



Jülich Supercomputing Centre


(Germany)

E.Suarez

Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich

Dual Architecture

Classic JSC approach: two production systems operated in parallel



JUQUEEN

2012 - 2018

- IBM Blue Gene/Q
- PowerPC[®] A2
 - 16x cores per node
 - 1.6 GHz
 - 16 Gbyte main memory
- 28,672 nodes
 - 458,752 cores
 - 28 racks
- Peak performance: **5,9 PFlop/s**
- Connected to JUST (GPFS cluster):
- 5D network



JUQUEEN:
Top500 Nov2017:
#22 worldwide
#6 in Europe

Planned decommissioning in April 2018

JURECA Cluster

Since 2015

- Dual-socket Intel Haswell (E5-2680 v3)
 - 12x cores/socket
 - 2.5 GHz
 - ≥ 128 GB main memory
- 1,884 compute nodes (45,216 cores)
 - 75 nodes: 2x K80 NVIDIA GPUs
 - 12 nodes: 2 x K40 NVIDIA GPUs
 - 512 GB main memory
- Peak performance: **2.2 Petaflop/s** (1.8 w/o GPUs)
- Mellanox InfiniBand EDR (full fat-tree)
- Connected to the central GPFS cluster JUST
 - 20 PByte online disk and
 - 100 PByte offline tape capacity



JURECA Booster


Since Nov 2017

- Intel Xeon Phi - KNL (7250-F)
 - 68x cores/socket
 - 1.4 GHz
 - 16 GiB MCDRAM, 96 GiB DDR4
- 1,640 compute nodes (111,520 cores)
- Intel OmniPath (non-blocking fat-tree)
- Connections
 - High-speed connection to Cluster
 - Shared GPFS file systems (JUST)
- Booster peak performance: 5 PFlop/s
- Cluster + Booster act as single system:
 - SLURM jobs can be distributed over both
 - Linpack over both sides: **3,78 PFlop/s**



JURECA
Cluster + Booster
Top500 Nov2017:
#29 worldwide
#10 Europe

Cluster Booster architecture

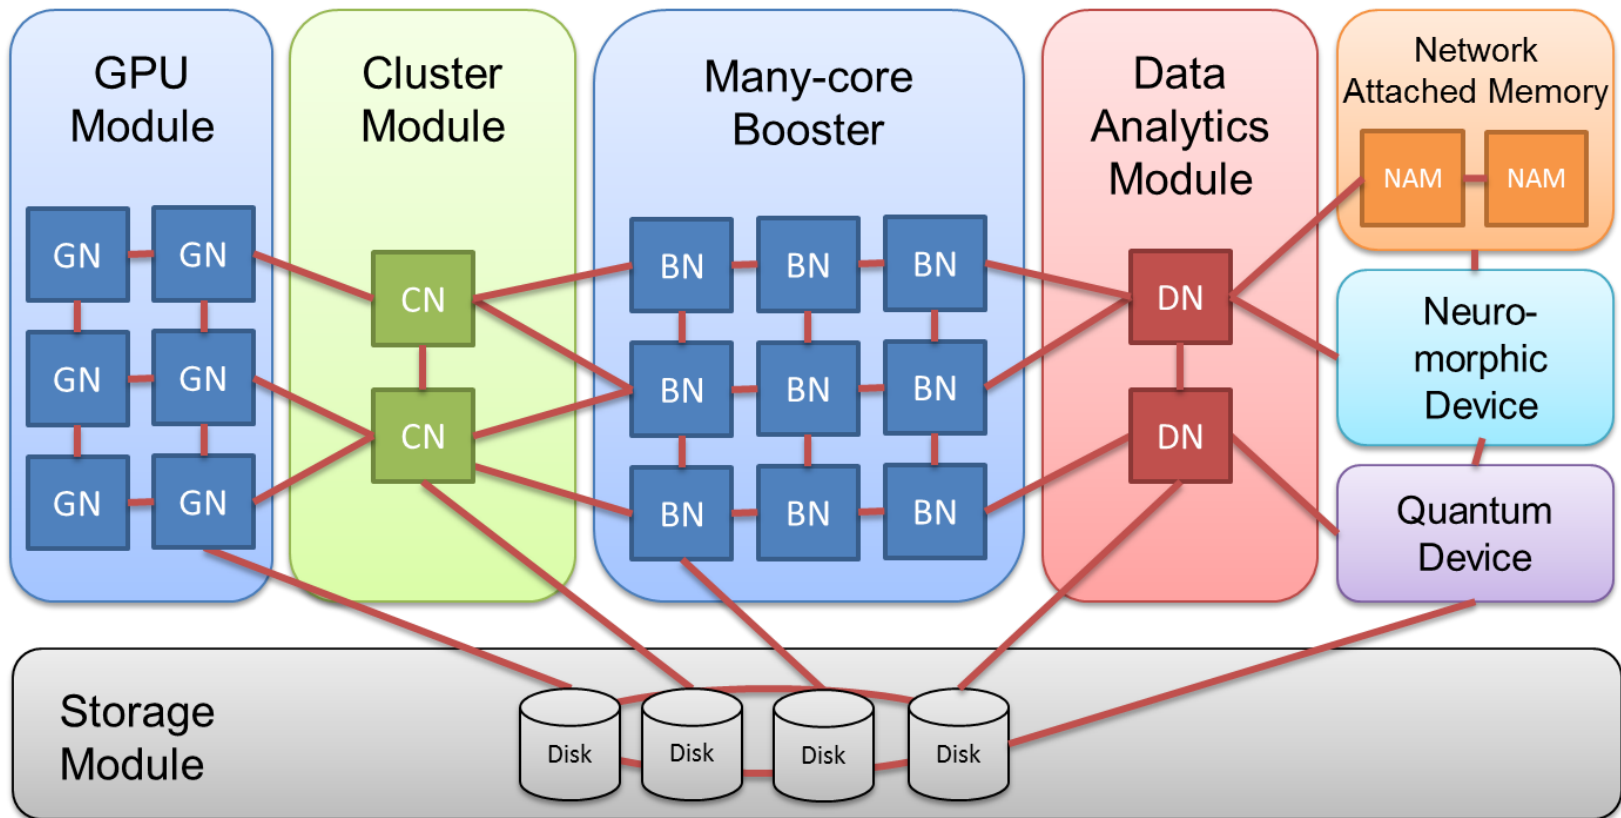


Low/Medium scalable code parts

Highly scalable code parts

Next step: Modular Supercomputing architecture

Generalization of the Cluster-Booster concept



- Different compute modules, each tailored to specific user needs
- Applications can access all and map their needs onto the HW

JUWELS:

First Tier-0/1 Modular Supercomputer

JUWELS Cluster (first module)

- 2559x nodes - Dual-socket Intel Platinum 8168
 - 24x cores/socket, 2.7 GHz
- Nodes Configuration:
 - 2271 standard (96 GB main memory)
 - 240 large memory (192 GB)
 - 48 with 4x NVIDIA V100 GPU each
- Infiniband EDR interconnect (fat tree)
- Peak performance:
 - **10.4 Pflop/s (CPU) + 1.6 Pflop/s (GPU)**



Installation planned in Q2/2018