



JUWELS BOOSTER

The last most remarkable stop (for now) of the NVIDIA application lab at JSC

21.06.2022 | D. ALVAREZ


BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC




BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC




BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC



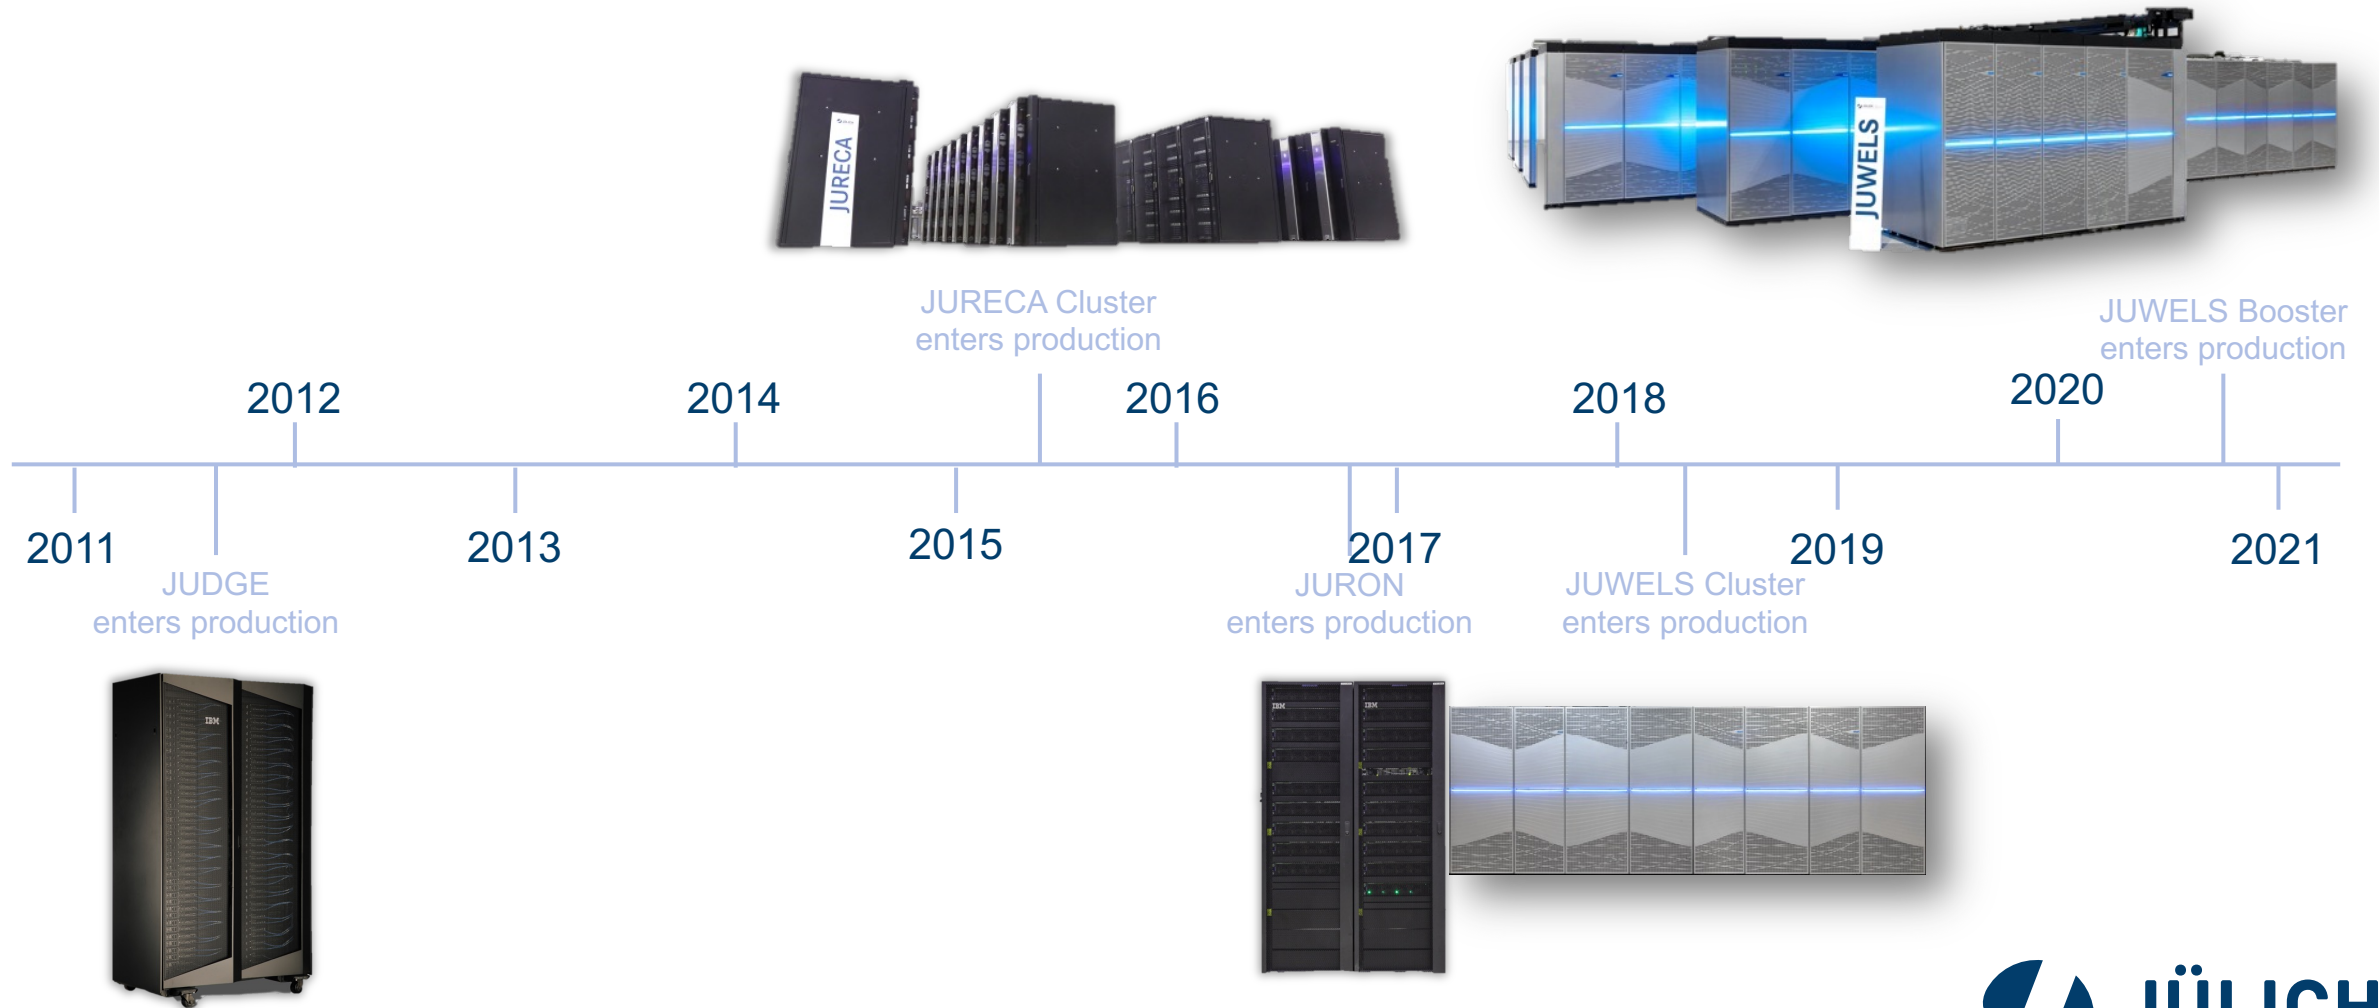
BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC



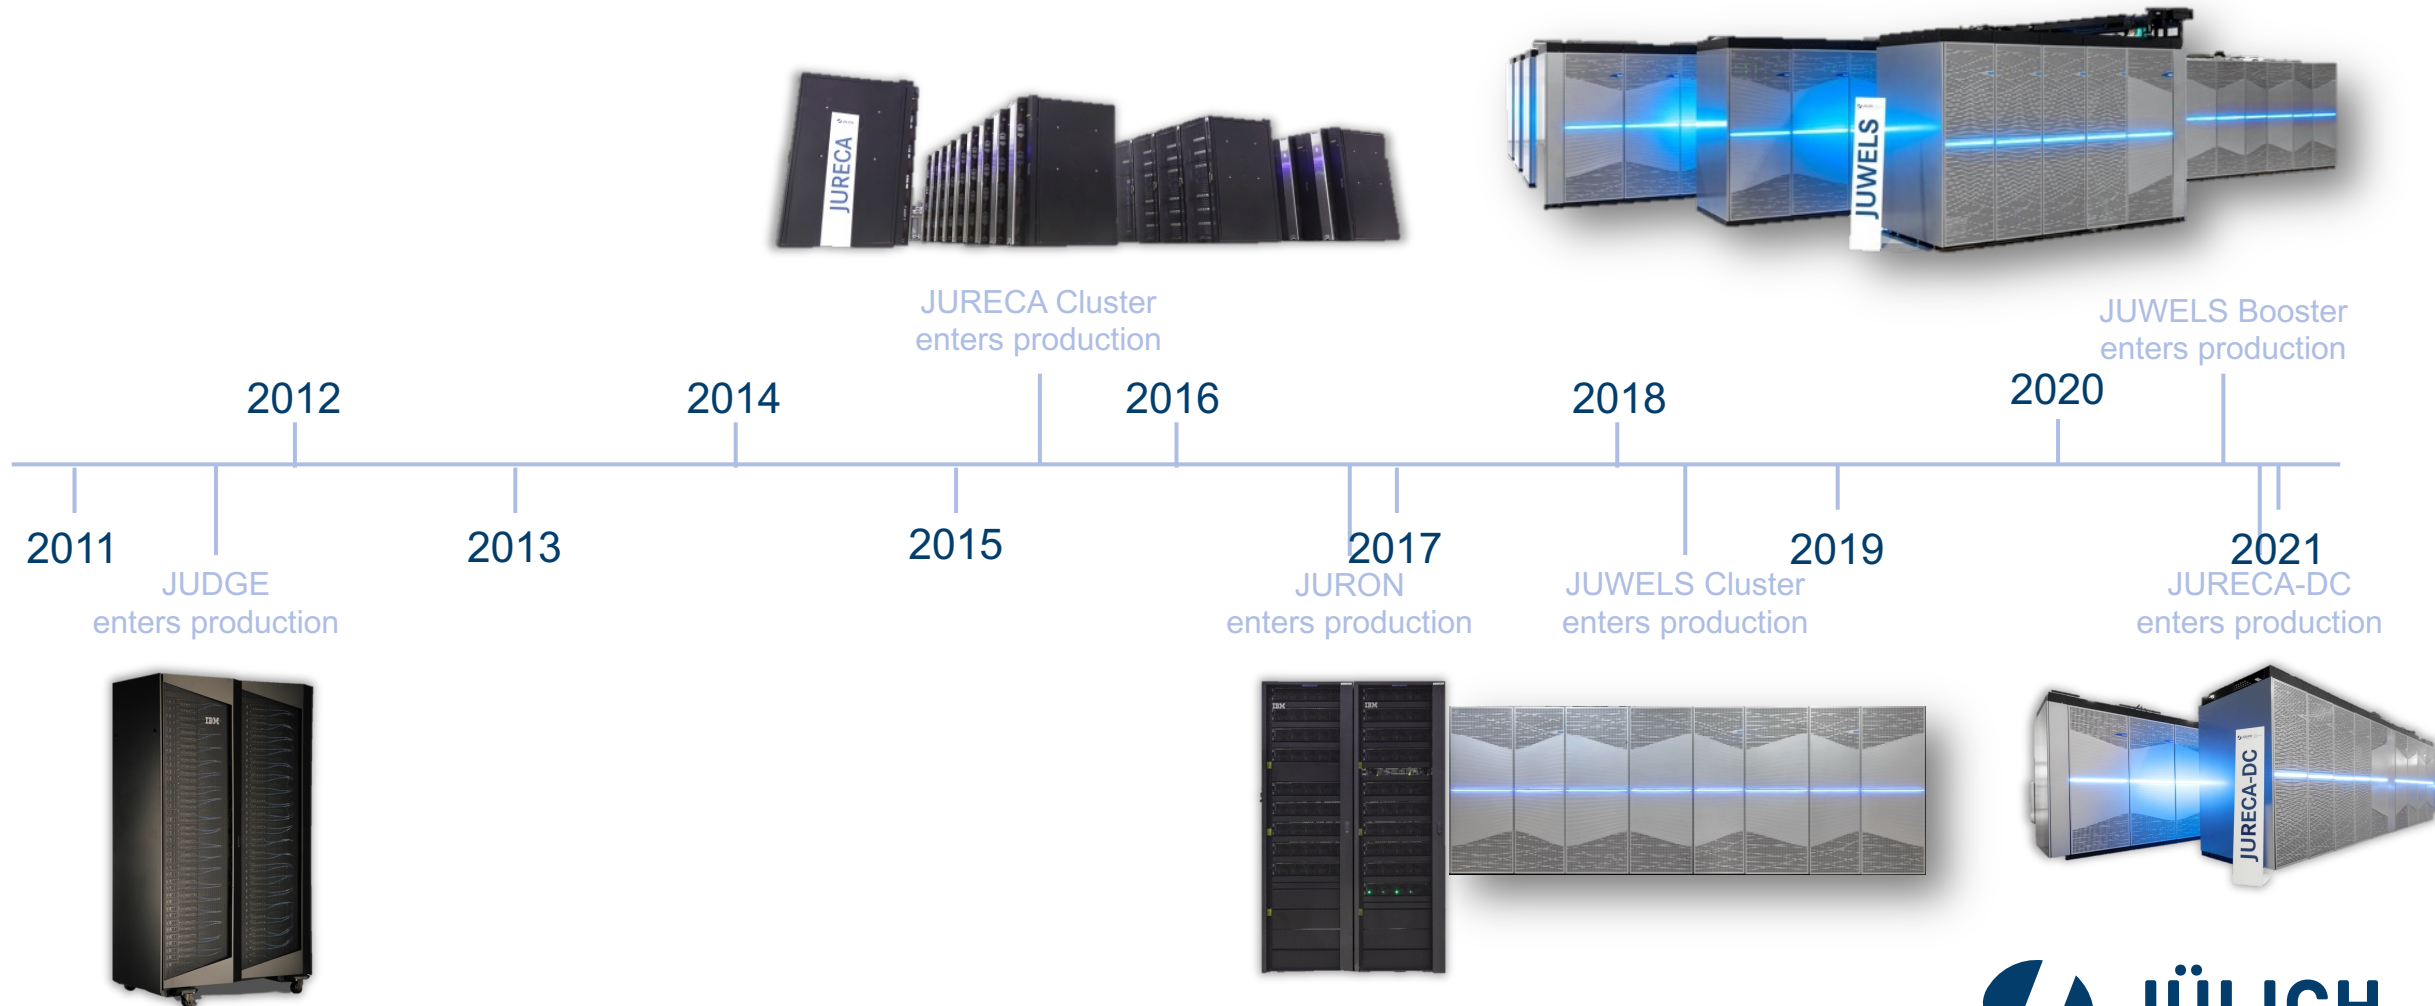
BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC



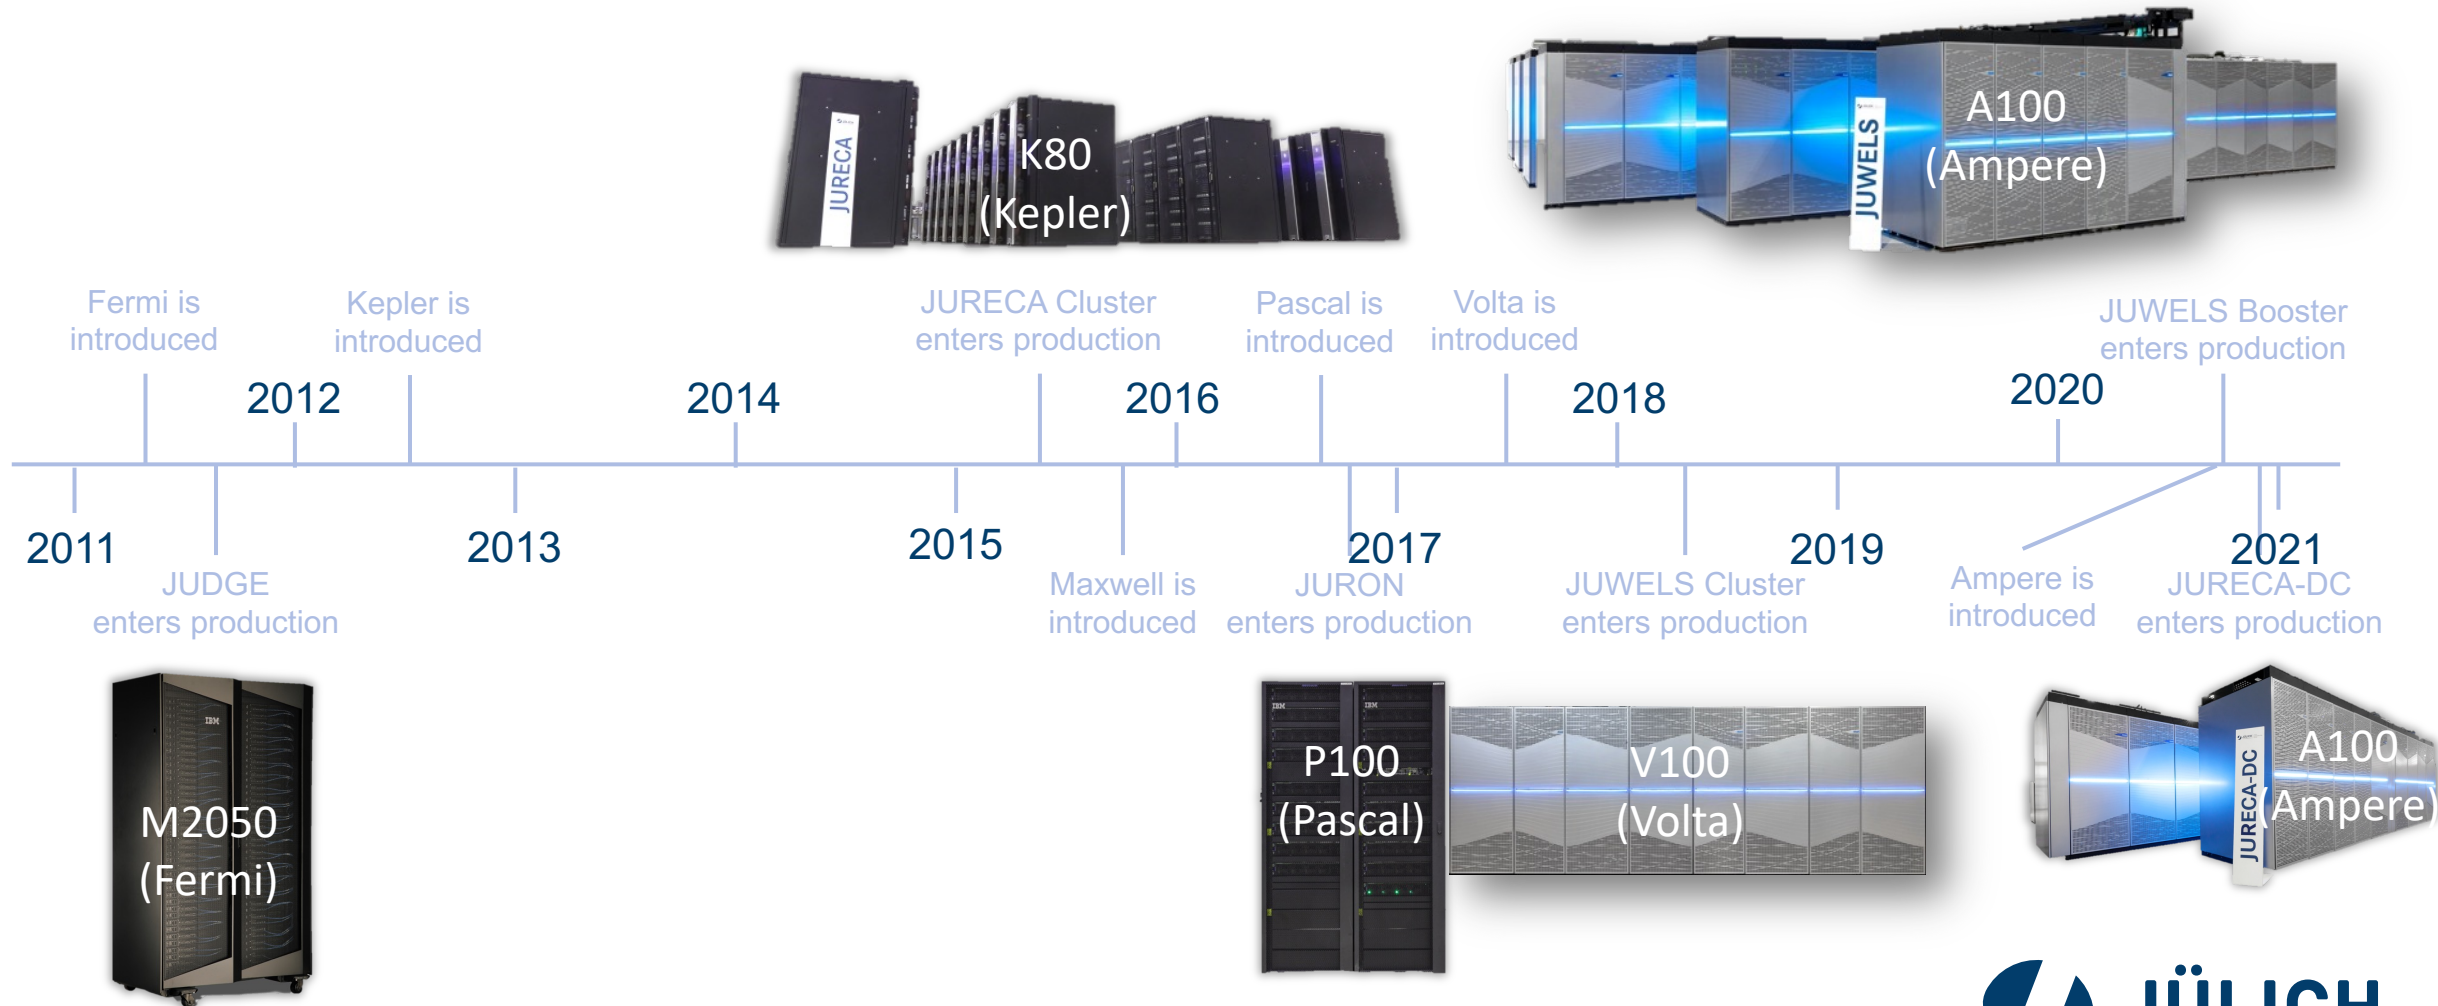
BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC



BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC



BRIEF TIMELINE OF SOME GPU PRODUCTION SYSTEMS AT JSC









#1 in Green



Top 100 in HPL

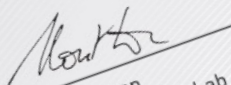
TOP 500 CERTIFICATE

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

is ranked
No. 11
among the World's TOP500 Supercomputers
with **44.12 PFlop/s Linpack Performance**
in the 59th TOP500 List published at the ISC22
Conference on June 01, 2022.
Congratulations from the TOP500 Editors


Erich Strohmaier
NERSC/Berkeley Lab


Jack Dongarra
University of Tennessee


Horst Simon
NERSC/Berkeley Lab


Martin Meuer
Prometeus


Wu-chun Feng
Virginia Tech

The GREEN 500 CERTIFICATE

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

is ranked
No. 18
among the World's TOP500 Supercomputers
with **308 GFlops/watts Performance**
in the Green500 List published at the ISC22
Conference on June 01, 2022.
Congratulations from the Green500 Editors


Kirk Cameron
Virginia Tech



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

is ranked
No. 11

among the World's TOP500 Supercomputers
with **44.12 PFlop/s Linpack Performance**
in the 59th TOP500 List published at the ISC22
Conference on June 01, 2022.
Congratulations from the TOP500 Editors

Erich Strohmaier
NERSC/Berkeley Lab

Jack Dongarra
University of Tennessee

Martin Meuer
Prometeus

Wu-chun Feng
Virginia Tech

Kirk Cameron
Virginia Tech


The GREEN 500 CERTIFICATE

the World's TOP500 Supercomputers
Green500 List published at the ISC22
Conference on June 01, 2022.
Congratulations from the Green500 Editors

is ranked
No. 18





with **308 GFlops/watts Performance**

#3 in TOP500 Europe
#7 HPCG500
#6 HPL-AI
Strong push back from AMD-based GPUs!




BOOSTER NODE DESIGN

JUWELS BOOSTER NODES


- 936 compute nodes 
 - 2x 24-core AMD Epyc Rome CPUs 
 - 2x 8 memory channels
 - 2x 256 GB DDR4 @ 3.2GHz
 - 96 PCIe Gen4 lanes
 - 512 GB DDR memory
 - 4x Nvidia A100 GPUs 
 - 9.7 / 19.5 TF/s peak
 - 40 GB HBM2
 - 1.5 TB/s memory performance
 - NVLink3 full mesh
 - 4 links (200GB/s) between GPU pairs
 - PCIe Gen4 x32 (64 GB/s)
 - 4x HDR200 InfiniBand adapter (1 per GPU) 



BOOSTER NODES



BOOSTER NODES




AMD 
2x EPYC 7402

BOOSTER NODES



4x A100 40GB
1x “Redstone”
board



AMD

2x EPYC 7402


BOOSTER NODES



4x A100 40GB
1x “Redstone”
board




2x boards
4x ConnectX-6




AMD

2x EPYC 7402

BOOSTER NODES




BOOSTER NODES



BOOSTER NODES


Past challenges



- Solved thanks to the close collaboration with some of the members of this application lab
- PCI switch
 - Insufficient bidirectional bandwidth
 - Triggered by a bug in the PCI switch FW
- Dead switches
 - Triggered by a bug in the VR for the PCI switch

BOOSTER NODES

Current challenges




- **Being solved** thanks to the close collaboration with some of the members of this application lab
- GPU overheating
 - Currently affecting node availability
 - Fix coming in the next 2-3 months
 - Implies changing the water blocks of the whole system
- Corrupt L3 caches
 - 1 code triggers data corruption under very specific circumstances
 - Still being analyzed

A photograph of a server room with rows of server racks. The racks are illuminated with a bright blue light, creating a strong visual theme. The text "NETWORK DESIGN" is overlaid in the center in a white, bold, sans-serif font. The racks have a textured, horizontal-slatted appearance. The perspective is from a low angle, looking down a row of racks into the distance.


NETWORK DESIGN

BOOSTER TOPOLOGY CHOICE


- 4 HCAs per node, 936 nodes = 3744 end points
- Fat tree, 40 ports switches, no pruning:
 - 199 L1 switches
 - 200 L2 switches
 - 100 Top level switches
 - ~7488 copper cables
 - ~3744 optical cables
 - Expensive
 - Best performance
- DragonFly+
 - 199 L1 switches
 - 200 L2 switches
 - ~7488 copper cables
 - ~400 copper cables
 - Cheaper
 - Good performance
 - Needs adaptive routing



VS

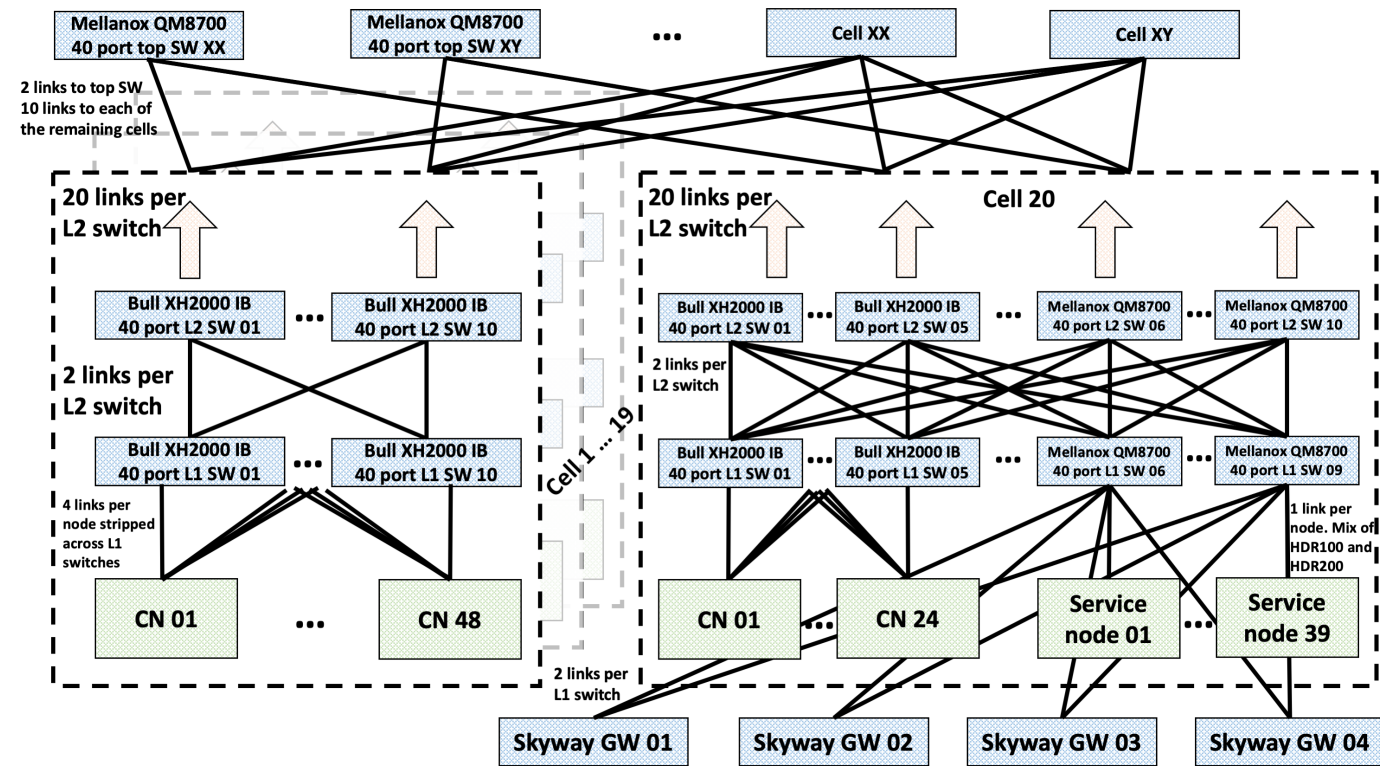


BOOSTER TOPOLOGY CHOICE



- 10 links between each cell
- 250 GB/s
- Non-minimal adaptive routing to avoid congestion
- 200 links (10 per cell) to Cluster
- 5 TB/s

- Fat tree as DragonFly group topology
- 48 nodes per cell (two racks)
- 4× links striped over L1 switches



JUWELS IB FABRIC

Welcome to the jungle

Cluster node

Cluster GPU node

Cluster switch

Cluster gateway

Top level switch

Booster node


Booster switch

Booster gateway

JUST-IME node

JUST-IME switch

Service node



JUWELS IB FABRIC

Welcome to the jungle

FDR/EDR/HDR100/HDR200

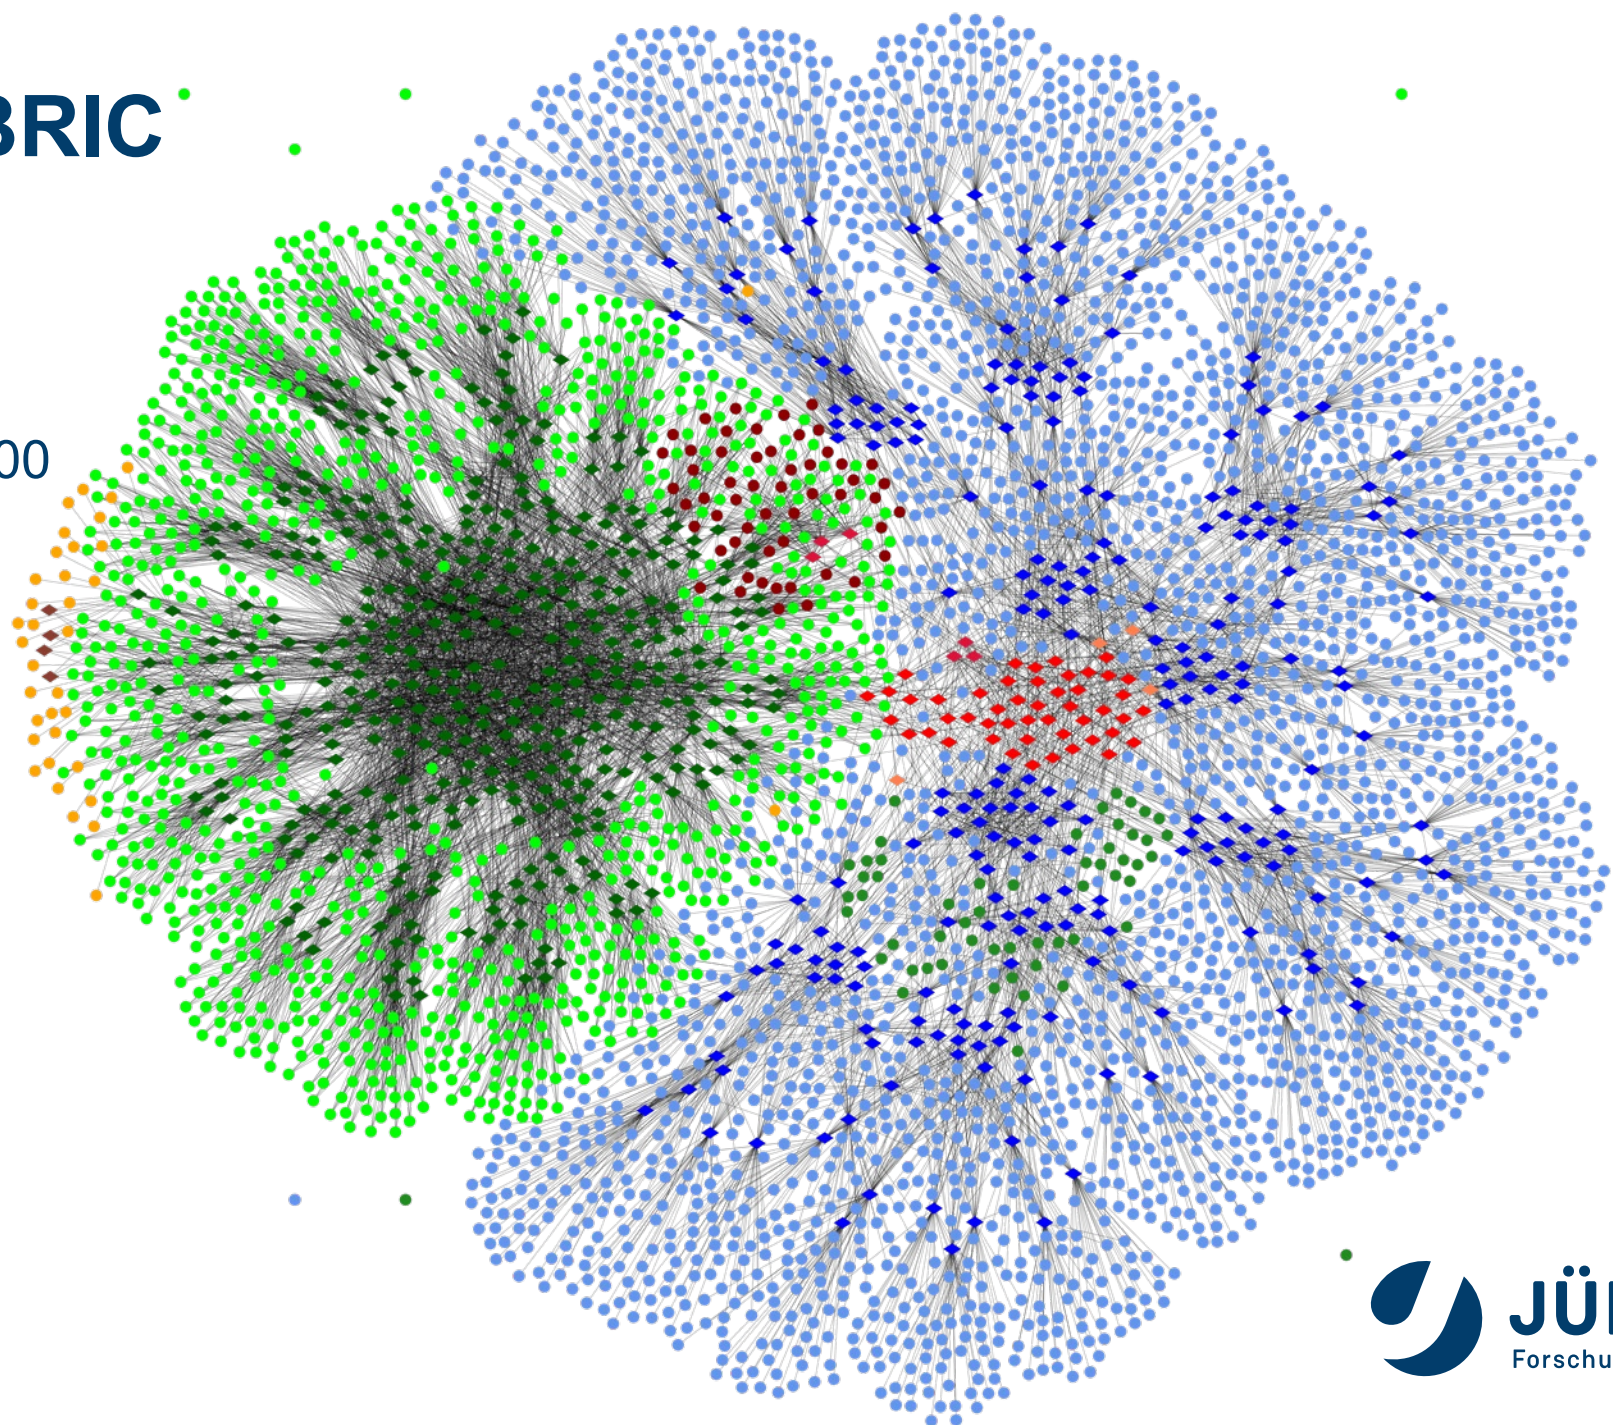
35000 ports

8 gateways

693 switches

6500+ endpoints


3 routing algorithms



JUWELS IB FABRIC

Near future plans

- Introduce new IB-Ethernet gateways
 - Better resiliency in case of GW failure
 - Side effect: 2x BW to storage
- Enable SHARP/HCOLL MPIs
 - In the wish list for some time
 - Should be stable enough now (multiple switch FW bugs in the past)
 - Need to be assessed




INFRASTRUCTURE



JUWELS INFRASTRUCTURE

- Booster:
 - 39 Sequana XH2000 racks
 - 6 rows
- Transformers 1 & 2 full utilized
- Direct warm-water cooling
 - 34 - 37 °C facility-side inlet temperature
 - Now with dry coolers!
 - Much more efficient
 - Sorry for the service interruptions!




```
*****
* Welcome to                                     *
*
*      _ _ _ _ _ \ \ / / _ _ _ _ _ / _ _ _ _ _ |   Juelich Wizard
*      | | | | | \ \ / / | | | | | \ _ _ _ _ _ \   for
*      | _ | | | _ | \ v v / | | | | | _ _ _ _ _ )   European Leadership
*      \ _ _ / \ _ _ / \ _ / \ _ / | | | | | _ _ _ _ _ /   Science
*
*****
```

SOFTWARE

KEY ENABLING SOFTWARE

"Power is nothing without control"

- Ansible as provisioning system
- SLURM as scheduler
- ParaStation Modulo as resource manager/health checking
- EasyBuild as scientific software package management
- RHEL (service nodes internal storage in Ceph)
- Rocky Linux 8



KEY ENABLING SOFTWARE

Near future plans

- Enable GPU frequency locking on job prologue, to a user specified frequency



KEY ENABLING SOFTWARE

Near future plans

- Enable GPU frequency locking on job prologue, to a user specified frequency
- Enable topology plugin
 - Help to keep nodes in a job together and minimize job fragmentation
- Not straight forward due to network topology/multiple HCAs per node



JUWELS BOOSTER

The last ~20 months in (some) numbers



JUWELS BOOSTER

The last ~20 months in (some) numbers

- 586523 successfully completed jobs
- Too many small scale jobs 😞
- 714 users
- Top 15 users responsible for 399798 jobs
- 3 Gordon Bell prize candidate jobs/users
- 4449 tickets opened
 - 3967 node HW tickets
 - 205 IB tickets

THANK YOU