JUPITER - THE ARRIVAL OF EXASCALE IN EUROPE

SC23 - Denver

2023-11 I JÜLICH SUPERCOMPUTING CENTRE

Member of the Helmholtz Association

EuroHPC Joint Undertaking

Bundesministerium für Bildung und Forschung

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

JÜLICH Forschungszentrum

Shaping Change

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 superco
CPUs, and the initial co
supercomputers in the v
on ARM.

That is a big disappointr
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 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

JUPITER – TIMELINE

Bundesministerium für Bildung und Forschung

EuroHPC Joint Undertaking

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

- 17.12.2021: Call for Expression of Interest (EoI) for Hosting Entity
- 14.02.2022: Deadline EoI Submission
- Q2 2022: Hearings & Hosting site decision and announcement
- Q1-Q3 2023: Procurement (Competitive Dialogue)
- **03.10.2023: Contract Signature**
- **Q3/Q4 2023: Installation Planning**
- Q1/2024: Start installation of JUPITER
- End of 2024: Put in operation JUPITER

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Digital Europe programme, as well as by Germany through the BMBF and the MKW.

JUPITER HIGH LEVEL ARCHITECTURE

93 ExaFLOPS of AI | 1.0 EF Delivered HPC | 24,000 GH200
Quantum-2 InfiniBand | 1.2PB/s Aggregate Bandwidth | 18.2 MW

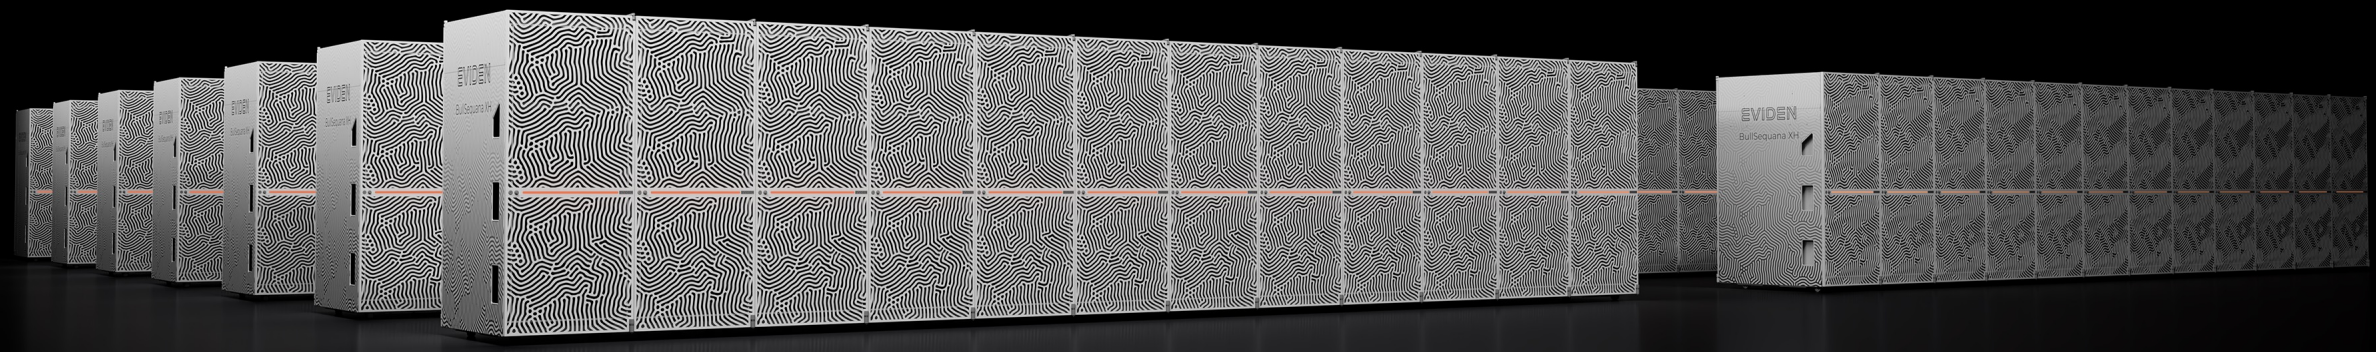

EuroHPC Joint Undertaking
EVIDEN an atos business
JÜLICH Forschungszentrum
NVIDIA
ParTec MODULAR SUPERCOMPUTING

JUPITER – HIGH-LEVEL ARCHITECTURE

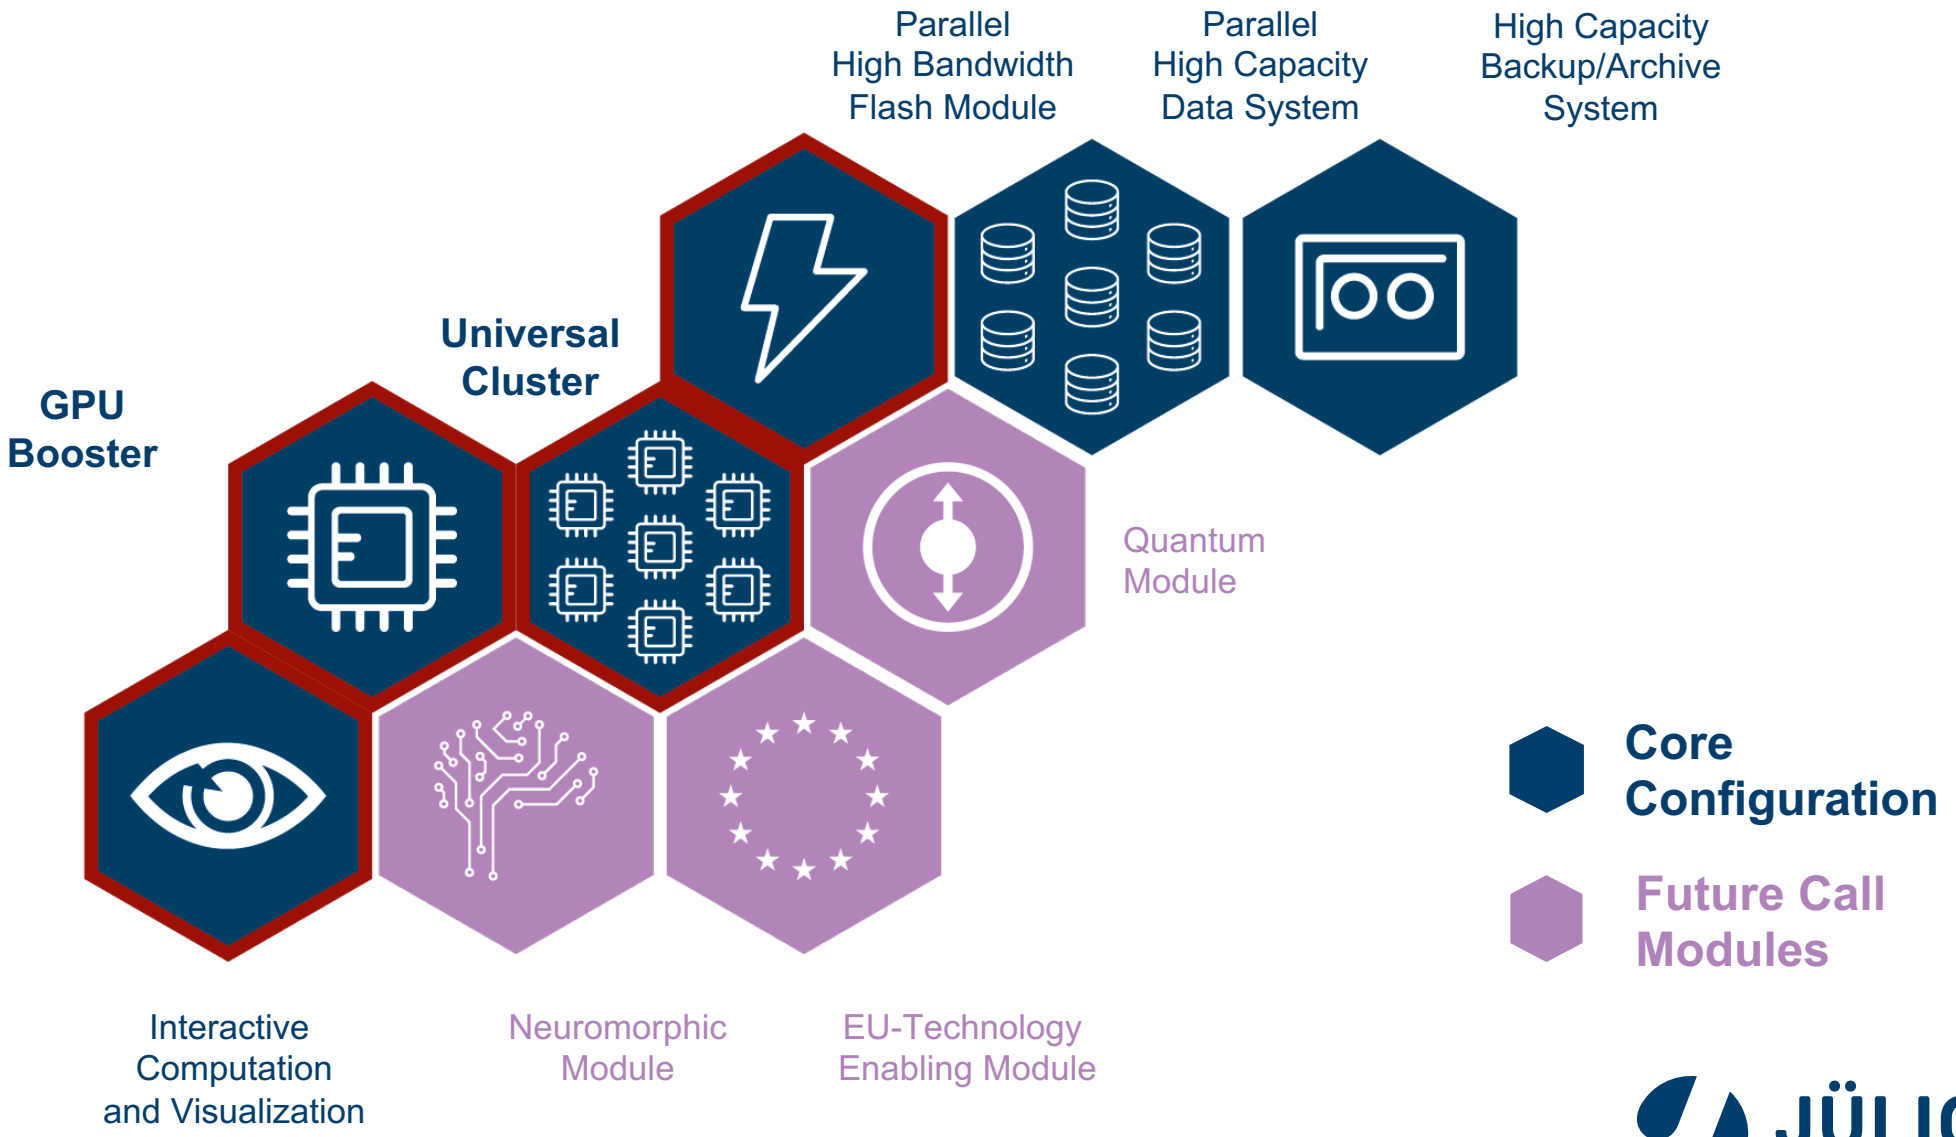

JUPITER – THE BOOSTER

Highly-Scalable Module for HPC and AI workloads

EVIDEN an atos business

NVIDIA

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

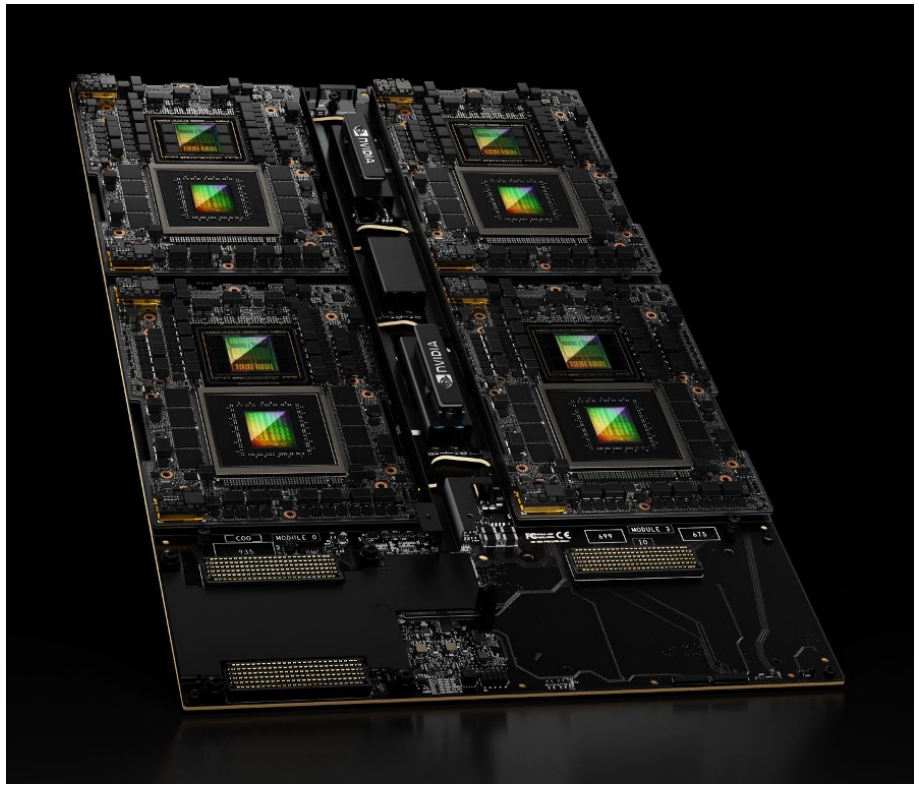

JUPITER – BOOSTER COMPUTE NODE ARCHITECTURE

Highly-Scalable Nodes for HPC and AI workloads

Node Specs

- 4× NVIDIA Grace-Hopper in SXM5 Board
- 4× NVIDIA InfiniBand NDR200
- 480 GB LPDDR5X / 384 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)

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
- 96 GB HBM3
  - ≥3600 GB/s
  - ~450 ns latency

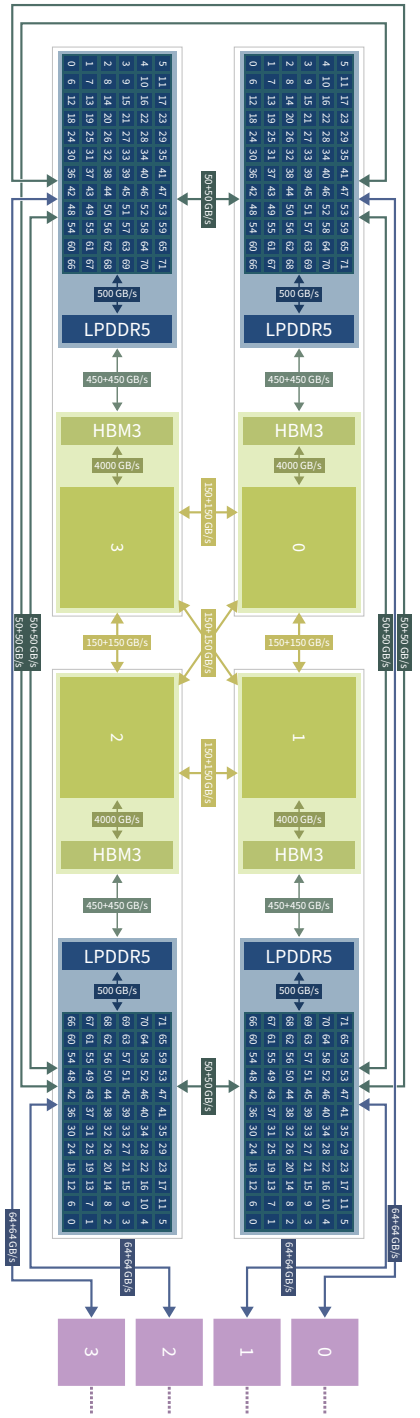

JUPITER – THE CLUSTER

General-Purpose Module for Mixed Workloads

- 5 >PetaFLOP/s (FP64, HPL)
- SiPearl Rhea1
  - 2 CPUs per node
  - >1300 nodes
- NVIDIA Mellanox NDR
  - 1 NDR200 NICs per compute node
- BullSequana XH3000
  - Direct Liquid Cooled blades
  - 3 compute nodes per blade

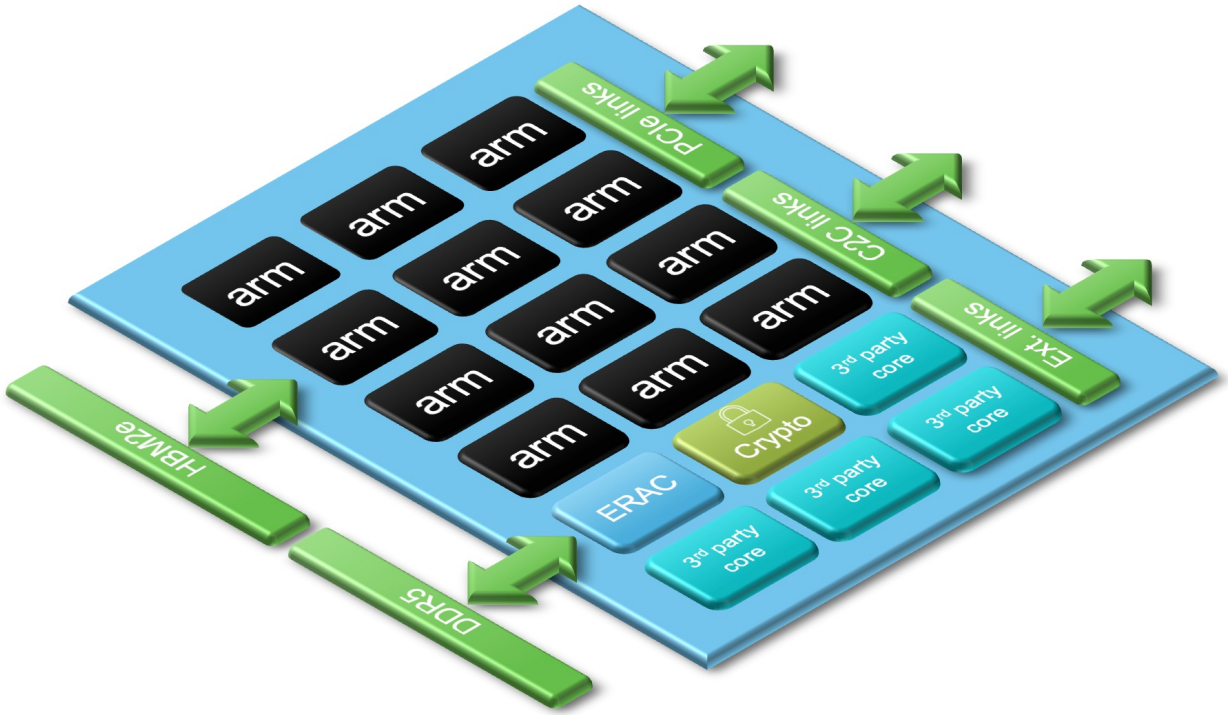

JUPITER – CLUSTER COMPUTE NODE ARCHITECTURE

General-Purpose Nodes for Mixed Workloads

Node Specs

- 2× SiPearl Rhea1
- 1× NVIDIA InfiniBand NDR200
- 512 GB DDR5 / 1024 GB

CPU Specs

- ARM Neoverse Zeus
  - 2 x 256 SVE per core
- >2.5 GHz
- HBM2e with 1.64 TB/s
- 256 GB DDR5
- PCIe Gen5

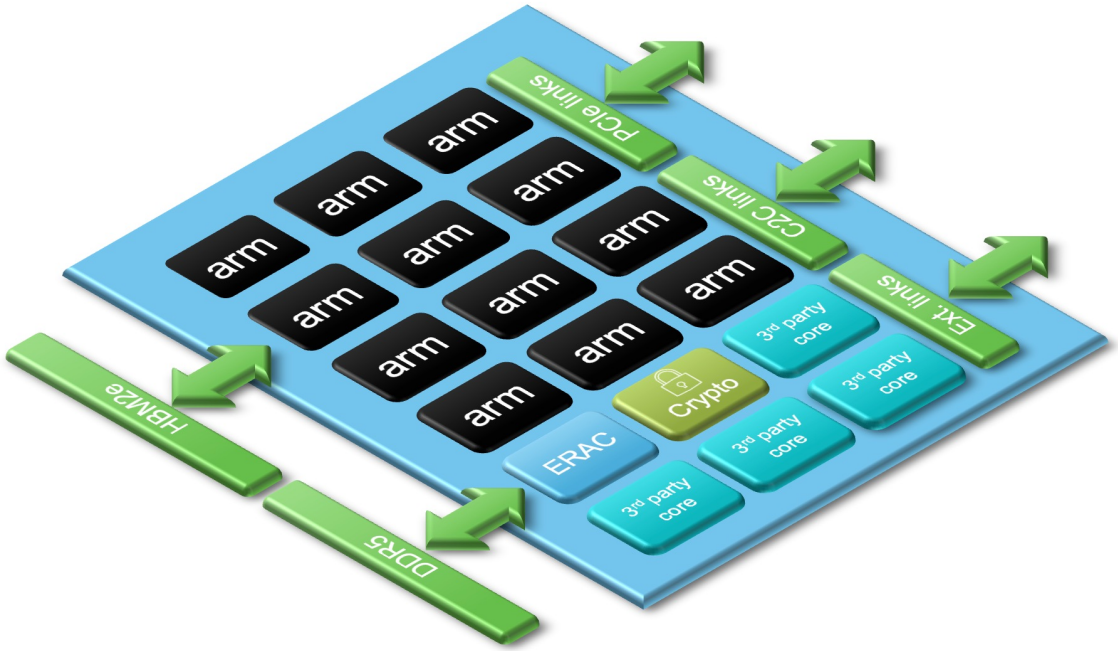

JUPITER – EXAFLASH

High-bandwidth low-latency SCRATCH

- Gross Capacity: 29 PB; Net Capacity: 21 PB
- Bandwidth: 2.1 TB/s Write, 3.1 TB/s Read
- 40× IBM ESS3500 Building Blocks (80 servers)
  - 4× NDR200 per server
  - 24× 30 TB NVMe drives per BB
  - IBM Storage Scale (aka Spectrum Scale/GPFS)
- Manager and Datamover Nodes
- Exclusive for JUPITER
  - Integrated into InfiniBand fabric

JUPITER – EXASTORE

High-bandwidth high-capacity HOME/PROJECT/DATA

- Target capacities
  - ~300 PB (2024, gross)
  - ~150 PB (2026 upgrade, gross)
- Exclusive for JUPITER
  - Integrated into the InfiniBand fabric
- Datamover Nodes
- Backup/Archive connected to existing facility
- Strong connection to JUST6 (central JSC storage)

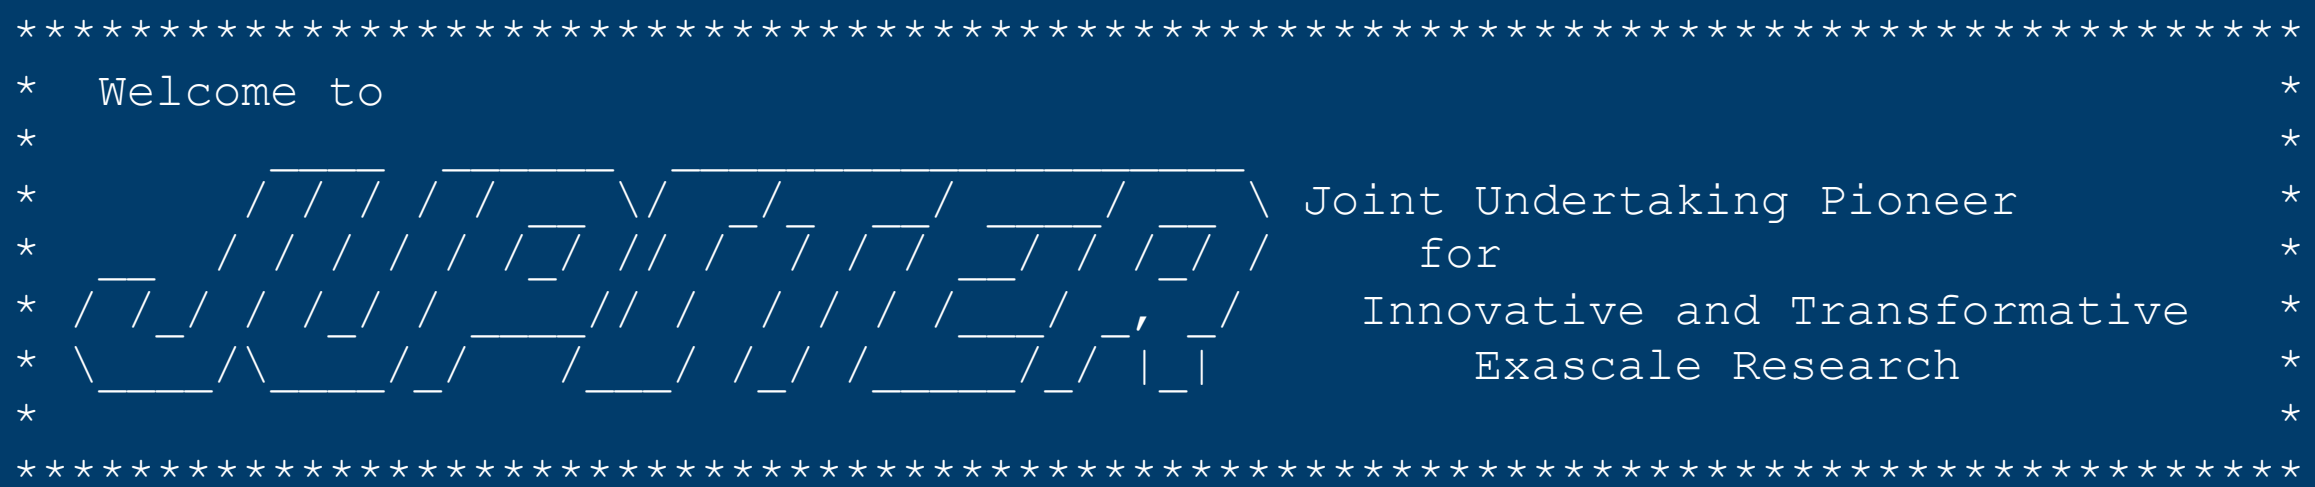

SYSTEM MANAGEMENT

JUPITER MANAGEMENT SUITE

”Power is nothing without control"

ParaStation MODULO

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

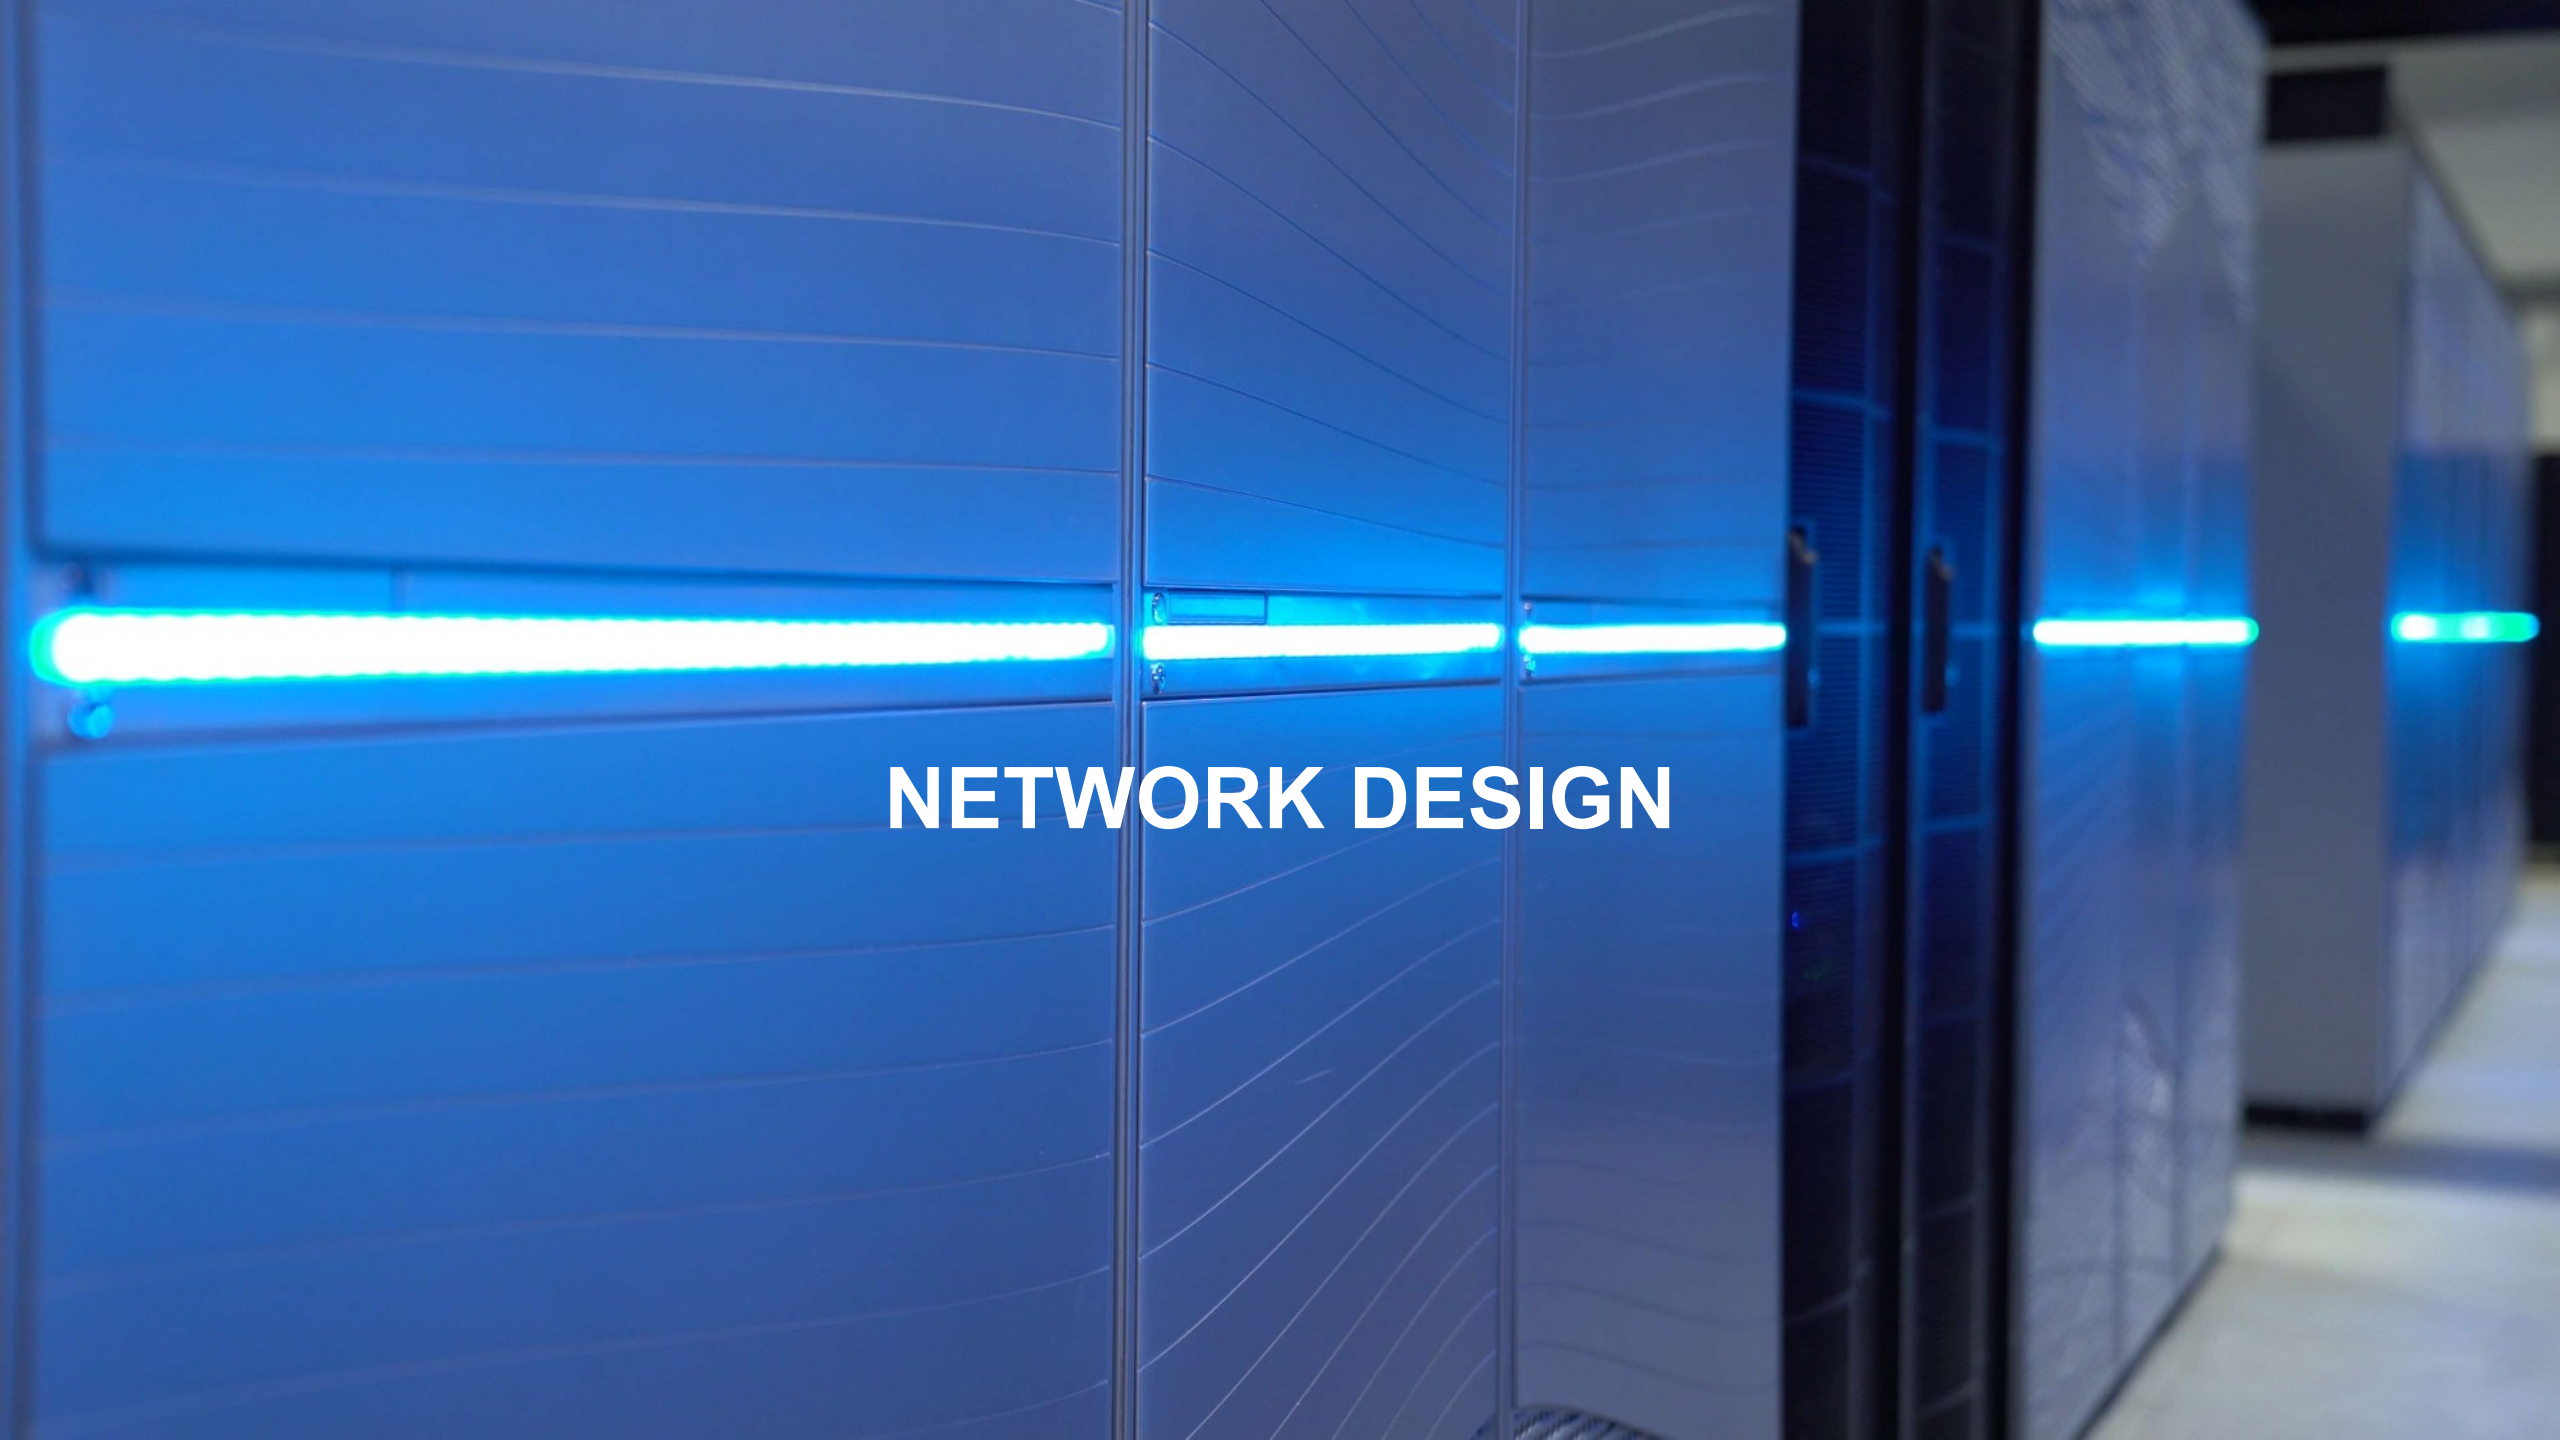
NETWORK DESIGN

JUPITER – INTERCONNECT

One Network to Rule Them All

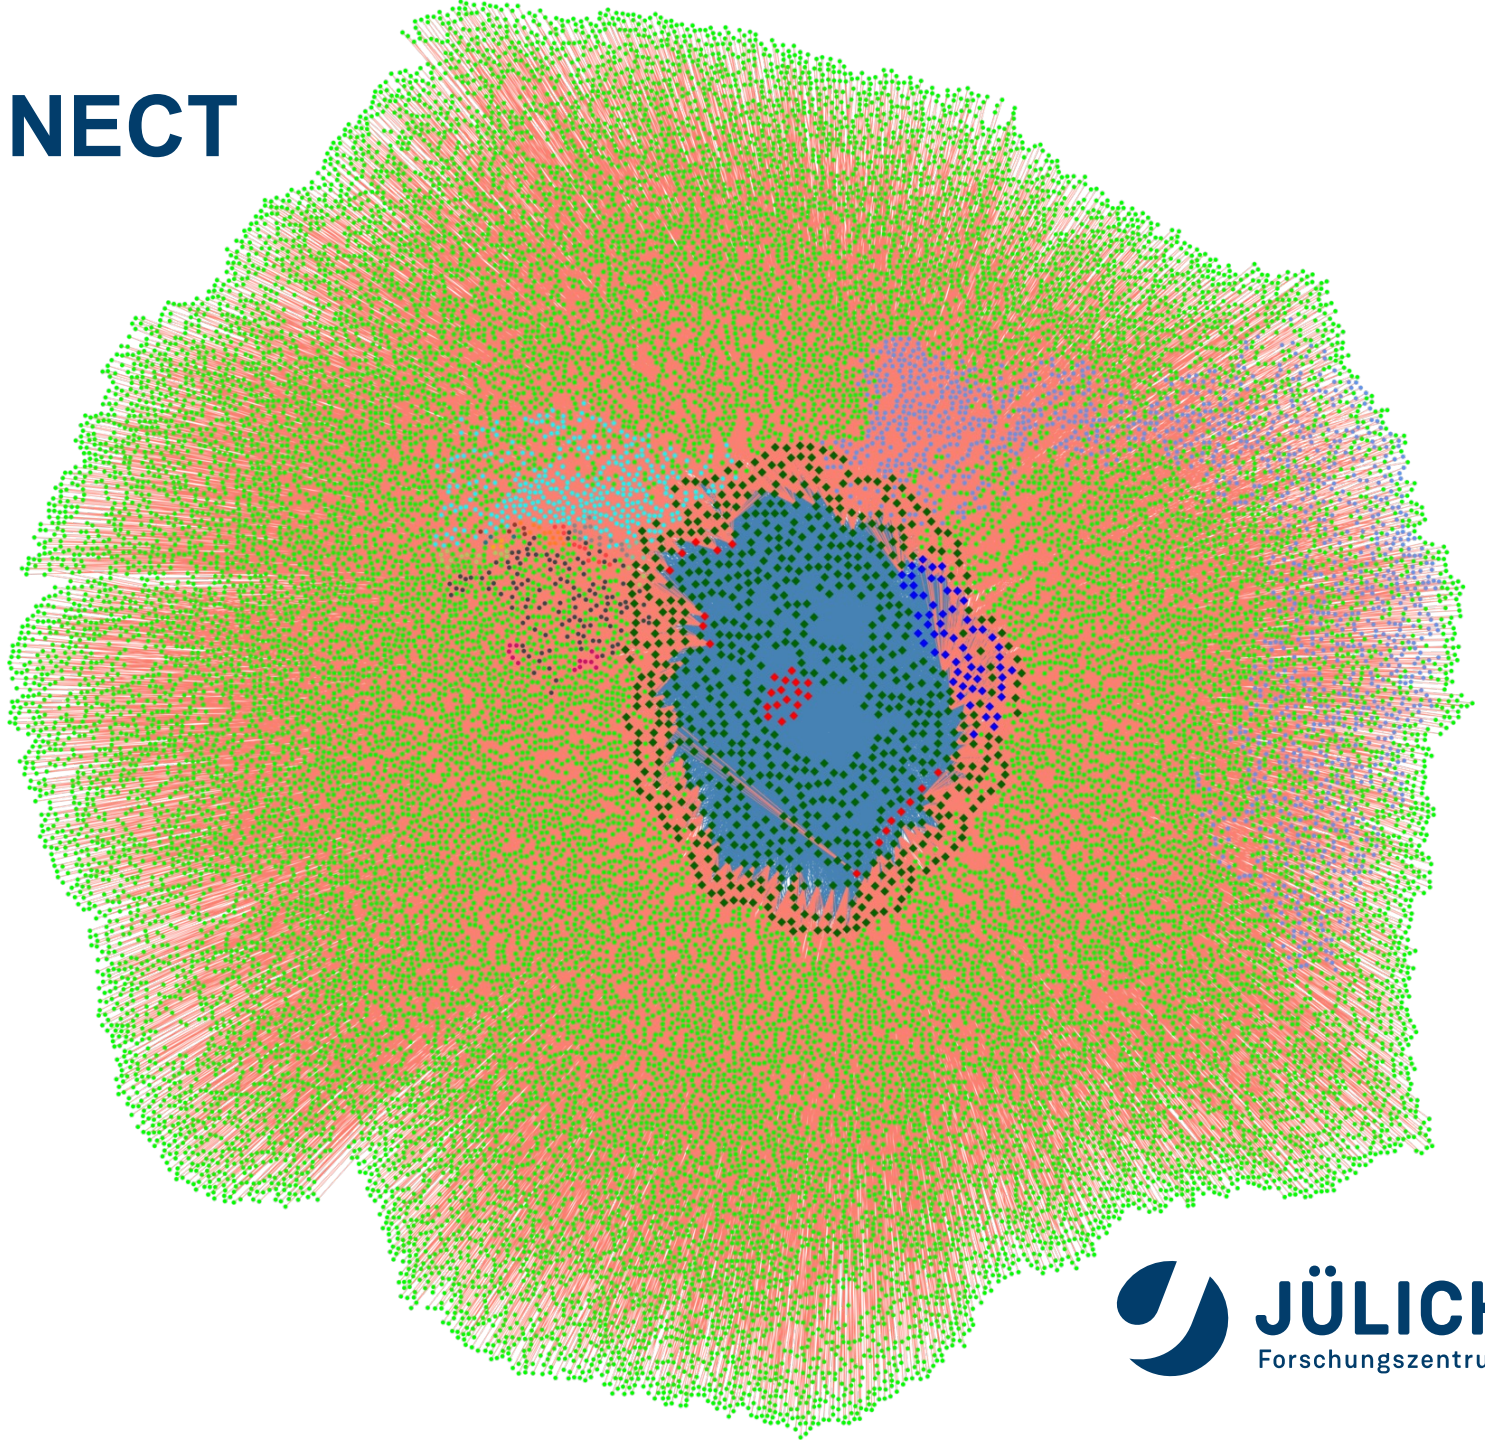

- Cluster module switch
- Booster module switch
- Booster node
- Flash storage node
- Misc node
- NDR link
- Admin module switch
- Cluster node
- Gateway node
- Storage node
- NDR200 link

JUPITER – INTERCONNECT

One Network to Rule Them All

EVIDEN an atos business

NVIDIA

- 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

JUPITER – INTERCONNECT

One Network to Rule Them All

EVIDEN
an atos business

NVIDIA

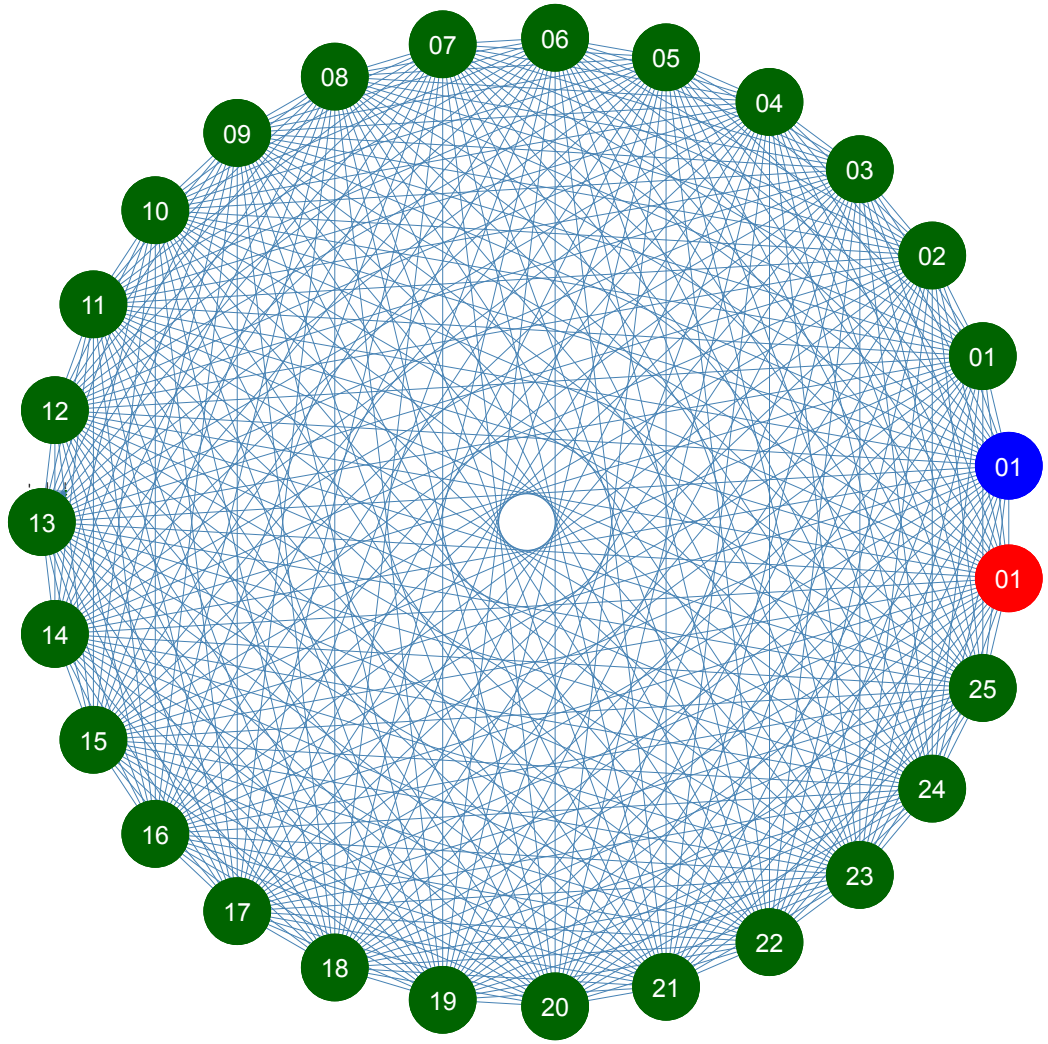

JÜLICH
Forschungszentrum

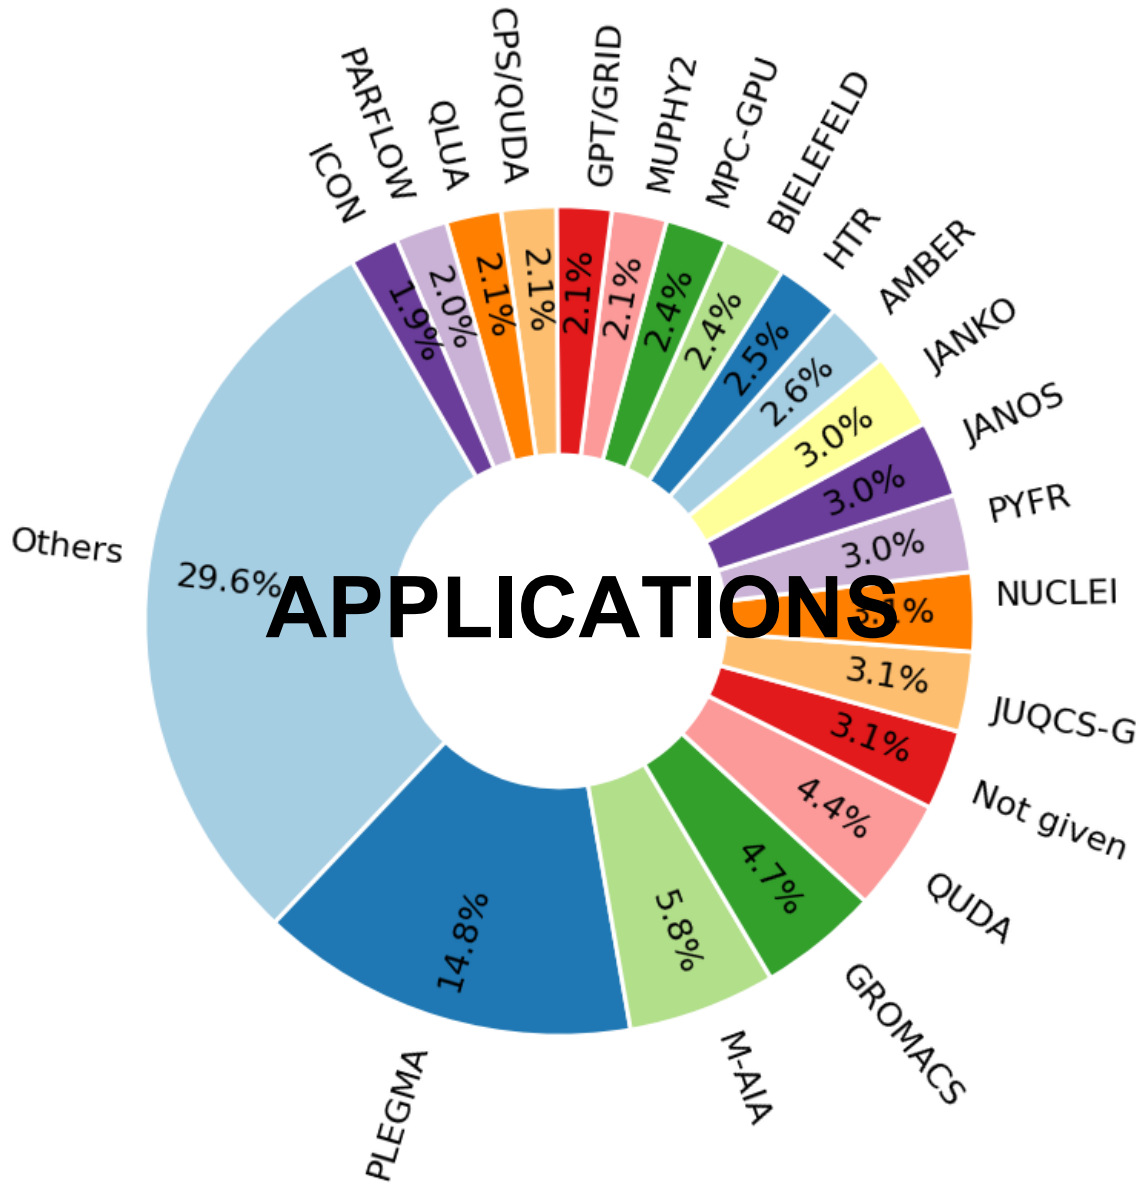
APPLICATIONS
Others
29.6%
PLEGMA
14.8%
M-AIA
5.8%
GROMACS
4.7%
QUDA
4.4%
Not given
3.1%
JUQCS-G
3.1%
NUCLEI
3.0%
PYFR
3.0%
JANOS
3.0%
JANKO
3.0%
AMBER
2.6%
HTR
2.5%
BIELEFELD
2.4%
MPC-GPU
2.4%
MUPHY2
2.1%
GPT/GRID
2.1%
CPS/QUDA
2.1%
QLUA
2.0%
PARFLOW
1.9%
ICON

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)*

JUREAP

Seeding Exascale in Europe!

More information and the first call will be announced soon!

Contact: jureap@fz-juelich.de

JUPITER Research and Early Access Program

JÜLICH Forschungszentrum

JUWELS VS JUPITER: KEY FACTS

JUWELS VS. JUPITER

	JUWELS	JUPITER
Cluster	**CPU:** Intel Xeon Platinum 8168 **GPU:** NVIDIA V100 **Peak:** 10 PFlop/s	**CPU:** SiPearl Rhea1 **GPU:** none **Mem. Bandwidth:** 0,51 Byte/Flop
Booster	**CPU:** 2* AMD Epyc Rome **GPU:** 4× NVIDIA A100 GPUs **Peak:** 73 PFlop/s	**CPU:** 4* NVIDIA Grace **GPU:** 4* NVIDIA Hopper **Peak:** **>1 EFlop/s**
Network topology	Fat tree and DragonFly+	DragonFly+
System access	GCS or PRACE proposals	GCS and EuroHPC JU proposals
User support	HLST, SDL, ATML, training courses, targeted early access program	**same**

SUMMARY

The Past

- DEEP: Blueprint for the MSA, research project
- JURECA: First production MSA System (2017)
  - Network bridging InfiniBand/OmniPath
  - KNL manycore Booster

The Present

- JUWELS: MSA on its way to Exascale
  - Massive HPC + AI capabilities
- JUNIQ: Quantum Computing Services
  - QC on its own and as Modules in the MSA

The Future

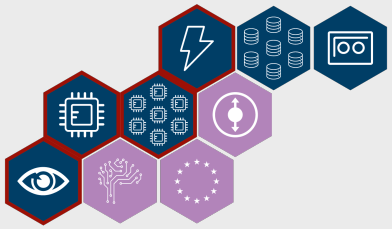

- JUPITER: MSA at Exascale
- DEEP-SEA: Developments for the next level
  - Make resource allocation (shrink, extend, distribute) more flexible to provide malleability
- HPC, AI and QC integration

JUPITER

The Arrival of Exascale in Europe

fz-juelich.de/jupiter | #exa_jupiter

Funding Agencies:

EuroHPC Joint Undertaking

Federal Ministry of Education and Research

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