



# FZJ OVERVIEW AND PLANS

## JUPITER - THE ARRIVAL OF EXASCALE IN EUROPE

Florian Berberich

31-01-2025 | CERN



Member of the Helmholtz Association



**EuroHPC**  
Joint Undertaking



Bundesministerium  
für Bildung  
und Forschung

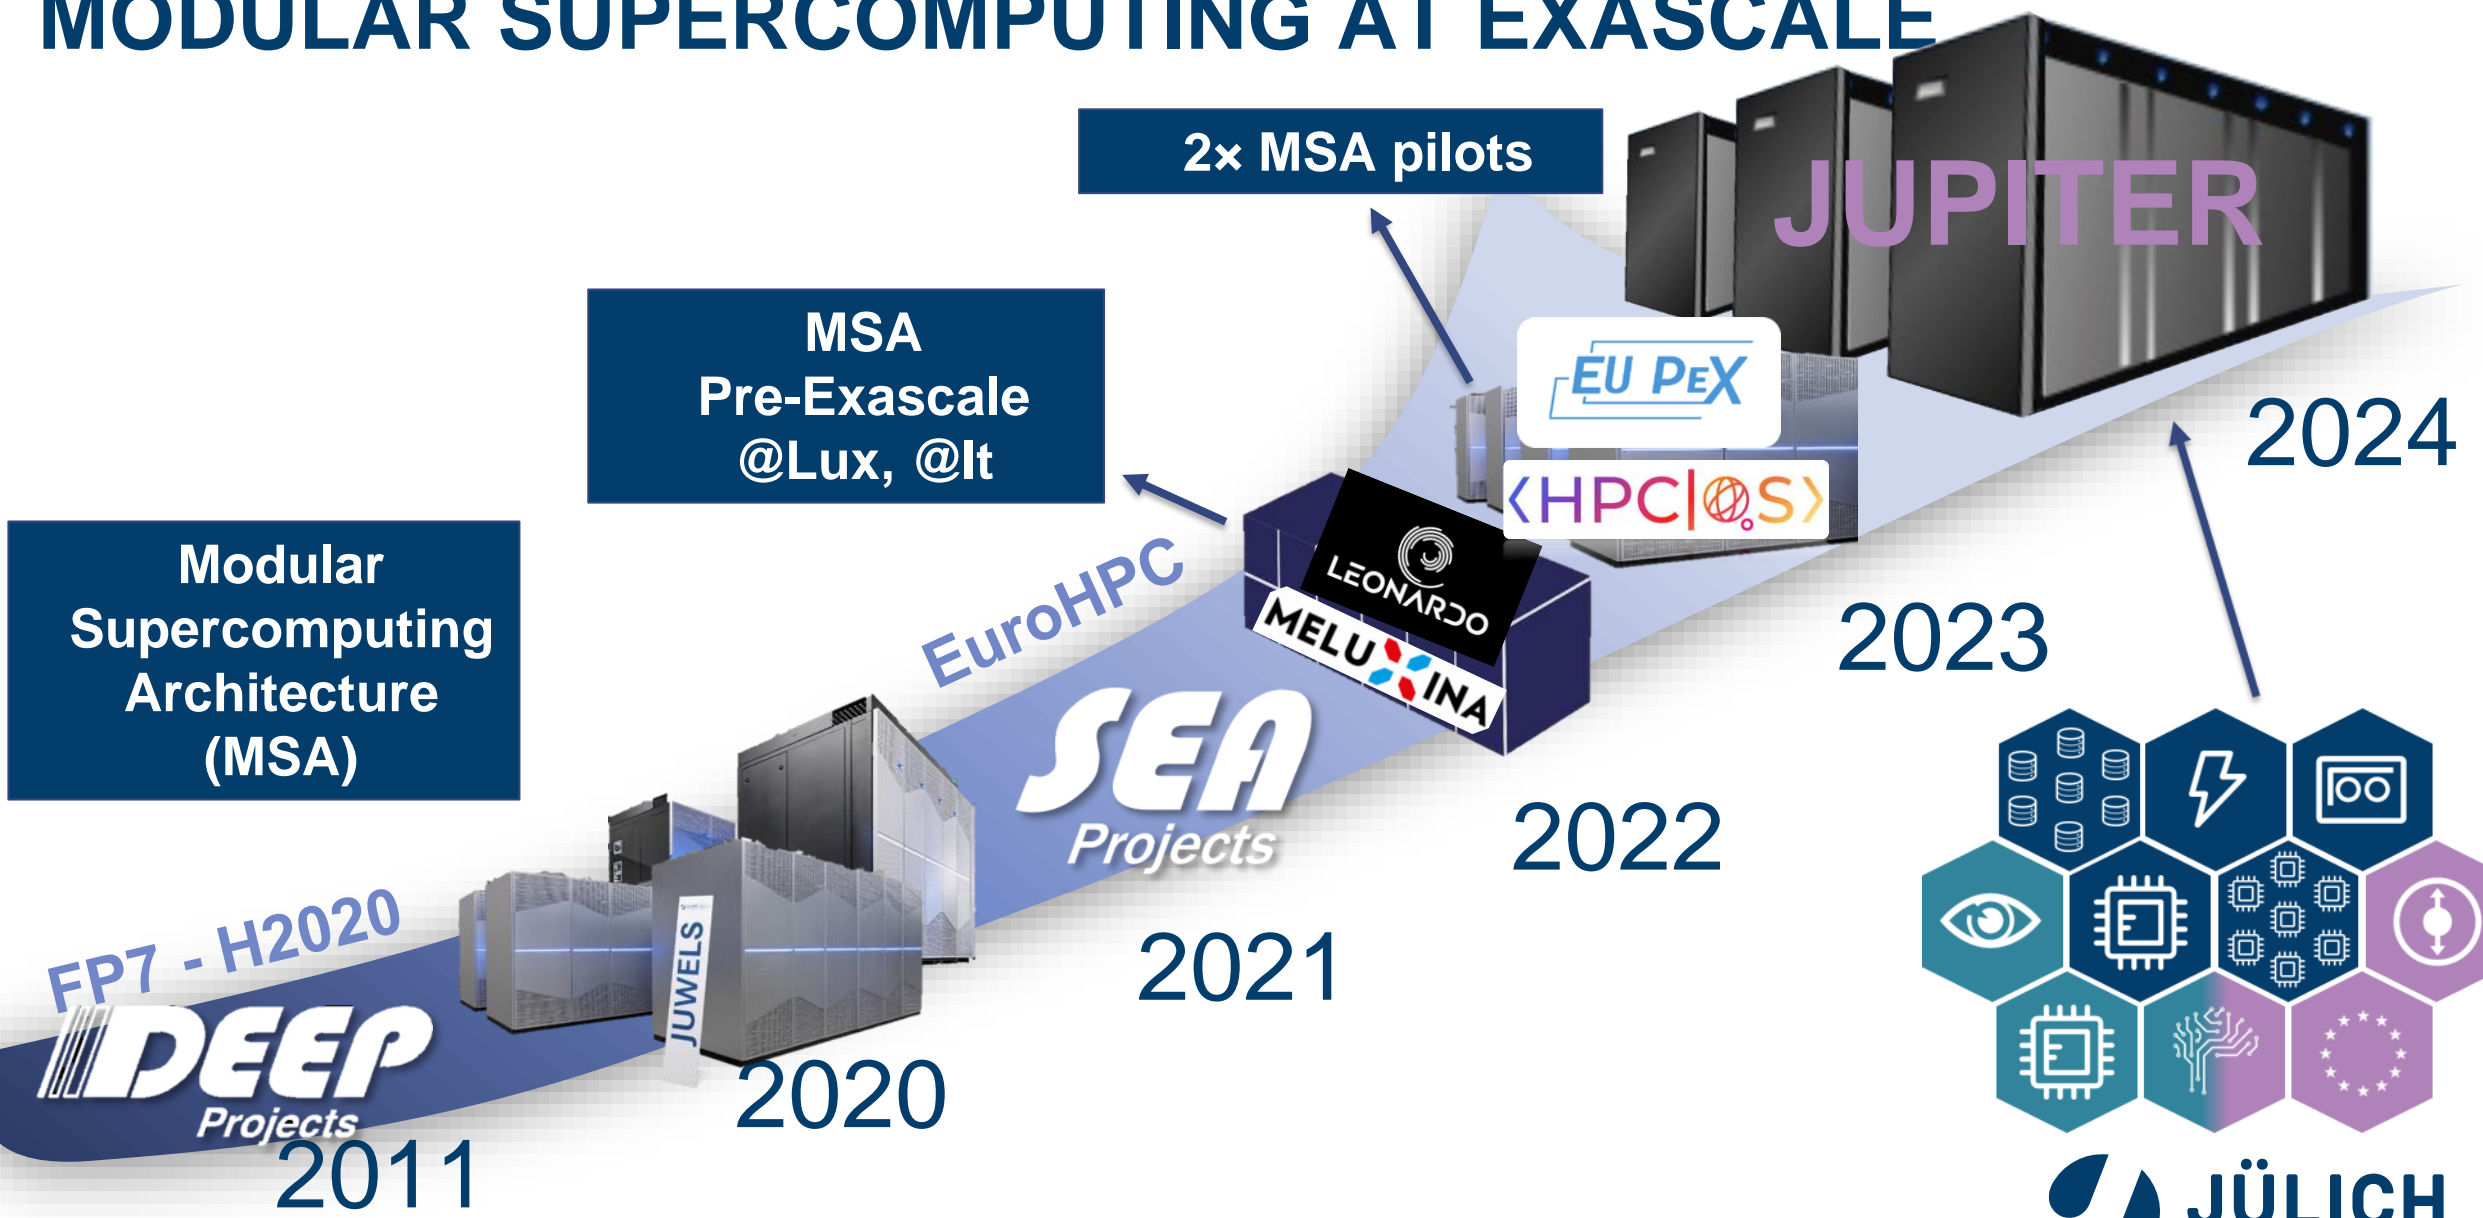
Ministerium für  
Kultur und Wissenschaft  
des Landes Nordrhein-Westfalen



**GCS**  
Gauss Centre for Supercomputing

**JÜLICH**  
Forschungszentrum  
*Shaping Change*

# MODULAR SUPERCOMPUTING AT EXASCALE

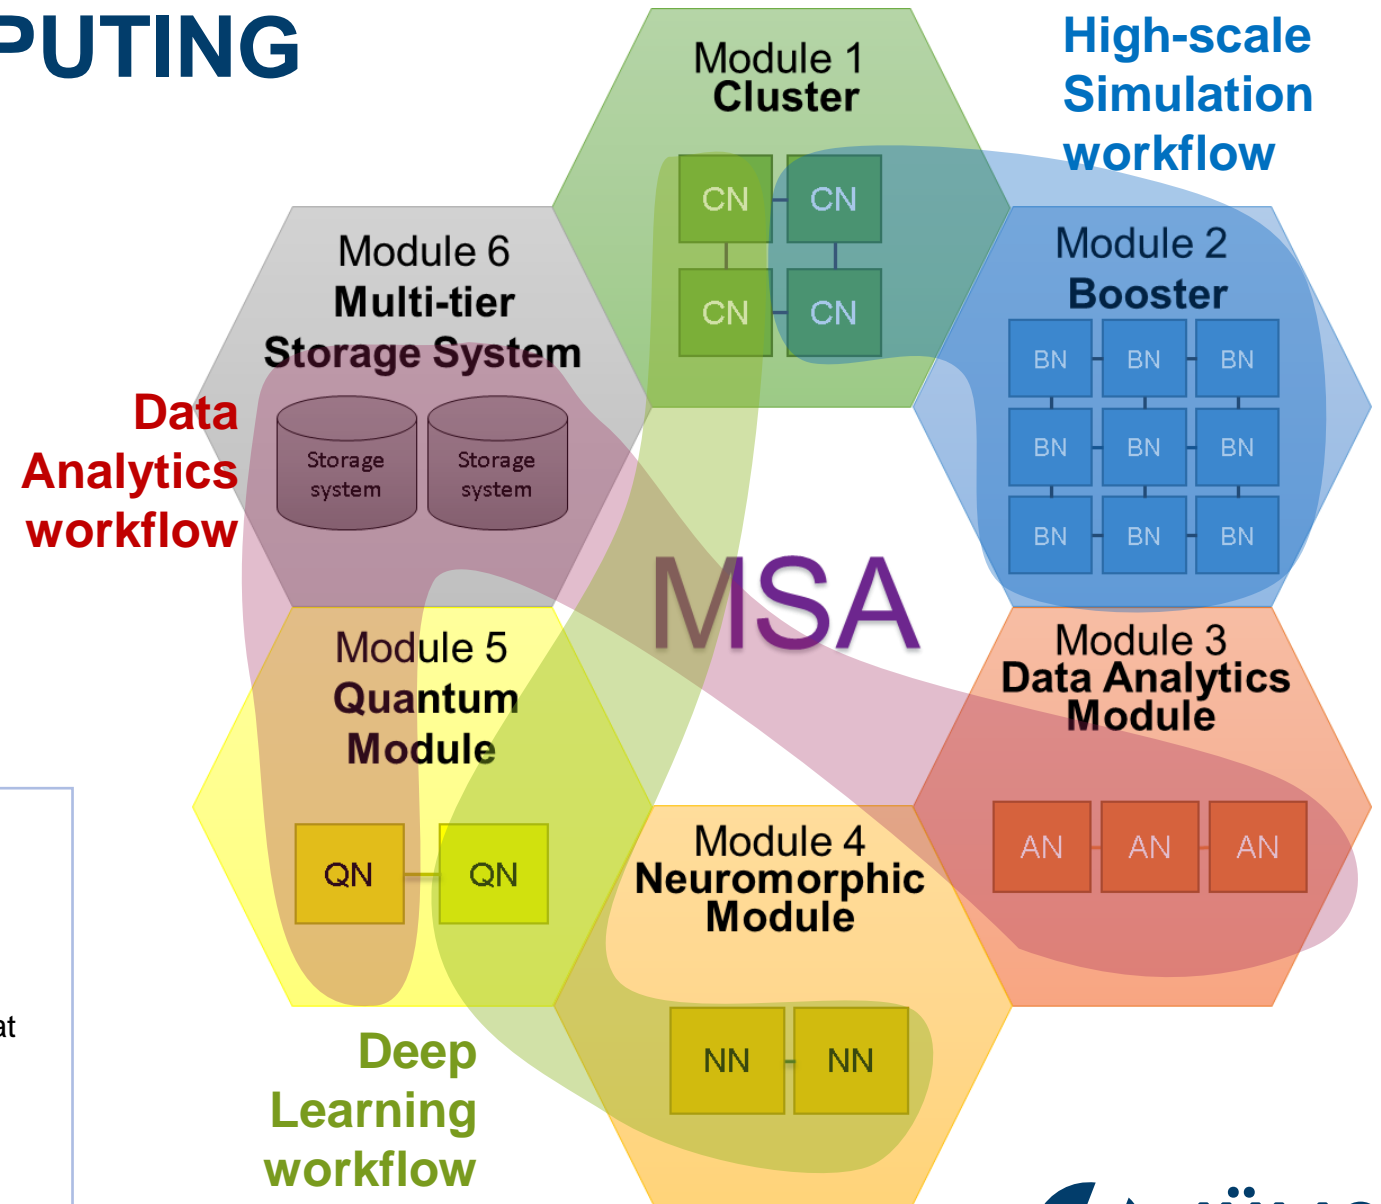


# MODULAR SUPERCOMPUTING ARCHITECTURE

## Composability of heterogeneous resources

- Cost-effective scaling
- Effective resource-sharing
- Match application diversity

- **E. Suarez**, N. Eicker, Th. Lippert, "Modular Supercomputing Architecture: from idea to production", Chapter 9 in Contemporary High Performance Computing: from Petascale toward Exascale, Volume 3, p 223-251, CRC Press. (2019)
- **E. Suarez**, N. Eicker, and Th. Lippert, "Supercomputer Evolution at JSC", Proceedings of the 2018 NIC Symposium, Vol.49, p.1-12, (2018)



# DISCOVERING JUPITER

- ParTec/Eviden Supercomputer Consortium
- Implementing Modular Supercomputing Architecture
- JUPITER Booster: High scalability; 1 EFLOP/s HPL, >70 EFLOP/s FP8
- JUPITER Cluster: High versatility; 0.5 B/FLOP balance
- Network: 200/400 Gigabit NVIDIA Mellanox InfiniBand NDR
- Storage: 29 PB Flash IBM Storage Scale 6000
- 17 Megawatt Linpack Power Consumption
- Direct Liquid Cooled (36 -> 4x degree) to enable heat-reuse




EVIDEN




# JUPITER – HIGH-LEVEL ARCHITECTURE



- JUPITER Management Stack
  - Modular System Architecture
  - ParTec ParaStation Modulo
  - Eviden SMC xScale
- Scheduler: Slurm
- Scientific Software: EasyBuild
- Jobreporting: LLview



- >1 ExaFLOP/s FP64 HPL
- >70 ExaFLOP/s FP8 AI
- ~6000 Nodes
- ~24000 NVIDIA Grace-Hopper
- NVIDIA Mellanox NDR
- BullSequana XH3000



- >5 PetaFLOP/s FP64 HPL
- ~1300 Nodes
- ~2600 SiPearl Rhea1
- BullSequana XH3000



- 210 PB useable SAS
- ~1,5 TB/s
- 22\* IBM SSS6000



- 23 Login Nodes
- NVIDIA Hopper/A40
- 200 Gbit/s connectivity



- 21 PB useable Flash
- ~2,5 TB/s
- 20\* IBM SSS6000

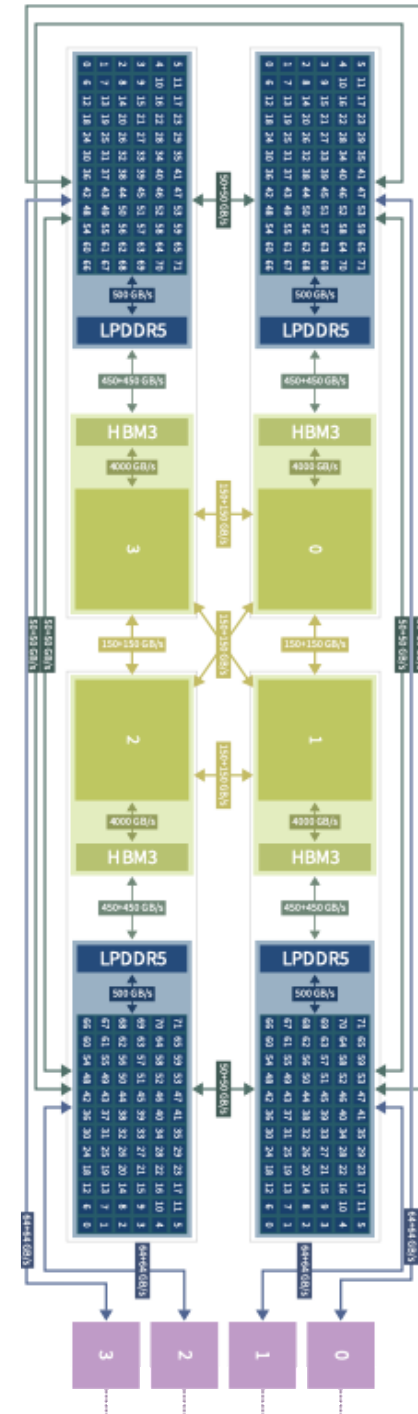
EU-Tech  
Enabling

Core Configuration

# JUPITER – THE BOOSTER

Highly-Scalable Module for HPC and AI workloads

- 1 ExaFLOP/s (FP64, HPL)
- NVIDIA Grace-Hopper CG1
  - ~5900 compute nodes
  - 4x CG1 chips per compute node
- NVIDIA Mellanox NDR
  - 4x NDR200 NICs per compute node
- BullSequana XH3000
  - Direct Liquid Cooled blades
  - 2x compute node per blade



# MODULAR DATA CENTER FOR JUPITER

EVIDEN



# MDC SHIPMENT START

10./11.9.2024





# JUREAP

## Seeding Exascale in Europe!




[jureap@fz-juelich.de](mailto:jureap@fz-juelich.de) • <https://www.fz-juelich.de/en/ias/jsc/jupiter/jureap>

JUPITER Research and Early Access Program

# JUPITER TEST RACK *JEDI*

- Actual JUPITER node design
  - 2x nodes per blade
  - 4x Grace-Hopper per node
- Usage
  - System management preparations
  - JUREAP
  - Application porting
- Also: GH200 COTS test nodes for some months





# BEYOND?

The future is now


AI

# CHALLENGES IN HPC "AFTER" EXASCALE

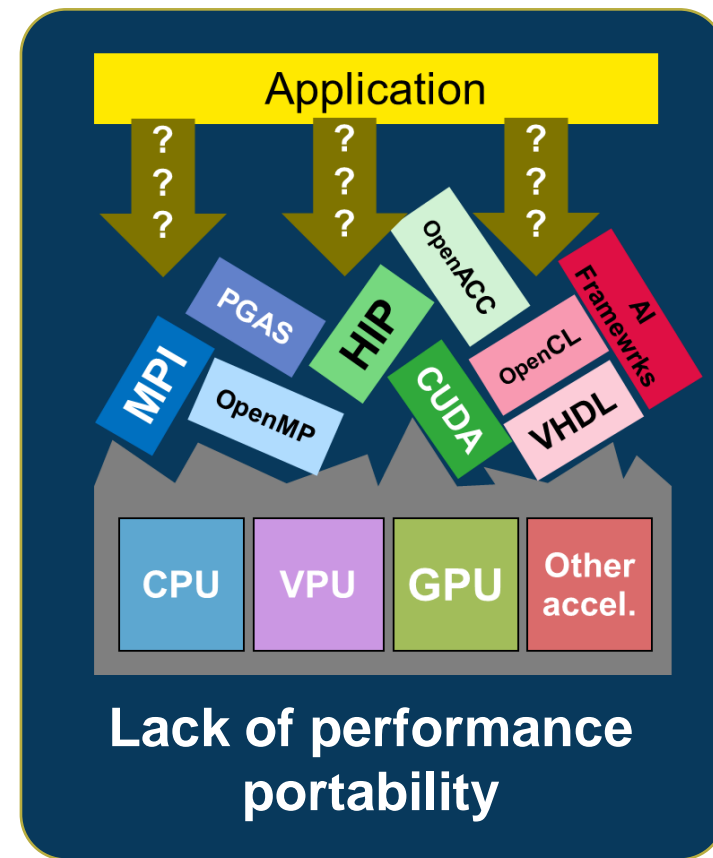
<https://woodsidepawprint.com/feature/2022/01/12/taiwan-vs-china-whats-at-stake/>



Technology dependency




HW heterogeneity



Lack of performance portability

Logos owned by respective companies

Graphic owned by Disney/Imgbin



HPC

HPC vs. Hyperscalars

<https://globalenergymonitor.org/projects/global-coal-plant-tracker/>



Justify use of energy and natural resources