



# JUPITER - THE PATH TO EXASCALE IN EUROPE


JU Pioneer for Innovative and Transformative Exascale Research

JUNE 2023 | BENEDIKT VON ST. VIETH

# HOW IT ALL STARTED

## DEEP Cluster-Booster Architecture

- The first ever CBA realized in the DEEP projects (Dynamical Exascale Entry Platform)
- Starting in 2011 – now running in its 4th generation (DEEP-SEA)



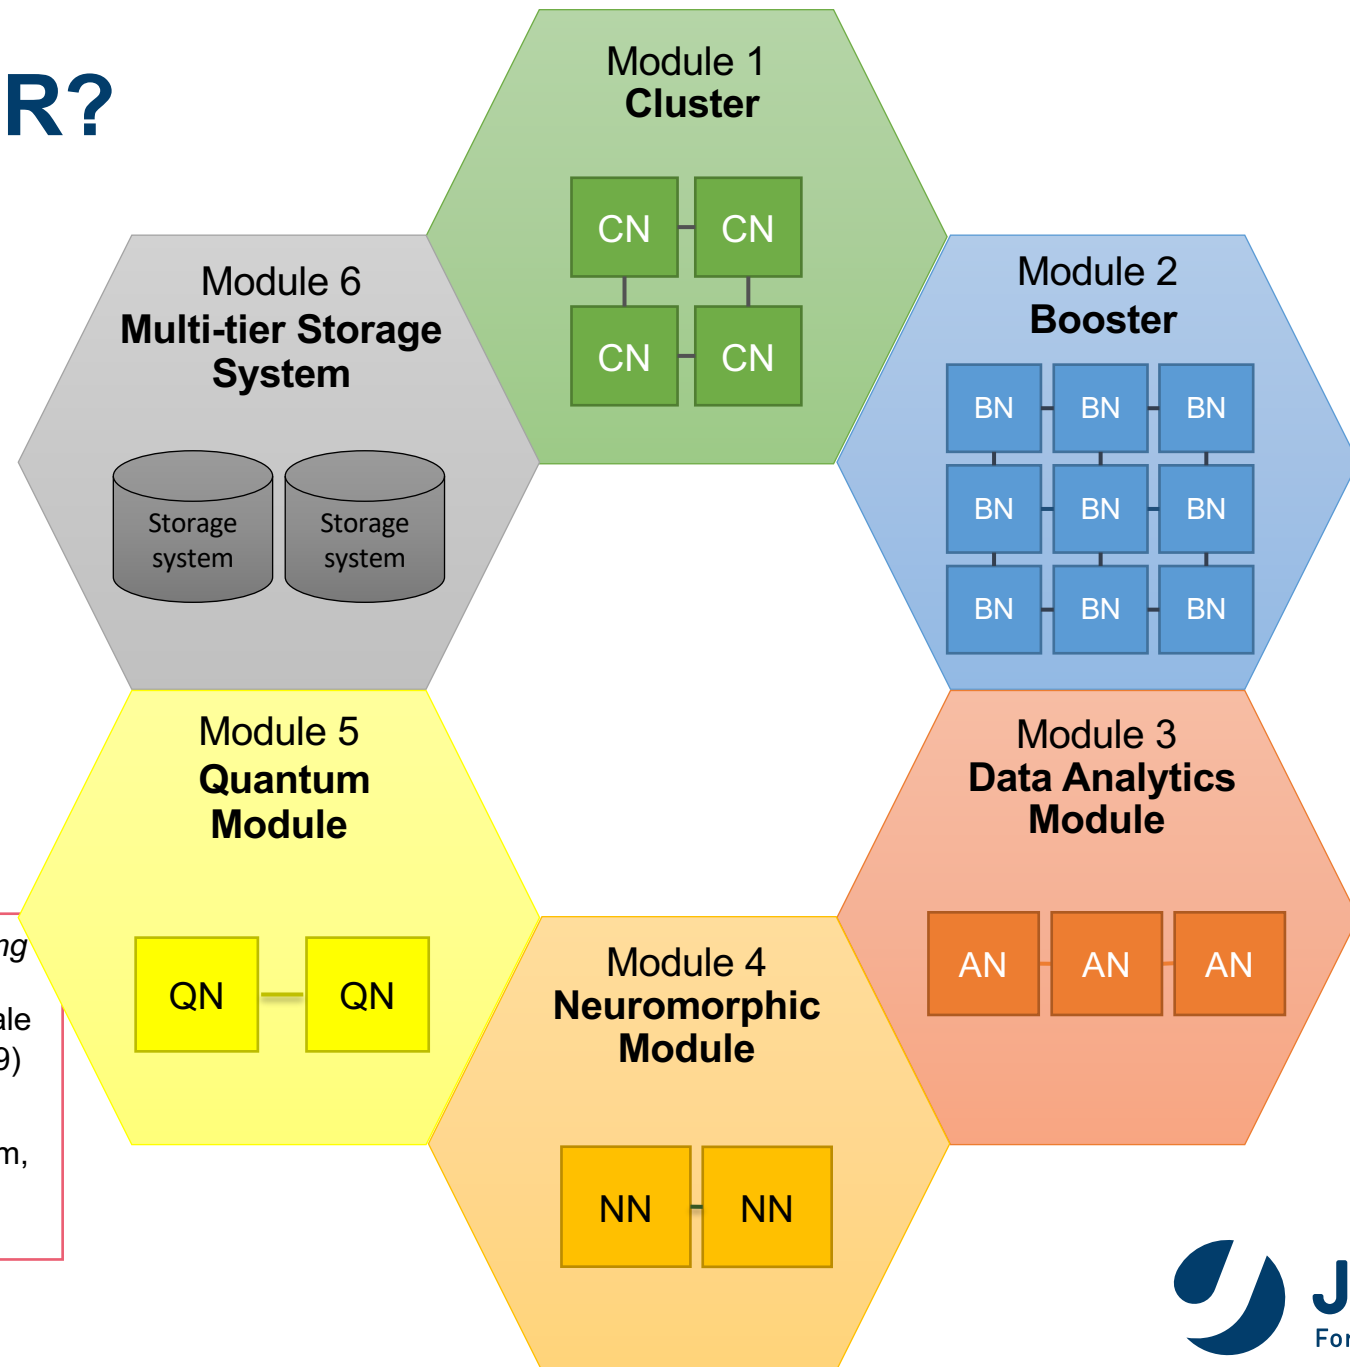
# CLUSTER? BOOSTER?

## Modular Supercomputing!

### Composability of heterogeneous resources

- Cost-efficient scaling
- Effective resource-sharing

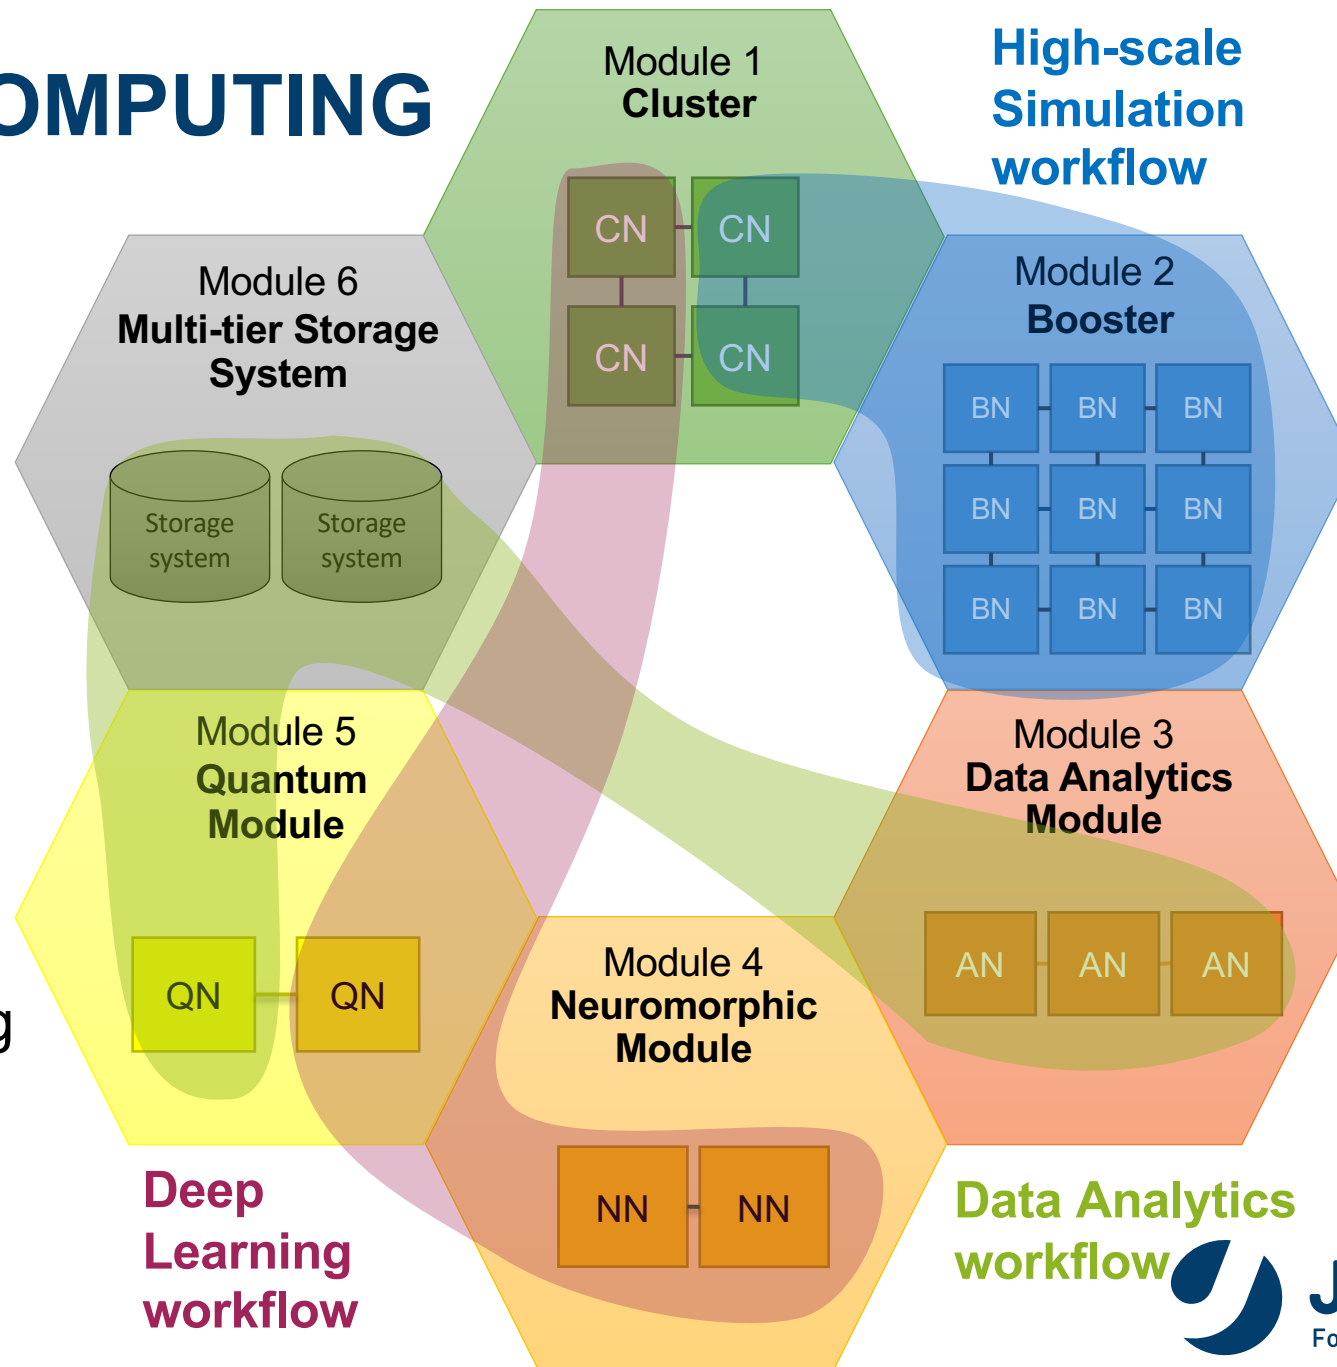
- **E. Suarez\***, N. Eicker, Th. Lippert, "Modular Supercomputing Architecture: from idea to production", Chapter 9 in Contemporary High Performance Computing: from Petascale toward Exascale, Volume 3, pp 223-251, CRC Press. (2019)
- **E. Suarez\***, N. Eicker, and Th. Lippert, "Supercomputer Evolution at JSC", Proceedings of the 2018 NIC Symposium, Vol.49, p.1-12, (2018) [online: <http://juser.fz-juelich.de/record/844072>].



# MODULAR SUPERCOMPUTING

## Composability of heterogeneous resources

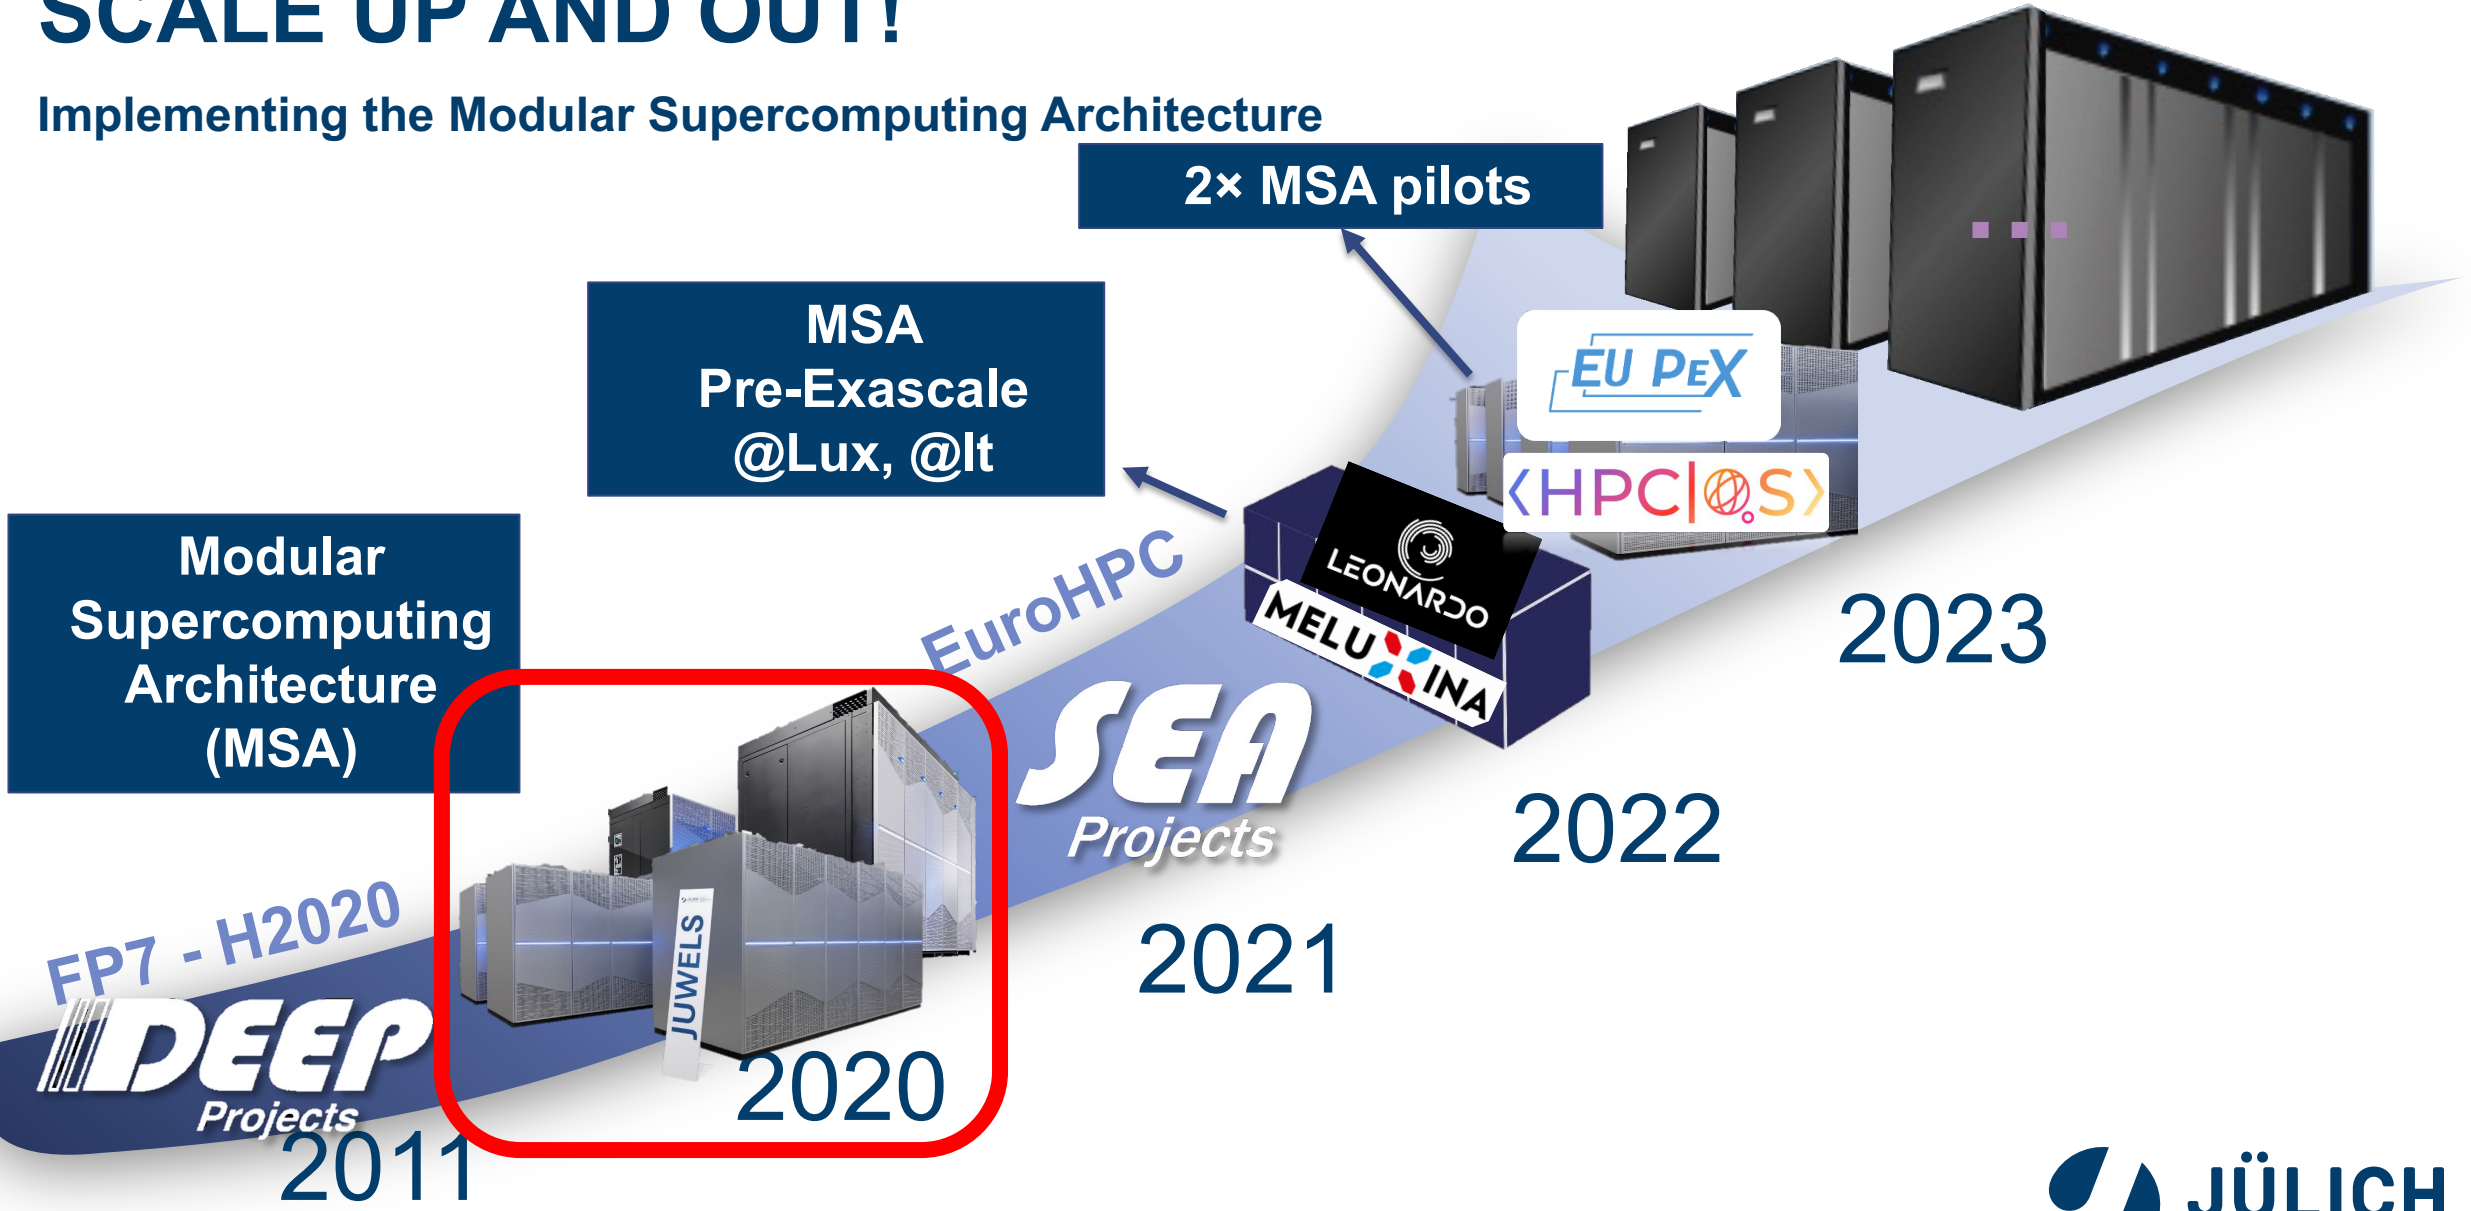
- Cost-efficient scaling
- Effective resource-sharing
- Fit application diversity
  - Large-scale Simulations
  - Machine- and Deep Learning
  - Artificial Intelligence
  - Quantum Computing





# SCALE UP AND OUT!

Implementing the Modular Supercomputing Architecture





# TOP 500 CERTIFICATE

The List.  
JUWELS Booster Module - Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100,  
Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite  
Forschungszentrum Juelich (FZJ), Germany

No. 7

is ranked  
among the World's TOP500 Supercomputers  
with 44.72 PFlops Linpack Performance  
in the 56th TOP500 Conference published at the SC2020 Virtual Conference on November 16, 2020.

Congratulations

Jack Dongarra  
University of Tennessee

Erich Strohmaier  
NERSC/Berkeley Lab

SC2020 Virtual  
TOP500 Editors

Horst Simon  
NERSC/Berkeley Lab

Martin Meuer  
Prometeus

## JUWELS Booster

#1 in TOP500 Europe (11/2020), #7 WW

#1 in Green500 among the top 100 in HPL

#5 HPCG500

#4 HPL-AI

The GREEN 500

# CERTIFICATE

JUWELS Booster Module - Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100,  
Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite  
Forschungszentrum Juelich (FZJ), Germany

No. 3

is ranked  
among the World's TOP500 Supercomputers  
with 25.1 PFlops Linpack Performance  
in the Green500 Conference published at the SC2020 Virtual Conference on November 16, 2020.

Congratulations from the Green500 Editors

Wu-chun Feng  
Virginia Tech

Kirk Cameron  
Virginia Tech

Wu-chun Feng  
Virginia Tech



# JUWELS – THE EXASCALE PATHFINDER

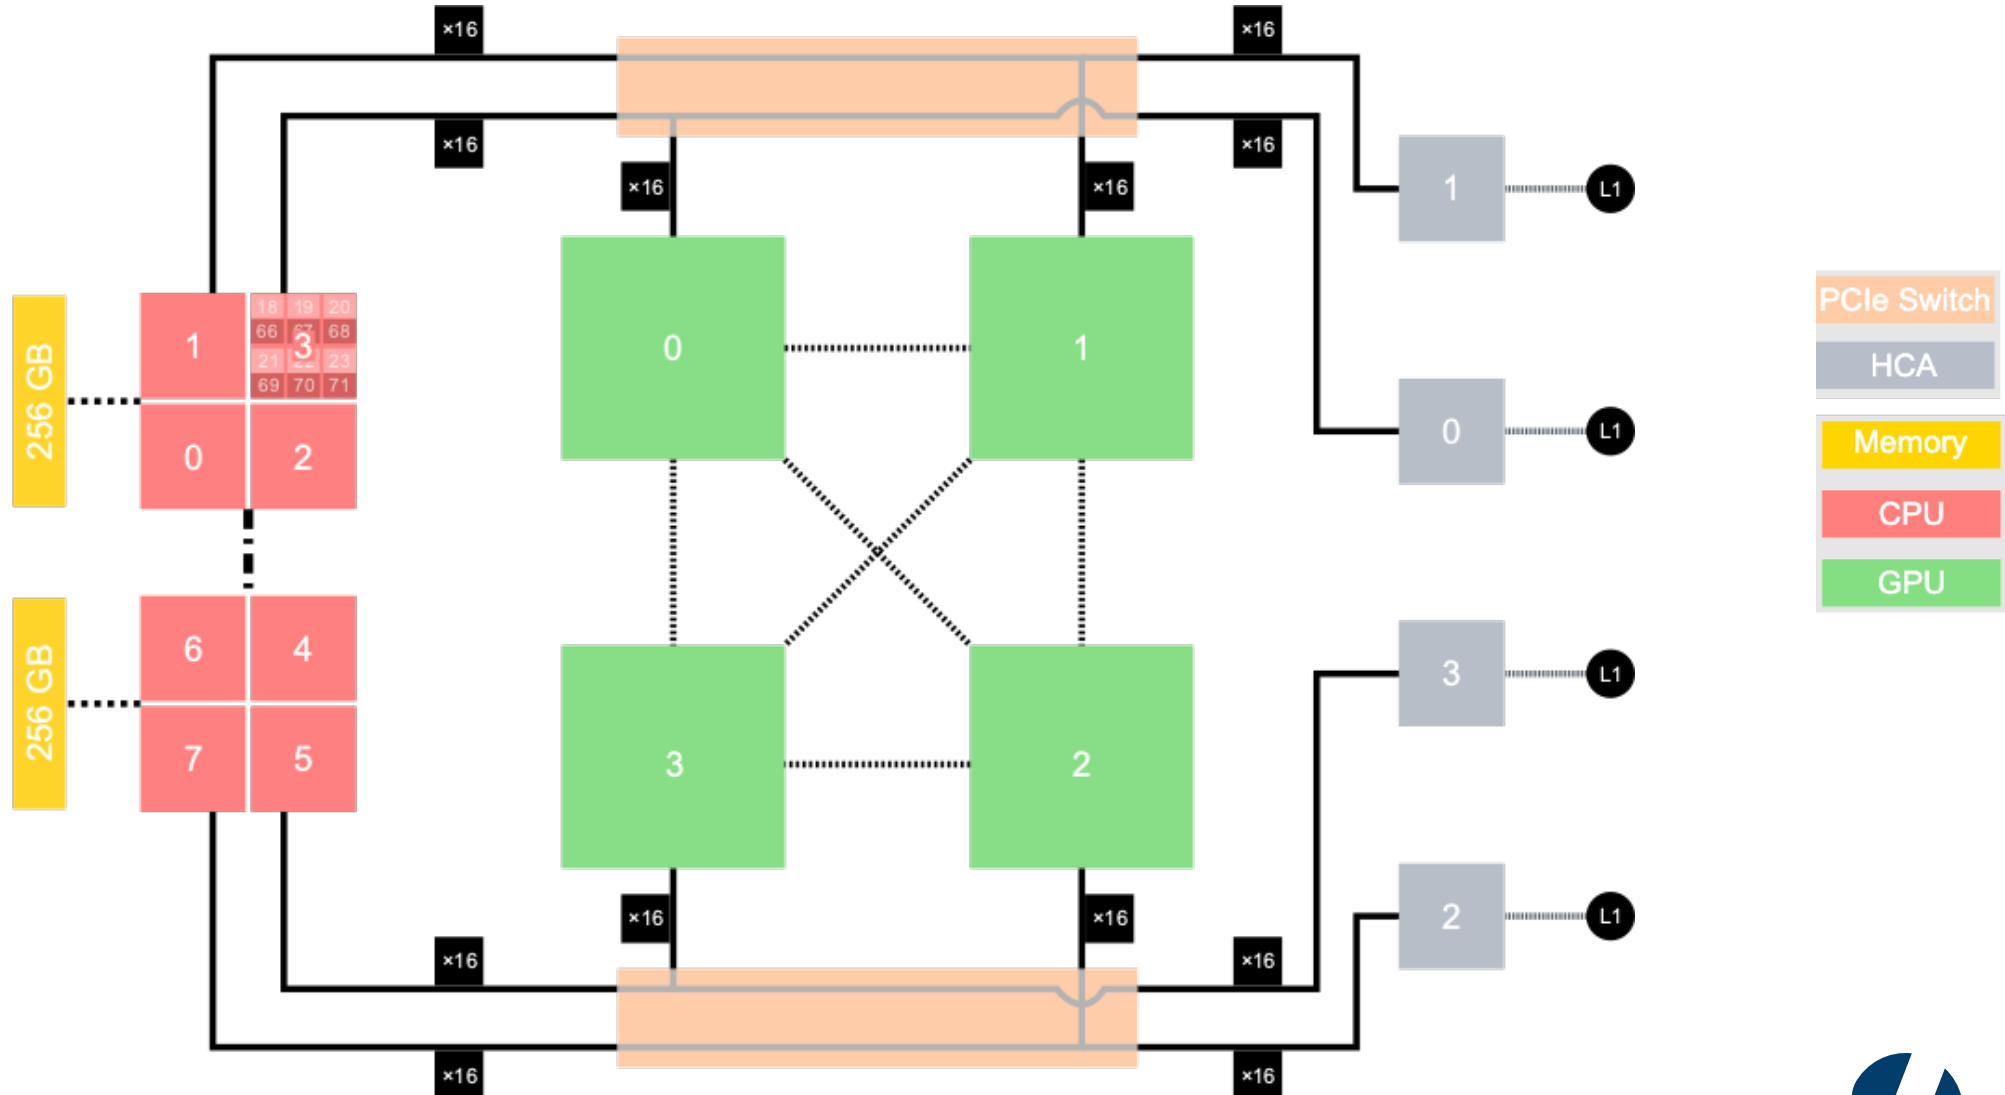
## JUWELS Cluster (w/o GPU nodes)

## JUWELS Booster

nodes	2511	x0.372	936	nodes
cores / node	48	x1	48	cores / node (CPU)
FP64 units (CPU)	384	x9	3456	FP64 units (GPU)
per node	4.15 TF	x19	78 TF	per node
peak performance	<b>10.6 PF</b>	x7	<b>73 PF</b>	peak performance
main memory	≥96 GB	x5.3	512 GB	main memory
high-bw memory	0 B	xNaN	160 GB	high-bw memory
memory bw per node	256 GB/s	x24	6 TB/s	memory bw per node
memory bw	0.6 PB/s	x9.3	5.6 PB/s	memory bw
link speed (EDR)	100 Gb/s	x2	200 Gb/s	link speed (HDR)
NIC per node	1	x4	4	NICs per node
topology	prun. FT		5 DF+	topology




# JUWELS BOOSTER – ITS ALL ABOUT ACCELERATORS



# ... STILL ALL ABOUT ACCELERATORS

FRONTIER – First known Exascale System (2022)



To Slingshot





# JUWELS – ENABLING SOFTWARE

”Power is nothing without control”

- ParaStation Modulo
  - Resource management
  - ParaStation MPI
  - Imaging

- |   |     |
|---|-----|
| <ul style="list-style-type: none"><li>• Ansible as provisioning system</li><li>• SLURM as scheduler</li><li>• EasyBuild as scientific software package management</li><li>• RHEL (service nodes internal storage in Ceph)</li><li>• Rocky Linux 8</li></ul> | JSC |
|---|-----|

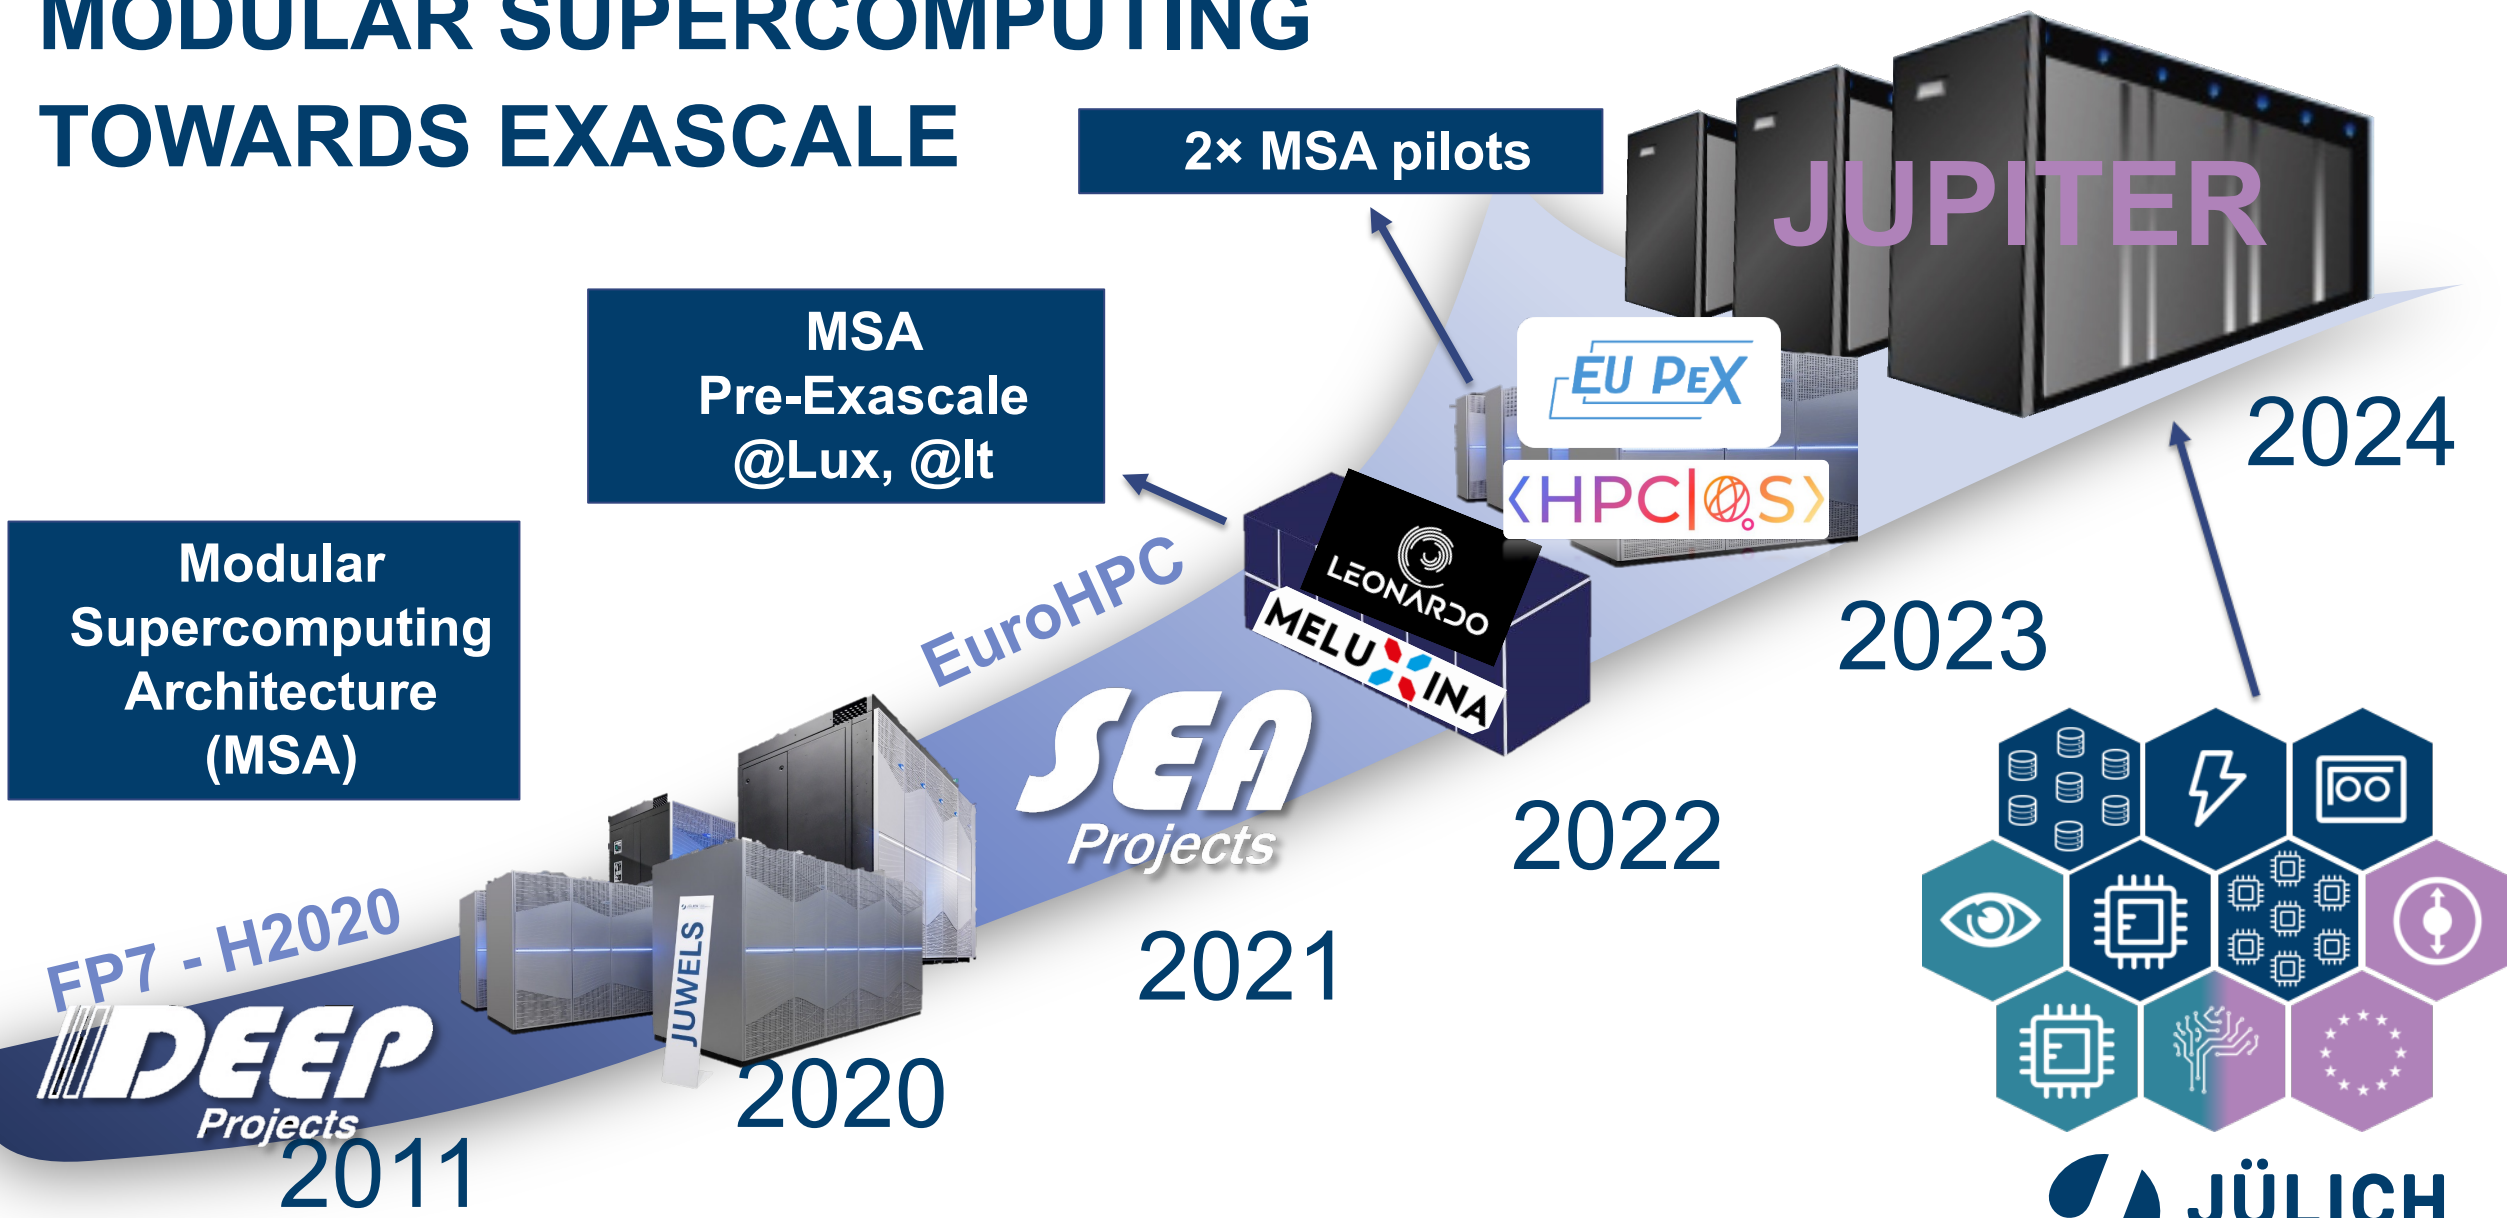
- Several Eviden (Atos) Tools



ANSIBLE



# MODULAR SUPERCOMPUTING TOWARDS EXASCALE



Modular  
Supercomputing  
Architecture  
(MSA)

MSA  
Pre-Exascale  
@Lux, @It

2x MSA pilots

JUPITER

2024

2023

2022

2021

2020

2011

FP7 - H2020  
**DEEP**  
Projects


EuroHPC

**SEA**  
Projects

LEONARDO  
MELU  
INA

EU PEX

<HPC|OS>



# JUPITER – BREAKING NEWS!

15.06.2022



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 Standort für die ersten 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



## Germany to host Europe's first exascale supercomputer

Jupiter added to HPC solar system

Dan Robinson

Thu 16 Jun 2022 // 07:38 UTC



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.



# JUPITER – TIMELINE



- 17.12.2021: Call for Expression of Interest (EoI) for Hosting Entity
- 14.02.2022: Deadline EoI Submission
- 16.05.2022: Hearings
- 15.06.2022: Hosting site decision and announcement
- 14.12.2022: Signature hosting agreement
- **Q1-Q3 2023: Procurement (Competitive Dialogue)**
- Q1/2024: Start installation of JUPITER
- End of 2024: Put in operation JUPITER



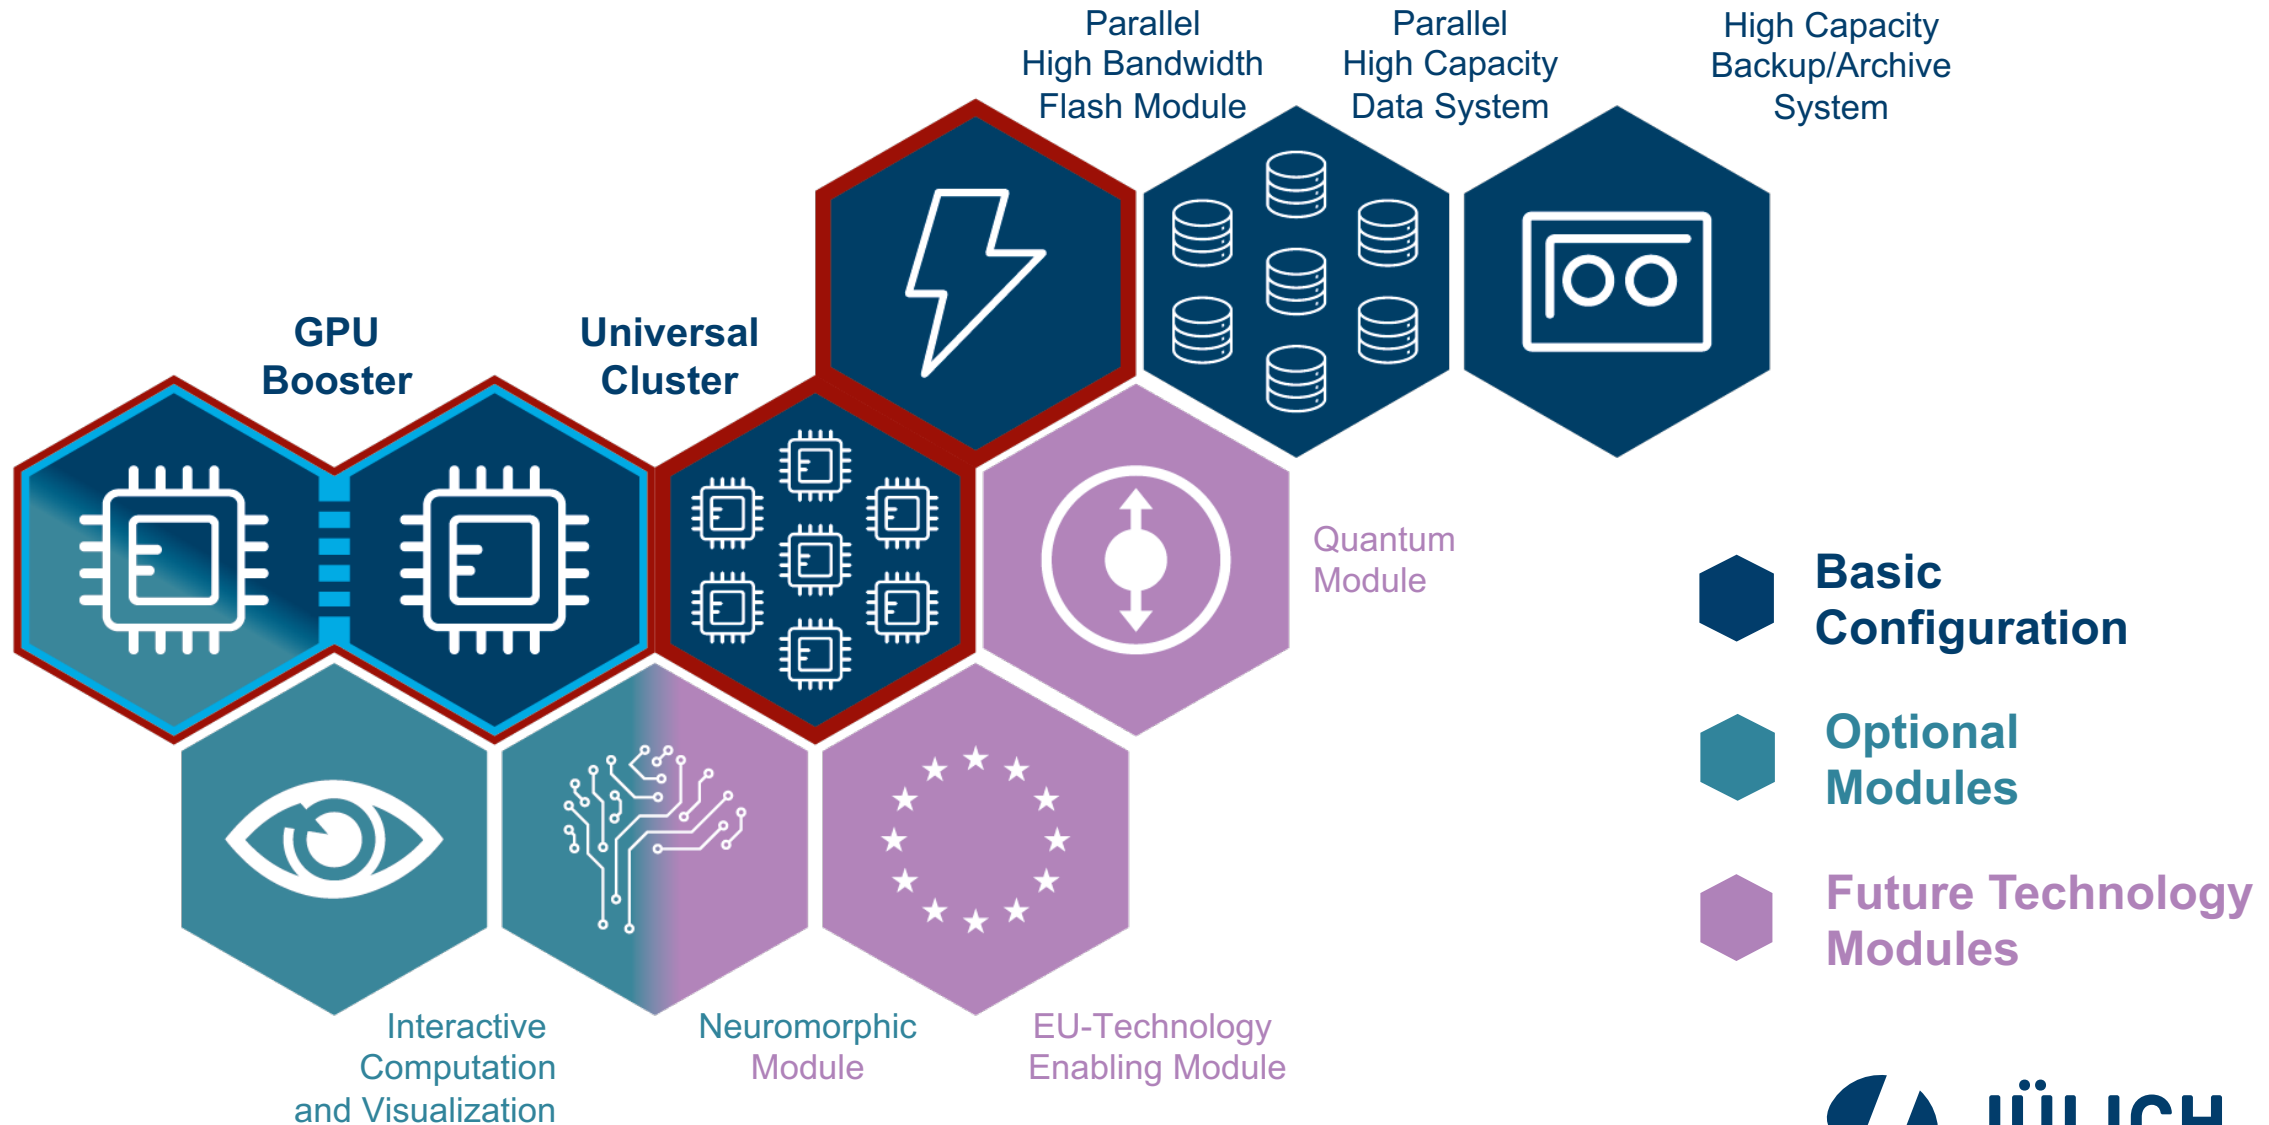
The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Digital Europe programme, as well as by Germany through the BMBWF and the MKW.

# JUPITER – WHERE WE CURRENTLY ARE

## Running the Competitive Dialogue

- Procurer: EuroHPC
- Technical Description: JSC
- Evaluation: JSC + EuroHPC
  
- Dialogue Sessions: JSC + EuroHPC
  
- Technical Specifications, Benchmarks: JSC
- Evaluation: 4 Experts

# JUPITER – ARCHITECTURE AND OPTIONS




# JUPITER – MODULES AT A HIGH LEVEL


- CPU-based general-purpose Cluster Module (Byte/Flop target > 0,4)
- GPU-based Booster Module (NVIDIA or AMD)
- Future Technology Modules: require additional funding/calls
  - EU-Technology Enabling – e.g. SiPearl Cronos
  - Quantum Module – e.g. European Quantum Flagship
- Storage
  - SCRATCH – based on flash
  - HOME, DATA, ARCHIVE – based on spinning disks and tape (>1EB)
- Infrastructure
  - 27 MW: up to 20 MW for Booster, ~6 MW for Cluster, ~1 MW service and storage
  - Warm water cooling

# GPU: XPU (CPU + GPU)

## AMD MI300



## NVIDIA Grace-Hopper




<https://www.hardwareluxx.de/index.php/news/hardware/grafikkarten/60485-isscc-2023-auf-cache-auf-compute-fohgt-compute-auf-cache-bei-amd.html>

<https://developer.nvidia.com/blog/inside-nvidia-grace-cpu-nvidia-amps-up-superchip-engineering-for-hpc-and-ai/>



# JUPITER – SOFTWARE

- User Applications
  - Large application benchmark set covering several major user communities and algorithmic patterns
  - Subset of benchmarks for assessing the scalability of the system
  - Early Access Program on JUPITER to increase usability from day 1
- Scientific Software Stack
  - EuroHPC JU calls for Exascale (ready) software
  - Exascale Computing Project results



# JUWELS VS. JUPITER (EXPECTATIONS)

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

# 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](https://fz-juelich.de/jupiter) | [#exa\\_jupiter](https://twitter.com/#!/exa_jupiter)



Funding Agencies:



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

