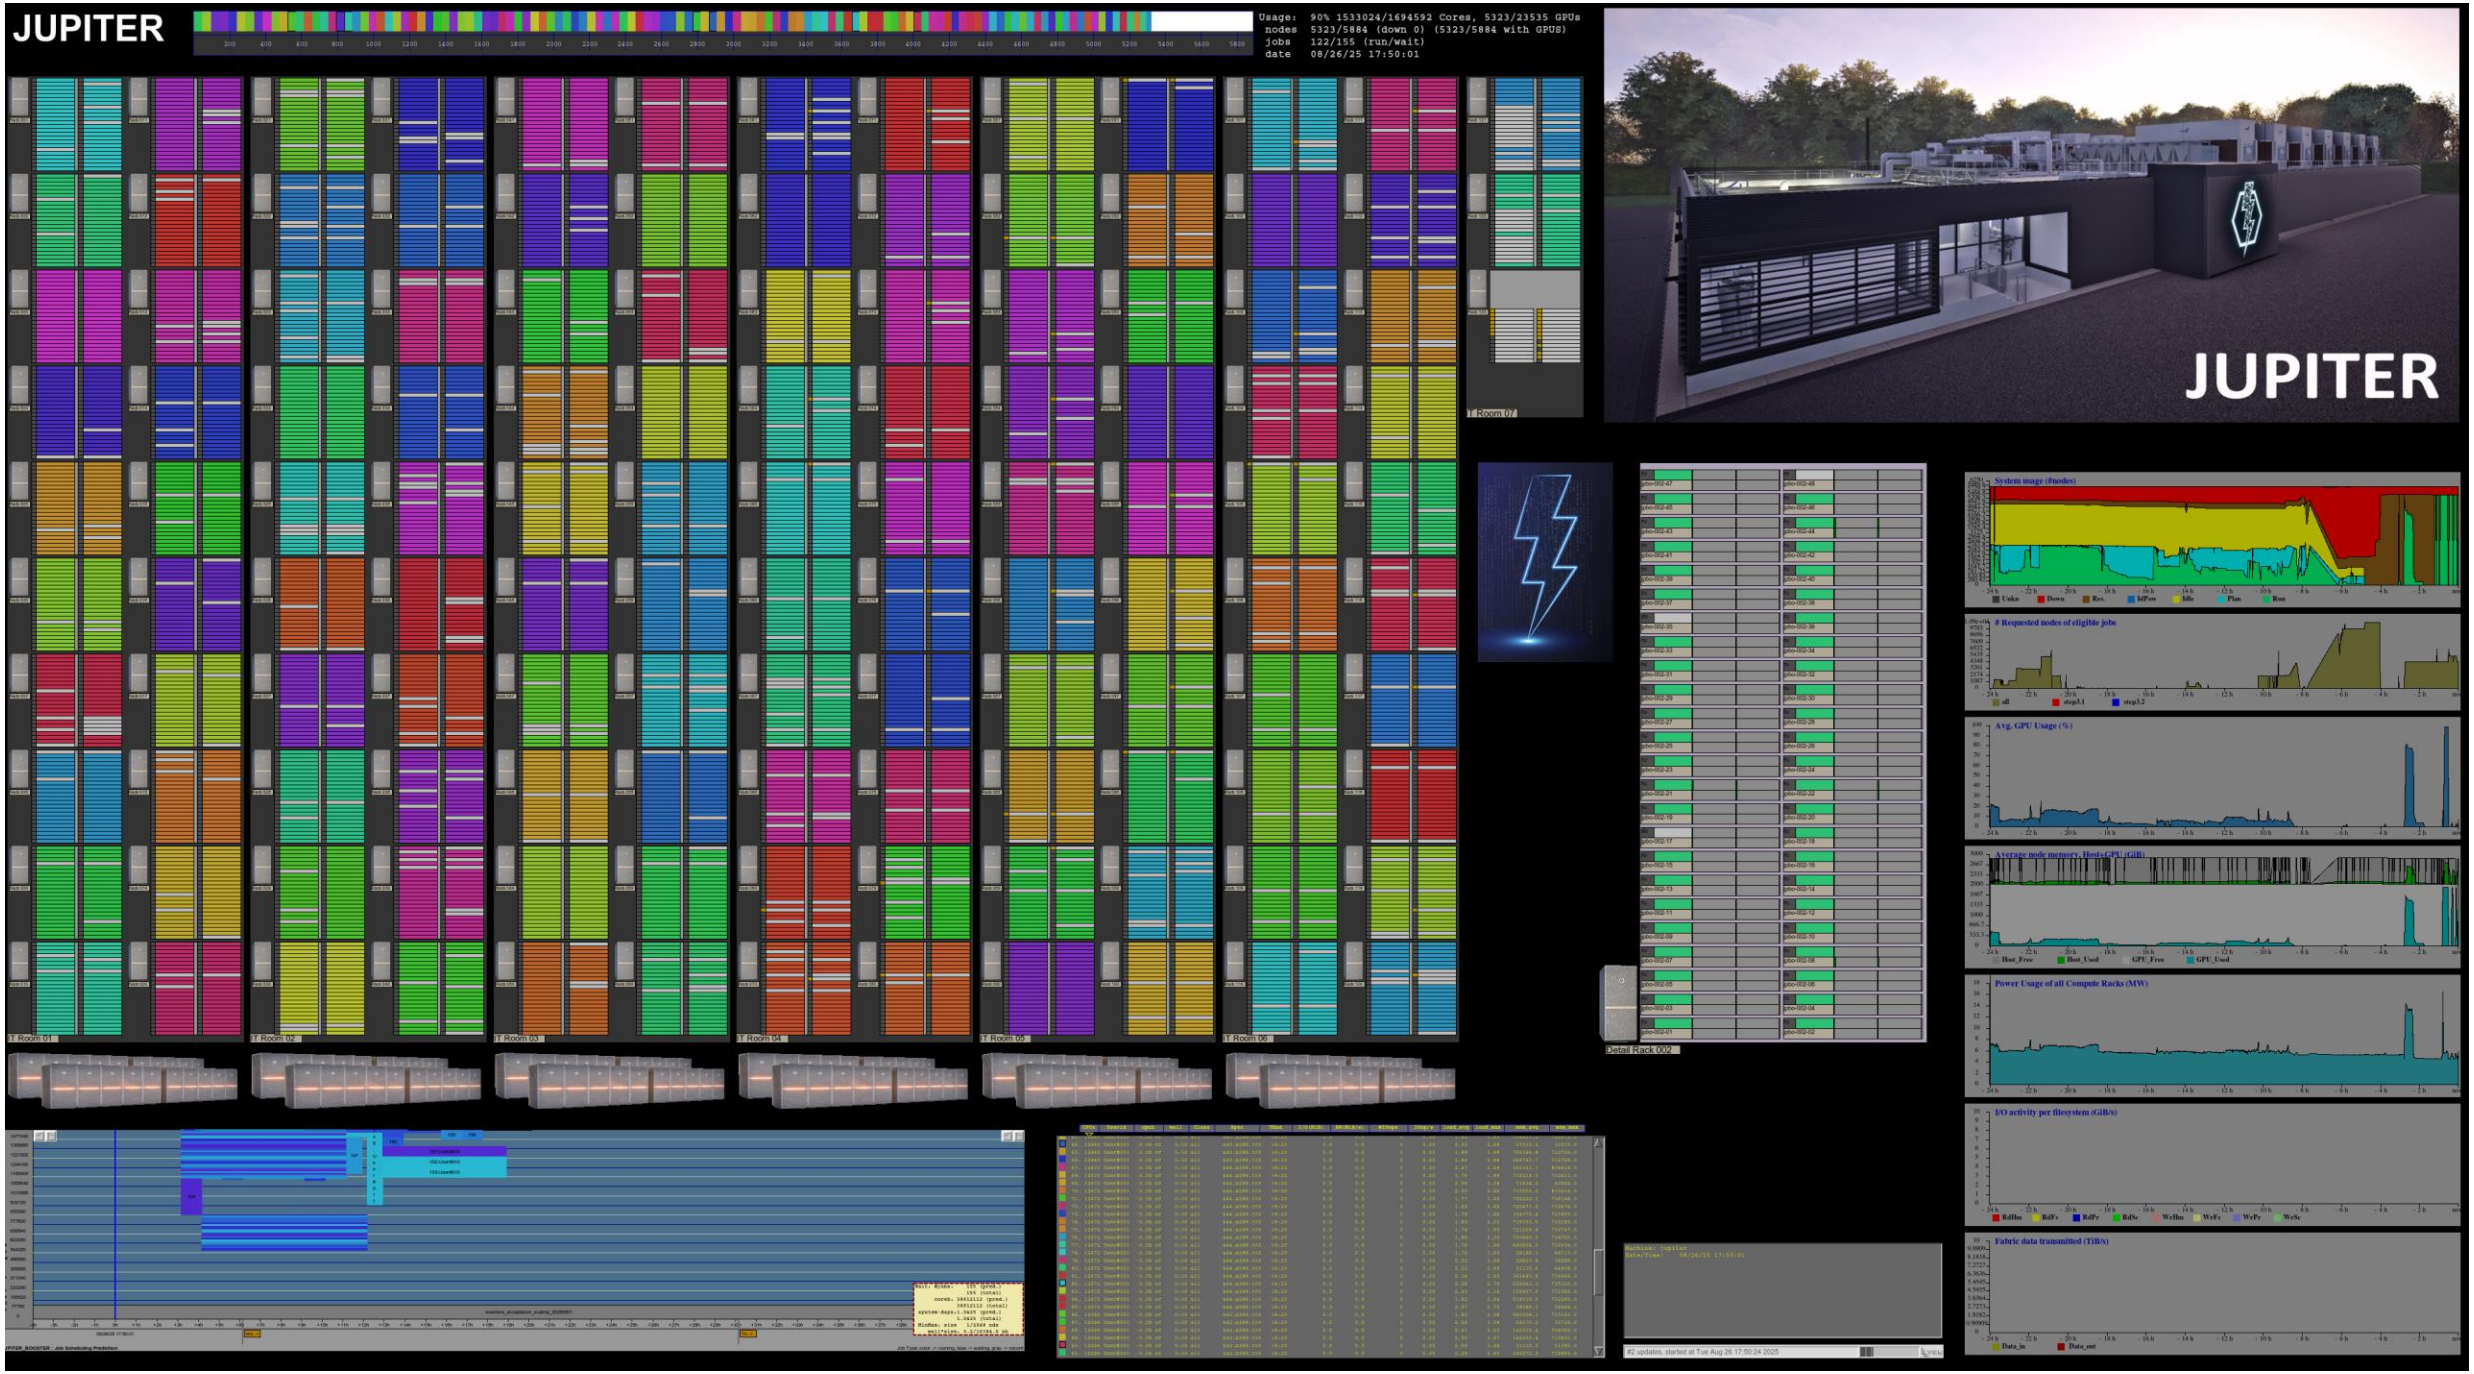
JUPITER
Usage: 90% 1533024/1694592 Cores, 5323/23535 GPUs
nodes 5323/5884 (down 0) (5323/5884 with GPUs)
jobs 122/155 (run/wait)
date 08/26/25 17:50:01
JUPITER
T Room 01
T Room 02
T Room 03
T Room 04
T Room 05
T Room 06
T Room 07
Detail Rack 002
System usage (#nodes)
Requested nodes of eligible jobs
Avg. GPU Usage (%)
Average node memory, Host&GPU (GiB)
Power Usage of all Compute Racks (MW)
I/O activity per filesystem (GiB/s)
Fabric data transmitted (TiB/s)

POPULATING JUPITER

Current Status and Next Steps

- JUPITER Research and Early Access Phase ongoing
  - > 100 Applications, 33 Lighthouses
- GCS AI Competition access started
- (Acceptance) Benchmarking running
  - TCO Application mix, High-Scaling
  - Storage Acceptance
- High-Performance Linpack
  - … otherwise it would be too easy
- **JUPITER Inauguration on 5.9.2025**
- Autumn/winter: EuroHPC/GCS calls

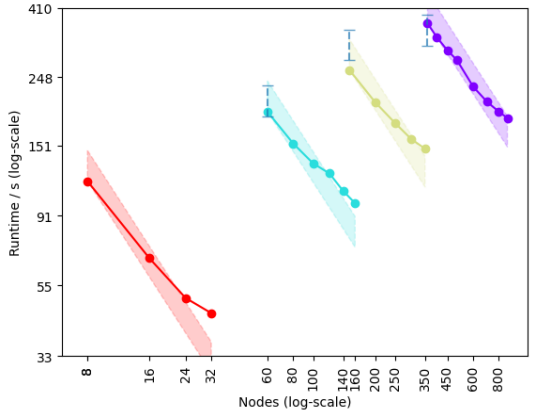

QUESTIONS?!


Welcome to
JUPITER
JU Pioneer for
Innovative and Transformative
Exascale Research

JOINING FORCES

Ministerium für
Kultur und Wissenschaft
des Landes Nordrhein-Westfalen

fz-juelich.de/jupiter

JUPITER – STORAGE (SCRATCH)

- Gross Capacity: 29 PB; Net Capacity: 21 PB
- Bandwidth: 2.1 TB/s Write, 3.1 TB/s Read
- 20× IBM SSS6000 Building Blocks (40 servers)
  - 2× NDR400 per server
  - 48× 30 TB NVMe drives per block
  - IBM Storage Scale (aka Spectrum Scale/GPFS)
- Manager and Datamover Nodes
- Exclusive for JUPITER
  - Integrated into InfiniBand fabric

Preliminary numbers, might change during installation

JUPITER – STORAGE (EXASTORE)

TCO contribution from JSC, not part of the JUPITER procurement

- Gross Capacity: 308 PB; Net Capacity: 210 PB
- Bandwidth: 1.1 TB/s Write, 1.4 TB/s Read
- 22× IBM SSS6000 Building Blocks (44 servers)
  - 2× NDR200 per server
  - 7× JBOD enclosures, each with 91x 22 TB Spinning Disks per Building Block (14014 disks)
  - IBM Storage Scale (aka Spectrum Scale/GPFS)
- Manager and Datamover Nodes
- Exclusive for JUPITER
  - Integrated into InfiniBand fabric

Preliminary numbers, might change during installation

JUPITER – STORAGE (EXATAPE)

TCO contribution from JSC, not part of the JUPITER procurement

- 369 Petabyte Tape Capacity
- Procurement Q3/Q4 2024, 2M€ Invest
- 2 x IBM TS4500 with LTO9
- TS1 in 16.4
  - 11 Frames
  - 15x LTO 9 tape drives
  - 10240 LTO9 media
- TS2 in 16.3
  - 11 Frames
  - 15x LTO9 tape drives
  - 10260 LTO9 media