JUPITER - FROM LAUNCH TO LANDING

Development Operations (DevOps) School for HPC – CARLA 2025

2025-09-24 I JÜLICH SUPERCOMPUTING CENTRE | BERND MOHR, DAMIAN ALVAREZ, 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

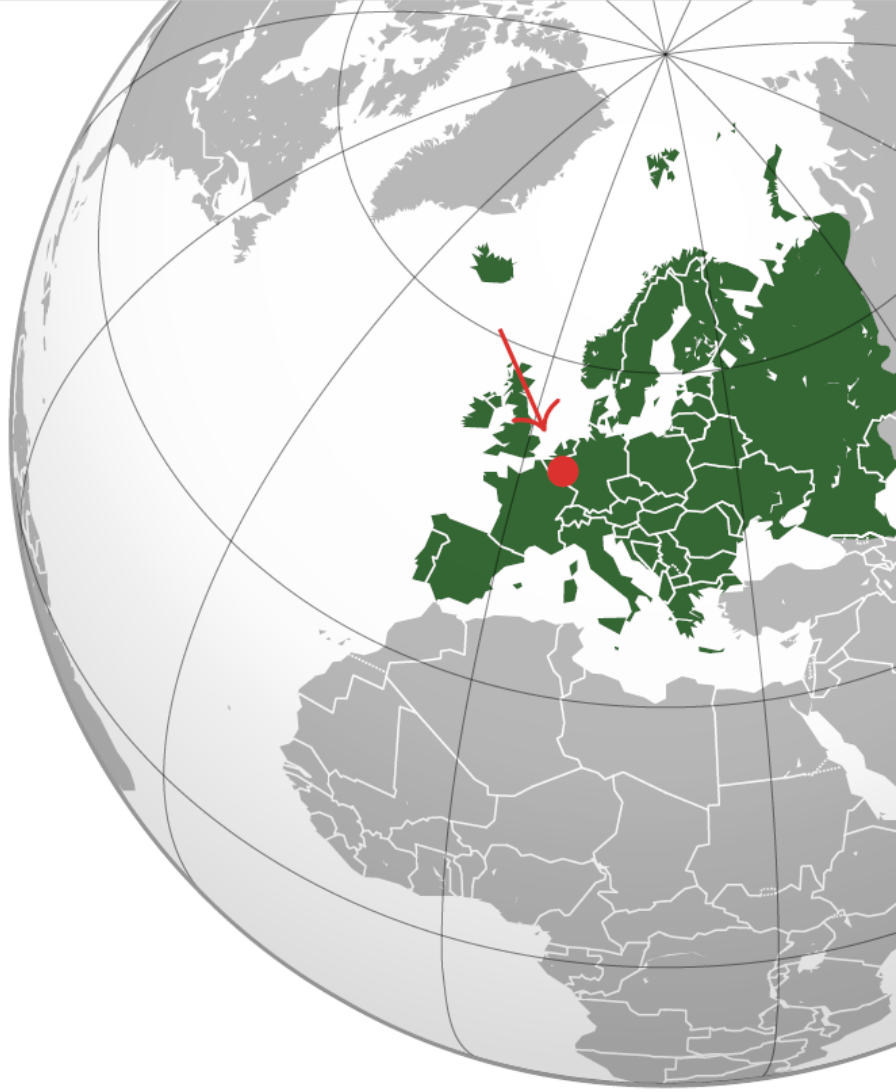

Forschungszentrum Jülich GmbH

JÜLICH SUPERCOMPUTING CENTRE

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

FORSCHUNGSZENTRUM JÜLICH GMBH

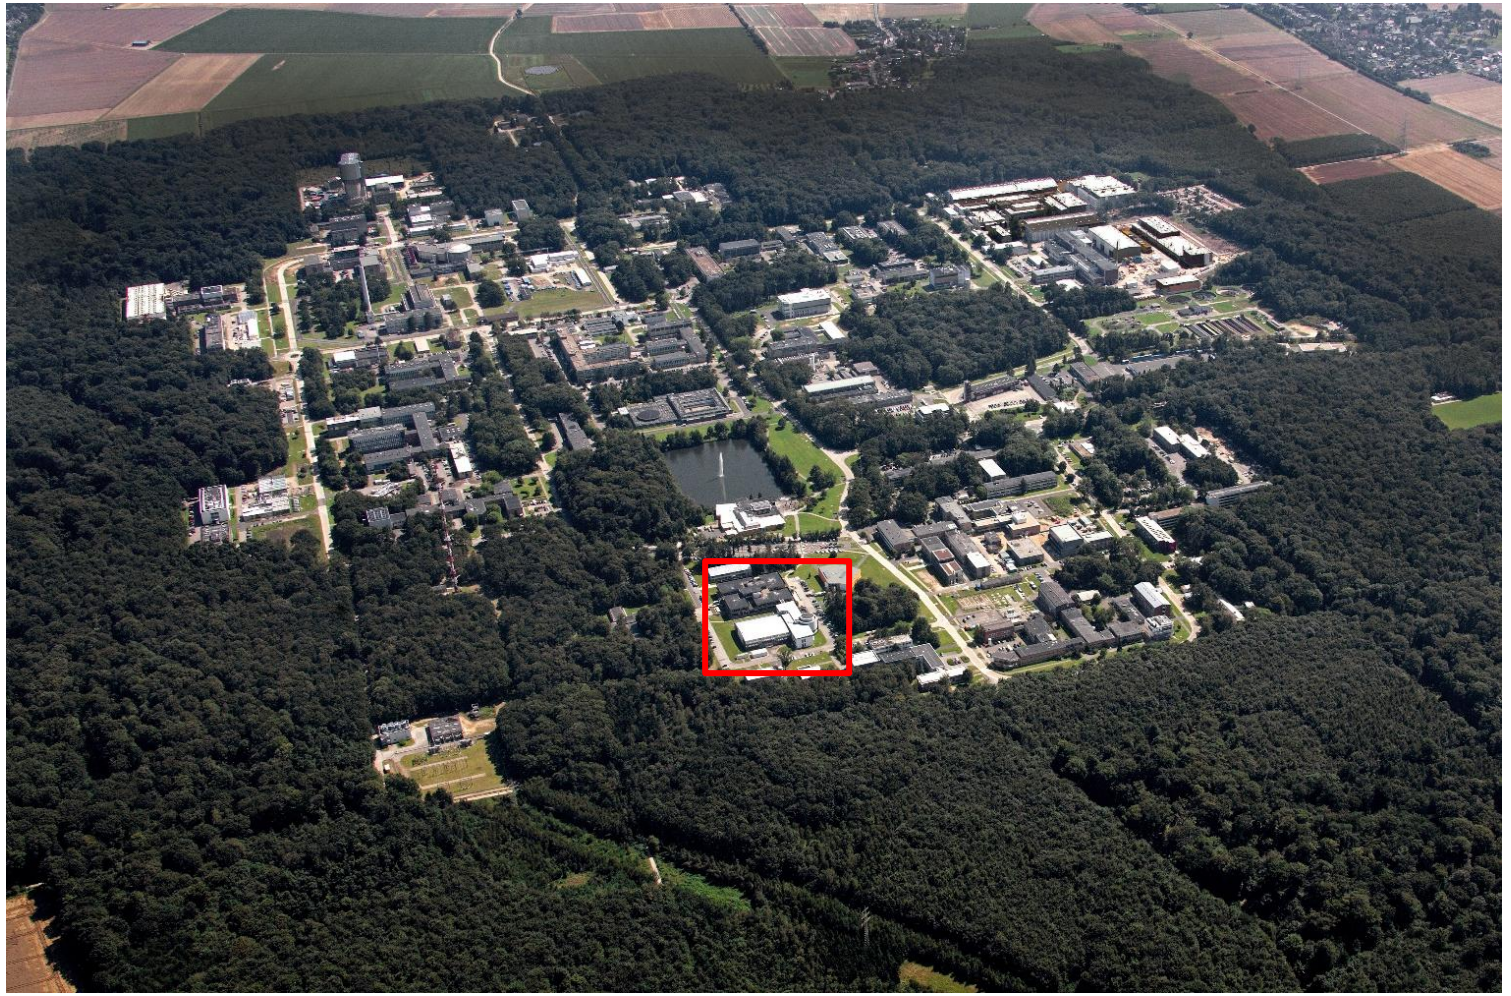

- Germany's largest national laboratory
- About 7500 employees
- Research areas
  - Information technology
  - Health (Neuroscience / brain research)
  - Energy
  - Atmosphere + Climate

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

JÜLICH SUPERCOMPUTING CENTRE (JSC)

HPC Centre for

- Forschungszentrum Jülich
- Jülich Aachen Research Alliance (JARA)
- Germany as GCS (1 of 3 German National Centres)
- Europe (EuroHPC JU) (1st European Centre inside PRACE)

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

The project still is like plasma in lightning:

- Hot
- Rapid changes
- Difficult to predict
- Want to touch it!
- Can kill me

Or like the planet:

- Solid core
- Nebulous surroundings
- Difficult to reach
- Fascinating
- Can also kill me!

J oint
U ndertaking
P ioneer for
I nnovative and
T ransformative
E xascale
R esearch

THE DISTANT PAST

BECOMING A HOSTING ENTITY/SITE

Call for Expression of Interest to host a high-end supercomputer by EuroHPC JU
17.12.2021

BECOMING A HOSTING ENTITY/SITE

CALL FOR EXPRESSION OF INTEREST for the selection of a Hosting Entity for a high-end Supercomputer

The objective of the call is to select hosting entities across the European Union that will support the acquisition and operation of the next generation of EuroHPC supercomputers.

PAGE CONTENTS

Details

Status	CLOSED
Reference	EUROHPC-2021-CEI-EXA-01
Publication date	17 December 2021
Opening date	17 December 2021
Deadline model	Single-stage
Deadline date	14 February 2022, 12:00 (CET)

Description

The EuroHPC JU will select a hosting entity for a high-end supercomputer and will conclude a hosting agreement, which will permit to establish a stable and structured partnership between the EuroHPC and the hosting entity for the acquisition and operation of the high-end supercomputer.

BECOMING A HOSTING ENTITY/SITE

Call for Expression of Interest to host a high-end supercomputer by EuroHPC JU
17.12.2021

Deadline for Application
14.02.2022

THE APPLICATION

- Opened 17. December 2021
- Closed 14. February 2022
- Contributions by all JSC divisions
- 86 pages
  - Proposed system architecture
  - Targeted user communities
  - Detailed cost calculations
    - **500m€ TCO**
  - Expected Infrastructure
  - …

Table of Contents

BECOMING A HOSTING ENTITY/SITE

Call for Expression of Interest to host a high-end supercomputer by EuroHPC JU
17.12.2021

Hearings
16.05.2022

Deadline for Application
14.02.2022

Hosting Entity Decision
15.06.2022

JUPITER – HOSTING ENTITY DECISION 15.06.2022

tagesschau

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

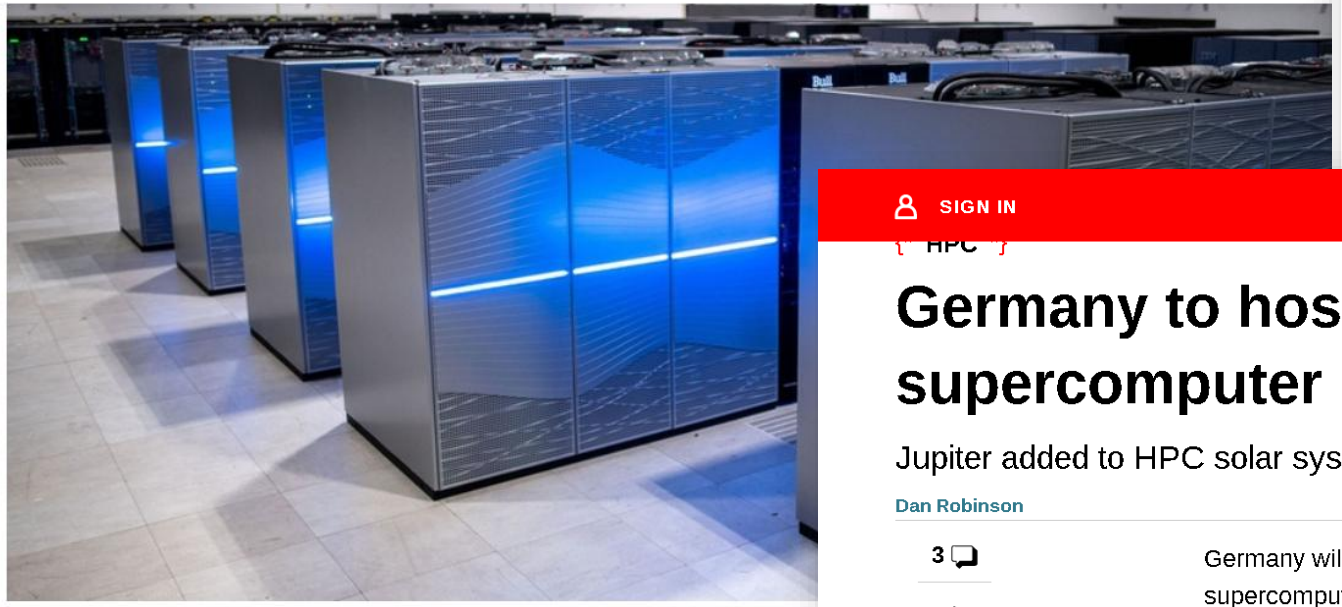

Hochleistungs-Rechner

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 A STARBASE

- **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)

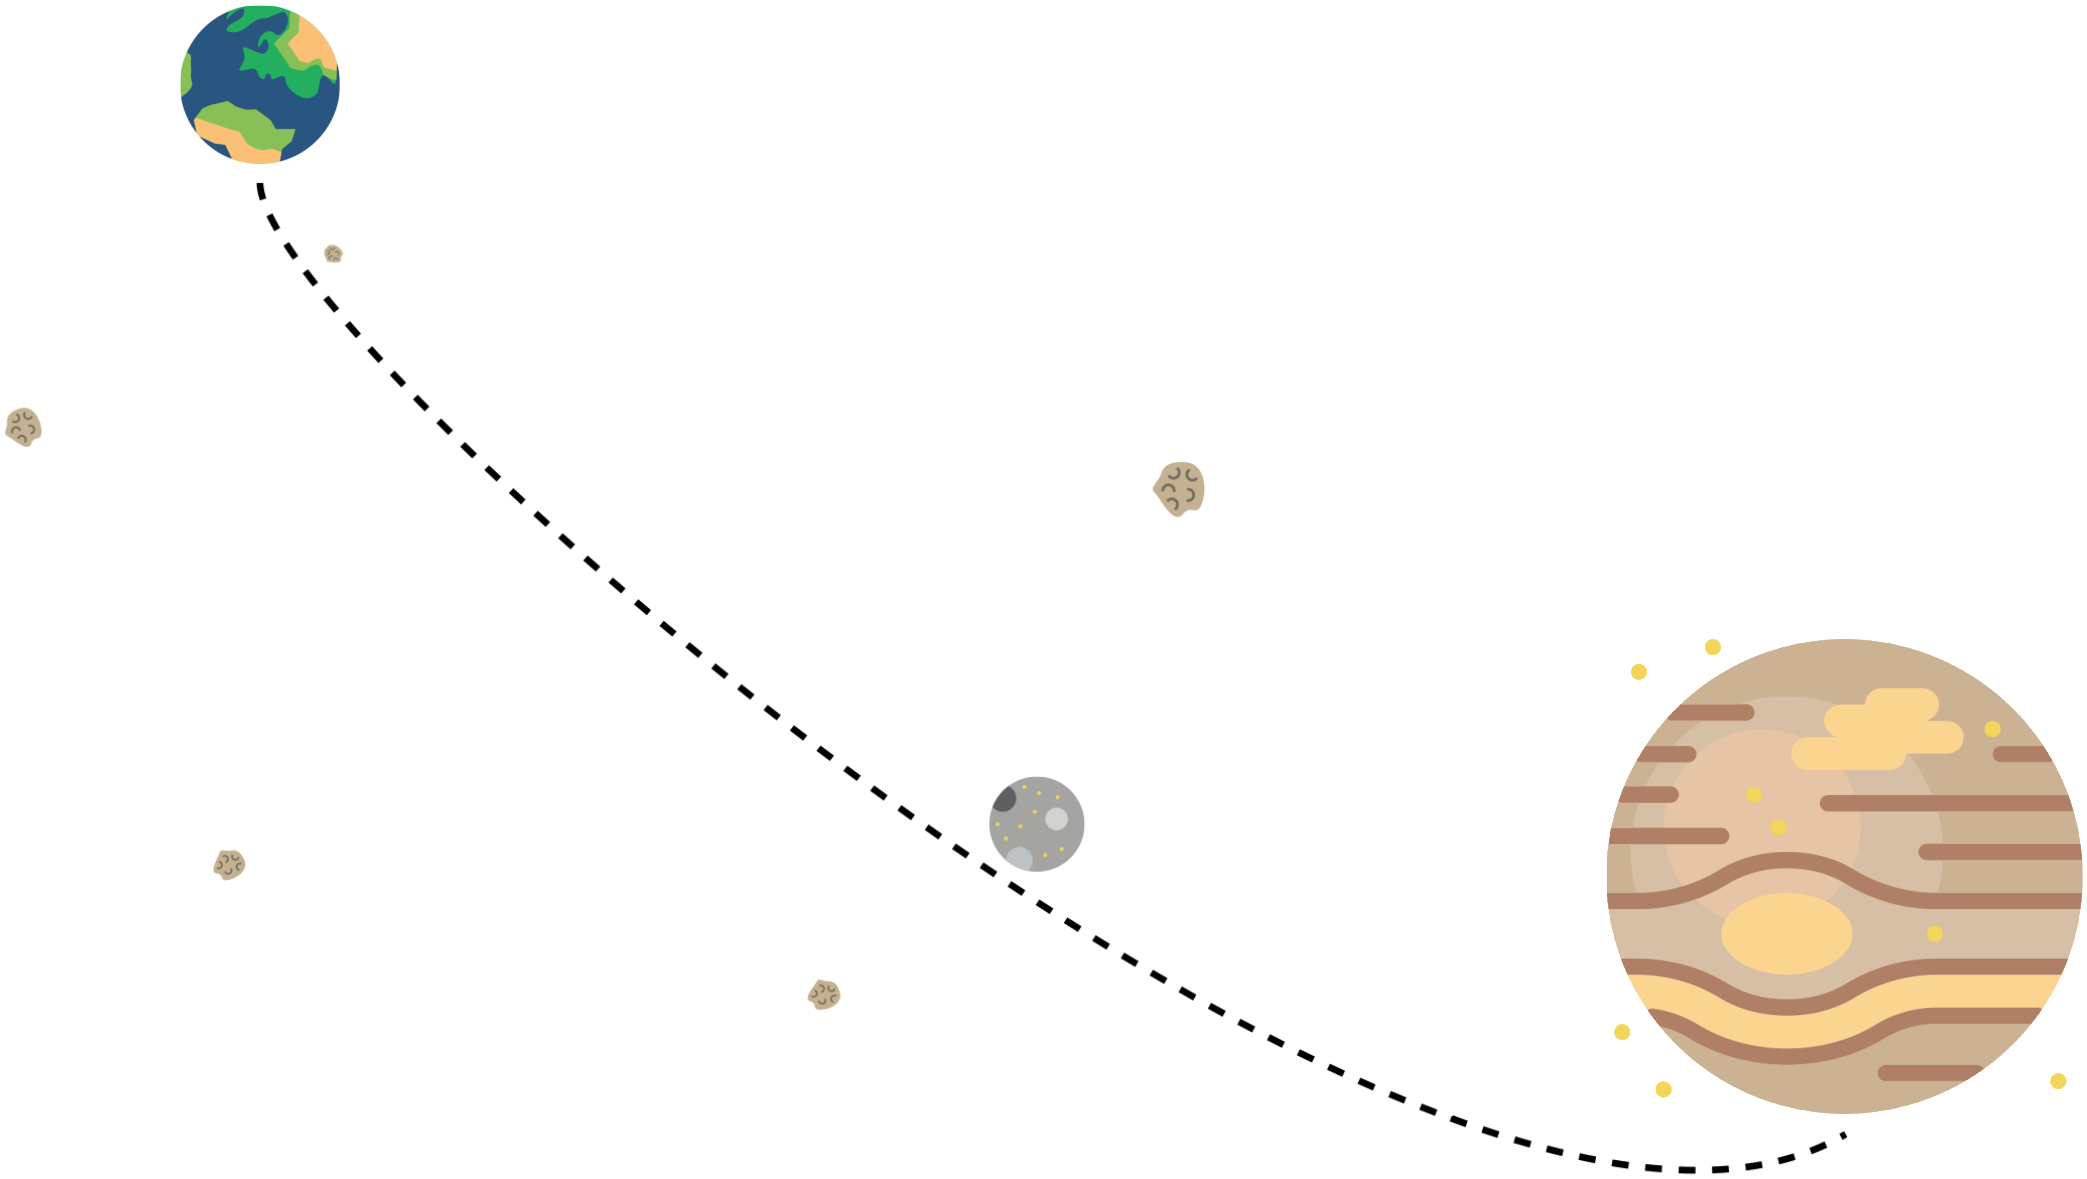

Ready for take off

READY FOR TAKE OFF

Competitive Dialogue - Descriptive Document

- Description of procurement procedure
- Overall budget, **273 M€**
- High-level description of targeted system
  - Implementing the MSA
  - Booster to achieve 1 EF
  - Cluster, preferably based on European IP
  - Flash storage module
  - Interconnect expectations
  - Login system sizing
  - System management

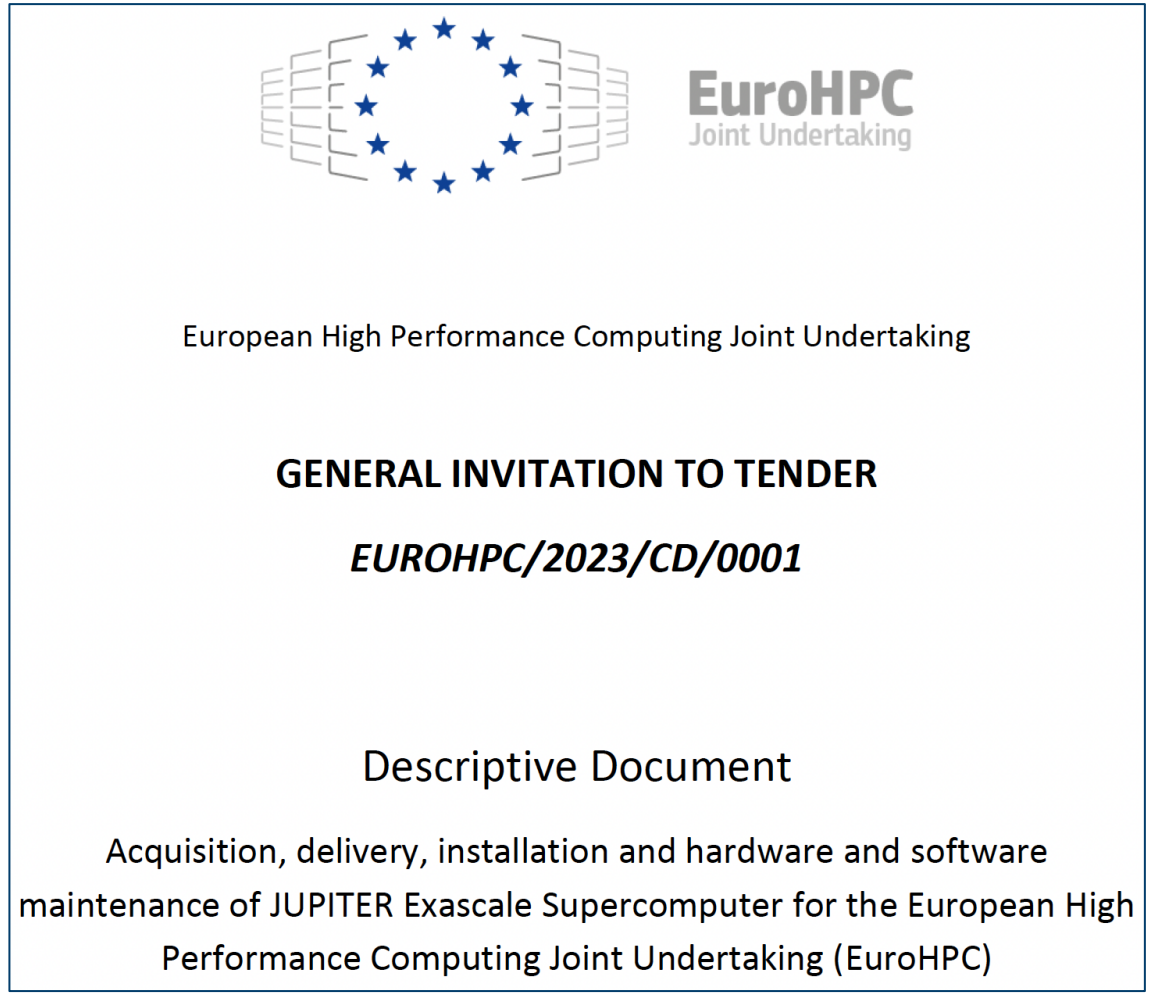

16. Jan 2023: Publish Call (Descriptive Document)

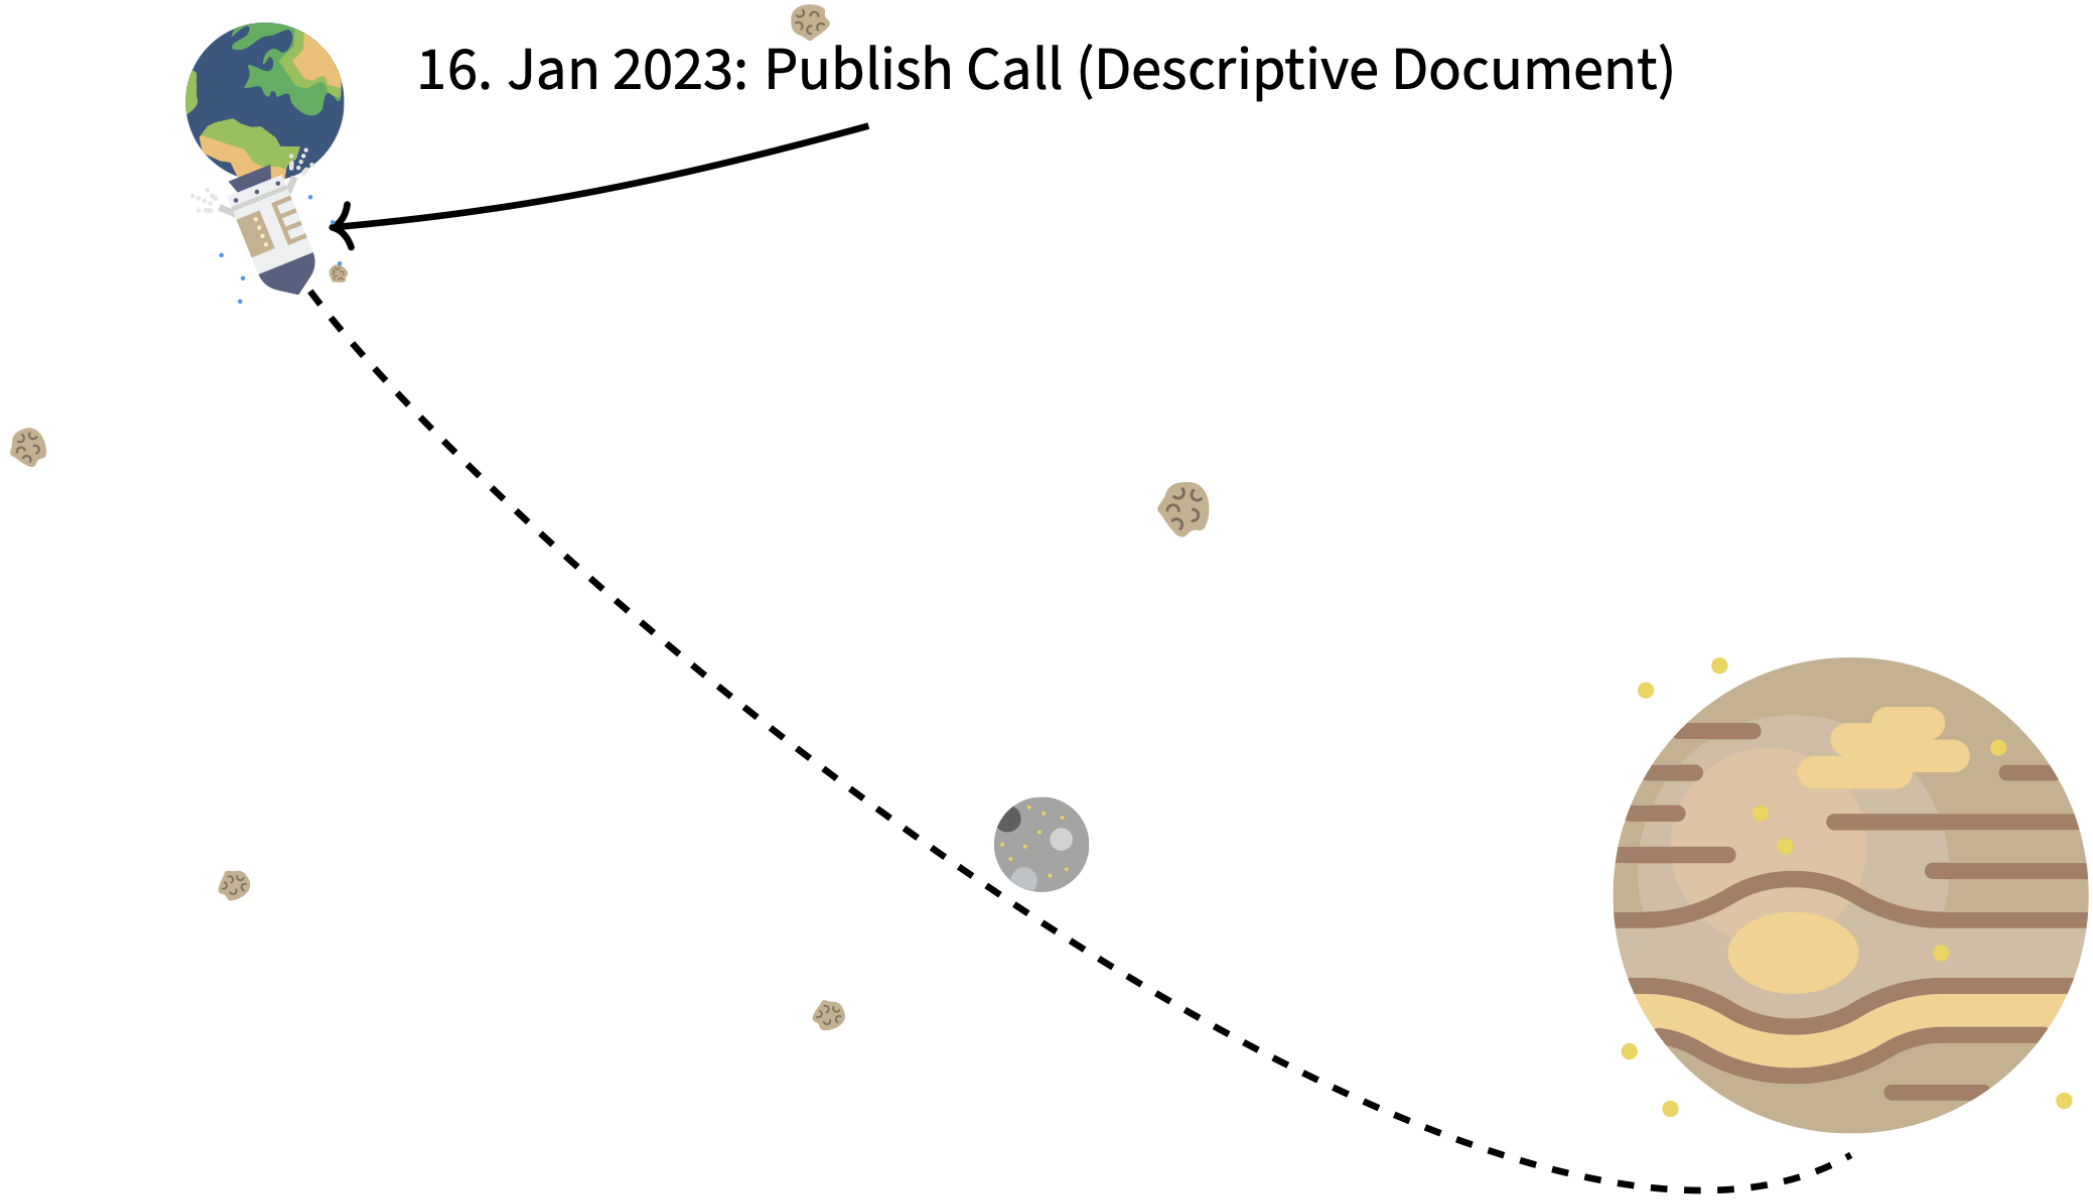

PROCUREMENT START – PUBLISHING THE CALL

CALL FOR PROPOSALS | Closed

Acquisition, Delivery, Installation and Hardware and Software Maintenance of JUPITER Exascale Supercomputer for the European High Performance Computing Joint Undertaking

The purpose of this call is to select one economic operator for the component acquisition, delivery, assembly, hardware and software installation and maintenance of JUPITER exascale supercomputer that will be owned by the EuroHPC JU.

PAGE CONTENTS

Details

Status	CLOSED
Reference	EUROHPC/2023/CD/0001
Publication date	16 January 2023
Opening date	16 January 2023
Deadline model	Single-stage
Deadline date	17 February 2023, 17:00 (CET)
Funding programme	Digital Europe Programme

16. Jan 2023: Publish Call (Descriptive Document)

17. Feb 2023: Deadline for Request for Participation

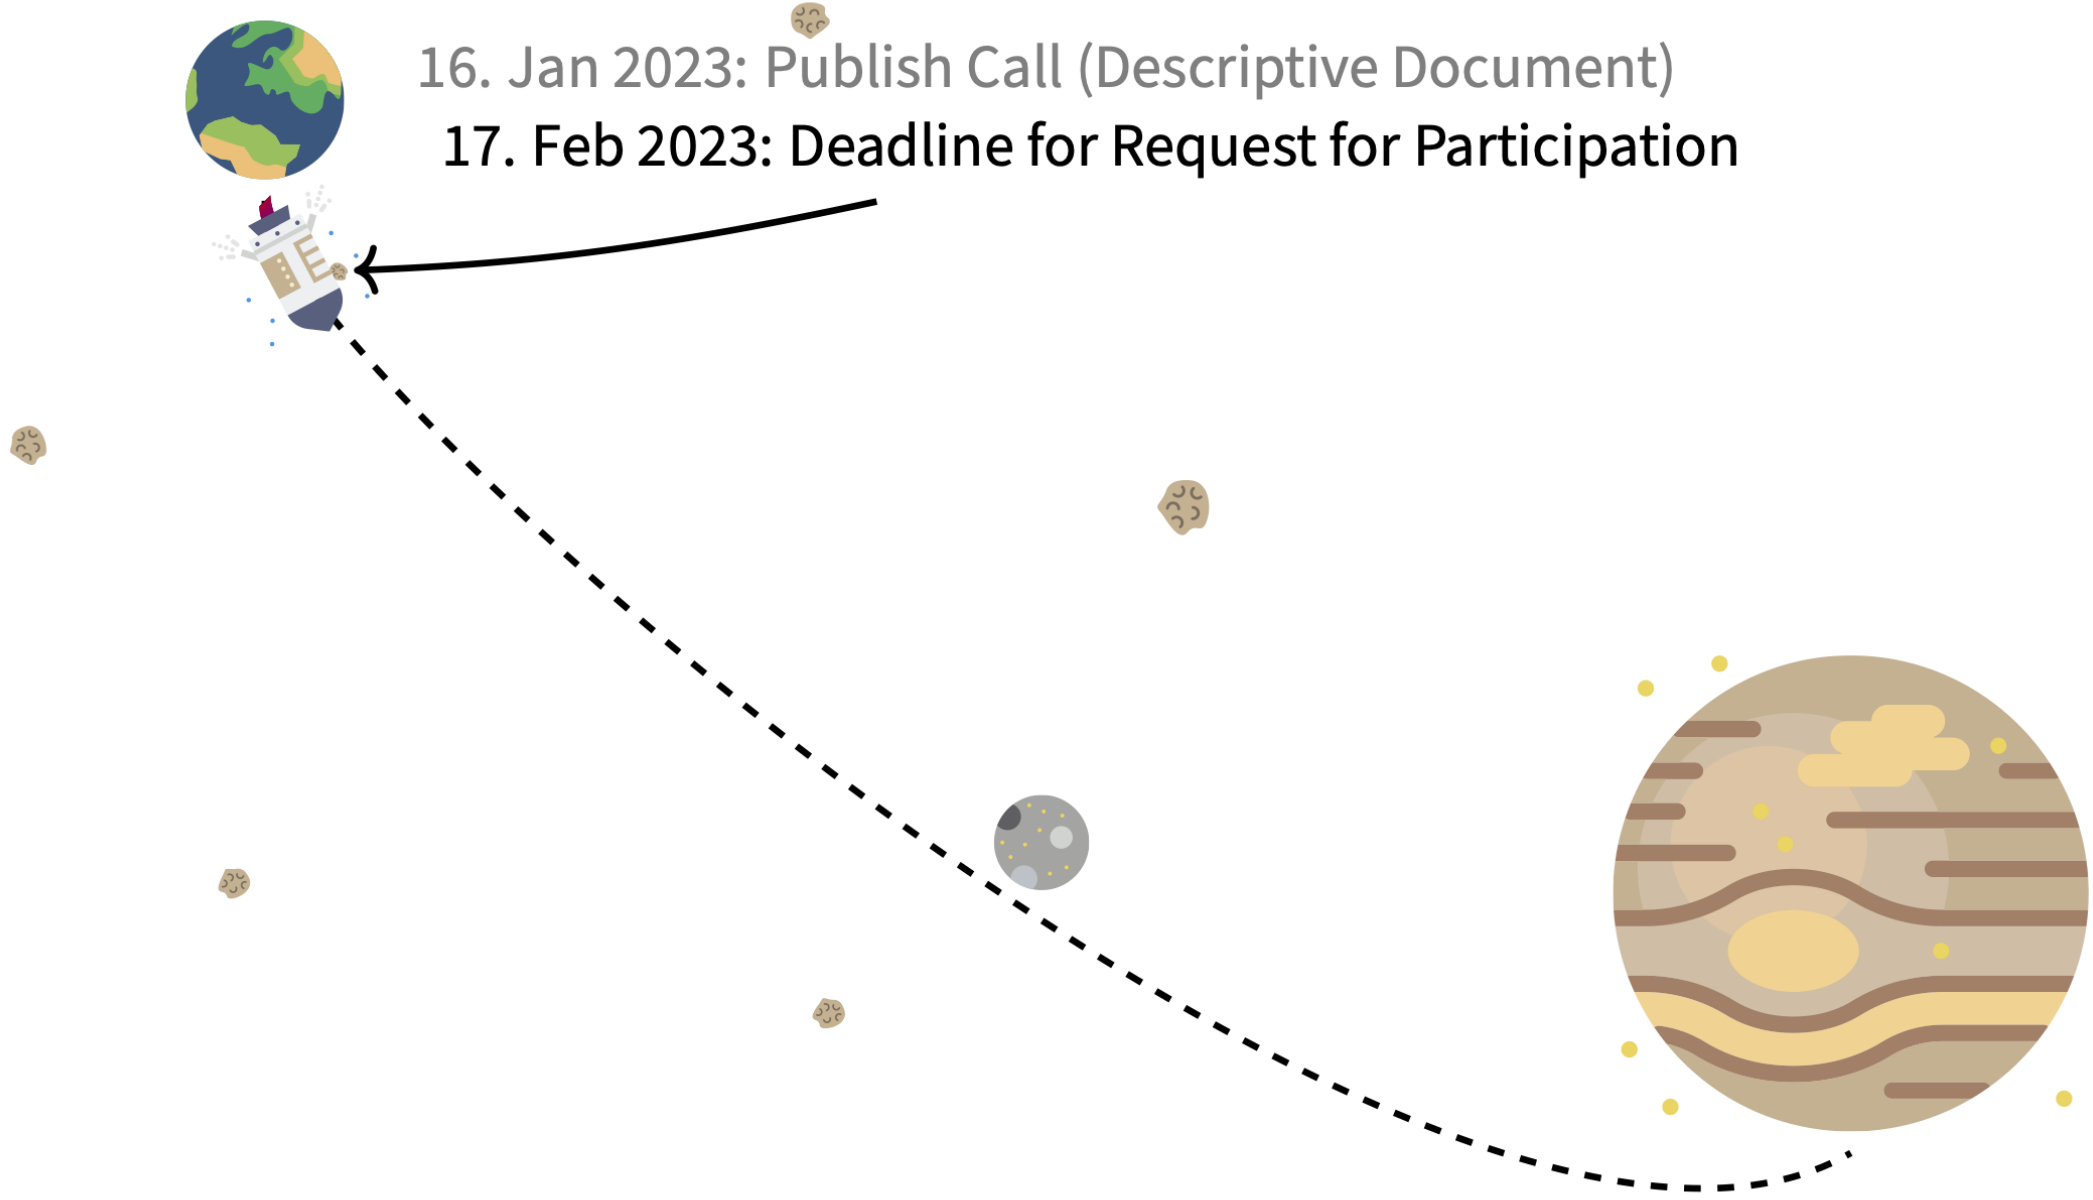

16. Jan 2023: Publish Call (Descriptive Document)

17. Feb 2023: Deadline for Request for Participation

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

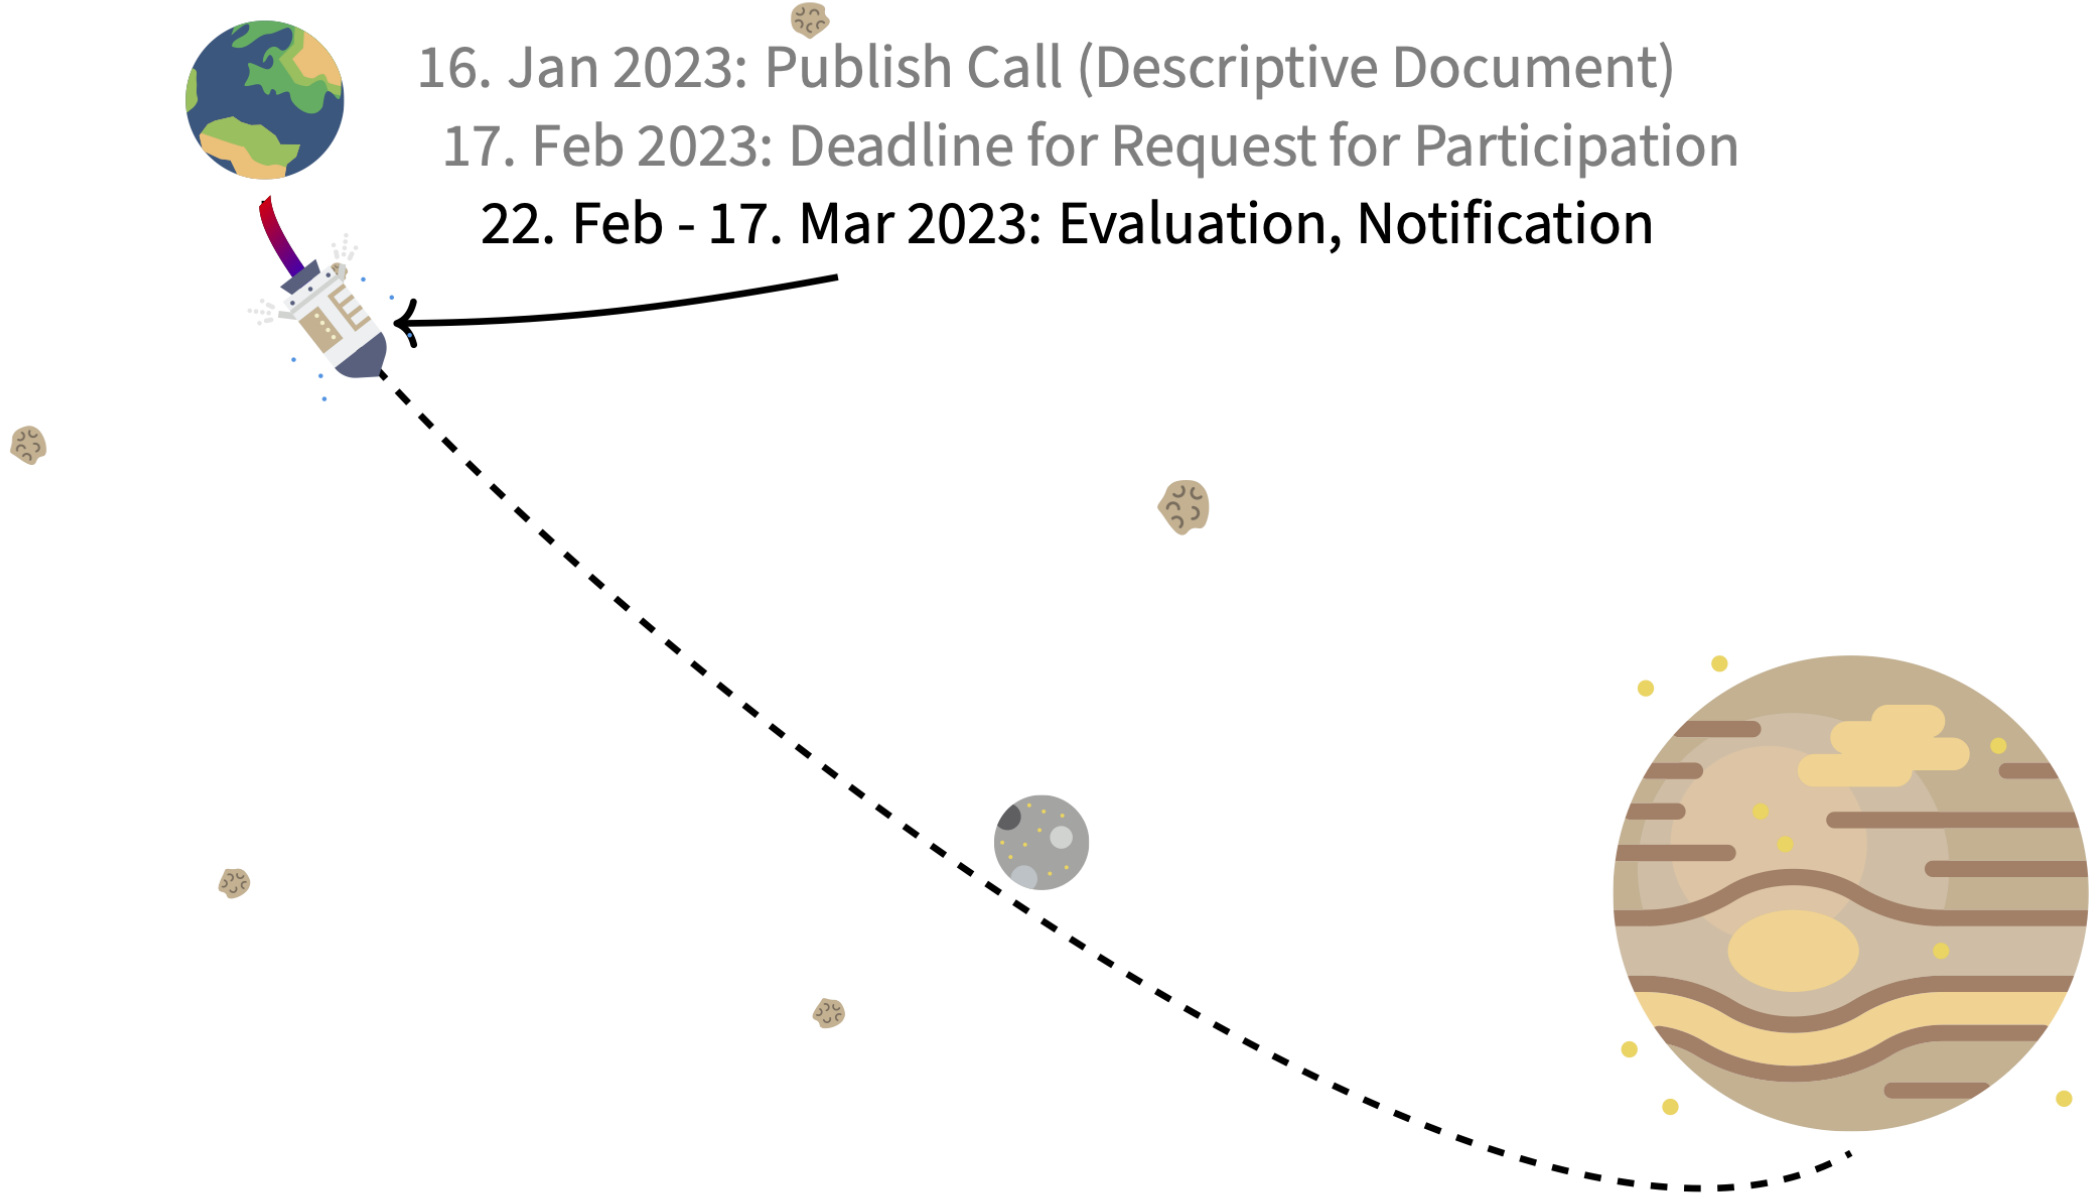

EVALUATION

6. March 2023

- Two proposals
  - HPE
  - ParTec-Eviden consortium
- Two times 80+ pages
  - Technical, legal, financial, CVs, blabla, …
- No digital copy (read paper!)
- One day in Luxembourg
- Two (or three) reviewers
  - Two by JSC
- Final evaluation report compiled by EuroHPC

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

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


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

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

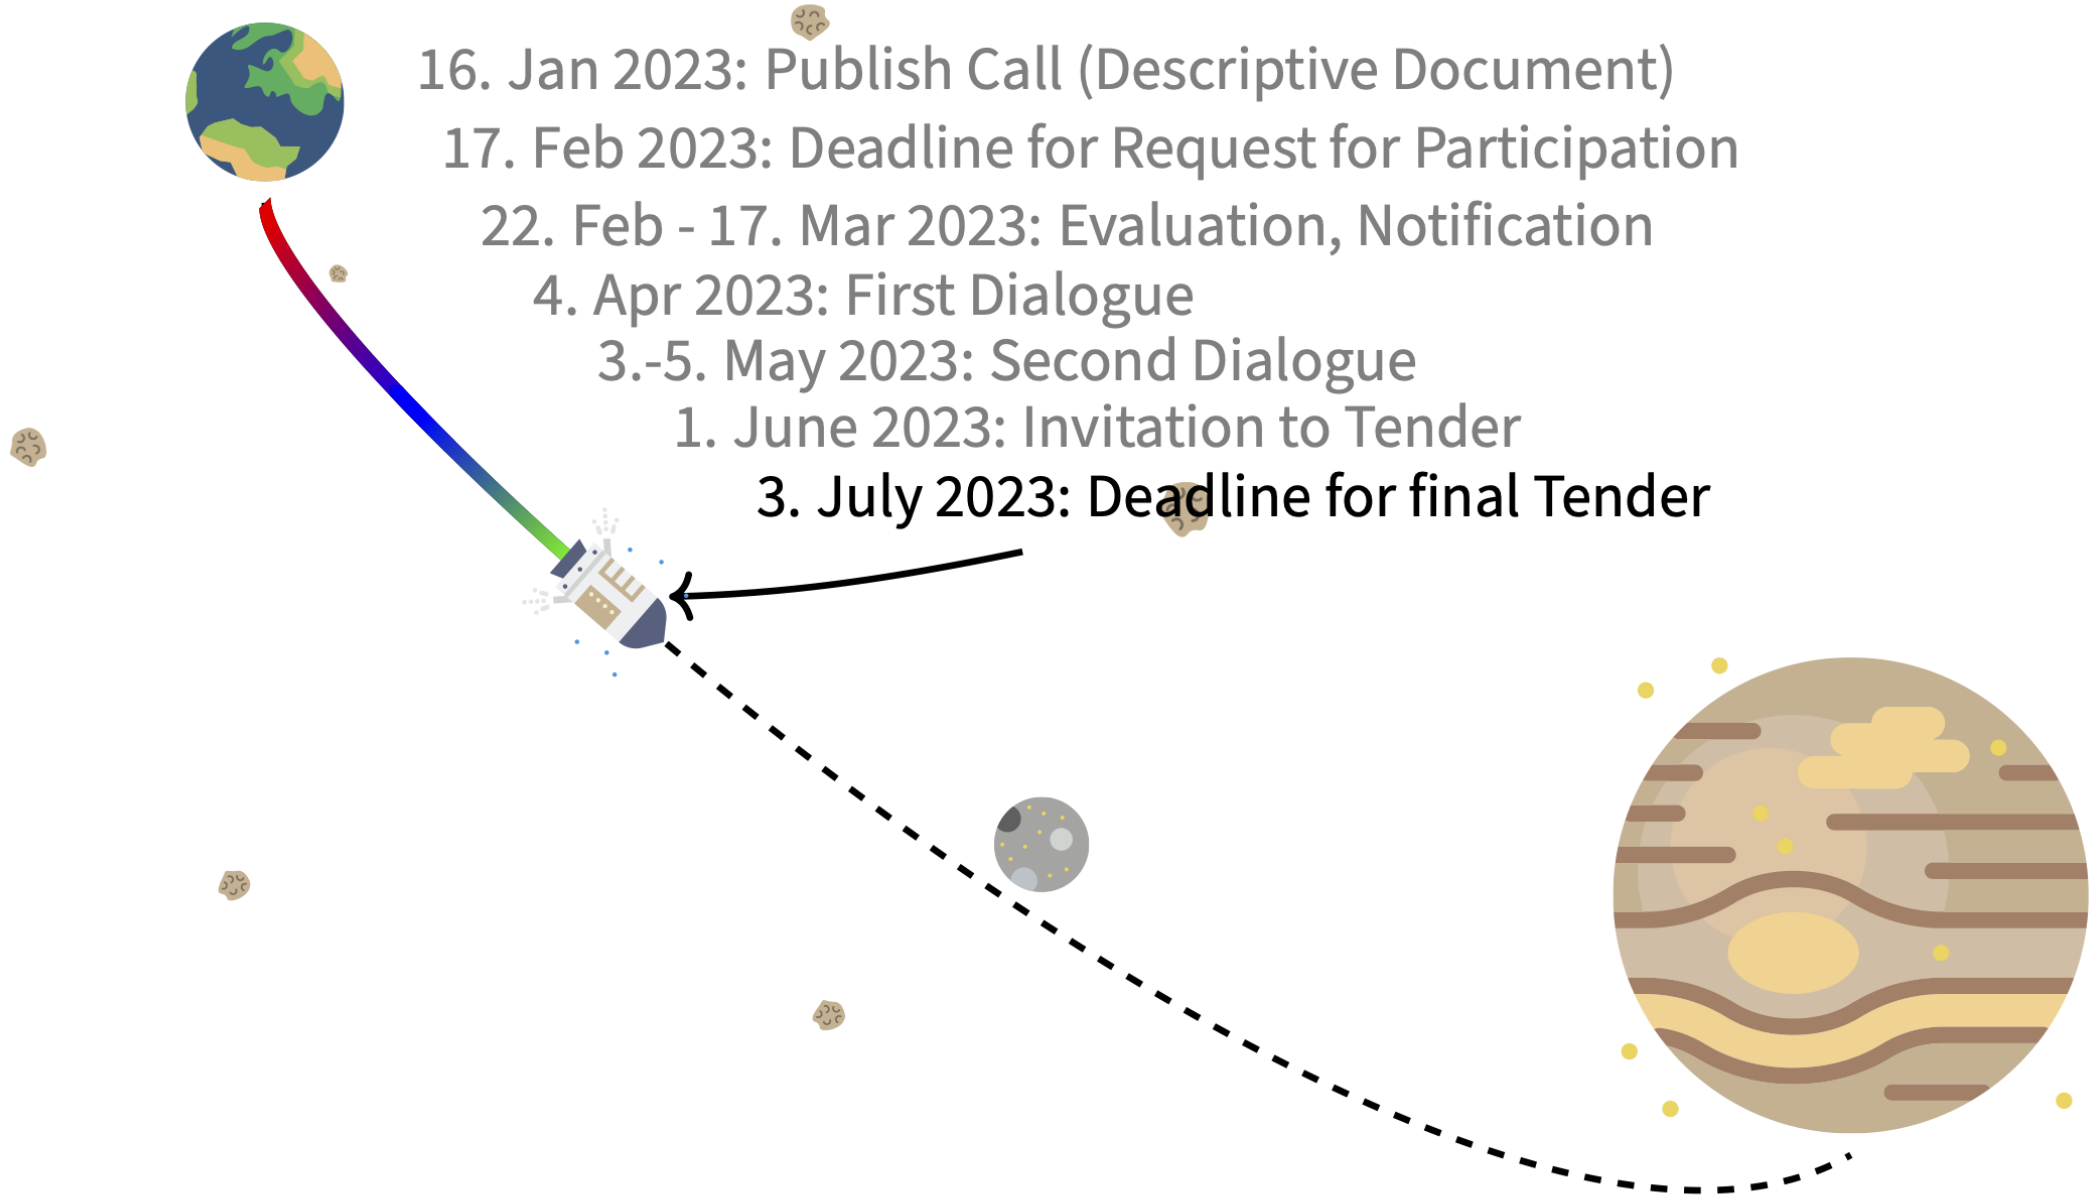

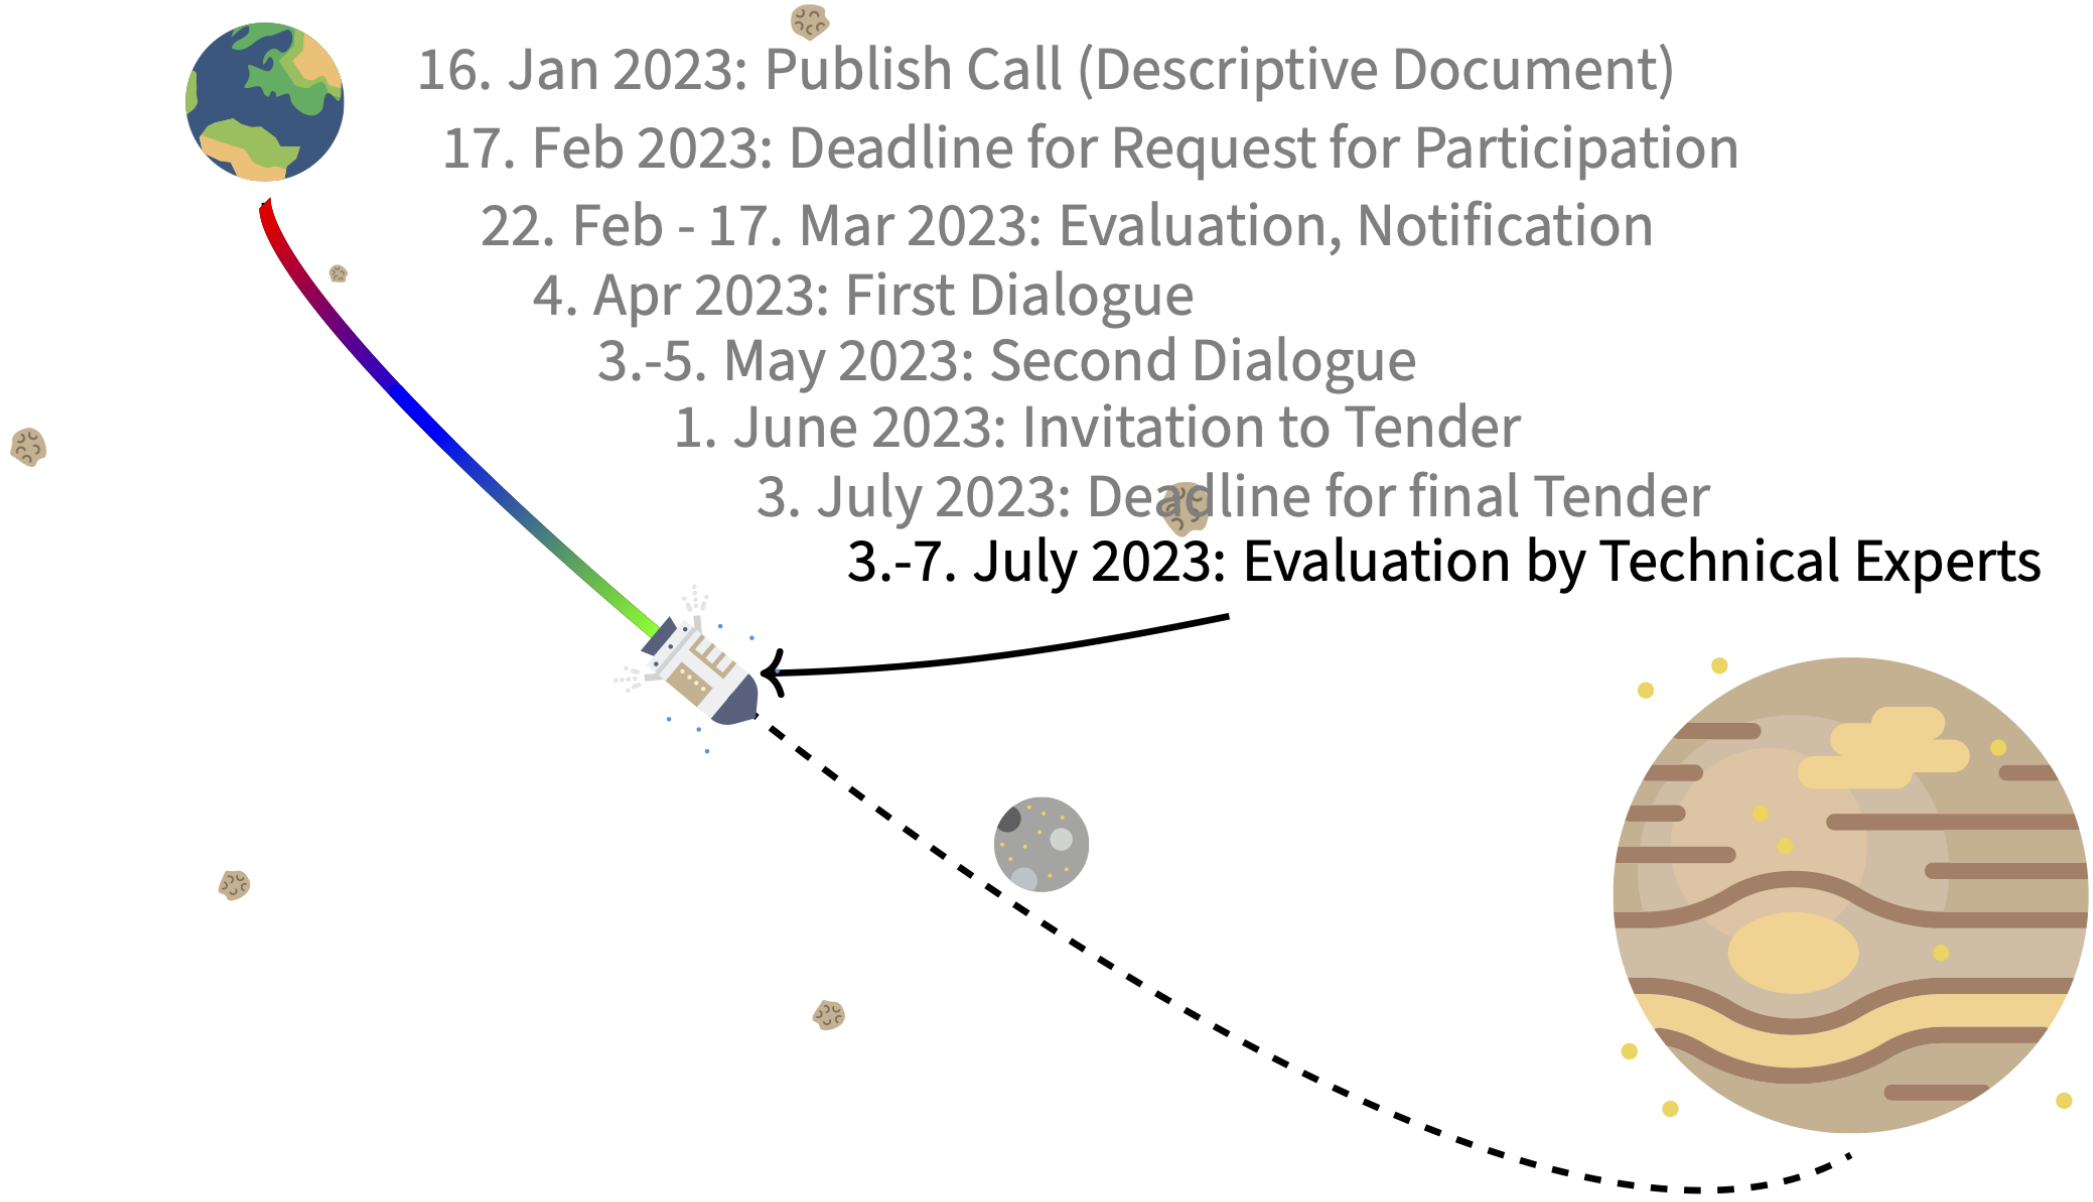
16. Jan 2023: Publish Call (Descriptive Document)
17. Feb 2023: Deadline for Request for Participation
22. Feb - 17. Mar 2023: Evaluation, Notification
4. Apr 2023: First Dialogue
3.-5. May 2023: Second Dialogue
1. June 2023: Invitation to Tender
3. July 2023: Deadline for final Tender
3.-7. July 2023: Evaluation by Technical Experts

FINAL EVALUATION

3. - 7. July 2023

- Three evaluators
  - Two by JSC
- Evaluation based on Technical Response Template (written by JSC)
  - Based on Technical Specification (written by JSC)

- Benchmark evaluation by JSC team
  - **Thank you benchmarkers! (EuroHPC mentioned they never saw something similar)**
- Throughout July additional questions to reviewer
- Final evaluation report compiled by EuroHPC


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

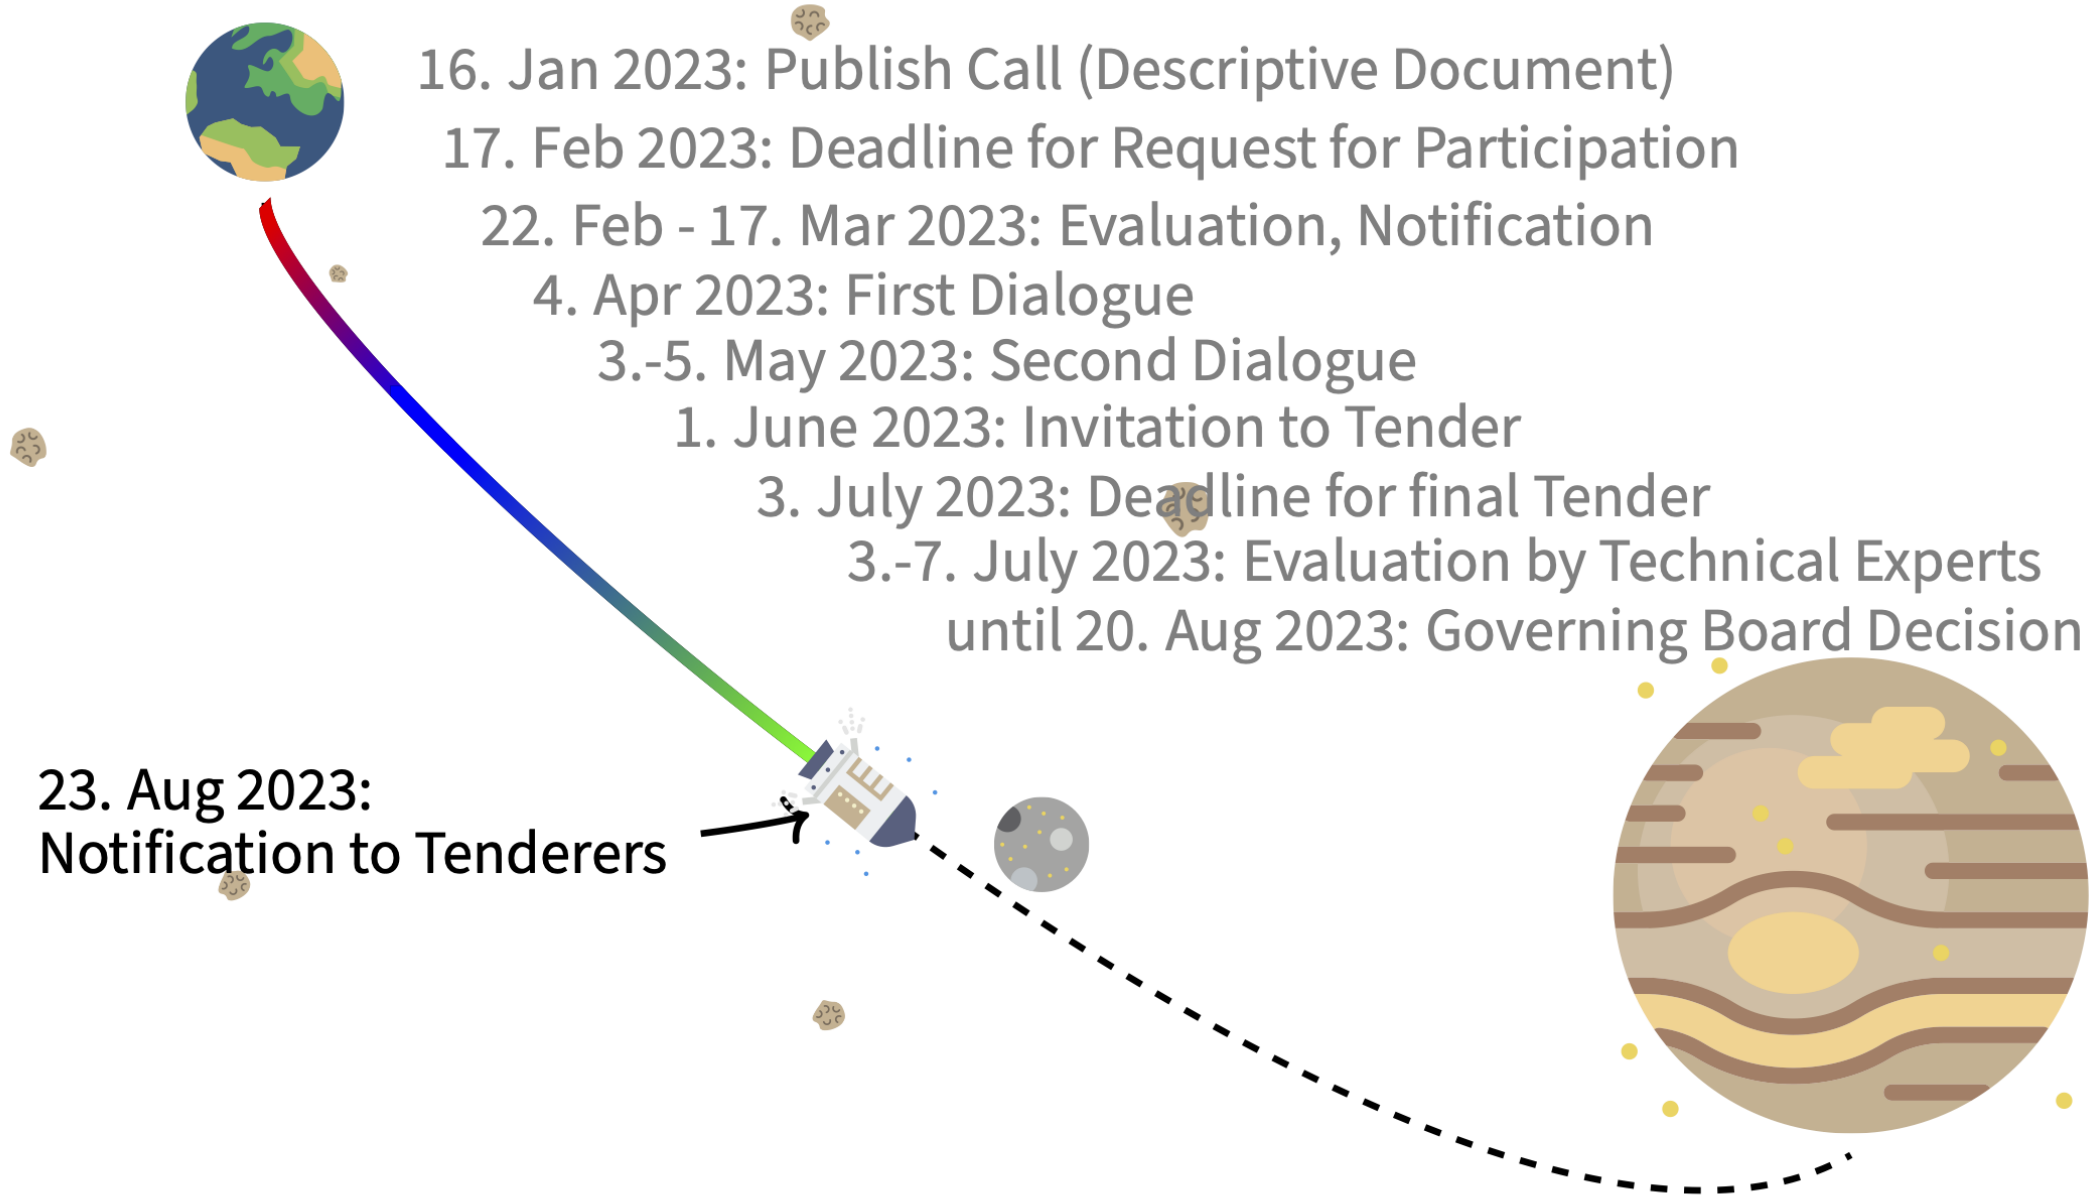
16. Jan 2023: Publish Call (Descriptive Document)
17. Feb 2023: Deadline for Request for Participation
22. Feb - 17. Mar 2023: Evaluation, Notification
4. Apr 2023: First Dialogue
3.-5. May 2023: Second Dialogue
1. June 2023: Invitation to Tender
3. July 2023: Deadline for final Tender
3.-7. July 2023: Evaluation by Technical Experts
until 20. Aug 2023: Governing Board Decision
23. Aug 2023:
Notification to Tenderers

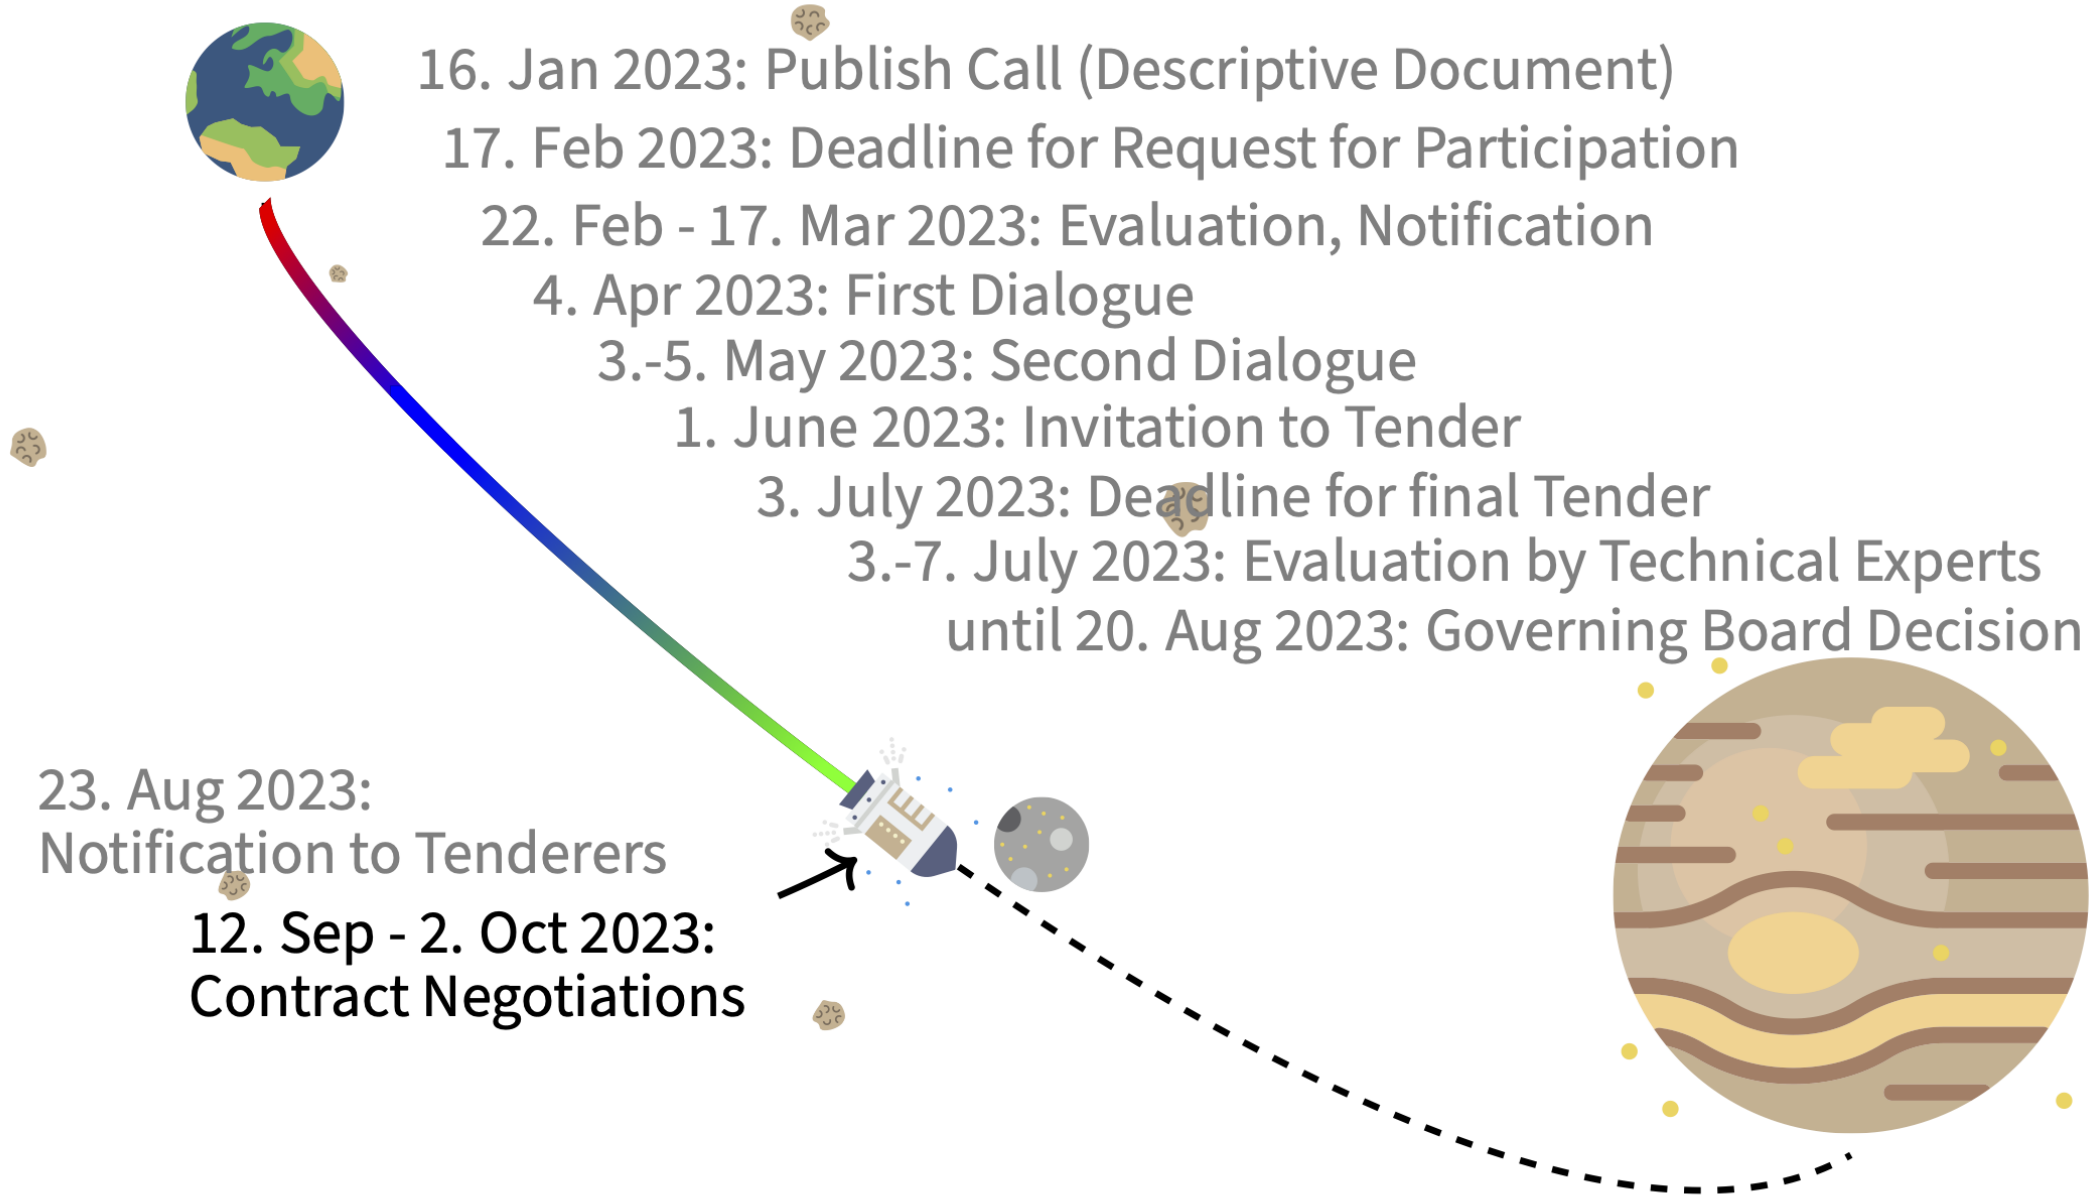
16. Jan 2023: Publish Call (Descriptive Document)
17. Feb 2023: Deadline for Request for Participation
22. Feb - 17. Mar 2023: Evaluation, Notification
4. Apr 2023: First Dialogue
3.-5. May 2023: Second Dialogue
1. June 2023: Invitation to Tender
3. July 2023: Deadline for final Tender
3.-7. July 2023: Evaluation by Technical Experts
until 20. Aug 2023: Governing Board Decision
23. Aug 2023:
Notification to Tenderers
12. Sep - 2. Oct 2023:
Contract Negotiations

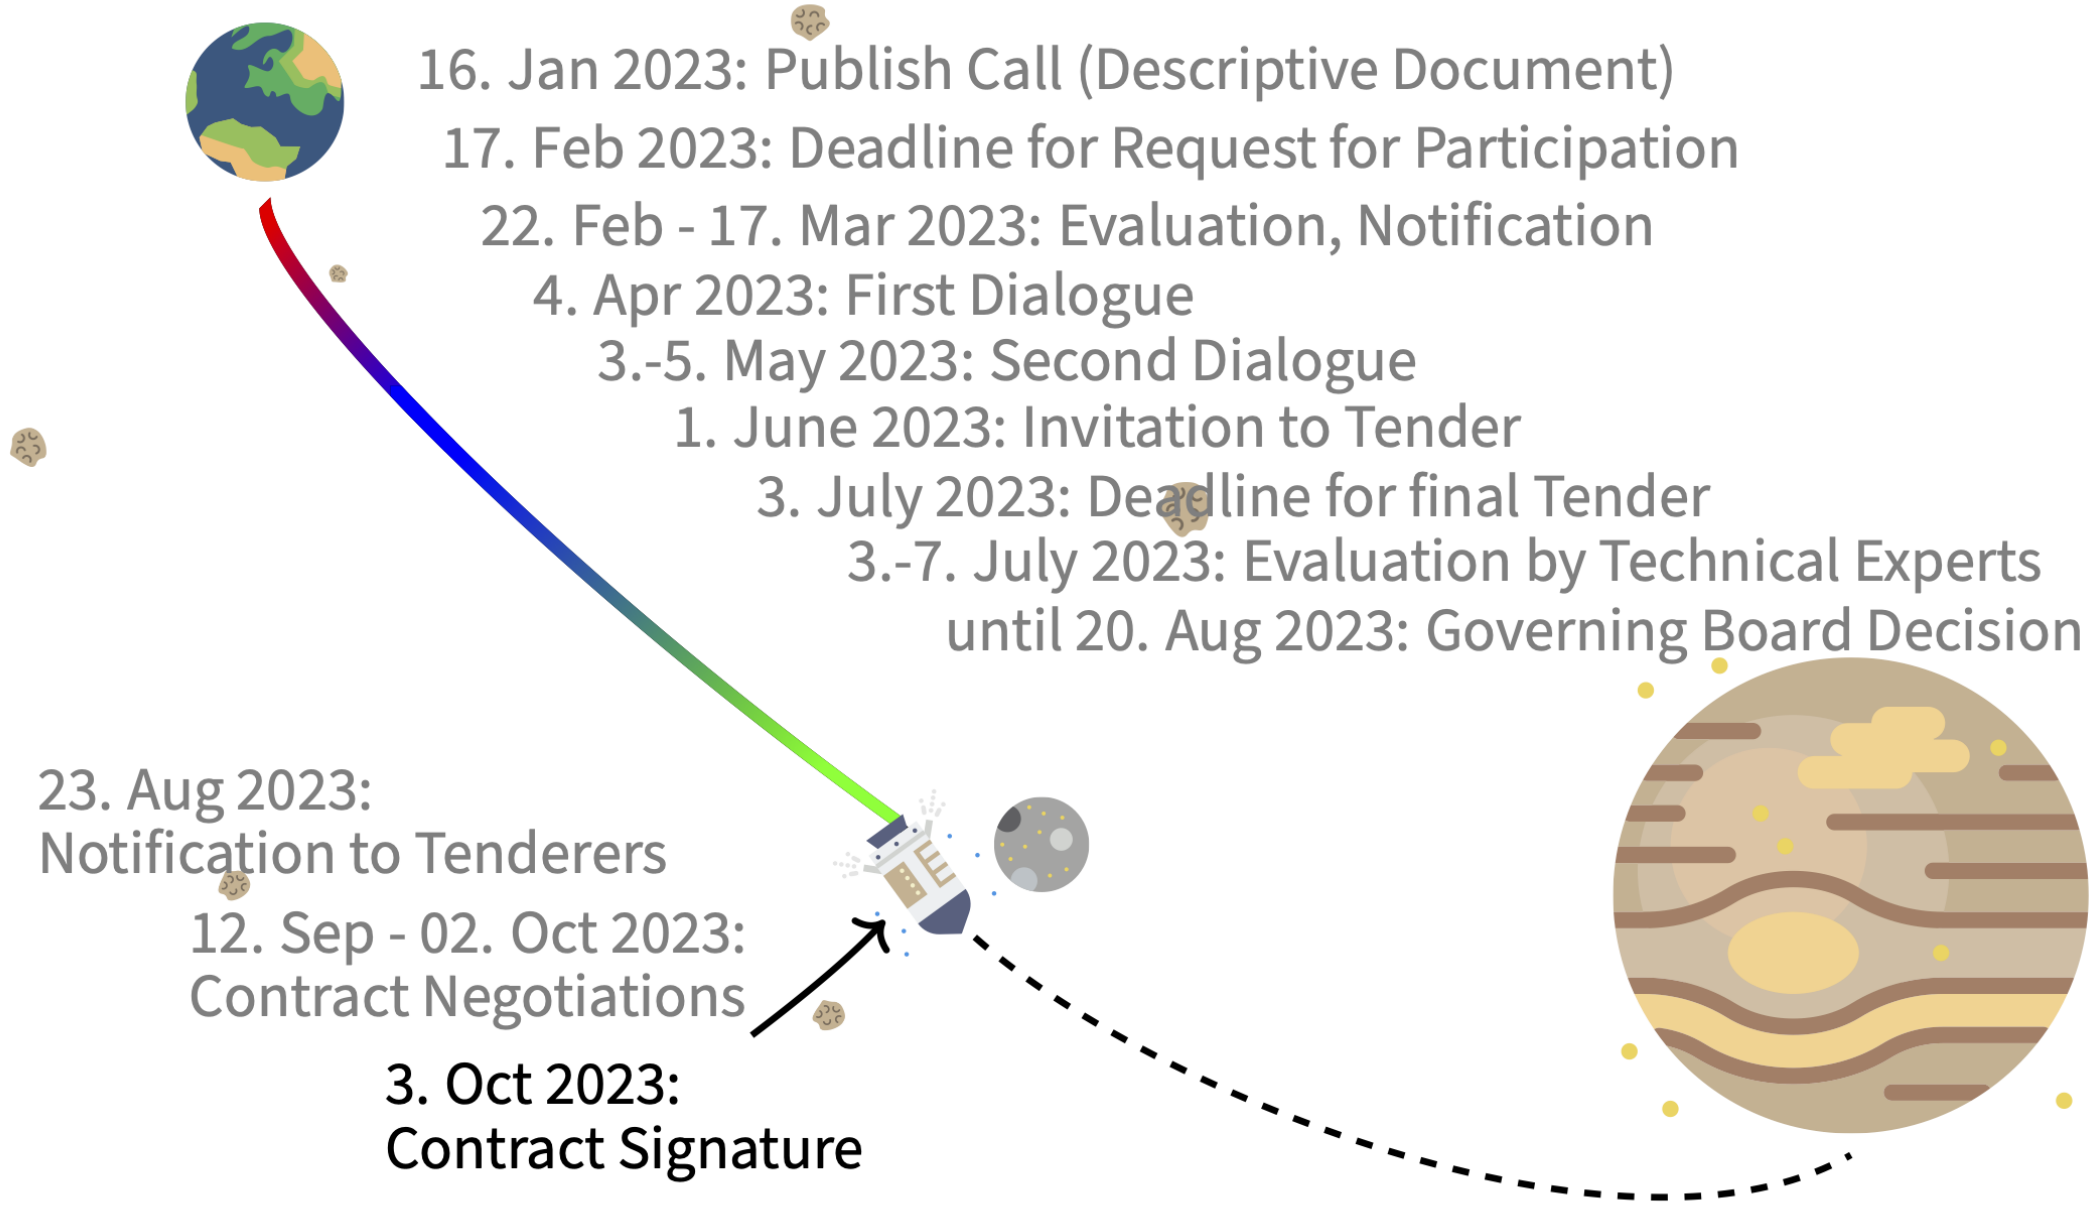
16. Jan 2023: Publish Call (Descriptive Document)
17. Feb 2023: Deadline for Request for Participation
22. Feb - 17. Mar 2023: Evaluation, Notification
4. Apr 2023: First Dialogue
3.-5. May 2023: Second Dialogue
1. June 2023: Invitation to Tender
3. July 2023: Deadline for final Tender
3.-7. July 2023: 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

HPCwire

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 ch Supercomputing Center in Germany in 2024. There has been a lot of speculation about what Jupiter l 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 ercomputers 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 ak 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

APPLICATIONS

ASSESSING WITH APPLICATIONS

- Theoretical FLOP/s and GB/s are nice; but building machines for users
- → **Applications** core of procurement assessment
- Define representative benchmarks, *ExaBench*
  1. Analyze JSC workload
  2. Select fitting applications
  3. Benchmark*ize* them
  4. Submit as part of specification

  $ Get best machine

JSC USERS

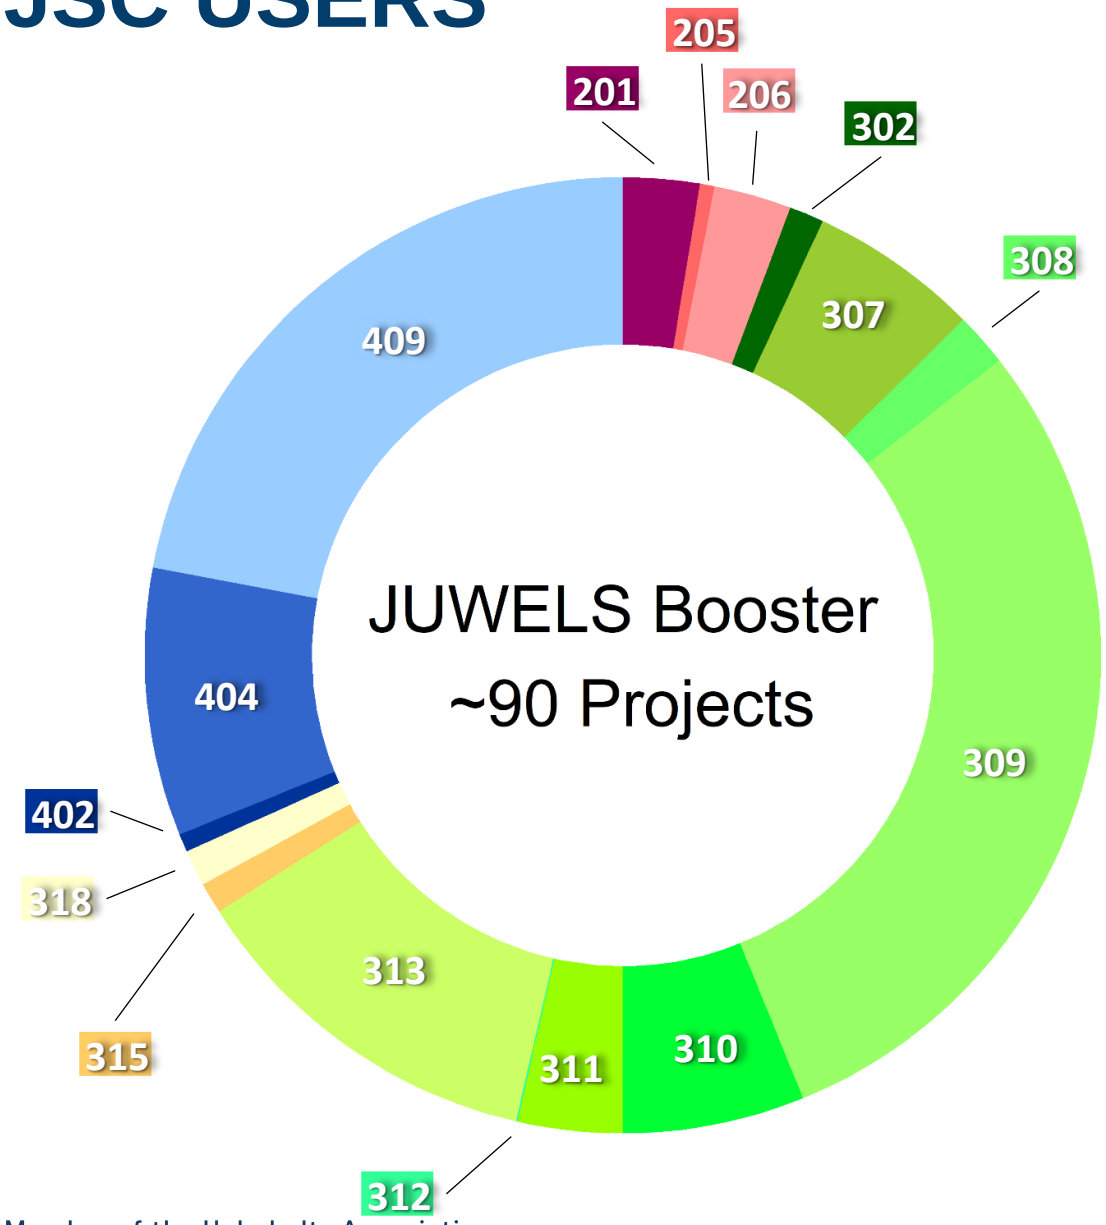

Reseach Fields

201	Basic Biological and Medical Research
205	Medicine
206	Neurosciences
302	Chemical Solid State and Surface Research
307	Condensed Matter Physics
308	Optics, Quantum Optics and Physics of Atoms, Molecules and Plasmas
309	Particles, Nuclei and Fields
310	Statistical Physics, Soft Matter, Biological Physics, Nonlinear Dynamics
311	Astrophysics and Astronomy
312	Mathematics
313	Atmospheric Science, Oceanography and Climate Research
315	Geophysics and Geodesy
318	Water Research
402	Mechanics and Constructive Mechanical Engineering
404	Heat Energy Technology, Thermal Machines, Fluid Mechanics
409	Computer Science

→ Define Benchmarks

APPLICATION SELECTION

- Selection criteria
  - Current workload
  - Future workload
  - Relevance
  - Balance with other applications
    - Domains
    - Programming models
    - Programming languages
    - Profile
  - Available PI/researcher

- Amber
- Arbor
- Chroma
- GROMACS
- ICON
- JUQCS
- nekRS
- ParFlow
- PIConGPU
- QuantumEspresso
- SOMA
- MMoCLIP
- NLP (Megatron)
- ResNet
- *DynQCD*
- *NAStJA*

FURTHER BENCHMARKS

- Augment application (complex) benchmarks with synthetic (simpler) benchmarks
- Application benchmarks: Test complex interplay of usage by real-world applications
- Synthetic benchmark: Test specific feature of system design

- OSU micro-benchmarks *(network/MPI)*
- STREAM CPU, GPU *(Memory)*
- Graph500 *(network)*
- HPCG *(memory, network)*
- HPL *(compute, network)*
- IOR *(storage)*
- Linktest *(network/MPI)*

BENCHMARKIZATION

bˌɛntʃmɑːɹkɪzˈeɪʃən, creating *benchmarks* of mere *applications*

- **Goal**: Version of application for vendors, for which we get (best) result back
- Implications: recipe, rules, verification pre-defined; only small corrections
- Steps
  - Define workload, metric (unit of time)
  - Create JUBE script for reproducibility, uniformity, abstraction
  - Add verification of results
  - *Benchmark* benchmark: Run, debug, scale
  - Add documentation, rules

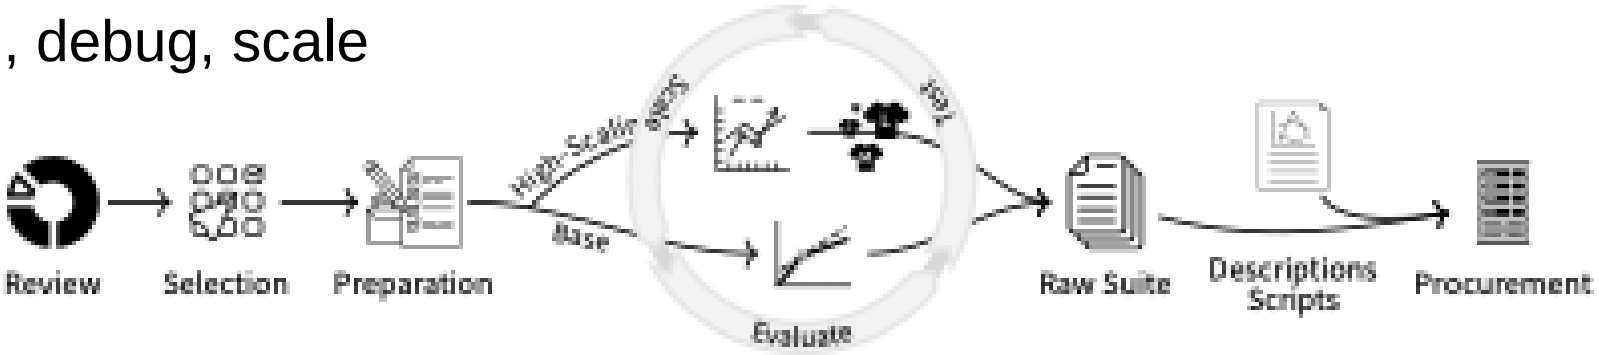

SUB-BENCHMARKS, VARIANTS

- Type of benchmarks
  - Applications benchmarks
  - Synthetic benchmarks
- Execution targets
  - JUPITER Booster (GPU, CPU)
  - JUPITER Cluster (CPU)
  - MSA
- Application benchmark categories
  - TCO
  - High-Scaling

	Booster			Cluster	MSA
Name	GPU	GPU High-Scale	CPU	CPU	
Amber	✓				
Arbor	✓	✓			
Chroma	✓	✓			
Gromacs	✓				
ICON	✓				
JUQCS	✓	✓			✓
nekRS	✓	✓			
ParFlow	✓				
PIConGPU	✓	✓			
Quantum ESPRESSO	✓				
SOMA	✓				
AI-MMoCLIP	✓				
AI-NLP	✓				
AI-ResNet	✓				
dynQCD				✓	
NAStJA				✓	
Graph500			✓		
HPCG	✓			✓	
HPL	✓			✓	
IOR			✓	✓	
LinkTest			✓	✓	✓
Multi-Flow IP			✓		
OSU	✓		✓	✓	
STREAM	✓			✓	

TCO

Total Cost of Ownership

- Traditional benchmark category
- ***How much of benchmark suite can be run in lifetime of system?*** Also: energy
- Key: same metric for each benchmark
  - Unit: time / s
  - Needed to convert rate → time
- One reference run for formula (e.g. 8 nodes); additional strong-scaled runs (e.g. 4, 16)
- Weights per individual benchmark
- Sophisticated formula for Cluster-Booster combination

HIGH-SCALING

- Give benchmarks a focus on large-scaleness of system
- Compare execution on full* JUWELS Booster to full* JUPITER Booster
  - *: Use 50 PFLOP/s$^{\text{th. peak}}$ part of JUWELS Booster

  → compare to 1000 PFLOP/s$^{\text{th. peak}}$ part of JUPITER Booster
- AKA **20×50 PF category**
- New challenge for us *(yay!)*
  - Design for unknown system, unknown device, unknown memory size
    Introduce 3 memory variants: small ($^2/_4$), medium ($^3/_4$), high ($^4/_4$ JWB A100 memory)
  - Hard to test on scale at JUWELS Booster
  - No way to test on scale required for JUPITER
  - Code issues at scale

- Arbor
  tiny ($^1/_4$), small, medium, large
- Chroma
  small, medium, large
- JUQCS
  small, large
- nekRS
  small, medium, large
- PIConGPU
  small, medium, large

FINAL BENCHMARK LISTS

		Booster			Cluster	MSA
Before Dialogue	After Dialogue	GPU	GPU High-Scale	CPU	CPU	
Amber	~~Amber~~	~~✓~~				
Arbor	Arbor	✓	✓			
Chroma	Chroma	✓	✓			
Gromacs	Gromacs (2)	✓				
ICON	ICON (2)	✓				
JUQCS	JUQCS	✓	✓			✓
nekRS	nekRS	✓	✓			
ParFlow	ParFlow	✓				
PIConGPU	PIConGPU	✓	✓			
Quantum ESPRESSO	Quantum ESPRESSO	✓				
SOMA	~~SOMA~~	~~✓~~				
AI-MMoCLIP	AI-MMoCLIP	✓				
AI-NLP	AI-NLP	✓				
AI-ResNet	~~AI-Resnet~~	~~✓~~				
dynQCD	dynQCD				✓	
NAStJA	NAStJA				✓	
Graph500	Graph500			✓		
HPCG	HPCG	✓			✓	
HPL	HPL	✓			✓	
IOR	IOR			✓	✓	
LinkTest	LinkTest			✓	✓	✓
Multi-Flow IP	~~Multi-Flow IP~~			~~✓~~		
OSU	OSU (2)	✓		✓	✓	
STREAM	STREAM	✓			✓	

SUBMITTED FILE, WEBSITE

Rolling release of benchmark (as-early-as-possible) via website; with hashes

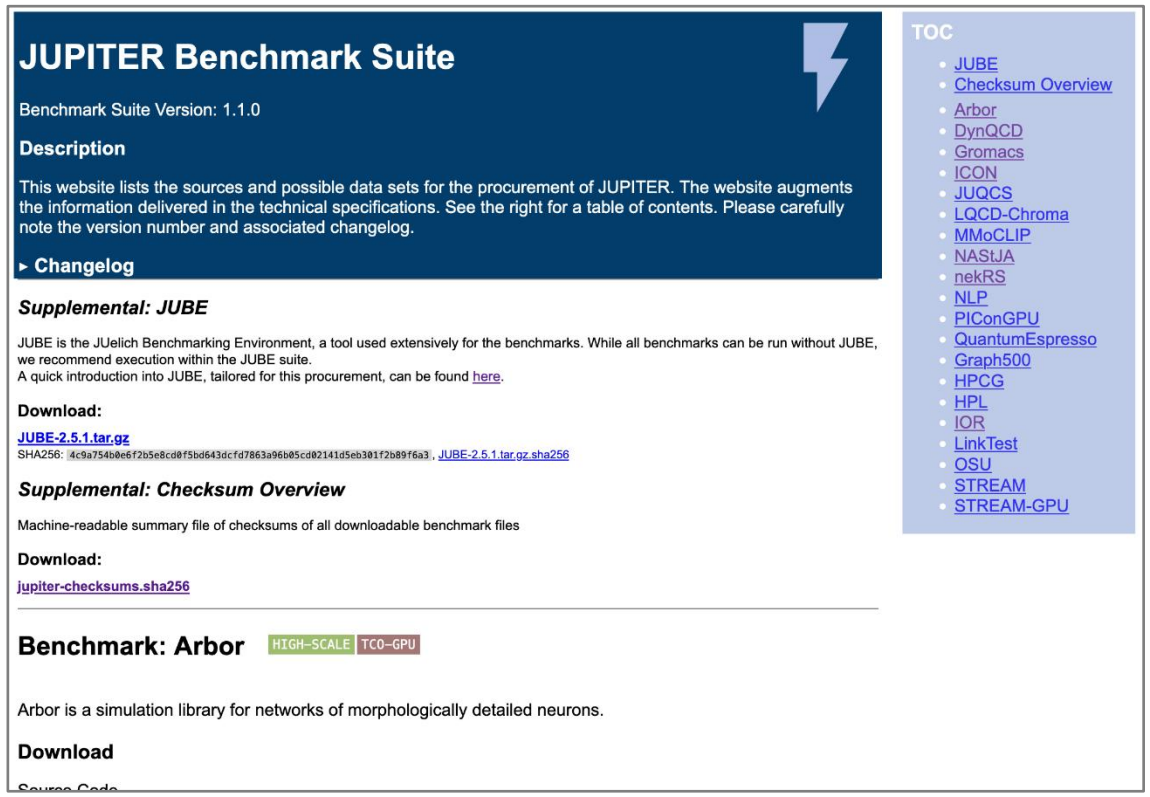

JUPITER Benchmark Suite

Benchmark Suite Version: 1.1.0

Description

This website lists the sources and possible data sets for the procurement of JUPITER. The website augments the information delivered in the technical specifications. See the right for a table of contents. Please carefully note the version number and associated changelog.

▸ **Changelog**

Supplemental: JUBE

JUBE is the JUelich Benchmarking Environment, a tool used extensively for the benchmarks. While all benchmarks can be run without JUBE, we recommend execution within the JUBE suite.
A quick introduction into JUBE, tailored for this procurement, can be found here.

Download:

JUBE-2.5.1.tar.gz
SHA256: 4c9a754b0e6f2b5e8cd0f5bd643dcfd7863a96b05cd02141d5eb301f2b89f6a3 , JUBE-2.5.1.tar.gz.sha256

Supplemental: Checksum Overview

Machine-readable summary file of checksums of all downloadable benchmark files

Download:

jupiter-checksums.sha256

Benchmark: Arbor HIGH-SCALE TCO-GPU

Arbor is a simulation library for networks of morphologically detailed neurons.

Download

Source Code

TOC

- JUBE
- Checksum Overview
- Arbor
- DynQCD
- Gromacs
- ICON
- JUQCS
- LQCD-Chroma
- MMoCLIP
- NAStJA
- nekRS
- NLP
- PIConGPU
- QuantumEspresso
- Graph500
- HPCG
- HPL
- IOR
- LinkTest
- OSU
- STREAM
- STREAM-GPU

Reference description, list of hashes, in attachment of Technical Specification

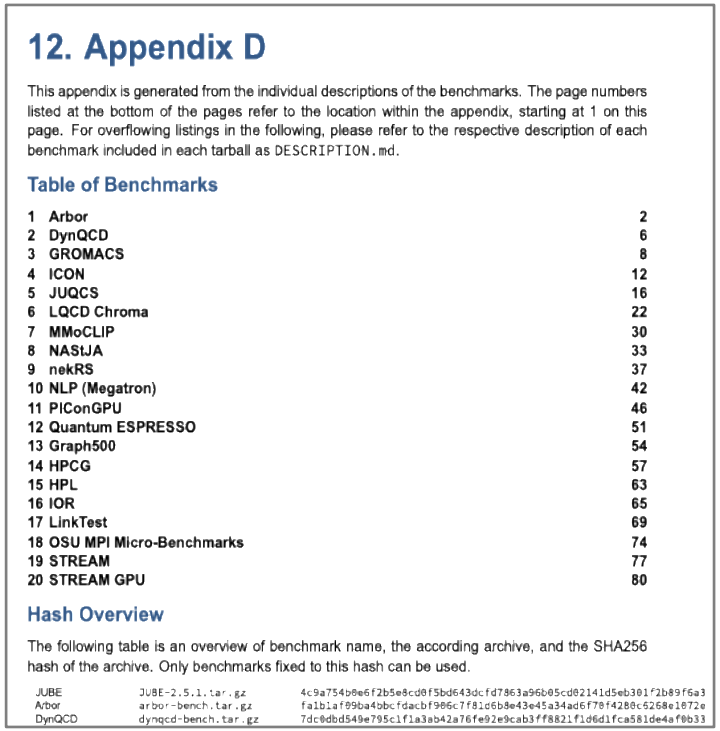

12. Appendix D

This appendix is generated from the individual descriptions of the benchmarks. The page numbers listed at the bottom of the pages refer to the location within the appendix, starting at 1 on this page. For overflowing listings in the following, please refer to the respective description of each benchmark included in each tarball as `DESCRIPTION.md`.

Table of Benchmarks

Hash Overview

The following table is an overview of benchmark name, the according archive, and the SHA256 hash of the archive. Only benchmarks fixed to this hash can be used.

JUBE	JUBE-2.5.1.tar.gz	4c9a754b0e6f2b5e8cd0f5bd643dcfd7863a96b05cd02141d5eb301f2b89f6a3
Arbor	arbor-bench.tar.gz	fa1b1af09ba4bbcfdacbf906c7f81d6b8e43e45a34ad6f70f4280c6268e1072e
DynQCD	dynqcd-bench.tar.gz	7dc0dbd549e795c1f1a3ab42a76fe92e9cab3ff8821f1d6d1fca581de4af0b33

SUBMITTED FILE, WEBSITE

Rolling release of benchmark (as-early-as-possible) via website; with hashes

Reference description, list of hashes, in attachment of Technical Specification

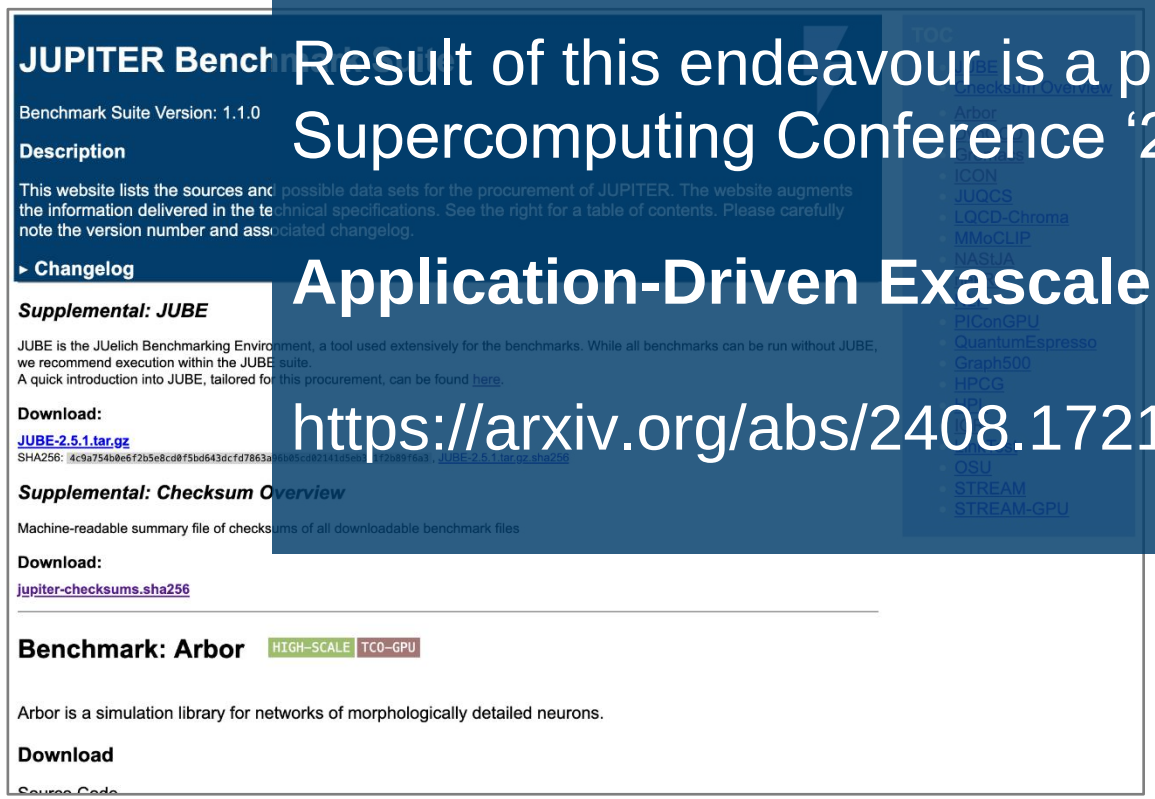

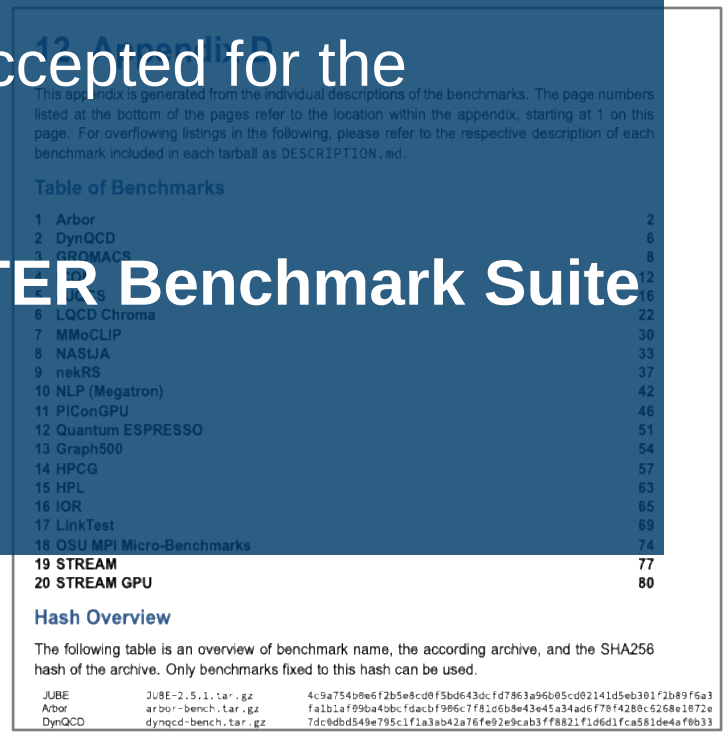

Result of this endeavour is a publication accepted for the Supercomputing Conference ‘24:

Application-Driven Exascale: The JUPITER Benchmark Suite

https://arxiv.org/abs/2408.17211

THE RESULT

DISCOVERING JUPITER

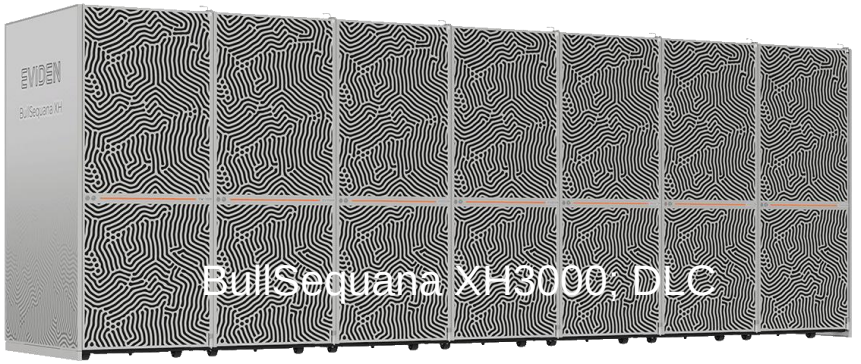

- First Exascale system in Europe (HPL); modular system
- JUPITER Booster: High scalablibty; 1 EFLOP/s HPL, >70 EFLOP/s FP8
  JUPITER Cluster: High versatility; 0.5 B/FLOP balance
- Network: InfiniBand NDR; Storage: 20 PB NVMe, 200 PB HDD
- Deployed in Modular Datacenter
- Building on: MSA (JUWELS); DEEP, EPI; ThunderX2, Ampere; …
- About 1.936.000 Arm cores

JUPITER – HIGH-LEVEL ARCHITECTURE

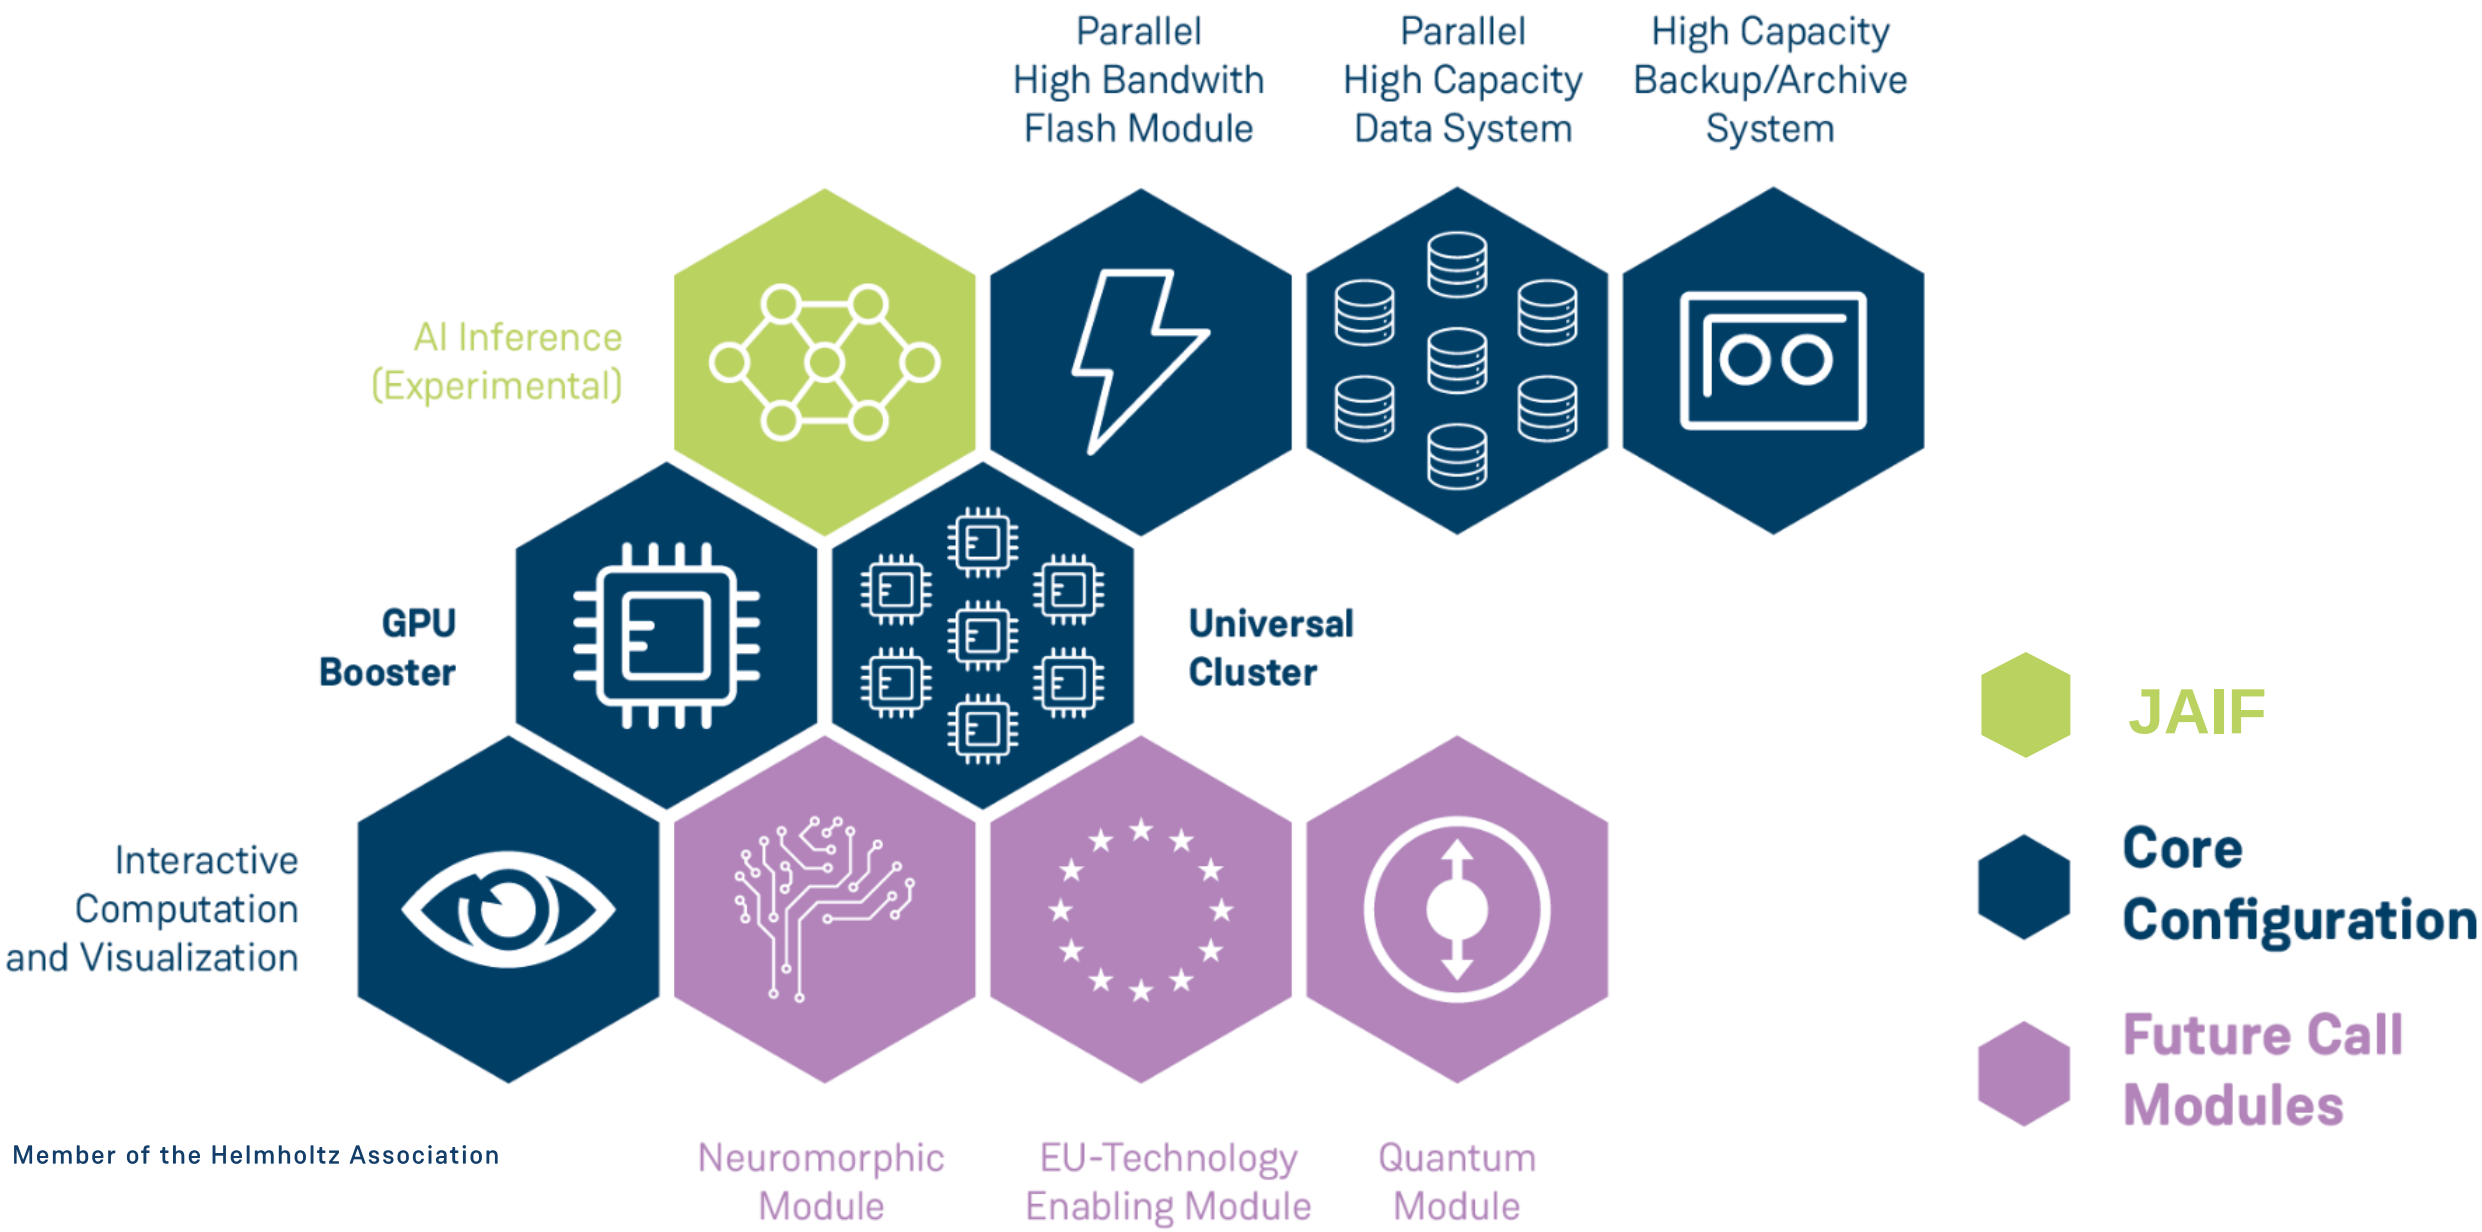

JUPITER MODULES

JUPITER Booster

- ~125 Racks BullSequana XH3000
  - 24 blades a rack, 2 nodes per blade
- Node design
  - ~6000 nodes; b
  - 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
  - 00 GB/s)

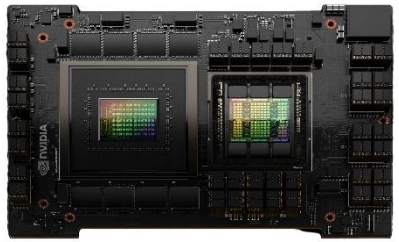

JUPITER Cluster

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

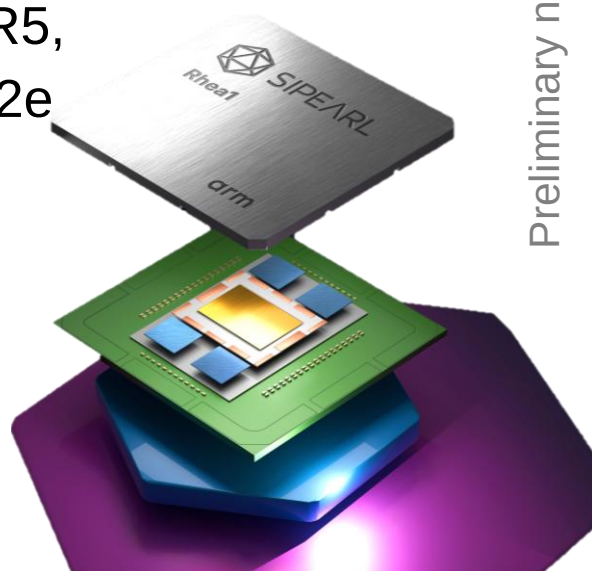

Preliminary numbers, might change during installation

JUPITER – THE BOOSTER

Highly-Scalable Module for HPC and AI workloads

- **1** ExaFLOP/s (FP64, HPL)
- NVIDIA Grace-Hopper CG1
  - ~6000 compute nodes
  - **4×** CG1 chips per compute node
- NVIDIA Mellanox NDR
  - **4×** NDR200 NICs per compute node
- BullSequana XH3000
  - Direct Liquid Cooled blades
  - **2× compute node per blade**

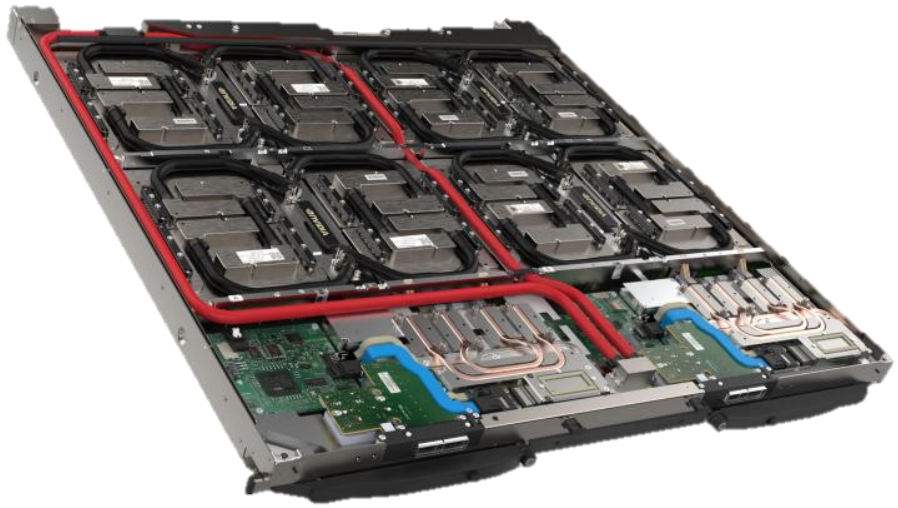

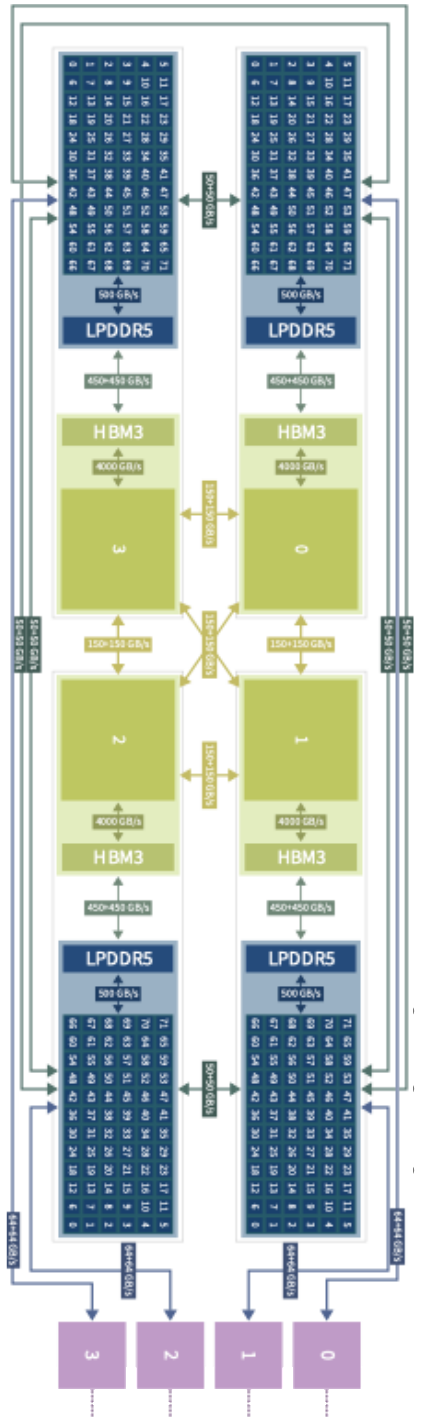

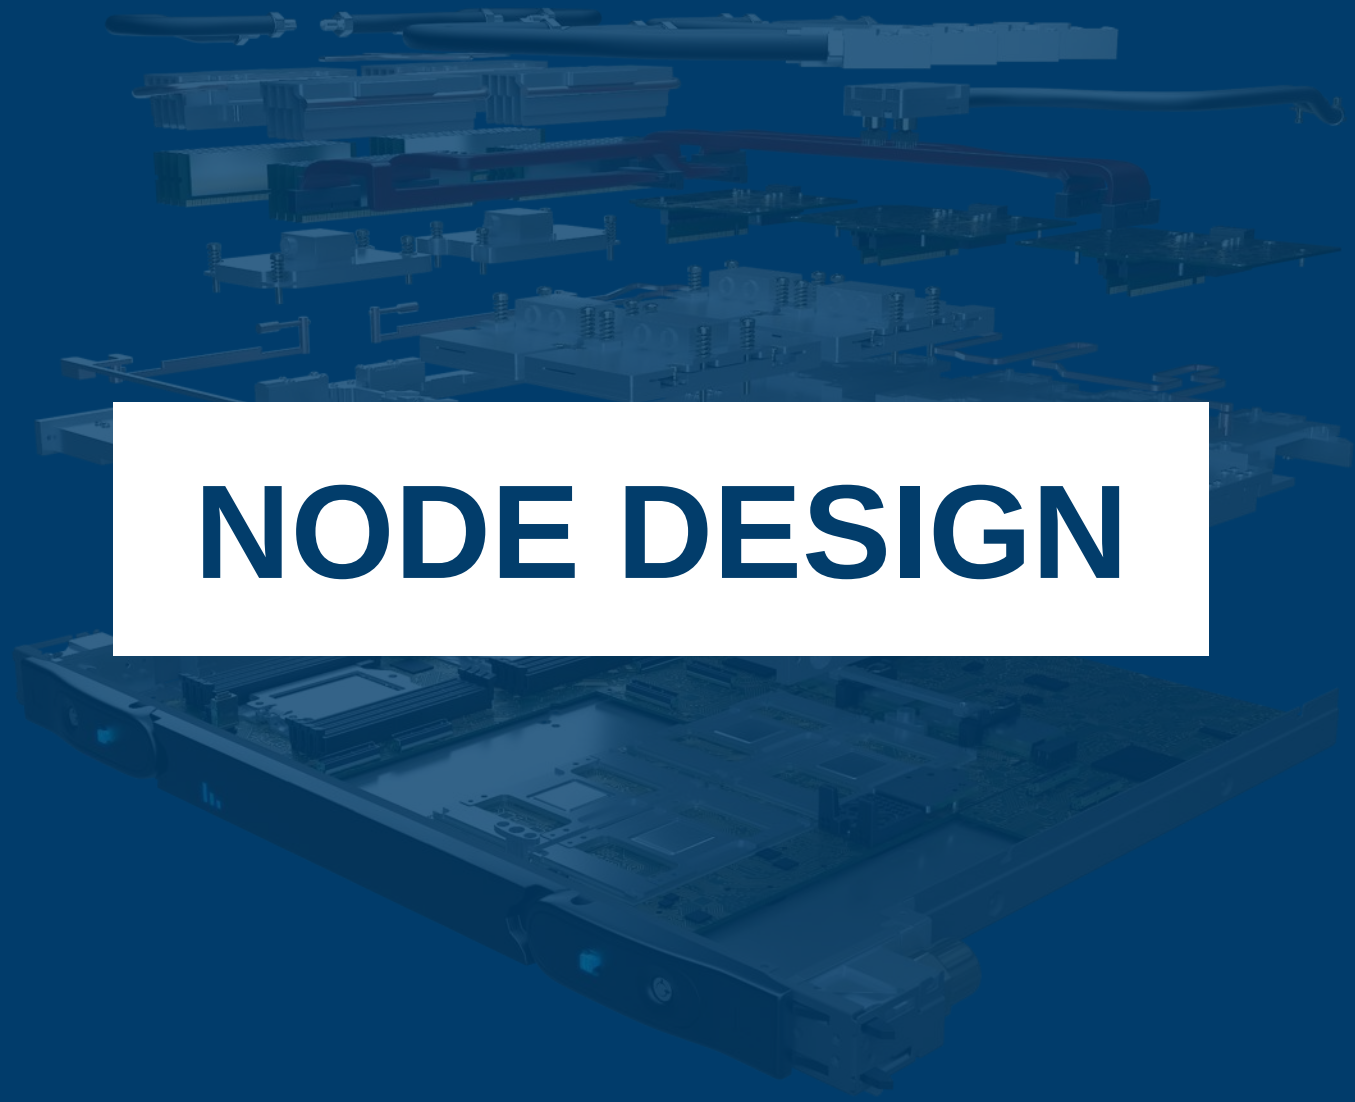

NODE DESIGN

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

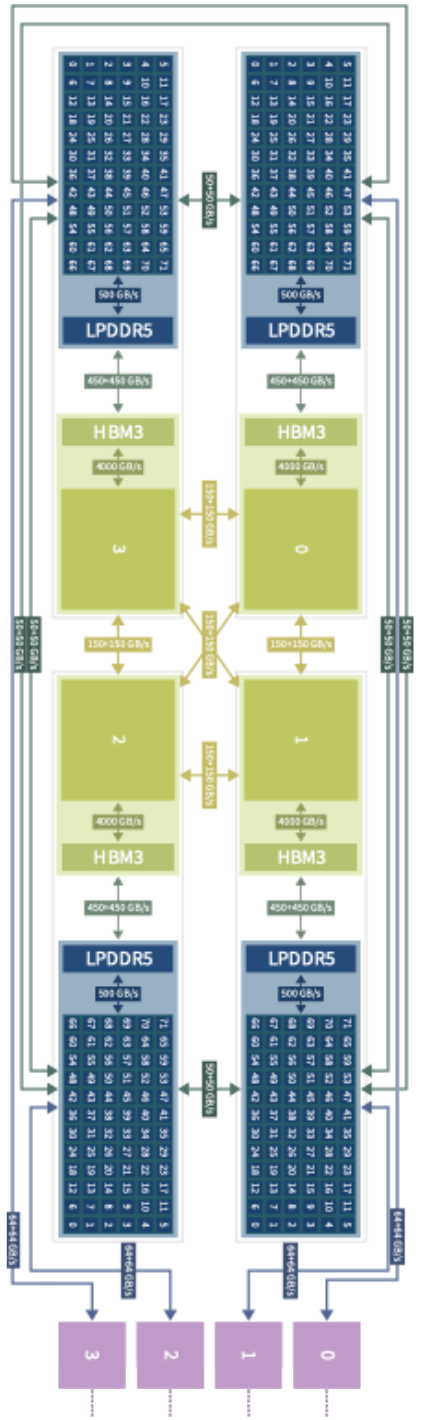

JUPITER BLADE OVERVIEW

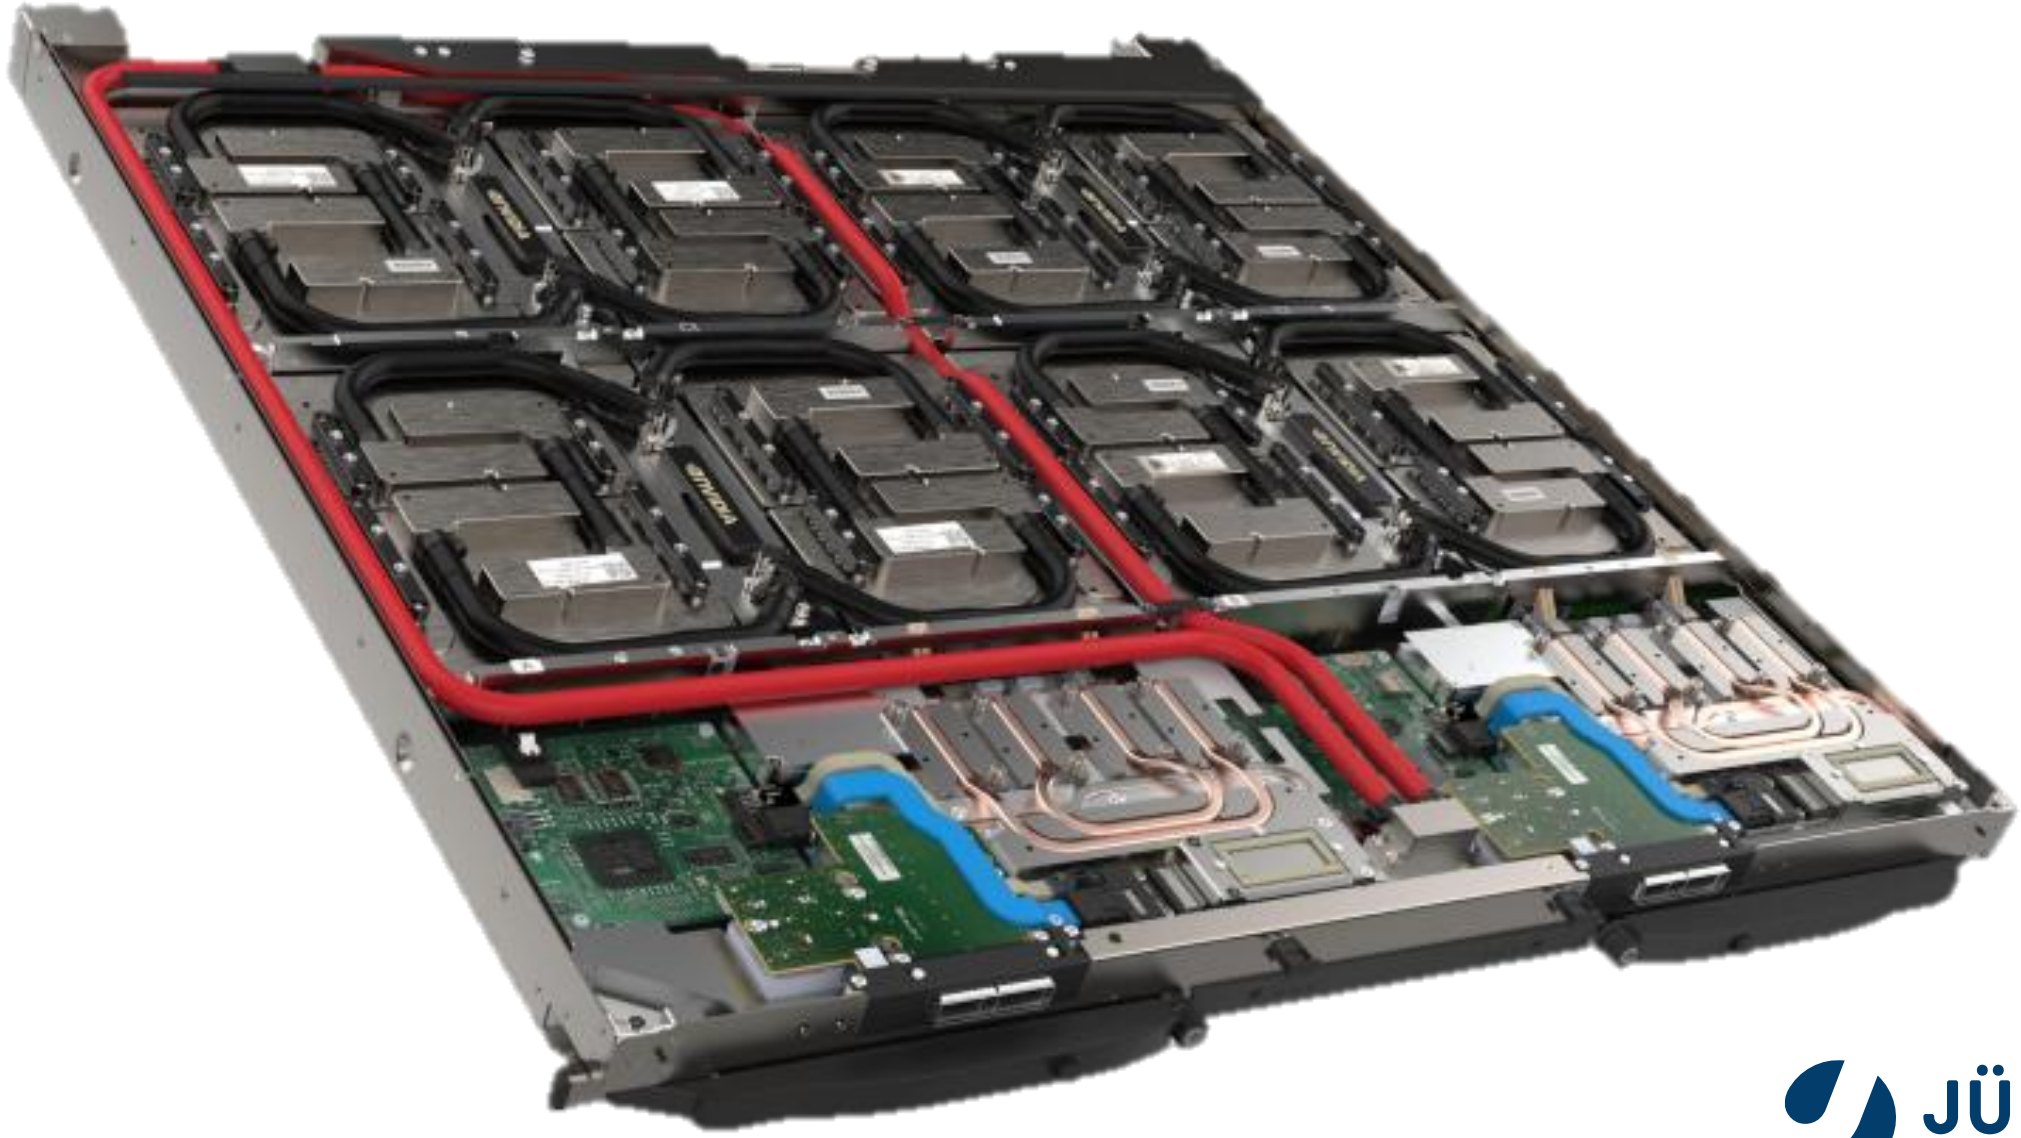

H100 GPU

GH200 Flavor (not PCIe, not HBM2e, or HBM3e)

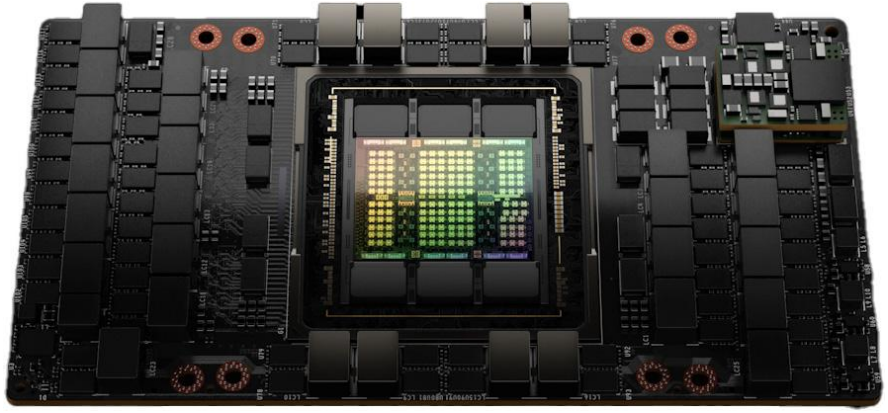

- 132 multiprocessors (*cores*)
  - 67 TFLOP/s $FP64^{TC}$
  - 989 TFLOP/s $FP16^{TC}$, 3958 TFLOP/s $FP8^{TC,Sp}$
- HBM3
  - 90 GB at 3.3 TB/s
- Notable new features
  - Thread-block clusters
  - Tensor memory accelerator
  - Transformer engine
- NVLink 4: 300 GB/s per GPU connection (900 GB/s)

→ Whitepaper

GRACE-HOPPER NODE OVERVIEW

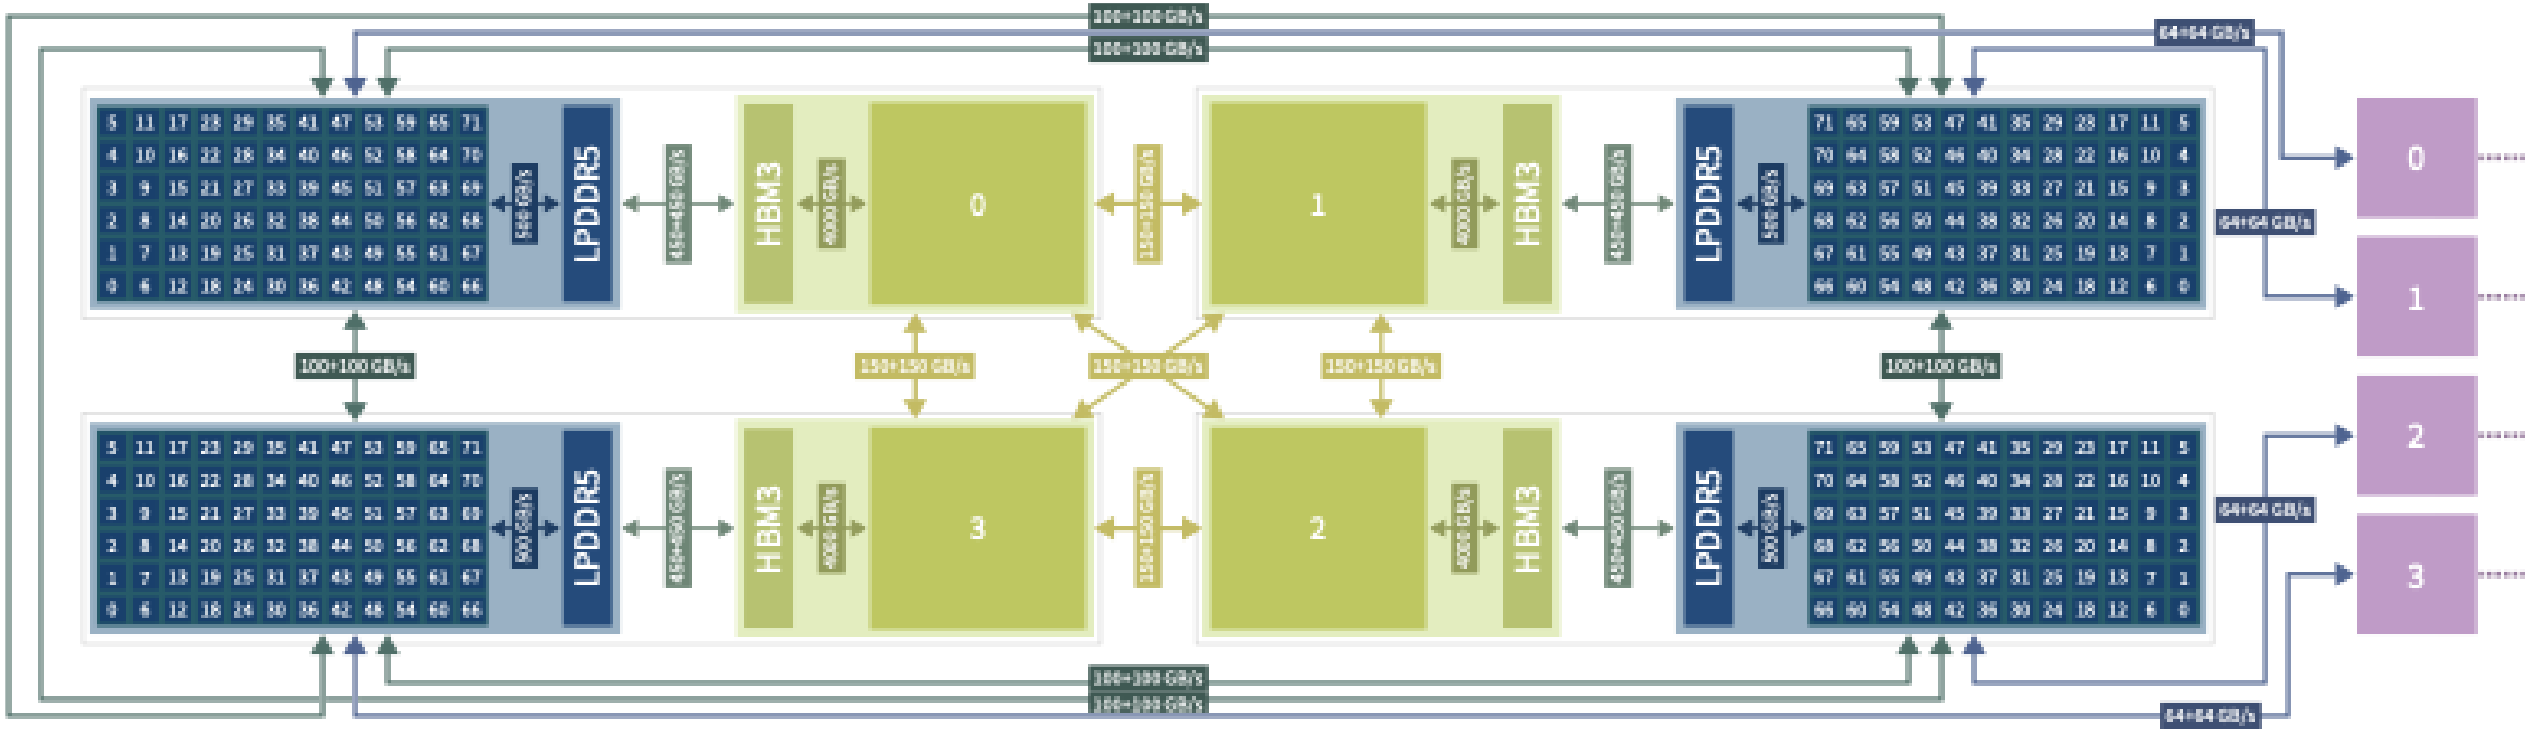

NODE COMPARISON

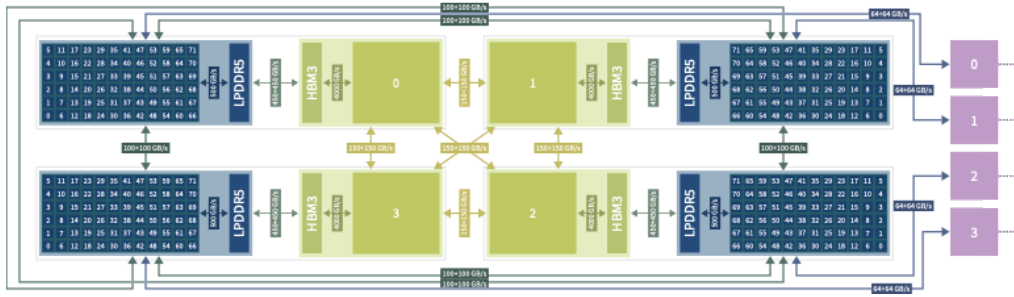

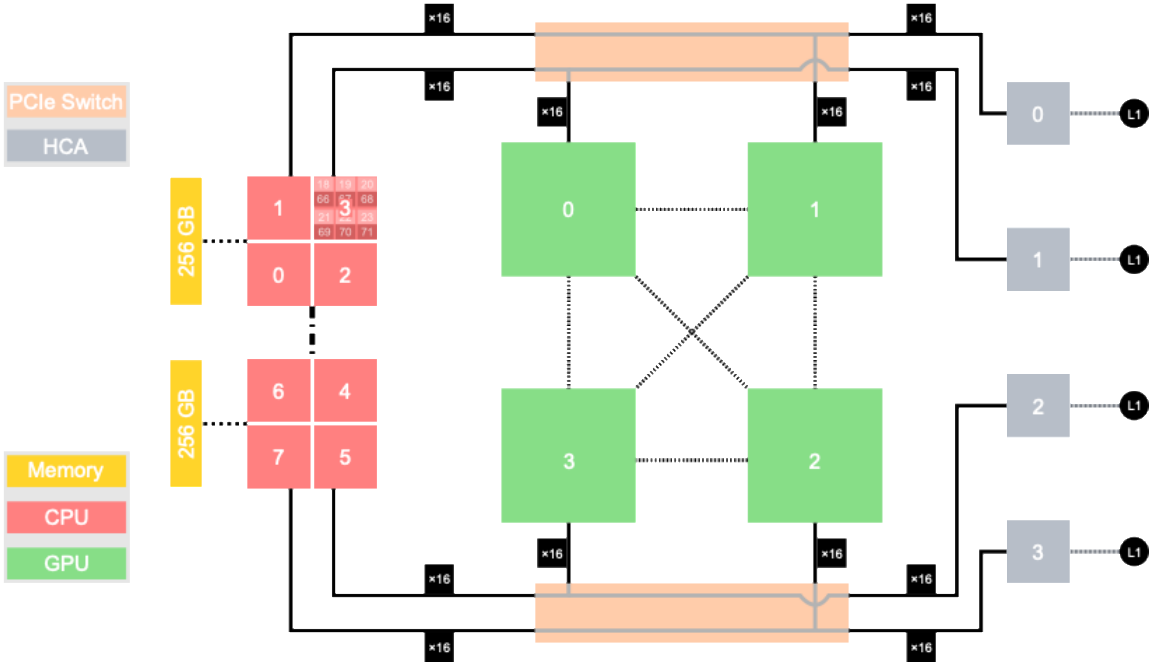

JUWELS Booster

NODE COMPARISON

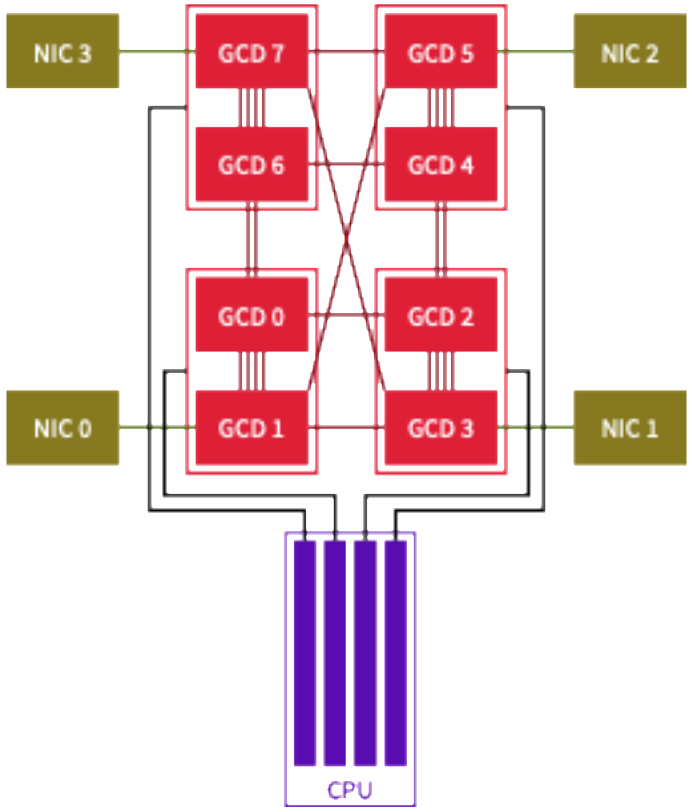

Frontier

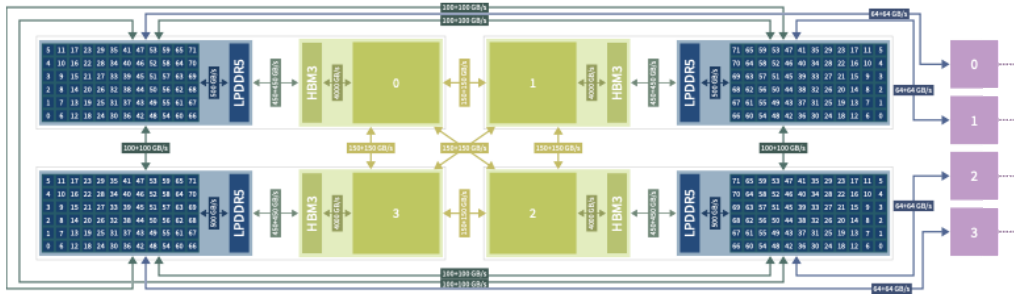

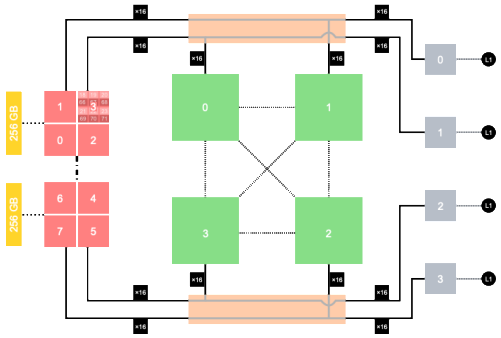

NODE COMPARISON

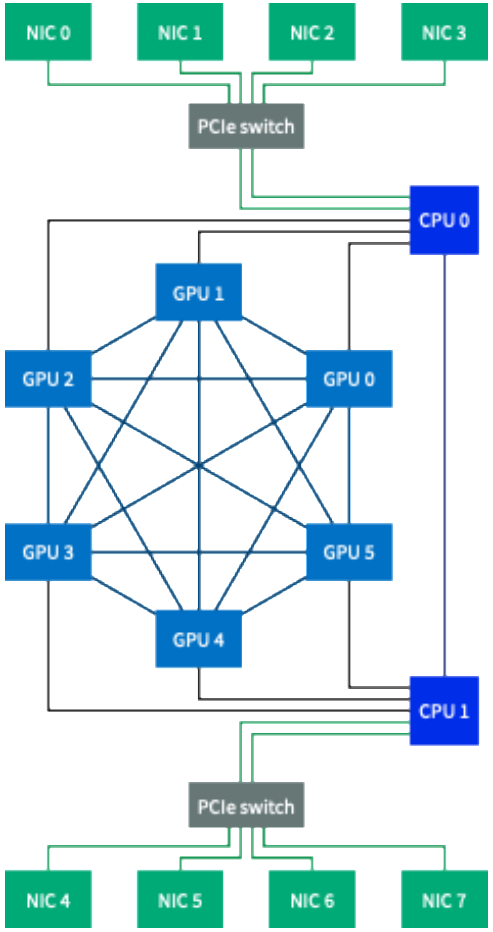

Aurora

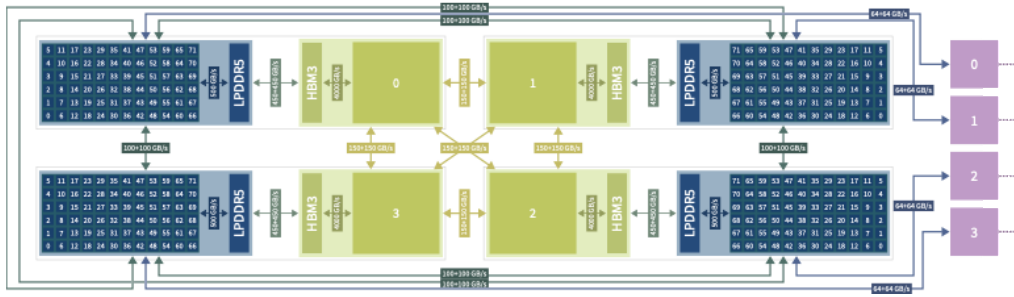

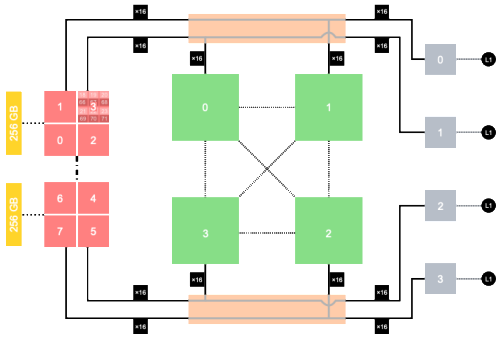

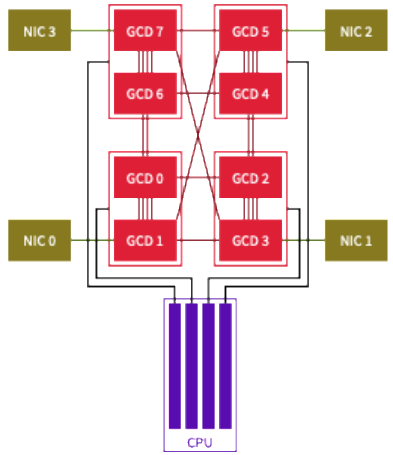

NODE COMPARISON

- JUWELS Booster: 2× CPU, 4× GPU, 4× IB
- JUPITER Booster: 4× CPU+GPU, 4× IB
- Frontier: 1× CPU, 4×(2× GPU), 4× Sling
- Aurora: 2× CPU, 6× GPU, 8× Sling
- El Capitan: 4× APU

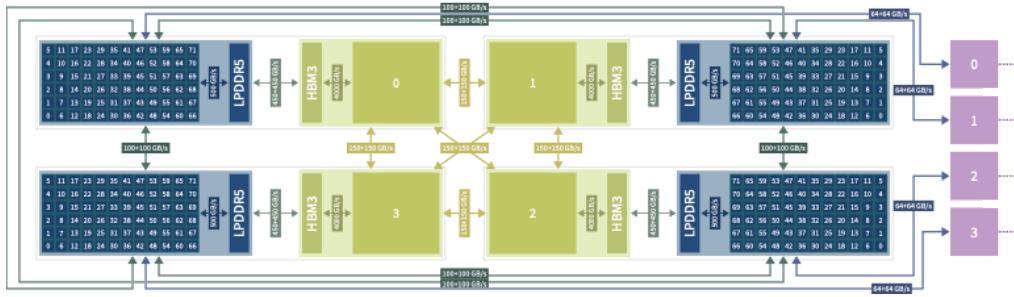

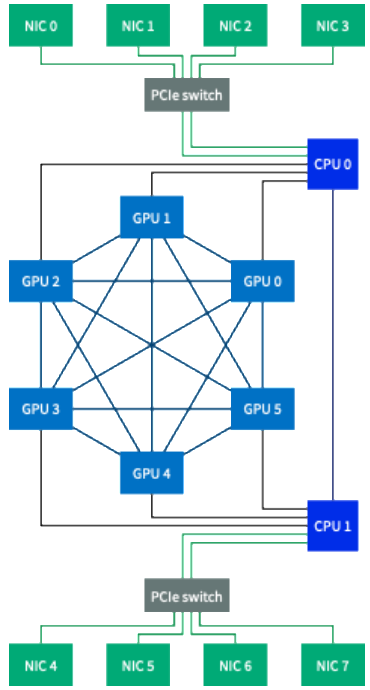

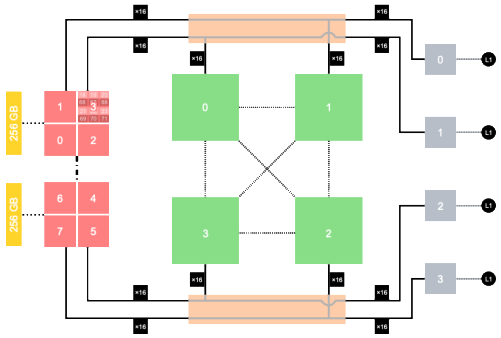

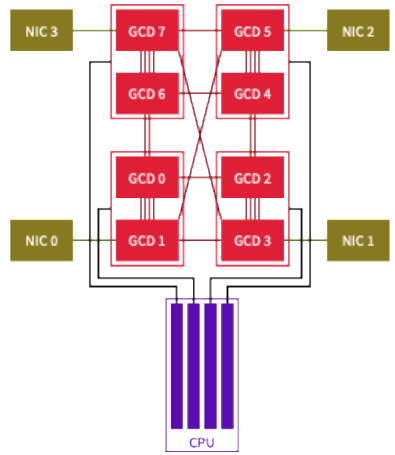

JÜLICH Forschungszentrum

NETWORK DESIGN

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)

JUPITER – INTERCONNECT

One Network to Rule Them All

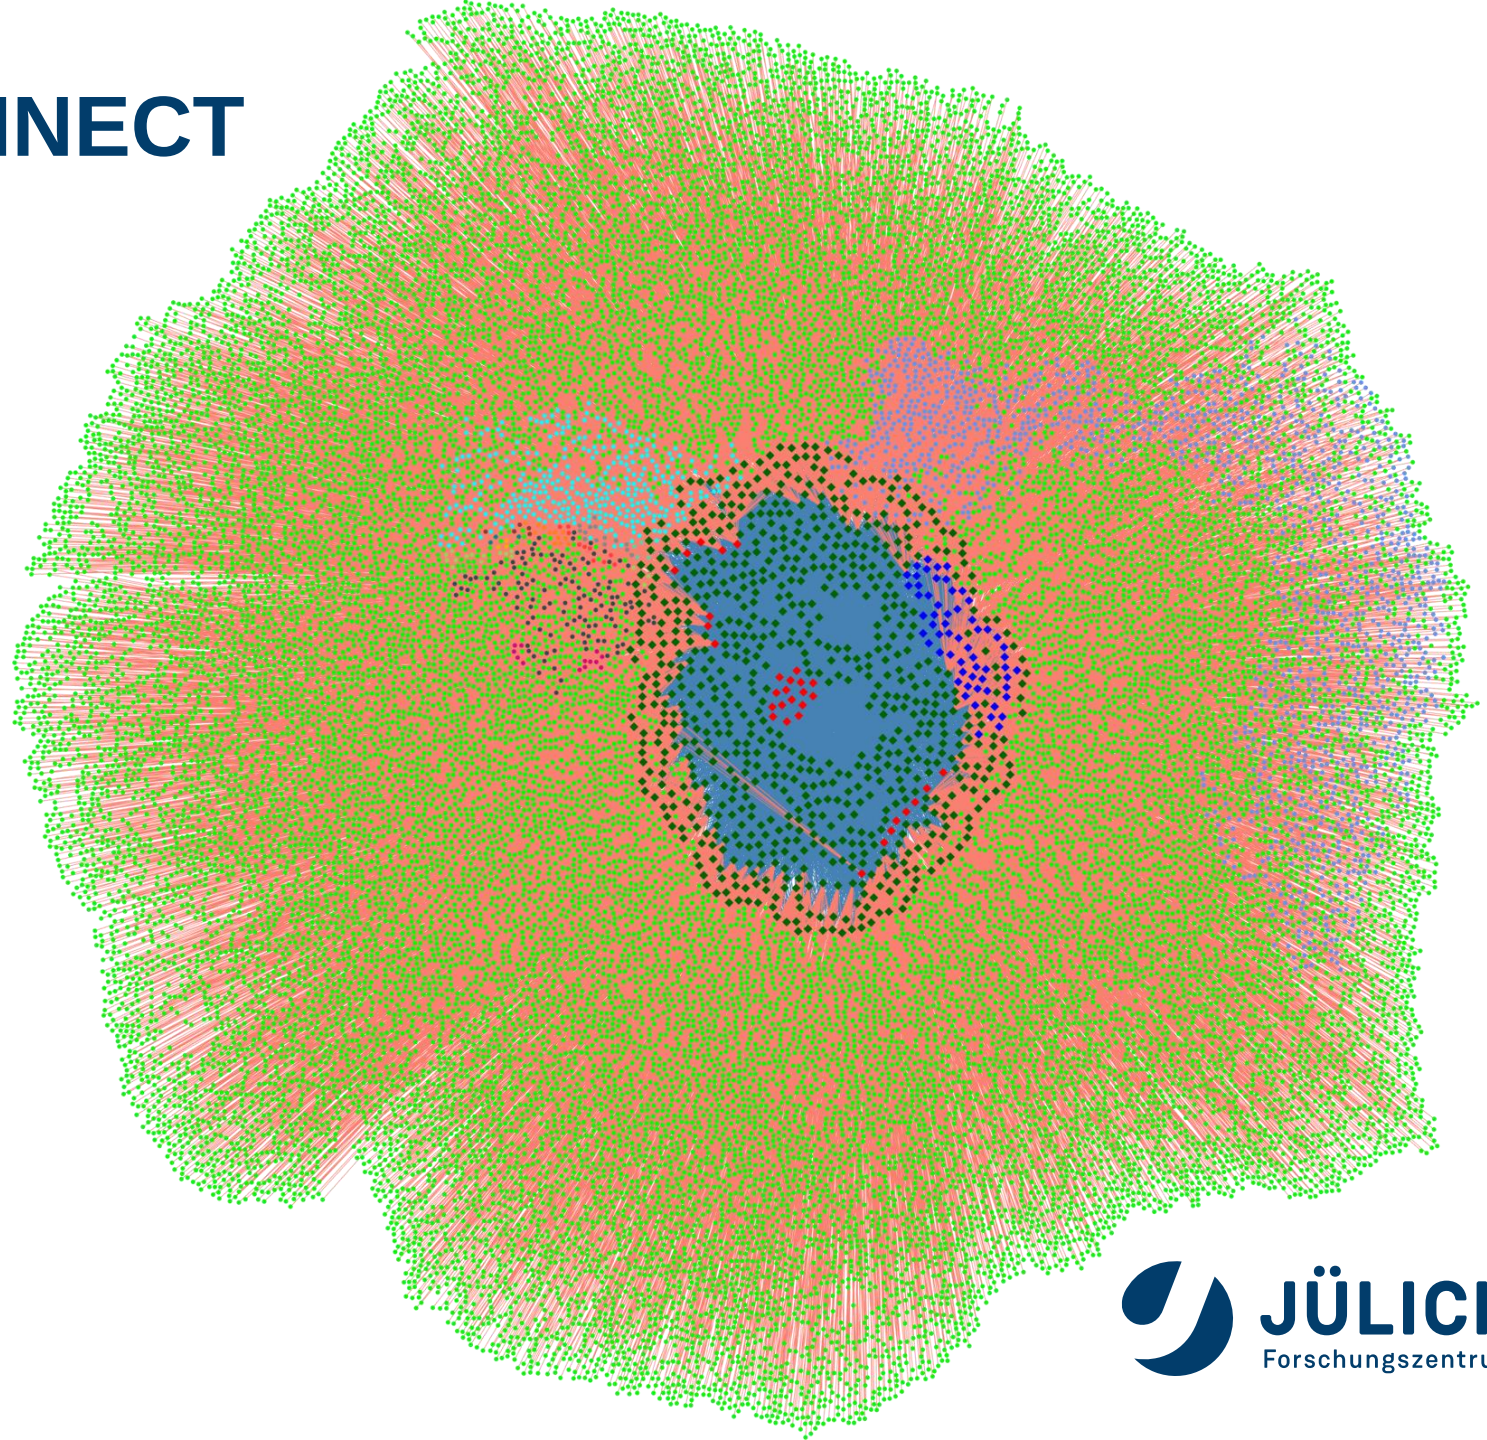

- Cluster module switch
- Booster module switch
- Booster node
- Flash storage node
- Misc node
- NDR link
- Admin module switch
- Cluster node
- Gateway node
- Storage node
- NDR200 link

JUPITER – INTERCONNECT

One Network to Rule Them All

EVIDEN
an atos business

NVIDIA

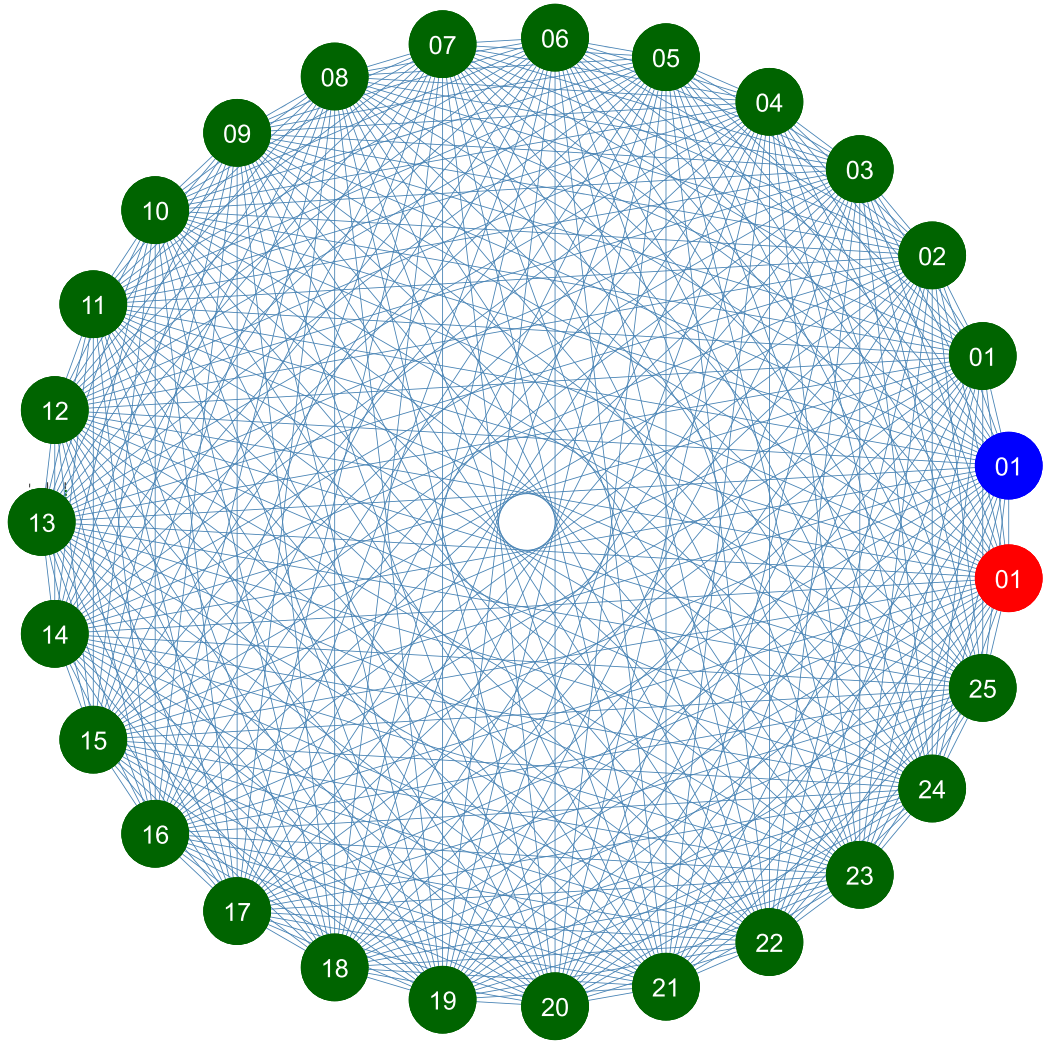

JÜLICH
Forschungszentrum

JUPITER – INTERCONNECT

One Network to Rule Them All

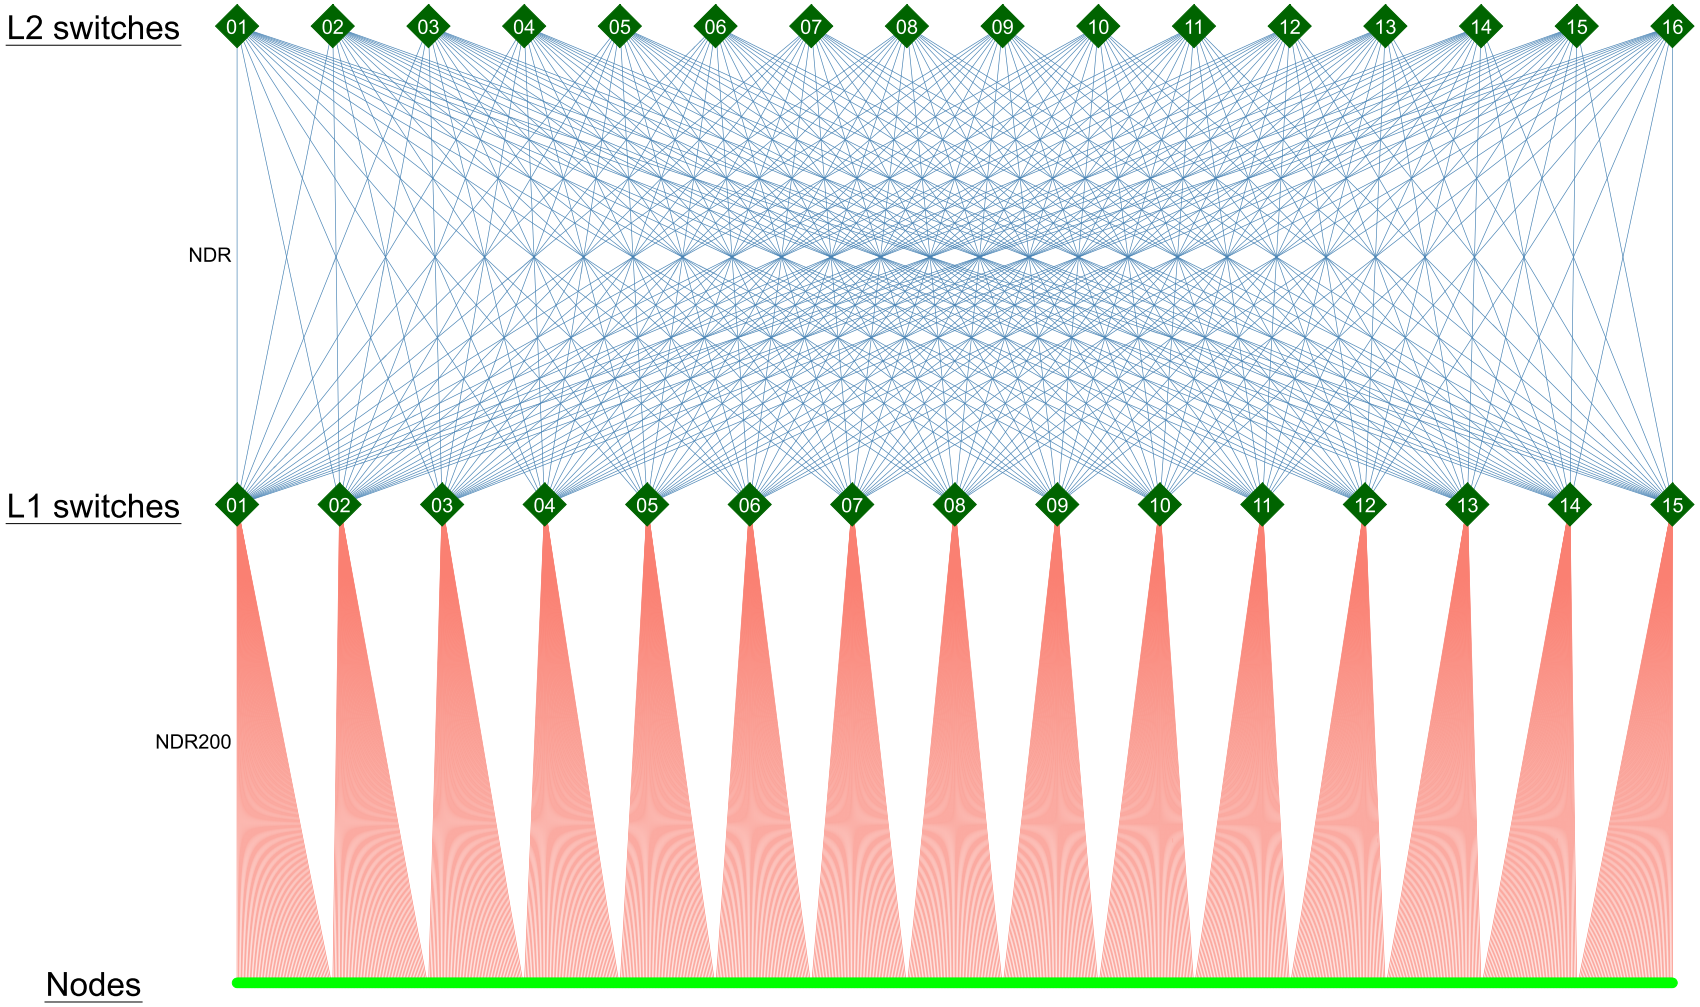

GPU cell

JUPITER – INTERCONNECT

One Network to Rule Them All

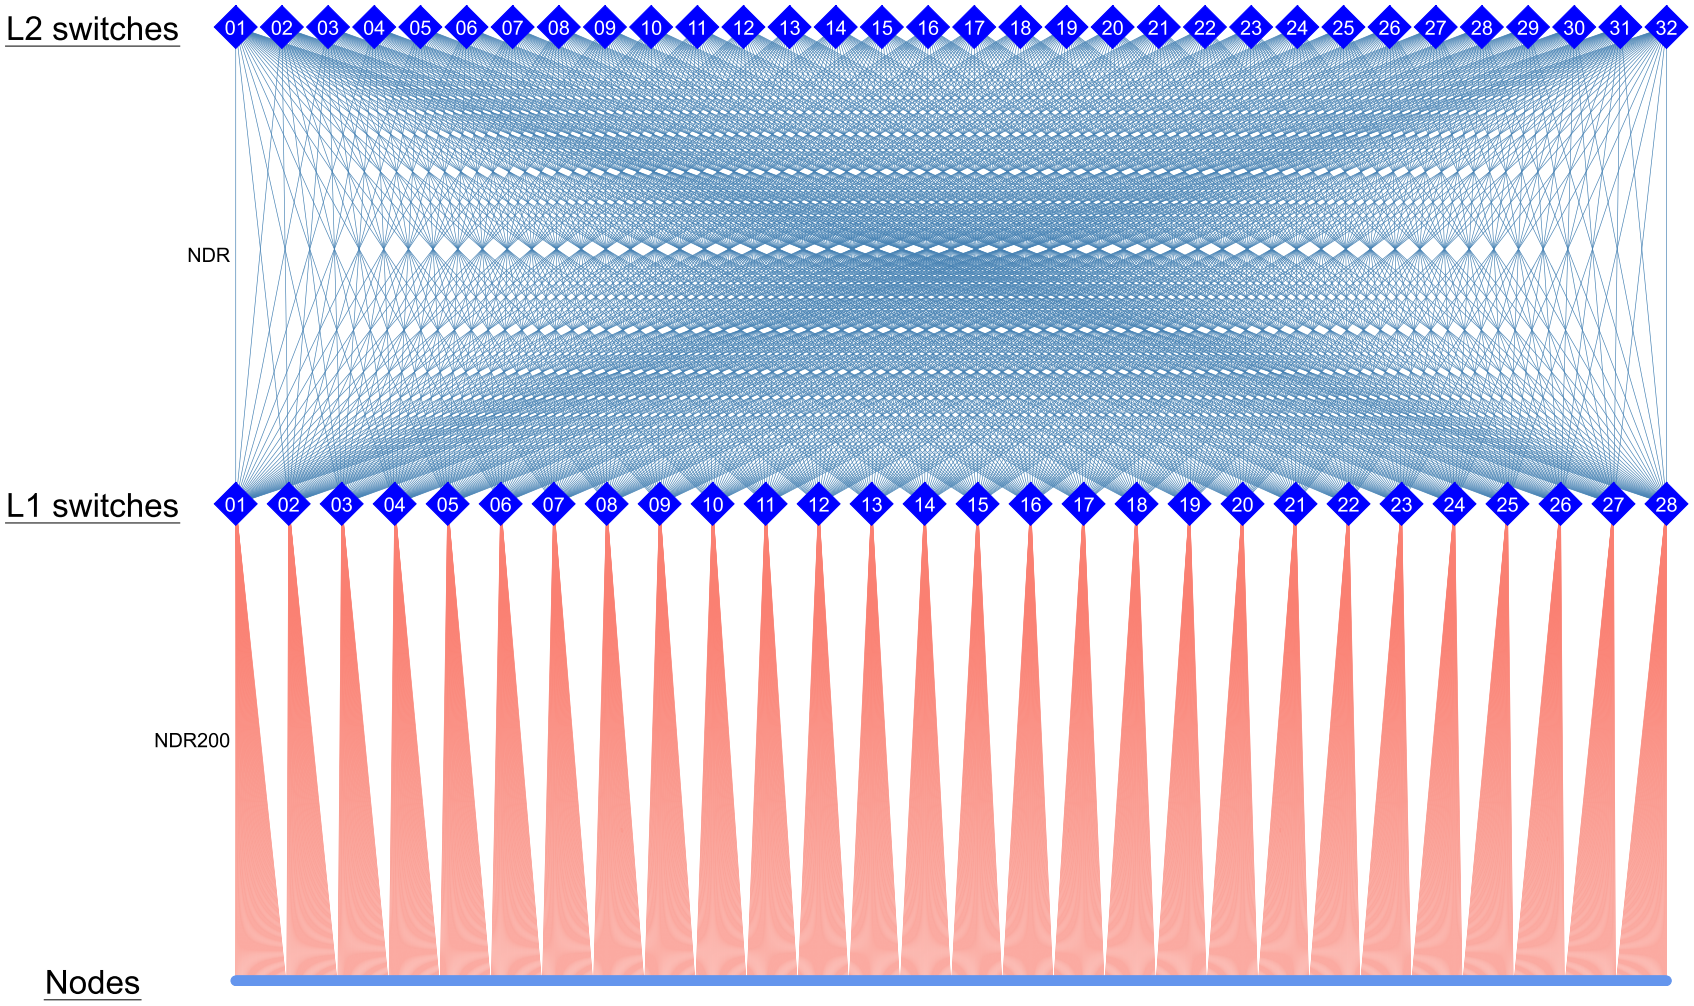

CPU cell

JUPITER – INTERCONNECT

One Network to Rule Them All

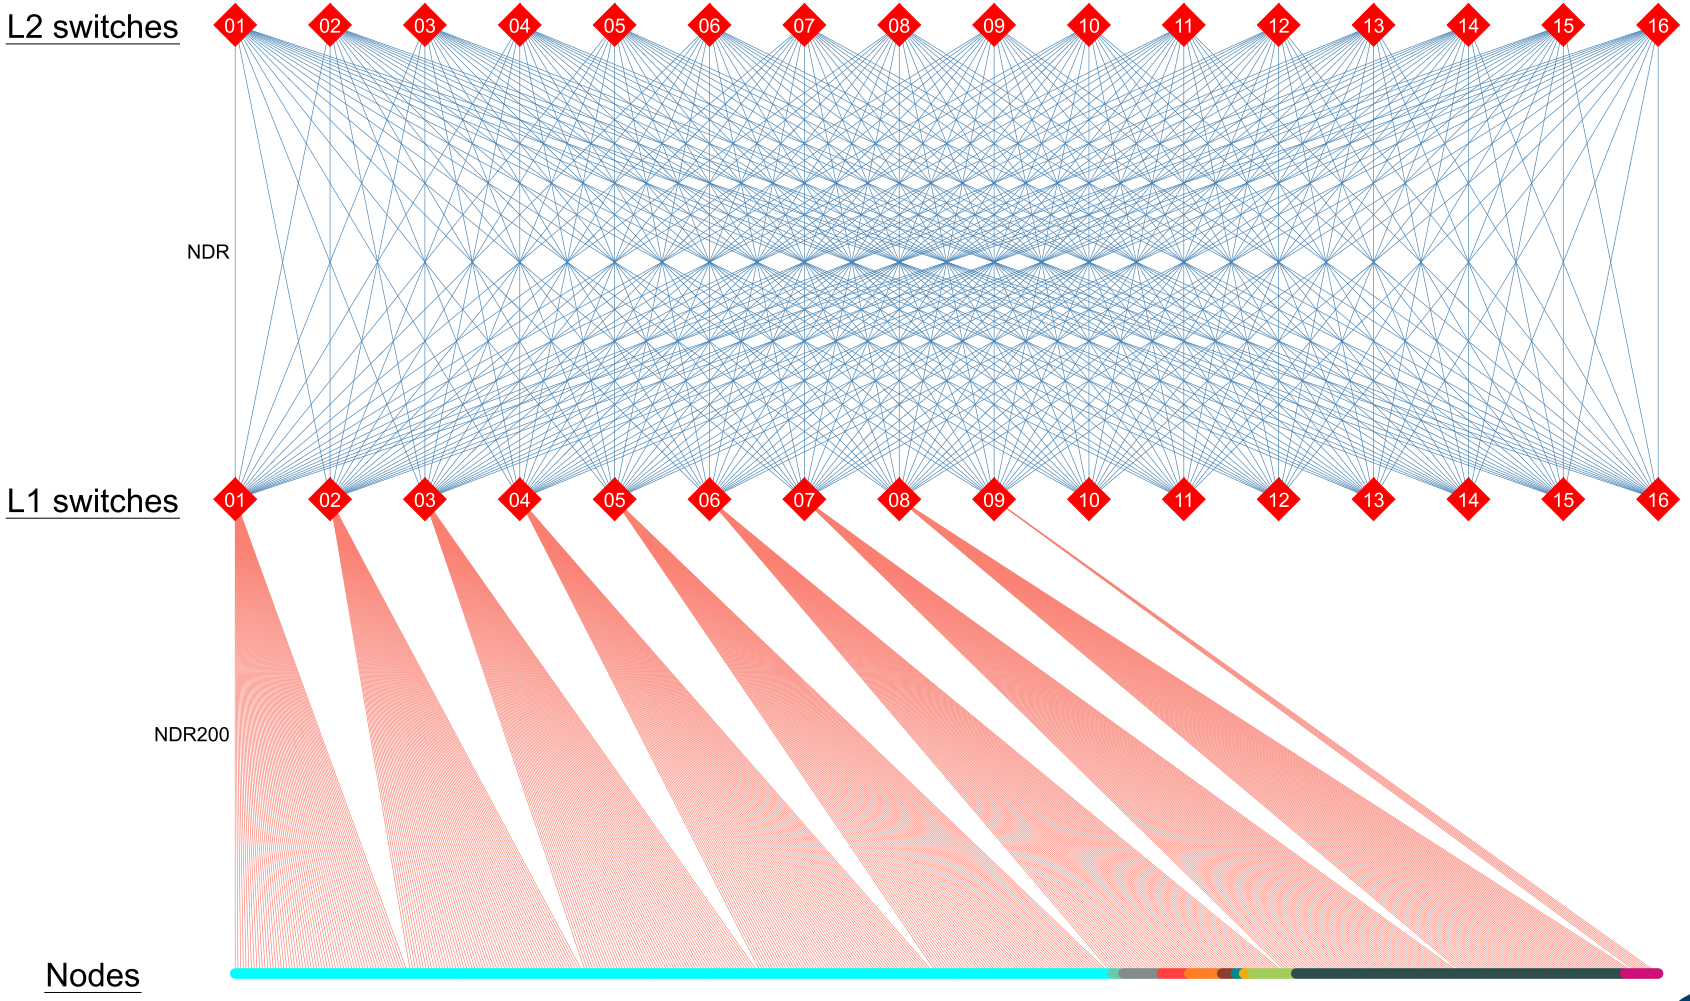

Admin cell

JUPITER – INTERCONNECT

One Network to Rule Them All

EVIDEN
an atos business

NVIDIA

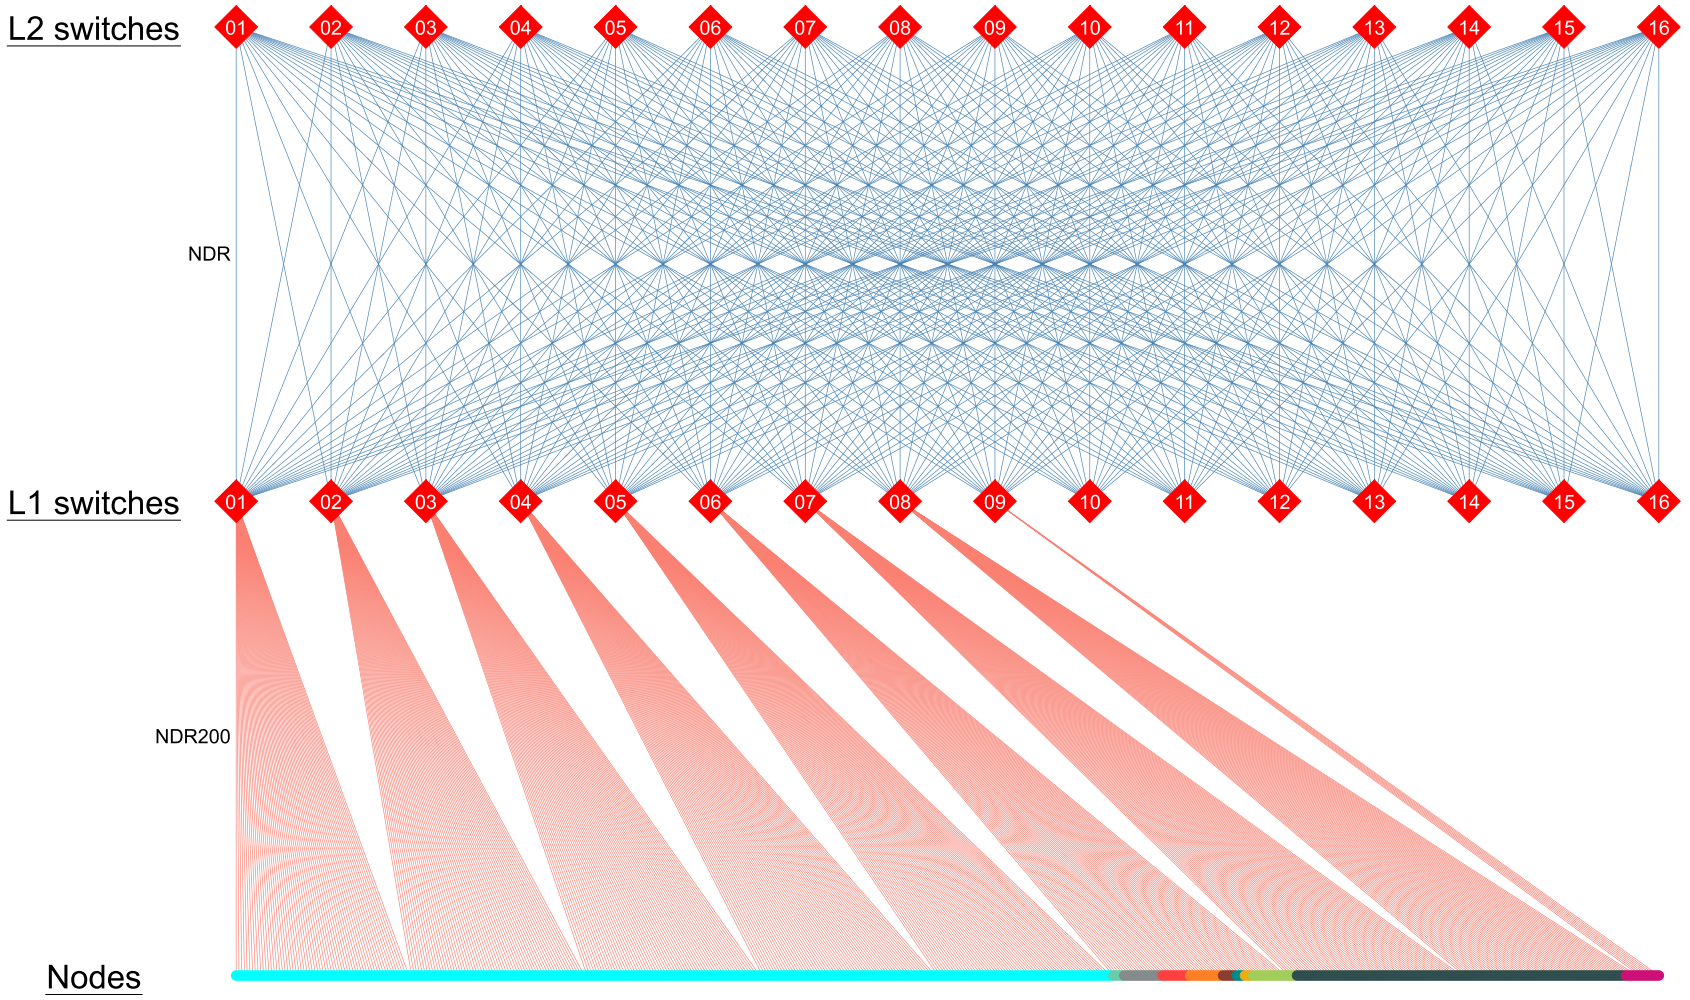

Old plot, topology and number of nodes not accurate

- 40x Flash nodes
- 44x Storage nodes
- 5x Datamover nodes
- 5x Cluster login nodes
- 12x Booster login nodes
- 3x Cluster vis nodes
- 3x Booster vis nodes
- 2x Gateways
- 22x management nodes

STORAGE

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

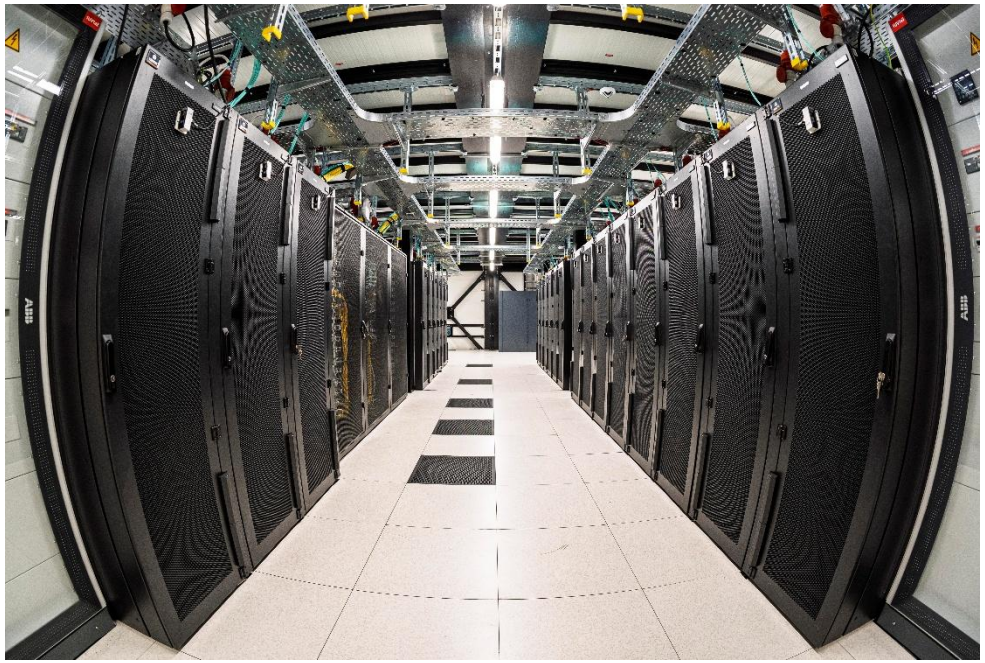

Preliminary numbers, might change during installation

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

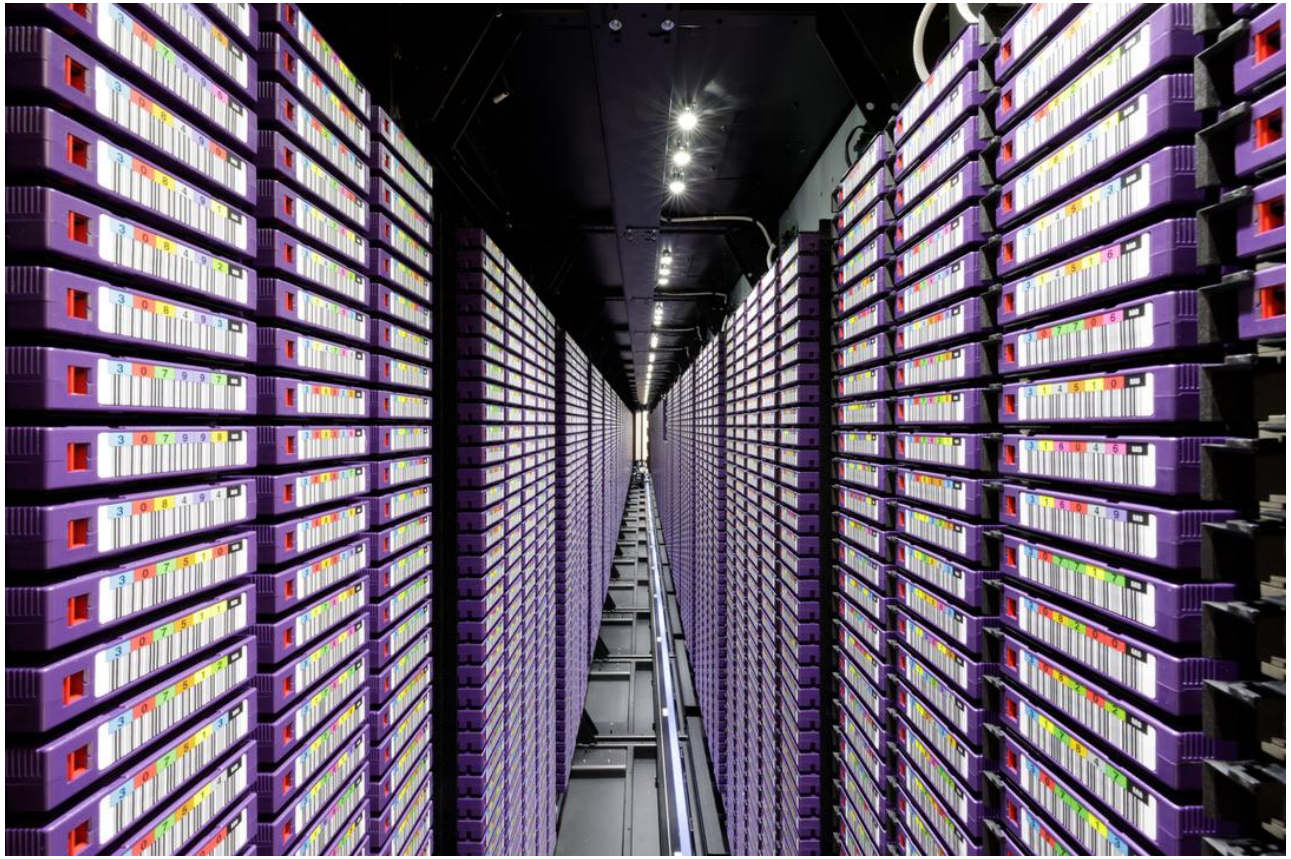

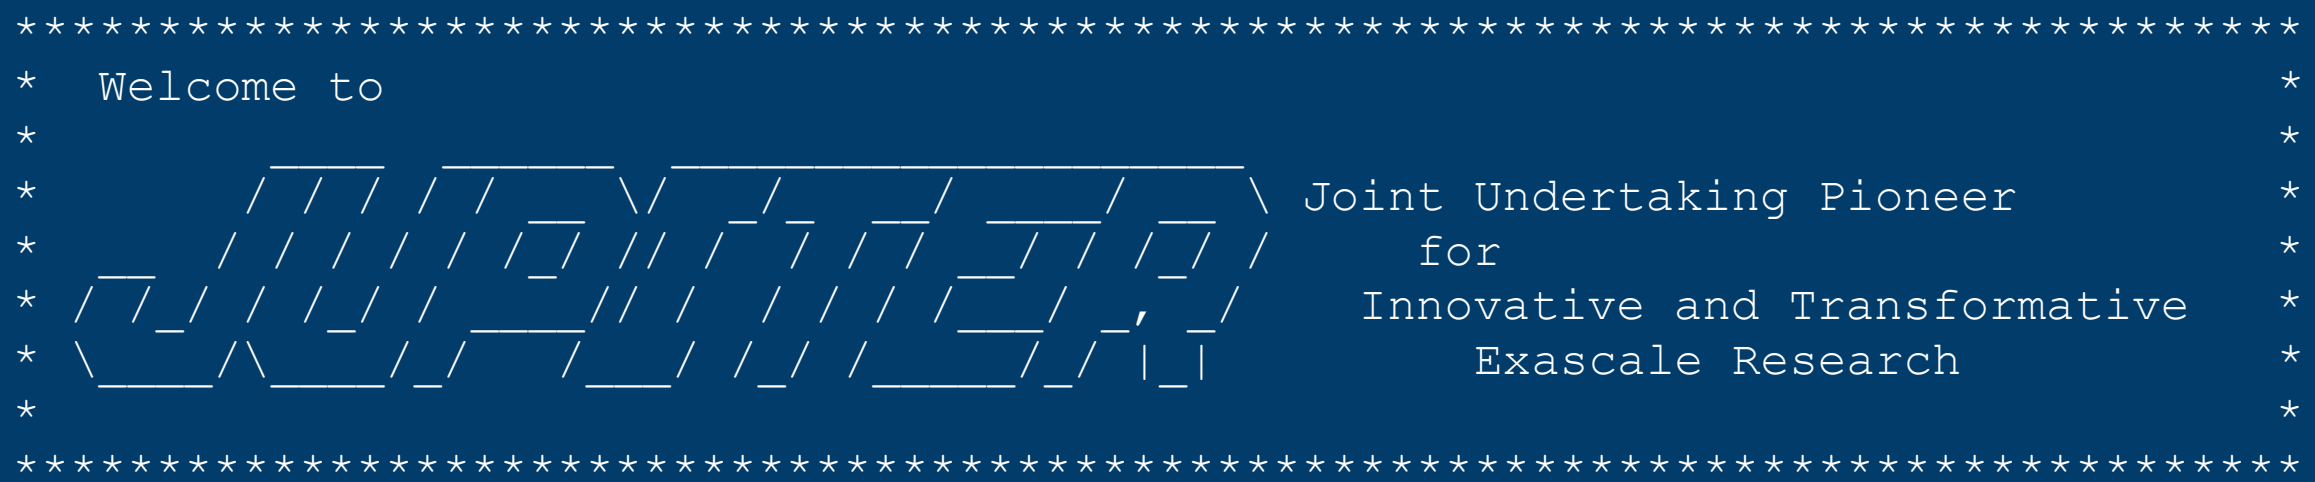

SYSTEM MANAGEMENT

JUPITER MANAGEMENT STACK

”Power is nothing without control"

- Eviden SMC xScale
- ParaStation Modulo
  - Resource management
  - ParaStation MPI
- Ansible as provisioning system
- SLURM as scheduler
- EasyBuild as scientific software package management
- RedHat Enterprise Linux 9

ParaStation MODULO

slurm workload manager

EasyBuild.io building software with ease

ANSIBLE

ceph

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	Red Hat
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	Red Hat
Boot Image(s) Management	ImageBuilder	ARM / x86 support, Tracking, Integration	Red Hat
Container(s)	UBI Universal Binary Images	Standard, Security, Consistency	Red Hat

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

JUPITER SOFTWARE STACK

JSC Software Team

- **Software Core Team**
  - 7 people
  - Responsible for core installation (GCCcore, compiler, MPI, Math)
  - Responsible for reviewing and merging PR into the JSC easybuild repository
- **Software Group**
  - Group of 25 contributors (60 in our JSC-internal EasyBuild chat)
  - Each module has has one responsible person which is applications/packages expert
  - Responsible for writing EasyConfig/EasyBlock, testing and validation as well as user questions

JUPITER SOFTWARE STACK

Stage Concept: Annual Curated Set of Software Versions

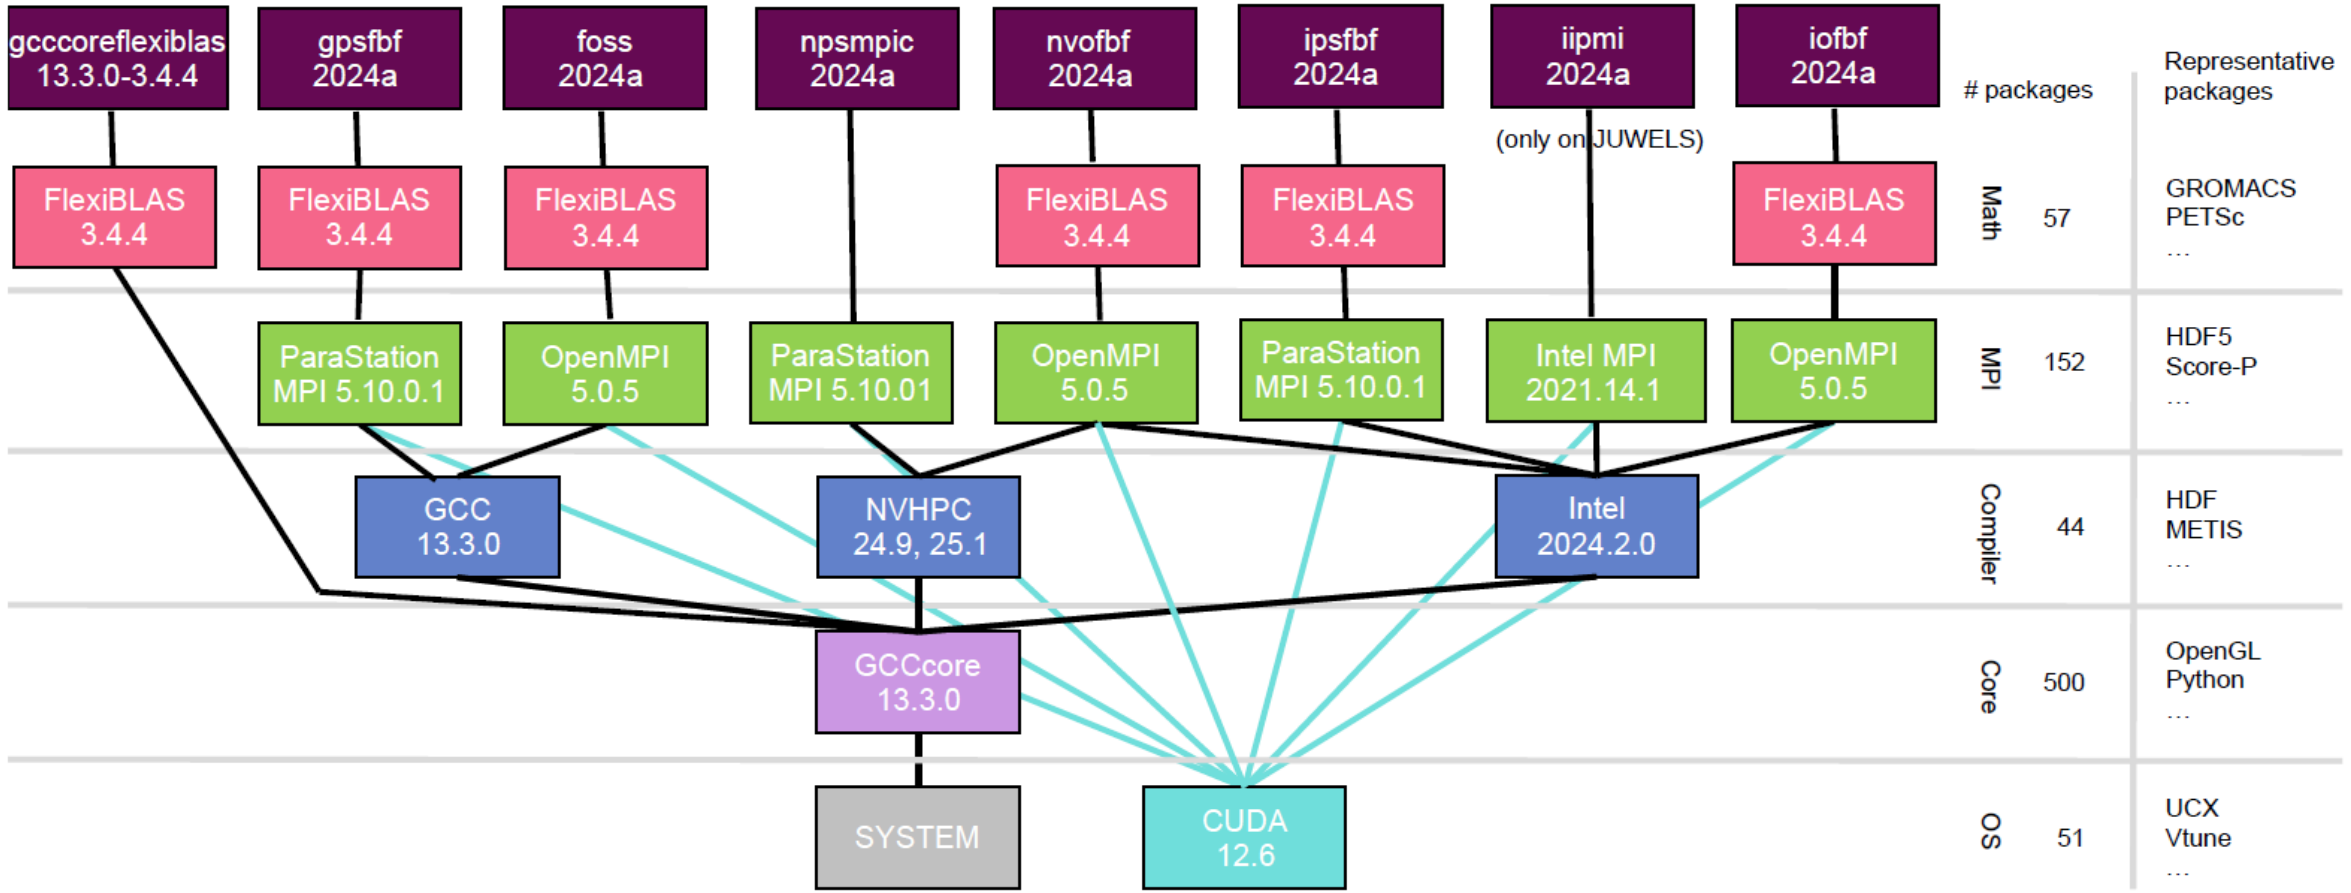

World-class Support

ALGORITHMS, TOOLS, METHODS LABS (ATML) + SIMULATION AND DATA LABS (SDL)

MULTI-TIER USER SUPPORT

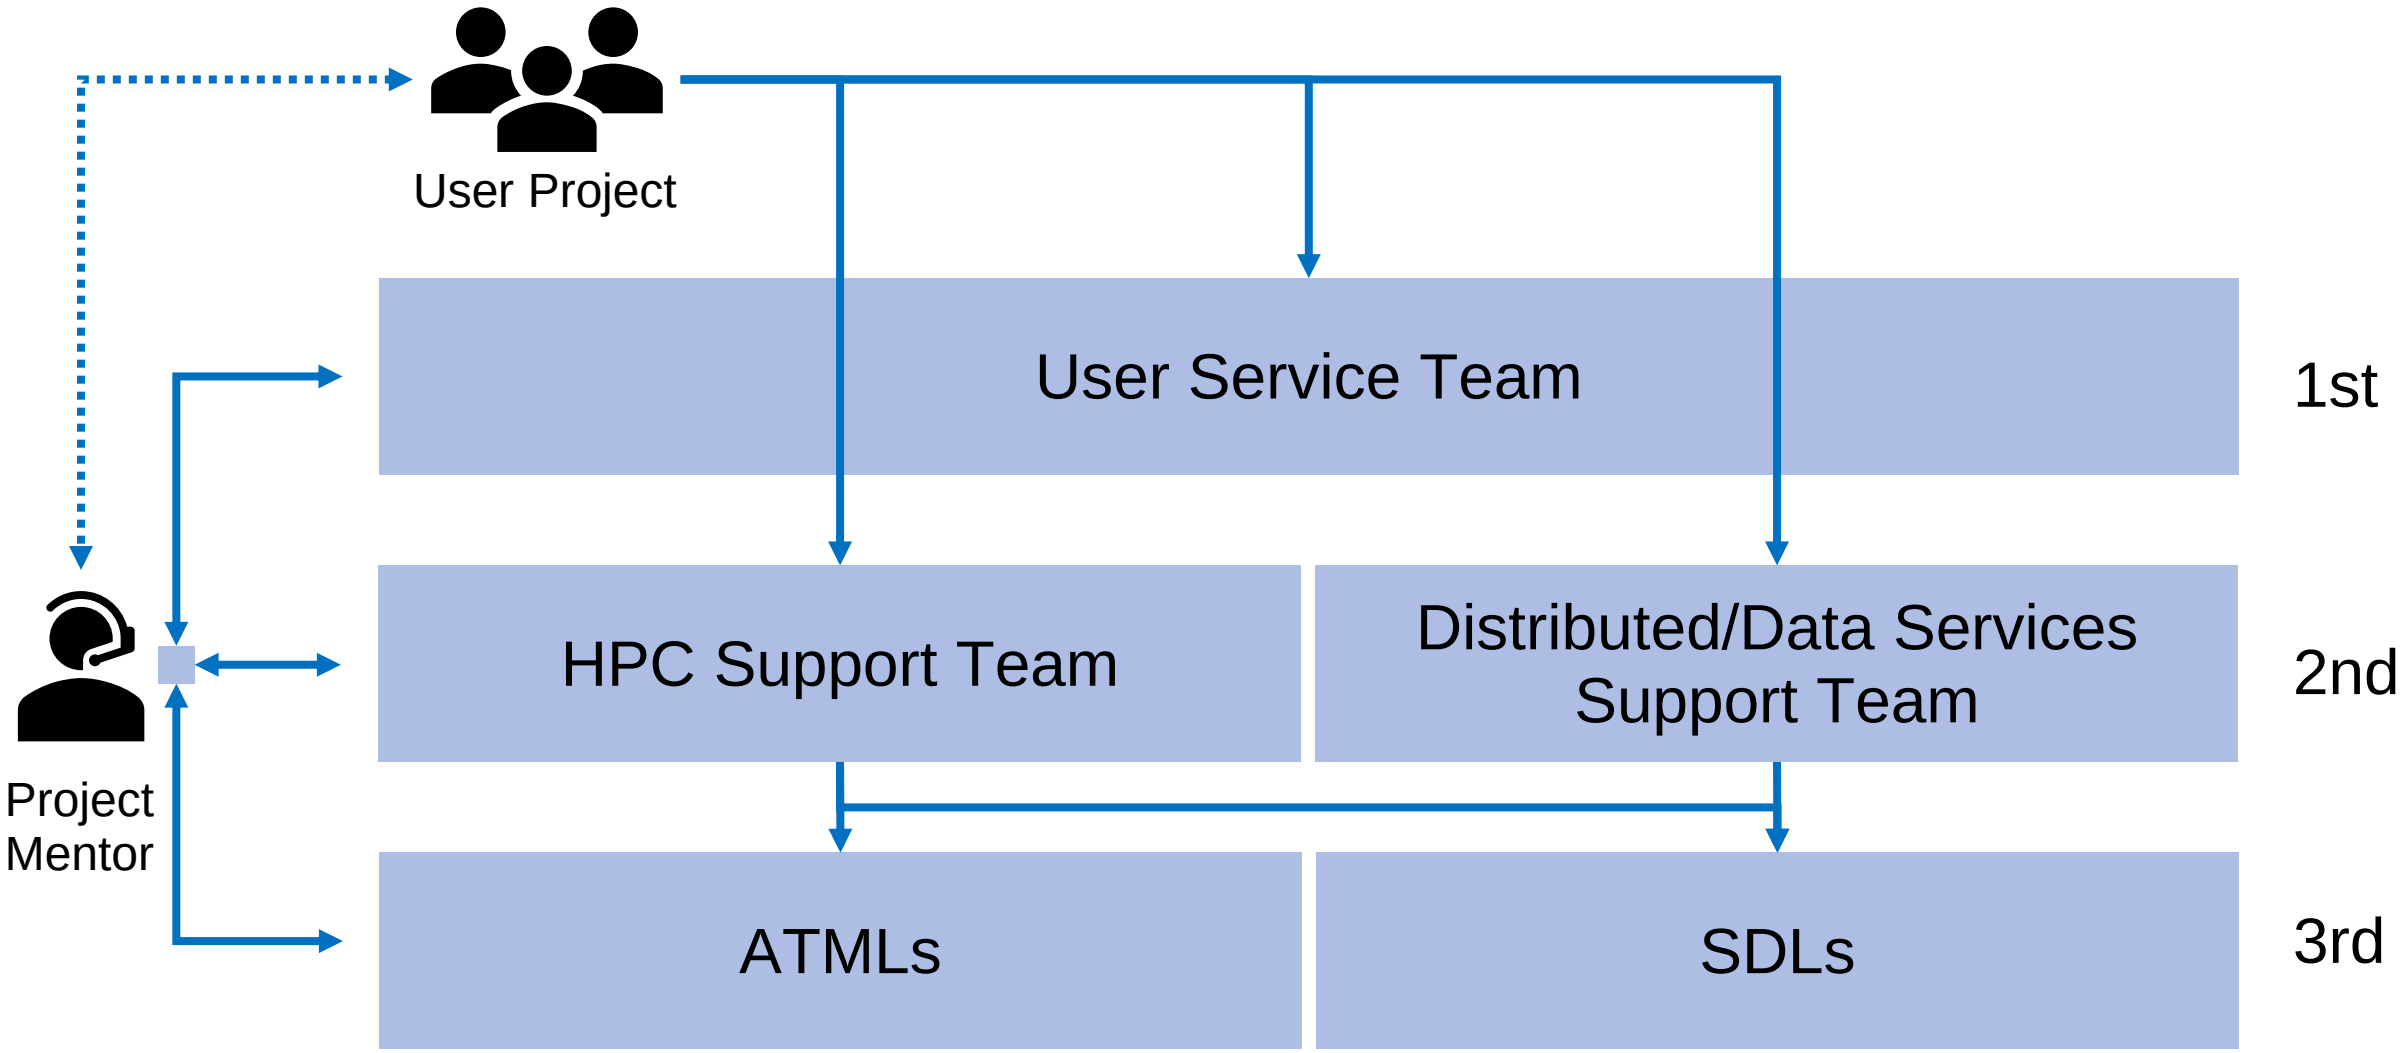

SUPPORT AND RESEARCH LANDSCAPE AT JSC

- Quantum Information Processing
- Earth System Data Exploration
- Computation Material Science
- Computational Structural Biology
- Next Generation Architectures
- Software for Modular
- Supercomputers
- RSE
- AI & ML for Healthcare

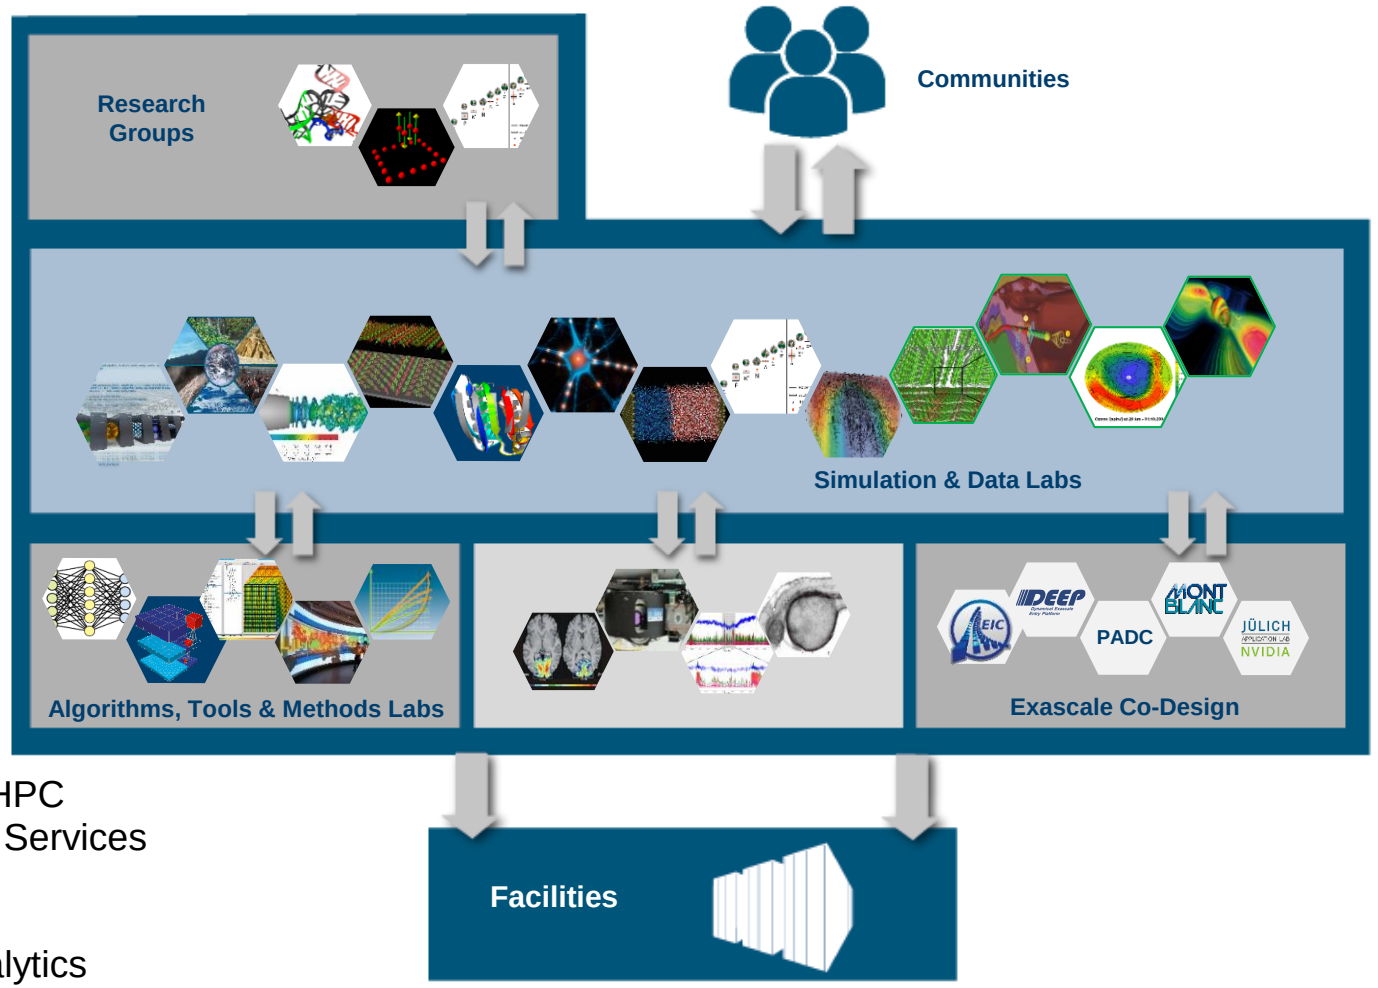

- Deep Learning
- Accelerating Devices
- Parallel Performance
- Application Optimization
- Applied Machine Learning
- Visualization & Interactive HPC
- Federation Technologies & Services
- Concurrency & Parallelism
- Advanced Time Integrators
- Data Management and Analytics
- Numerical & Statistical Methods

- Complex Particle Systems
- Quantum Materials
- Electrons and Neutrons
- Biology
- Neuroscience
- Fluids & Solids Engineering
- Plasma Physics
- Numerical Quantum Field Theory
- Astronomy & Astrophysics
- Climate Science
- Terrestrial Systems
- AI and ML for Remote Sensing

THE LESS DISTANT PAST

2023/2024 - THE PRESENT - SLAB, MDC, STORAGE, JEDI

… it is not only about waiting for JUPITER

ExaSTORE Procurement
300PB+ Spinning Disk Storage for JUPITER

2023/2024 - THE PRESENT - SLAB, MDC, 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: ~2300m²
- 1x Datahall (Storage, Management)
- 7x IT Modules (20 Racks per module)
- UPS, Generator
- Entrance area
- Workshop, Warehouse
- 15x 2,5 Megawatt Power Stations

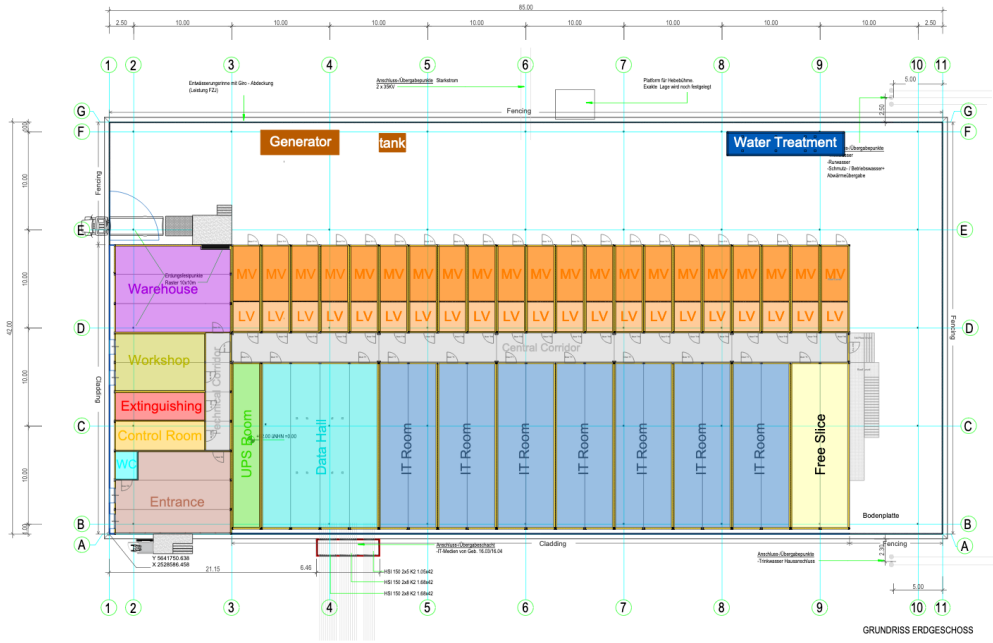

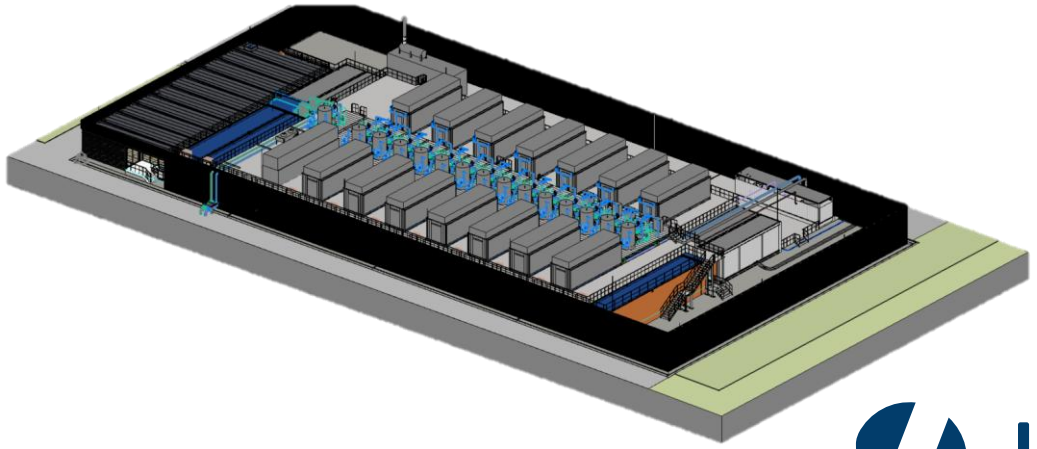

JÜLICH Forschungszentrum

JUPITER MODULAR DATA (MDC) CENTER

EVIDEN
an atos business

MODULAR DATA CENTER FOR JUPITER

EVIDEN

Virtual Tour: https://cloud.3dvista.com//hosting/8444525/0/

2023/2024 - THE PRESENT - SLAB, MDC, STORAGE, JEDI

… it is not only about waiting for JUPITER

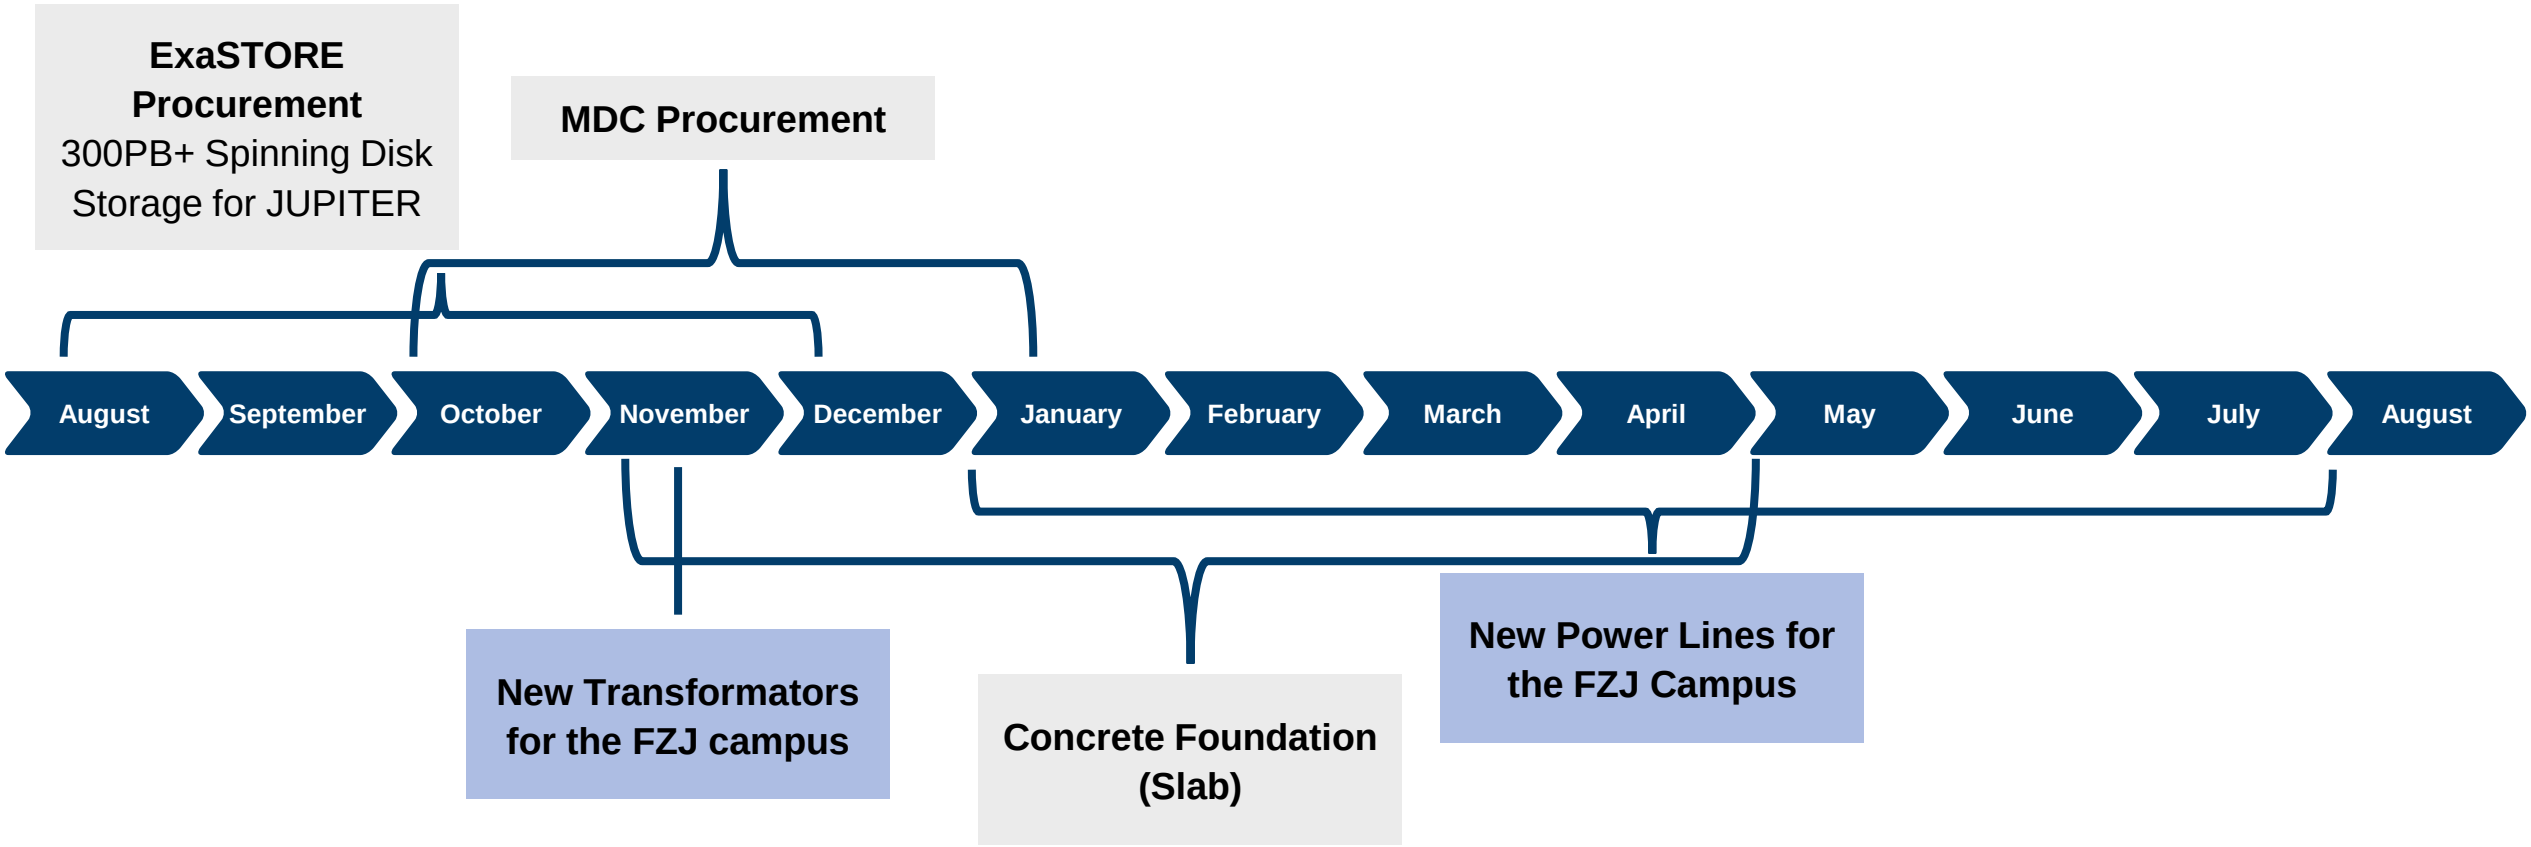

CONCRETE FOUNDATION

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

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

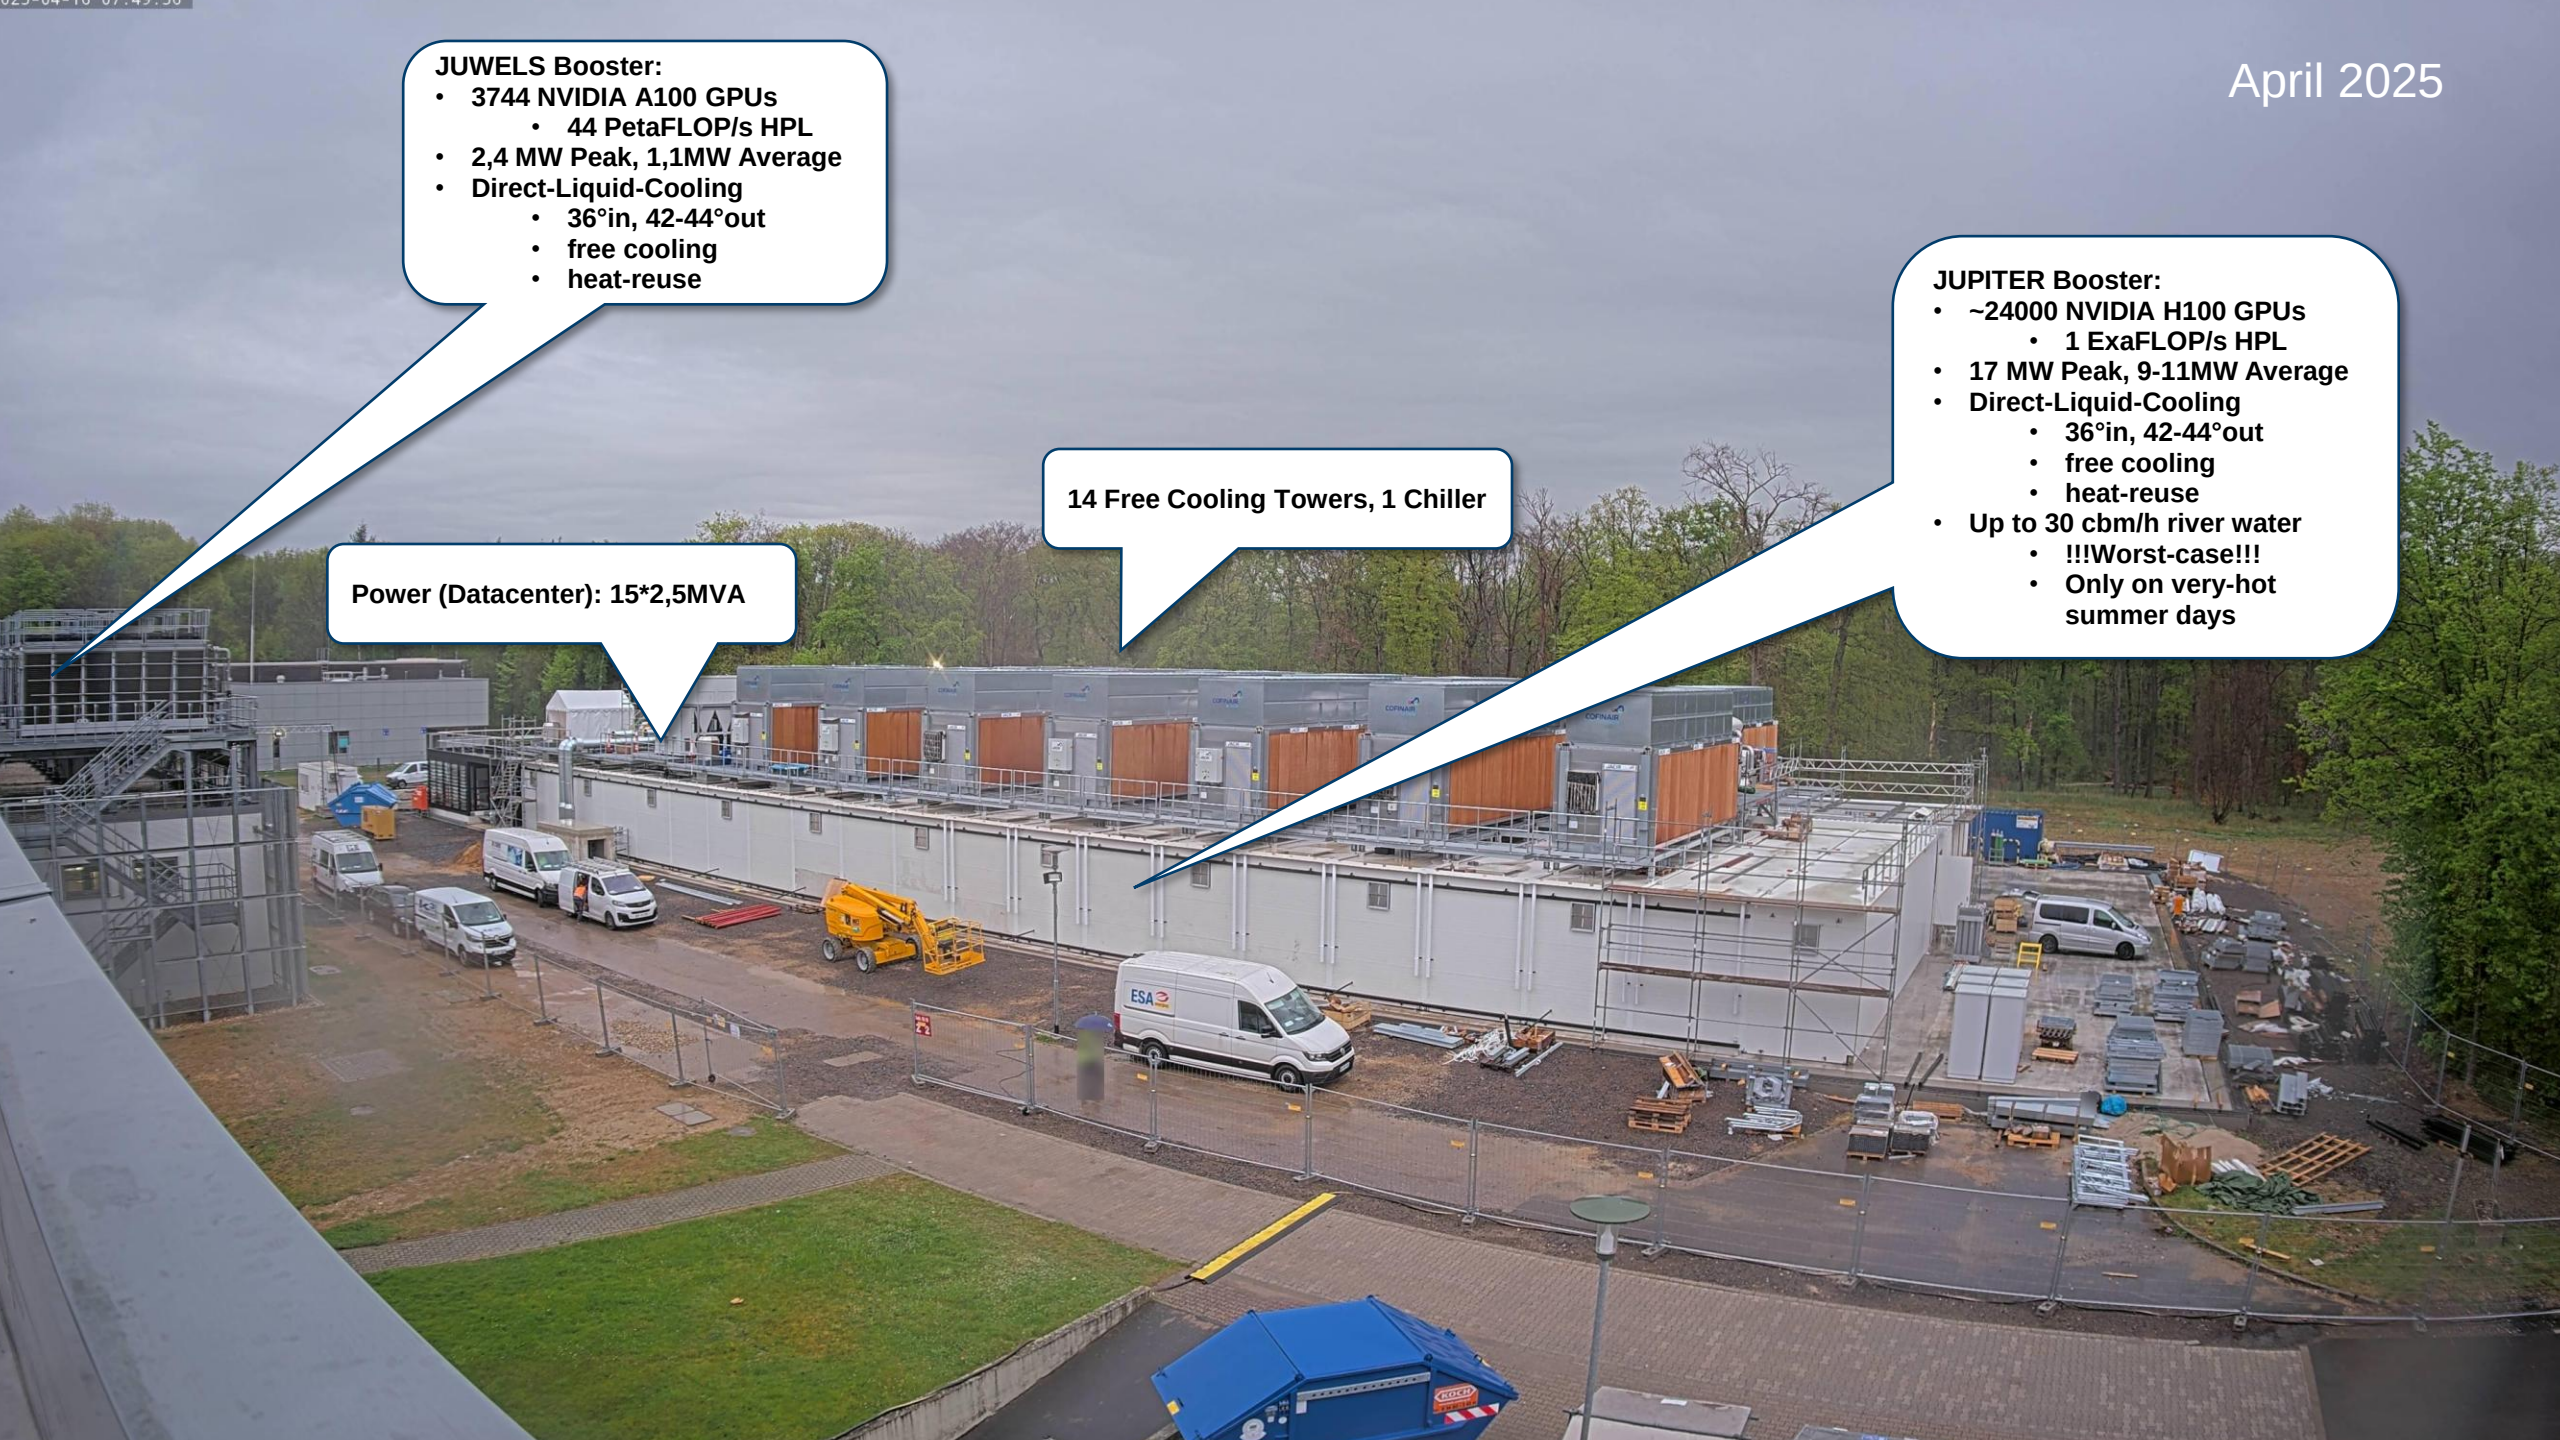
JUWELS Booster:
• 3744 NVIDIA A100 GPUs
• 44 PetaFLOP/s HPL
• 2,4 MW Peak, 1,1MW Average
• Direct-Liquid-Cooling
• 36°in, 42-44°out
• free cooling
• heat-reuse
Power (Datacenter): 15*2,5MVA
14 Free Cooling Towers, 1 Chiller
JUPITER Booster:
• ~24000 NVIDIA H100 GPUs
• 1 ExaFLOP/s HPL
• 17 MW Peak, 9-11MW 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

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

2023/2024 - THE PRESENT - SLAB, MDC, STORAGE, JEDI

… it is not only about waiting for JUPITER

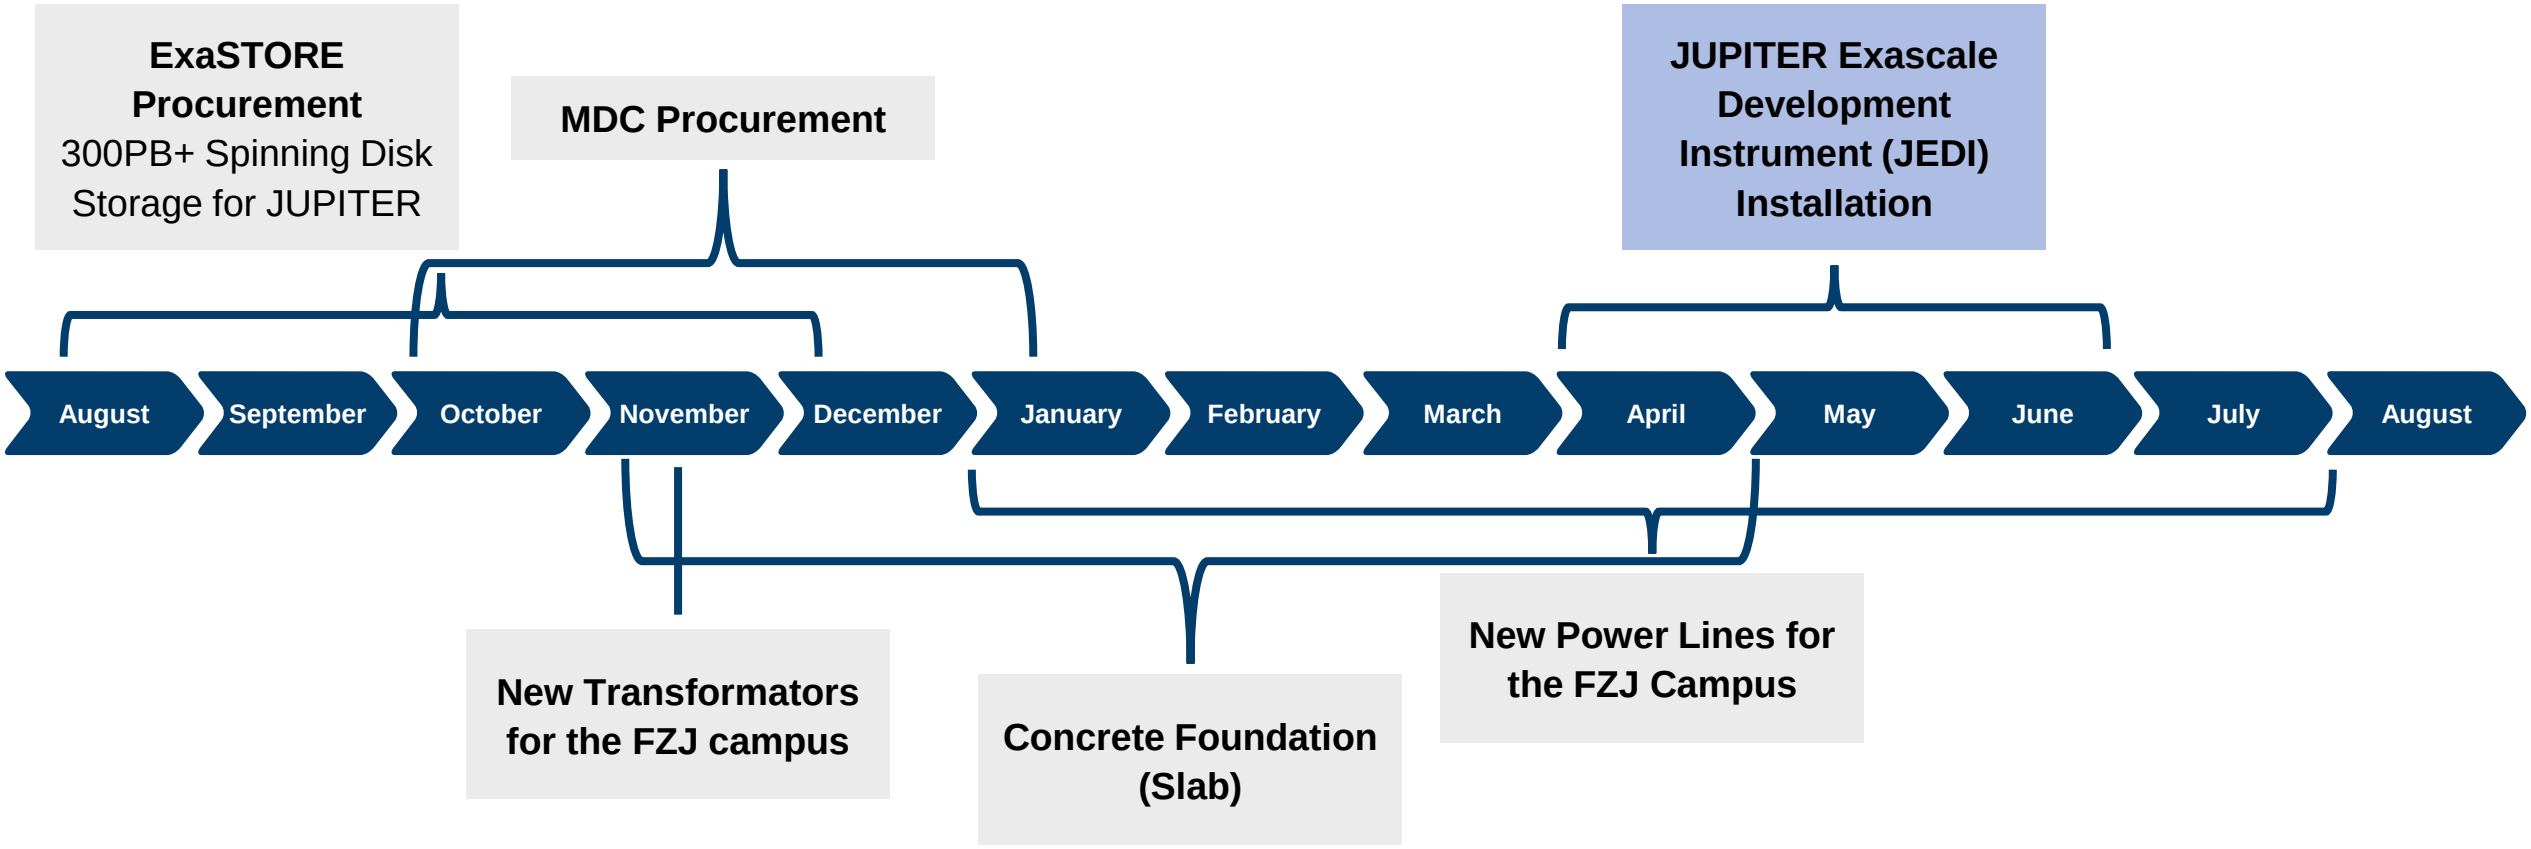

2024/2025 - MDC AND JUPITER INSTALLATION

Always in motion is the future…

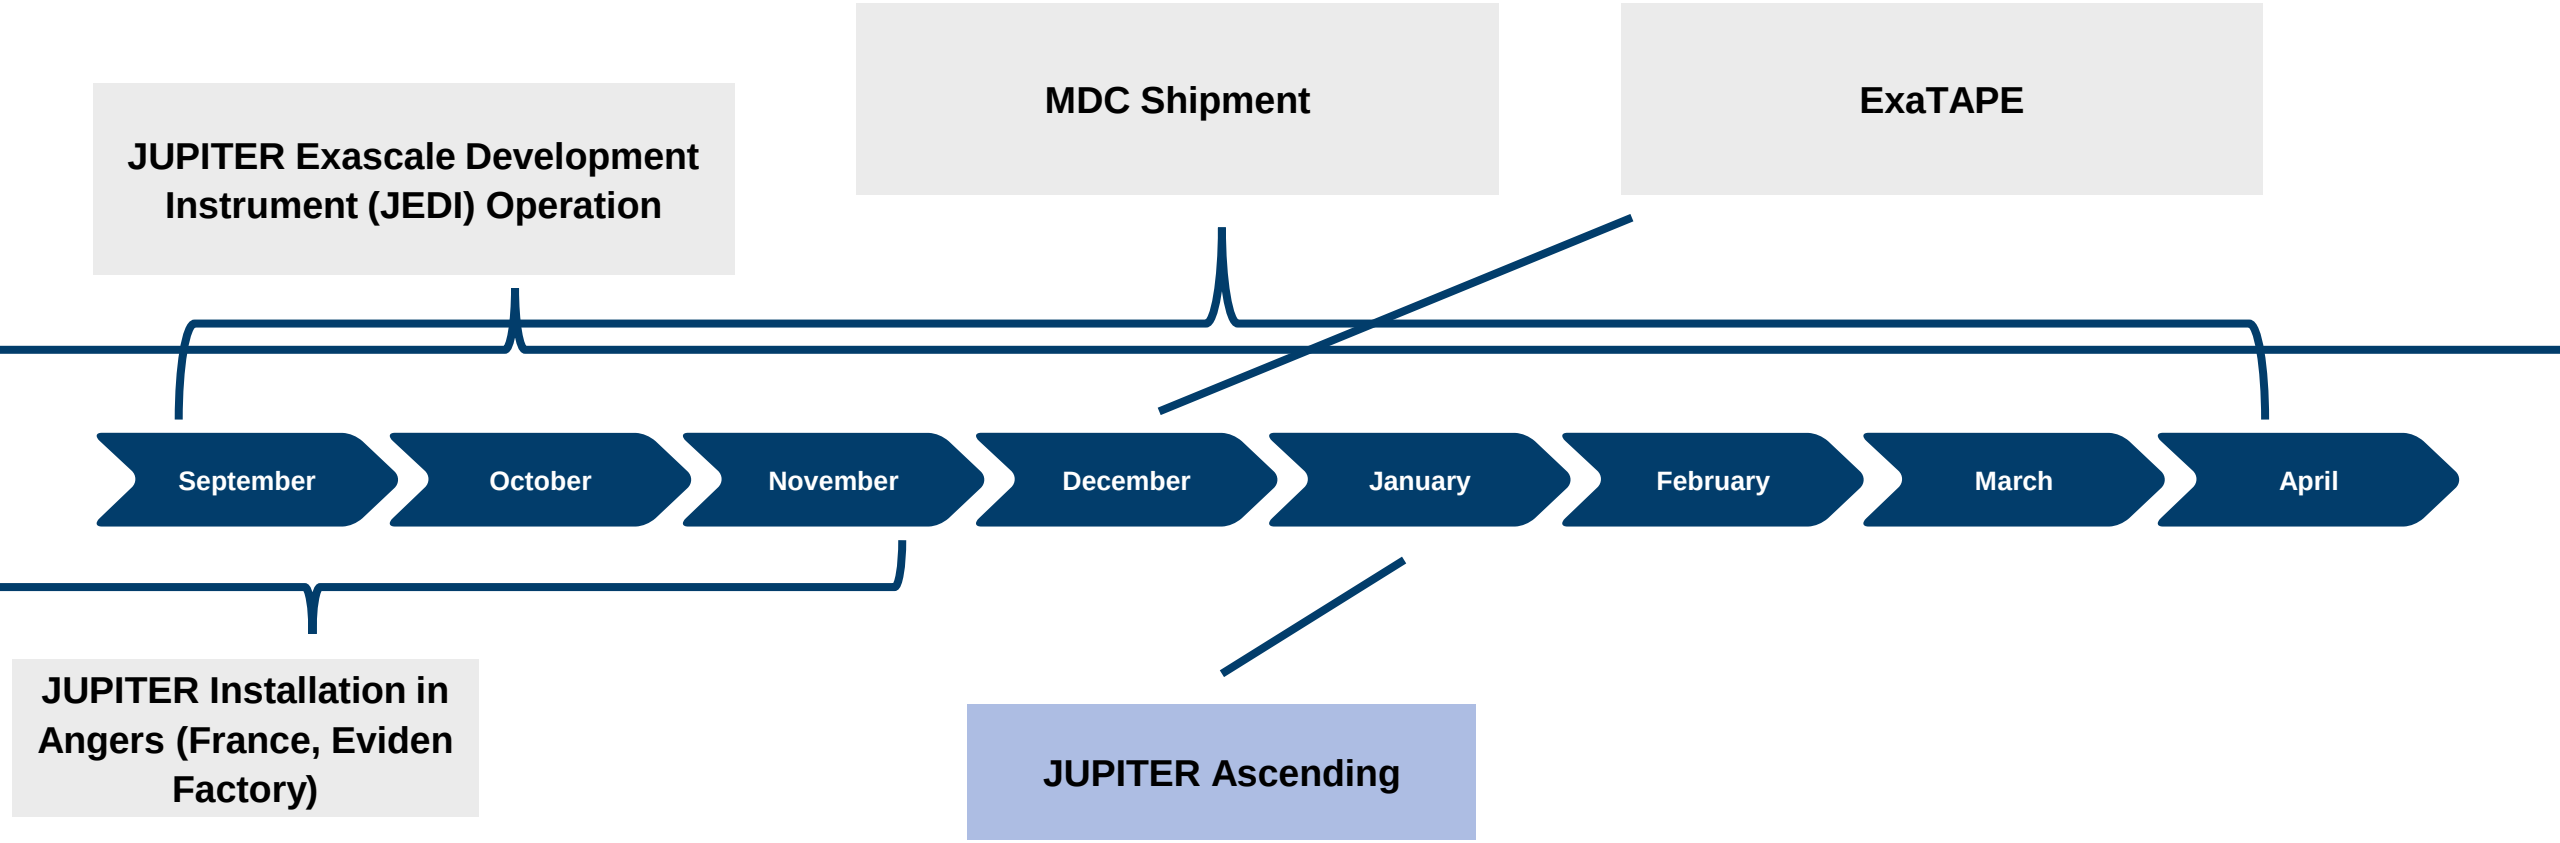

JUPITER ASCENDING

Since January 2025

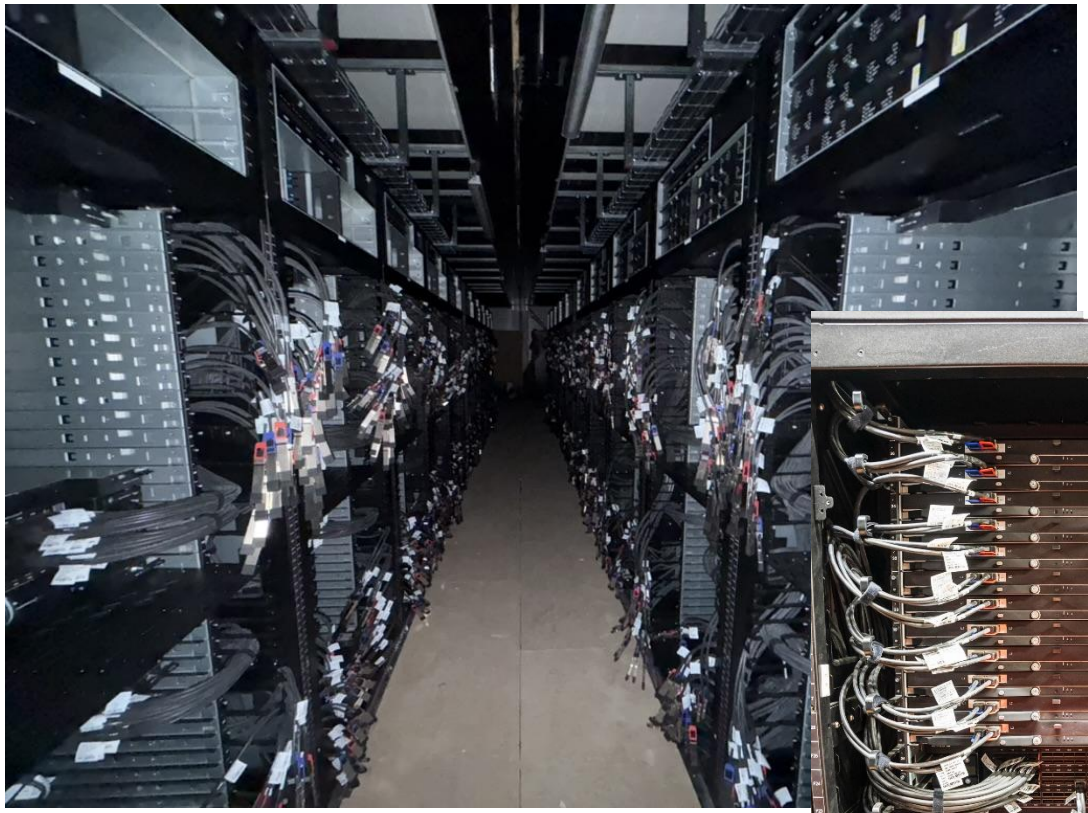

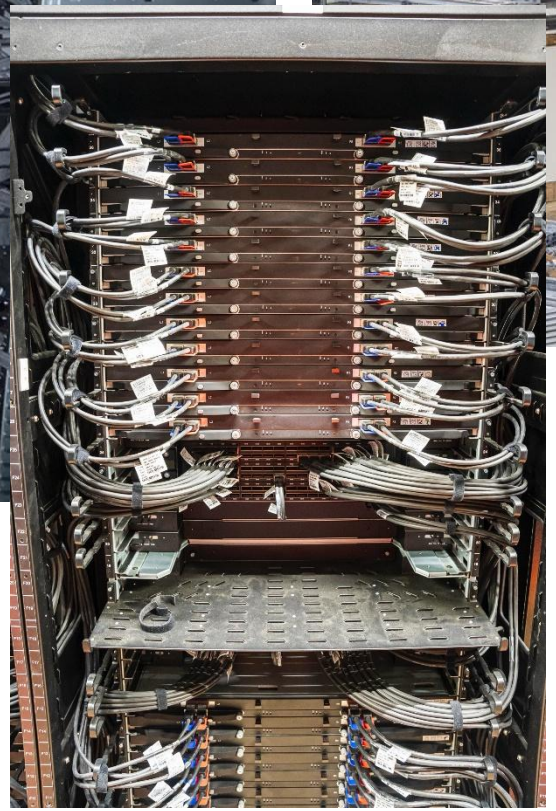


0221 / 956 86-0
www.COLONIA.info

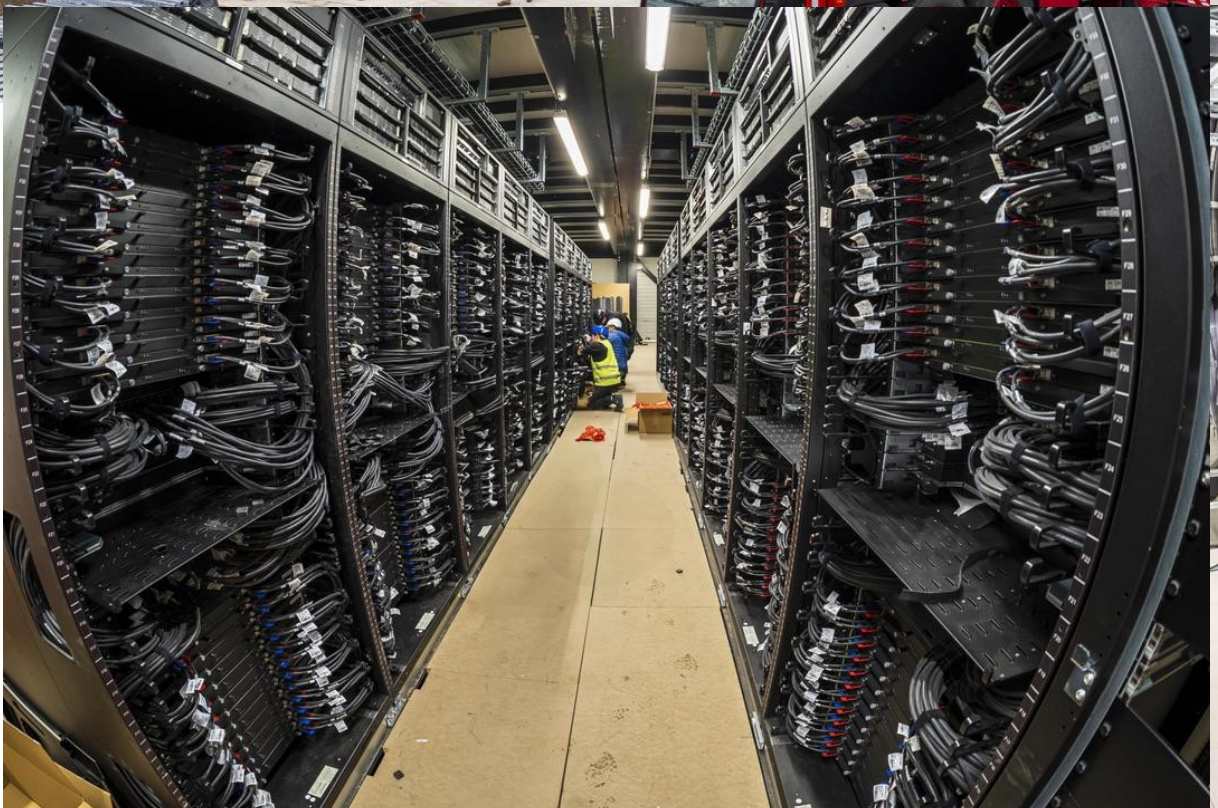

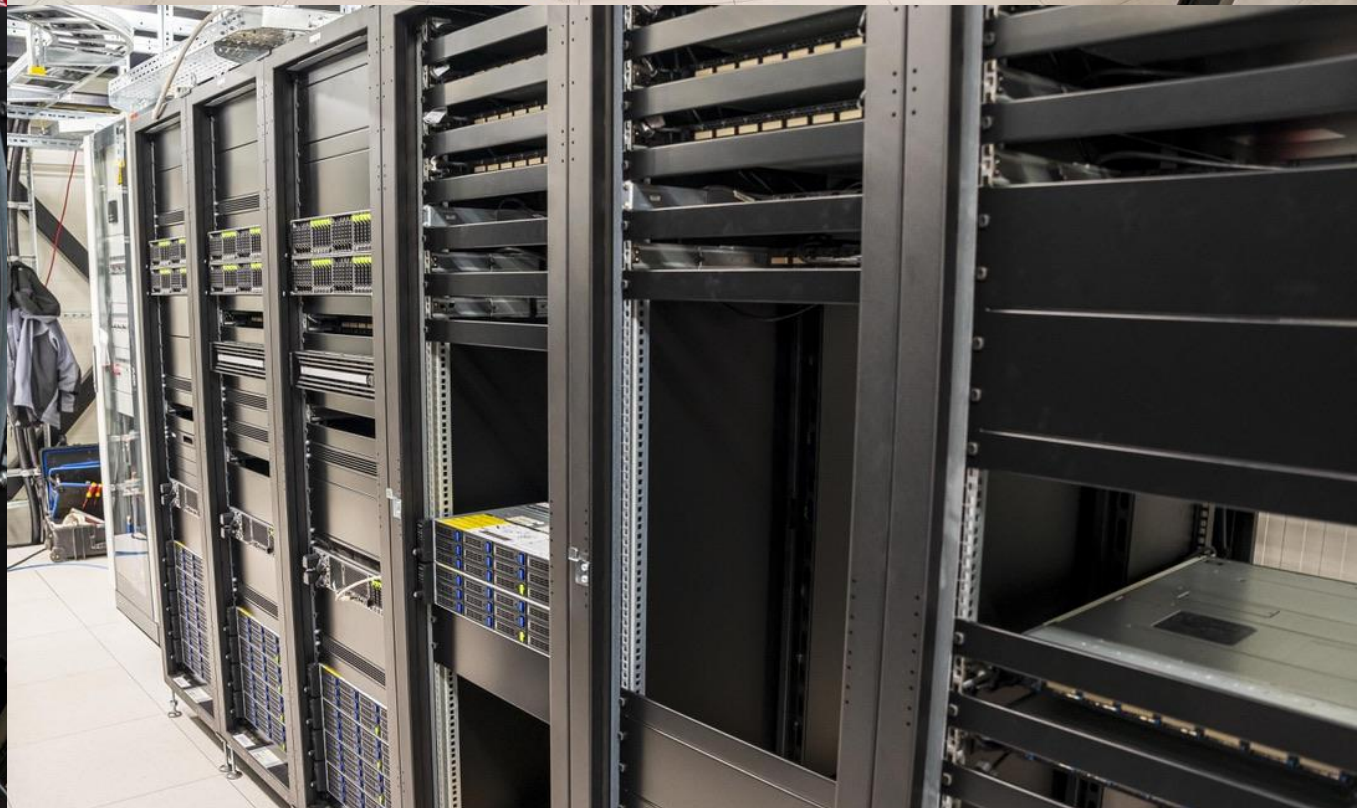

FIRST PUBLIC ACHIEVEMENTS

JUPITER EXASCALE DEVELOPMENT INSTRUMENT

EuroHPC / Forschungszentrum Jülich

- Eviden BullSequana XH3000
  - 24x Compute nodes (12x Blades)
    - NVIDIA quad-GH200 96GB Grace Hopper Superchip
    - Memory: 480GB on CPUs + 384GB on GPUs
    - NVIDIA quad-rail InfiniBand NDR200
  - 1x Network switch:
    - NVIDIA Quantum-2 NDR InfiniBand switch
  - All components are Direct Liquid Cooled

JEDI

#1 in Green500 (05/2024)
#189 in TOP500

CERTIFICATE

JEDI - BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail NVIDIA InfiniBand NDR200

EuroHPC/FZJ, Germany

is ranked

No. 189

among the World's TOP500 Supercomputers

with 4.50 PFlop/s Linpack Performance

in the 63rd TOP500 List published at the ISC24

Conference on June 01, 2024.

Congratulations from the TOP500 Editors

Erich Strohmaier
NERSC/Berkeley Lab

Jack Dongarra
University of Tennessee

Horst Simon
NERSC/Berkeley Lab

Martin Meuer
Prometeus

The GREEN 500

CERTIFICATE

JEDI - BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail NVIDIA InfiniBand NDR200

EuroHPC/FZJ, Germany

is ranked

No. 1

among the World's TOP500 Supercomputers

with 72.733 GFlops/watts Performance

in the Green500 List published at the ISC24

Conference on June 01, 2024.

Congratulations from the Green500 Editors

Wu-chun Feng
Virginia Tech

Kirk Cameron
Virginia Tech

TOP500 – JUNE 2025

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

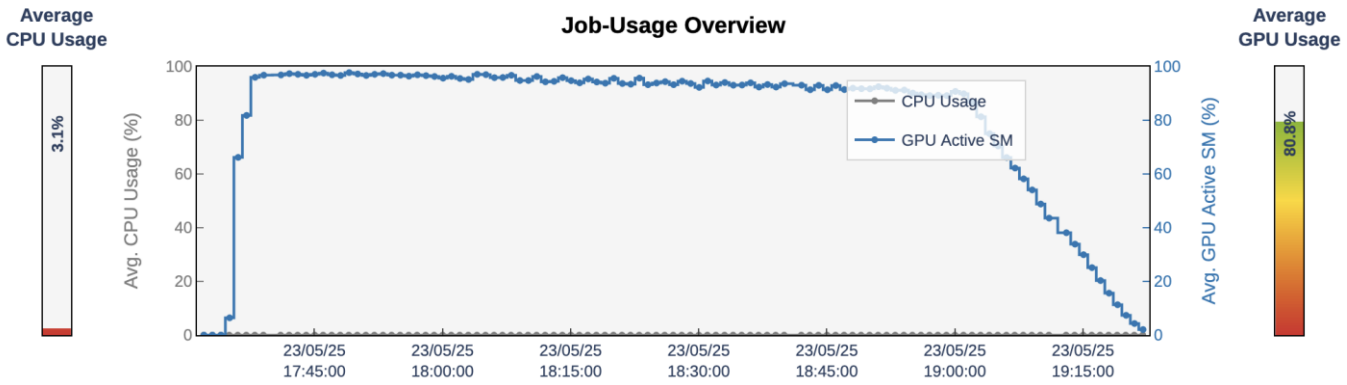

LLview job report of (a) HPL run

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
```

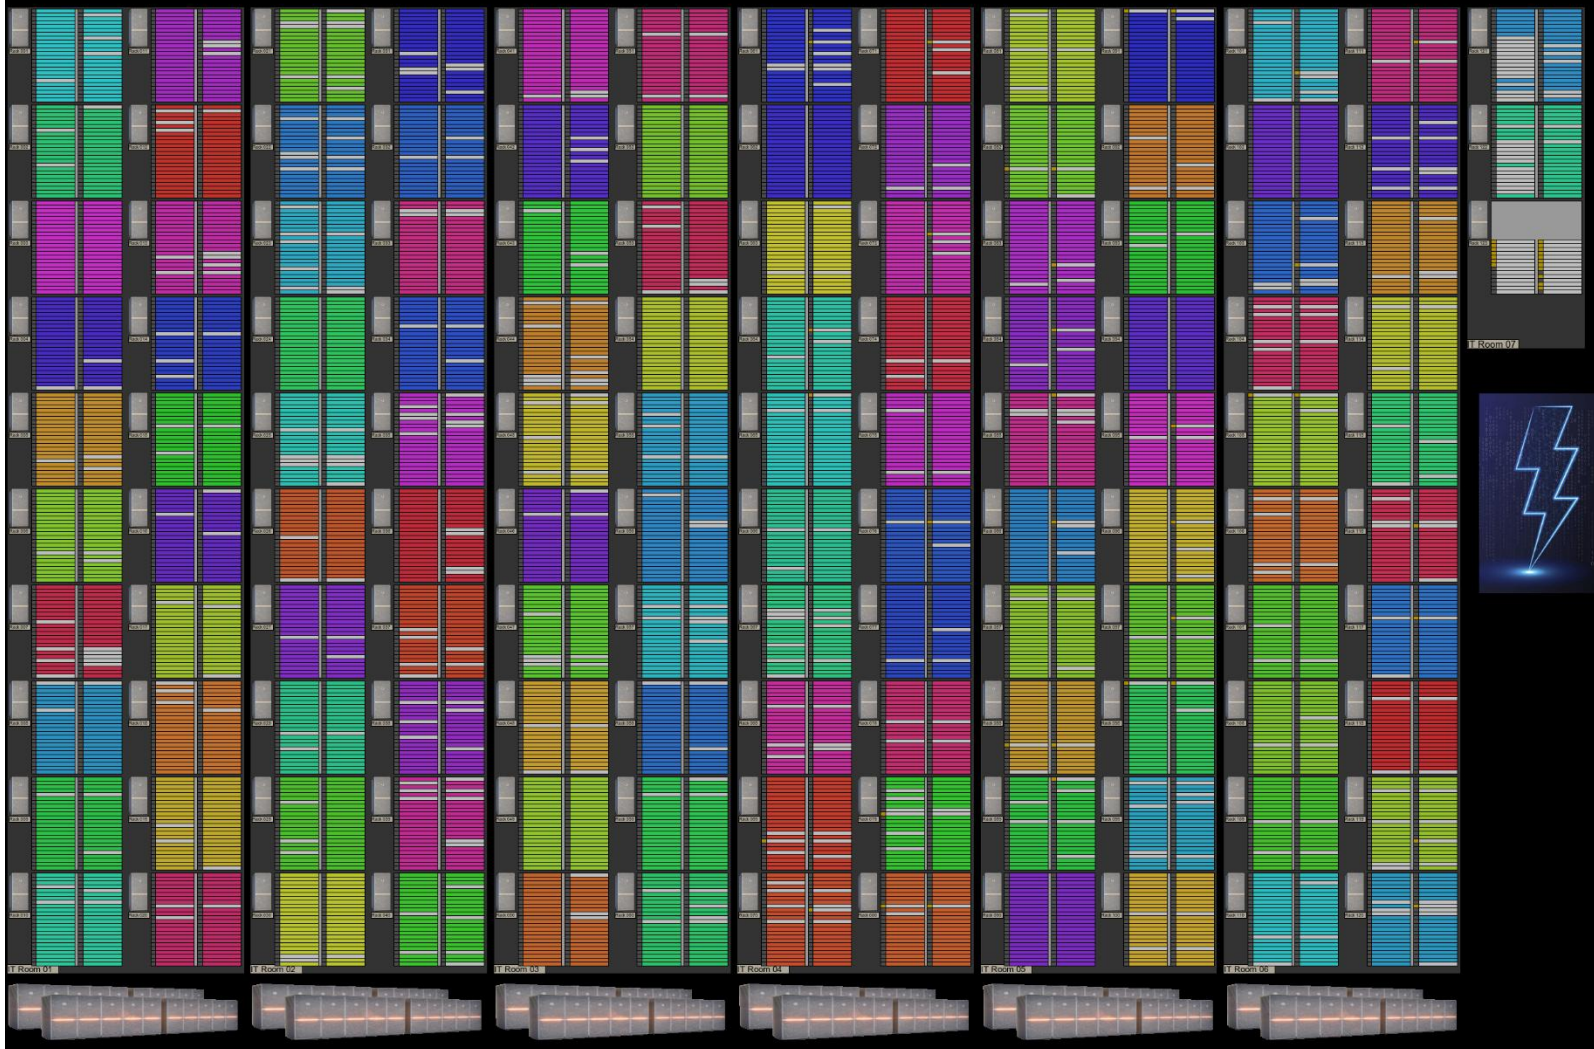

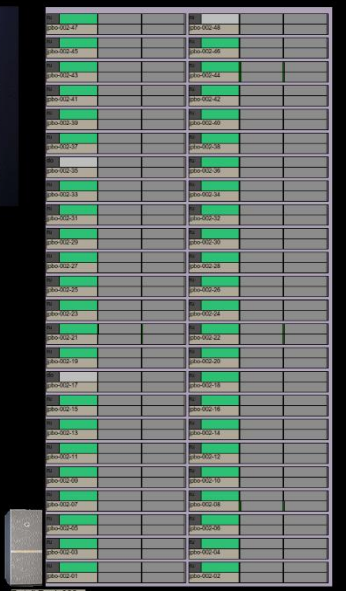

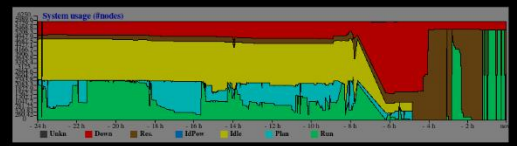

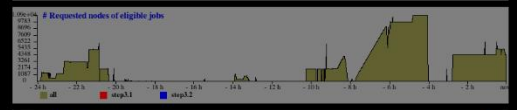

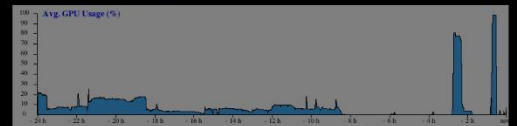

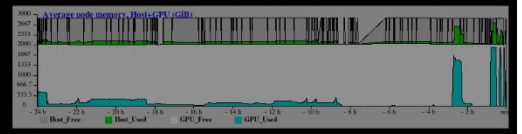

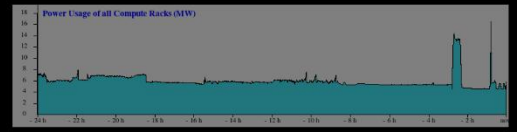

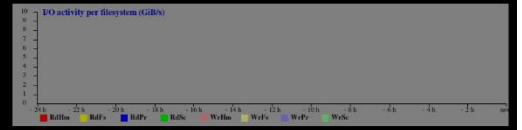

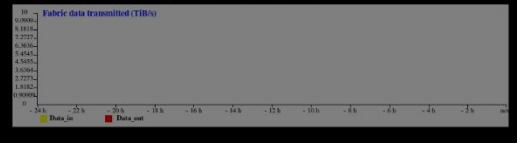

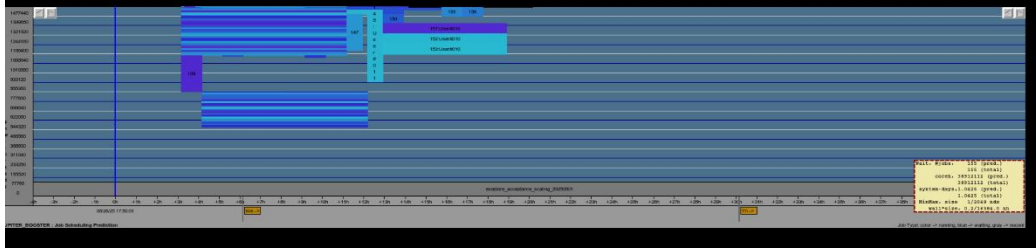

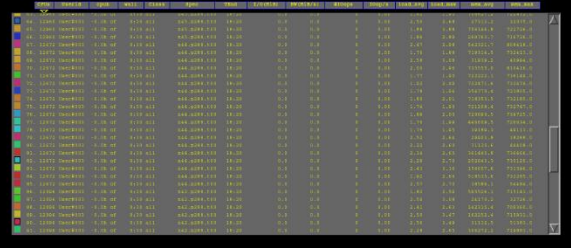

THE PRESENT


2025-09-10 06:59:56


2025-09-12 11:19:56


JUPITER
THE FIRST EXASCALE SYSTEM IN EUROPE
EuroHPC
Joint Undertaking
JÜLICH
Forschungszentrum
JUPITER
JOINING FORCES
EuroHPC
GCS
JÜLICH
ParTec
EVIDEN
NVIDIA
SIPEARL
IBM
jupiter.fz-juelich.de


EVIDEN


Inauguration: 5. Sep 2025
JÜLICH
Forschungszentrum

THE [NEAR] FUTURE

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

JÜLICH Forschungszentrum

Member of the Helmholtz Association

THE JAIF ONE STOP SHOP

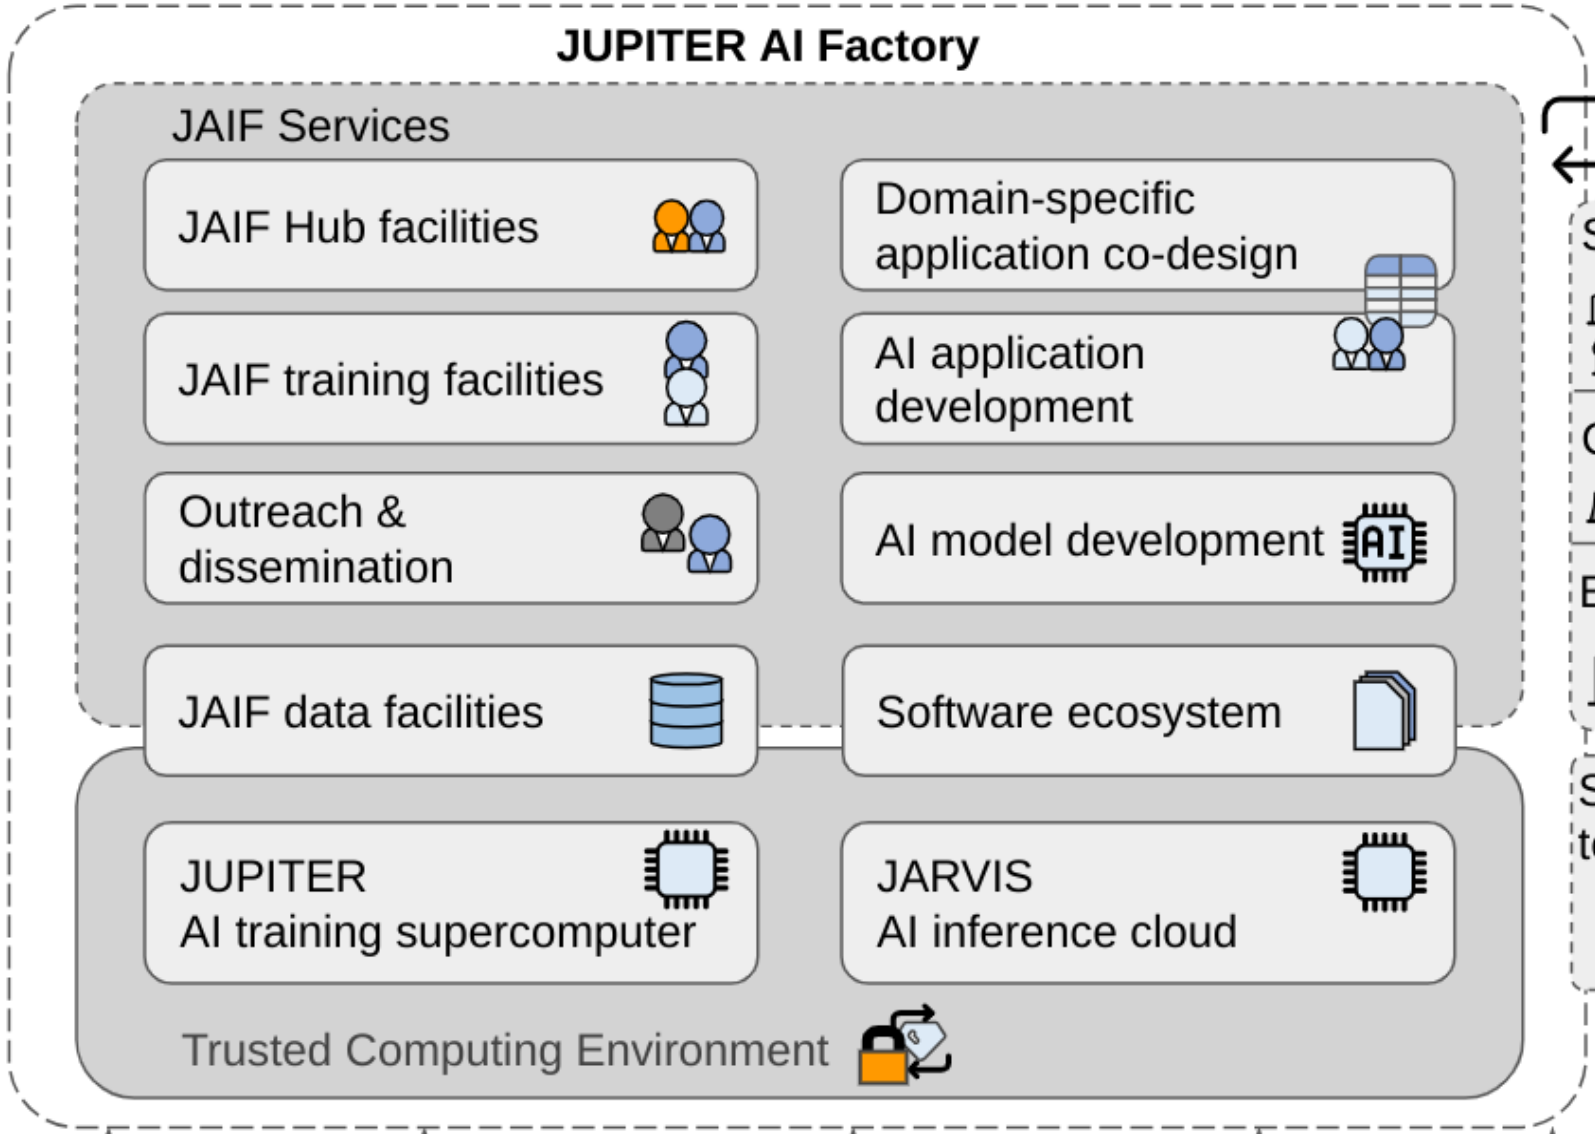

MODULAR JUPITER
HYBRID TRAINING/INFERENCE AI

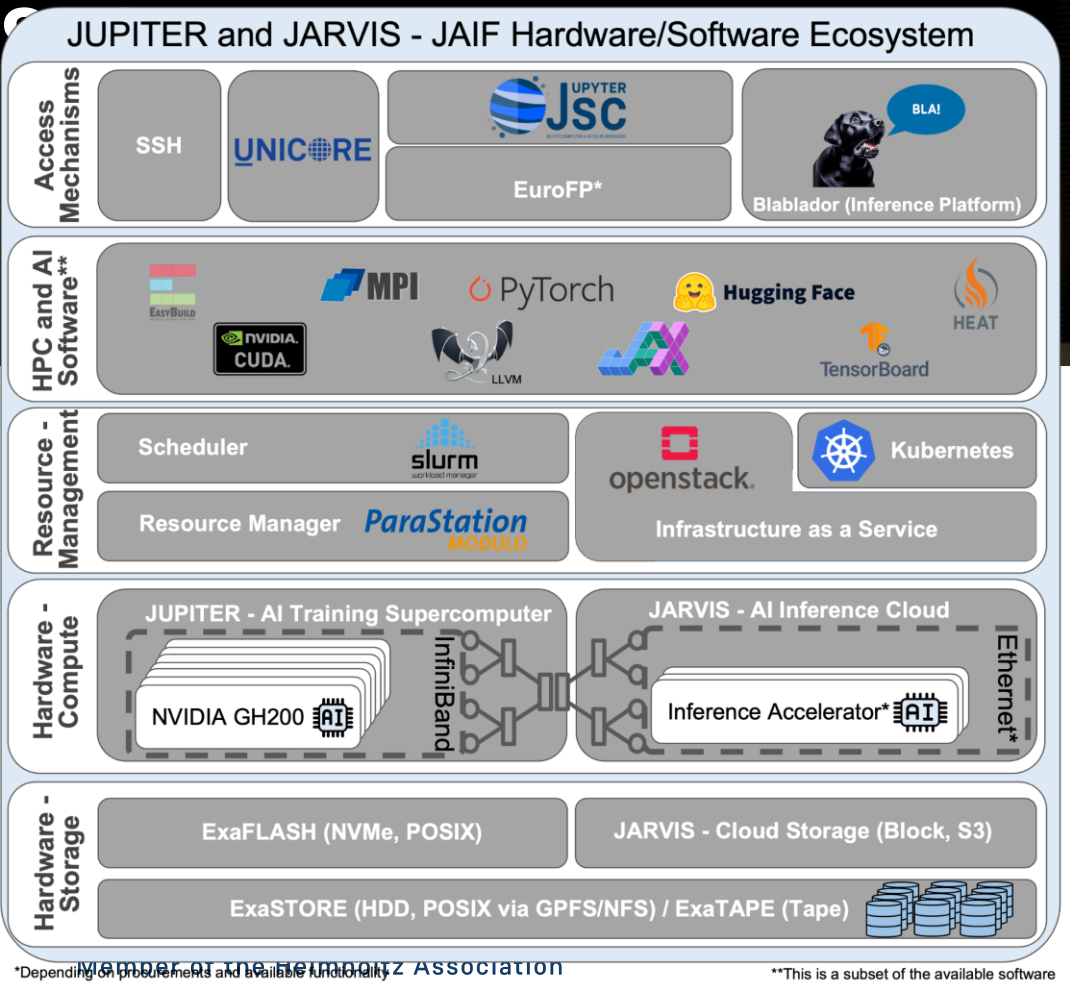

*Depending on procurements and available functionality

**This is a subset of the available software

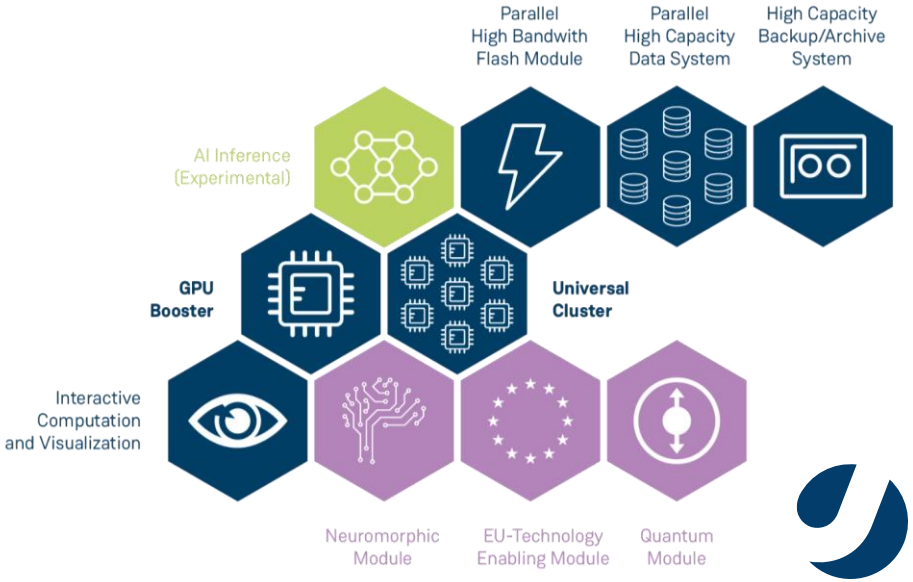

Forschungszentrum JUPITER AI FACTORY

POPULATING JUPITER

Current Status and Next Steps

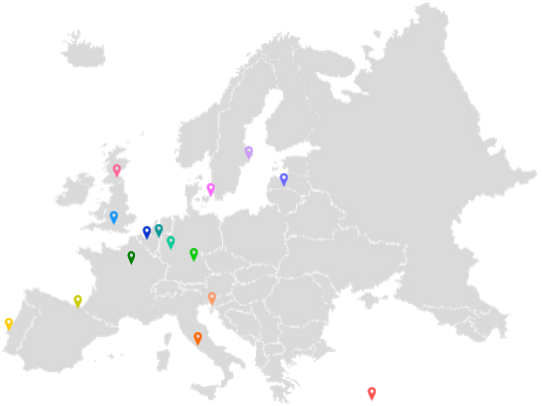

- JUPITER Research and Early Access Phase ongoing
  - > 100 Applications, 33 Lighthouses
- GCS AI Competition access started
- (Acceptance) Benchmarking running
  - TCO Application mix, High-Scaling
  - Storage Acceptance
- High-Performance Linpack
  - … otherwise it would be too easy
- **JUPITER Inauguration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

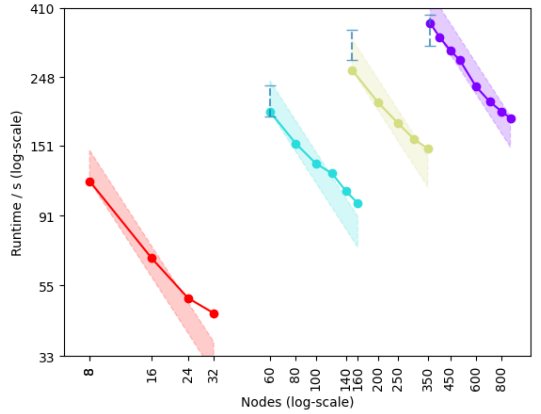

MISSION BRIEFING OVERVIEW

- En route to JUPITER: EuroHPC JU system hosted at JSC
- Launched with focus on applications
- ~6000 nodes,
  24 000 H100 GPUs, 1 728 000 Arm cores, 24 000 NDR200 endpoints
- Landing in Modular Data Center
- Preparing for descent:
  - JUREAP

→ jupiter.fz-juelich.de

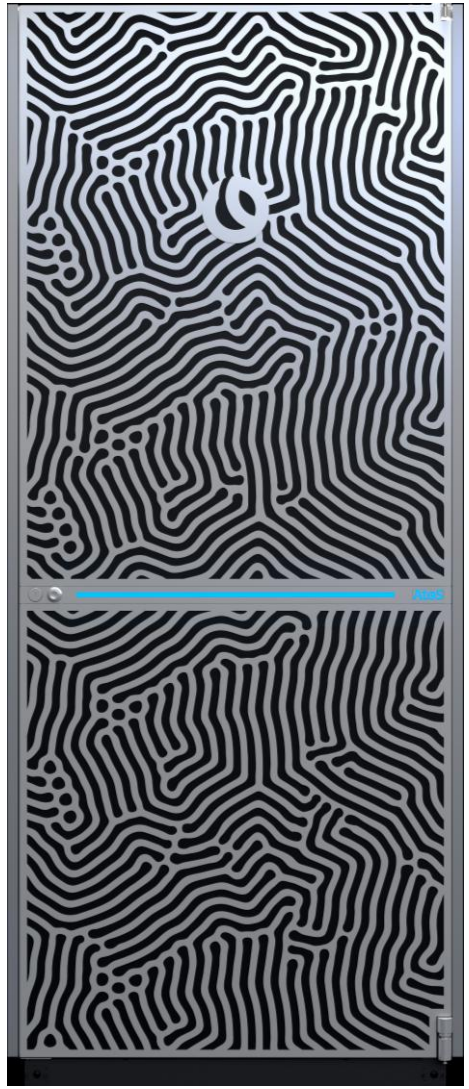

JUPITER

The Arrival of Exascale in Europe

fz-juelich.de/jupiter | #exa_jupiter

EuroHPC Joint Undertaking

Federal Ministry of Education and Research

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

GCS Gauss Centre for Supercomputing

JOINING FORCES

Ministerium für
Kultur und Wissenschaft
des Landes Nordrhein-Westfalen

fz-juelich.de/jupiter