Table of Contents

System overview

This page is supposed to give a short overview on the available systems from a hardware point of view. All hardware can be reached through a login node via SSH: **deep@fz-juelich.de**. The login node is implemented as virtual machine hosted by the master nodes (in a failover mode). Please, see also information about getting an account and using the batch system.

DEEP-EST Modular Supercomputer (prototype system)

The DEEP-EST system is a prototype of Modular Supercomputing Architecture (MSA) consisting of the following modules:

- Cluster Module (CM)
- Extreme Scale Booster (ESB)
- Data Analytics Module (DAM)

In addition to the previous compute modules, a Scalable Storage Service Module (SSSM) provides shared storage infrastructure for the DEEP-EST prototype (`/usr/local`) and is accompanied by the All Flash Storage Module (AFSM) leveraging a fast local work filesystem (mounted to `/work` on the compute nodes).

The modules are connected together by the Network Federation (NF), composed by different types of interconnects and briefly described below. **The setup will change into an "all IB EDR network" in the next months.**

Cluster Module

It is composed of 50 nodes with the following hardware specifications:

Cluster [50 nodes]: `dp-cn[01-50]`:	
• 2 Intel Xeon 'Skylake' Gold 6146 (12 cores (24 threads), 3.2GHz) • 192 GB RAM • 1 x 400GB NVMe SSD • network: InfiniBand EDR (100 Gb/s)	

Extreme Scale Booster

It is composed of 75 nodes with the following hardware specifications:

Extreme Scale Booster [75 nodes]: `dp-esb[01-75]`	
• 1 x Intel Xeon 'Cascade Lake' Silver 4215 CPU @ 2.50GHz • 1 x Nvidia V100 Tesla GPU (32 GB HBM2) • 48 GB RAM • 1 x 512 GB SSD • network: IB EDR 100 (Gb/s) (nodes `dp-esb[01-25]` to be converted from Extoll to IB EDR)	

Data Analytics Module

It is composed of 16 nodes with the following hardware specifications:

Data Analytics Module [16 nodes]: `dp-dam[01-16]`	
• 2 x Intel Xeon 'Cascade Lake' Platinum 8260M CPU @ 2.40GHz • 1 x Nvidia V100 Tesla GPU (32 GB HBM2) • 1 x Intel STRATIX10 FPGA (32 GB DDR4) • 384 GB RAM + 3 TB non-volatile memory (14 nodes with 2, 2 nodes with 3) • 2 x 1.5 TB Intel Optane SSD (1 for local scratch, 1 for BeeOND) • 1 x 240 GB SSD (for boot and OS) • network: EXTOLL (100 Gb/s) + 40 Gb Ethernet (to be converted to IB EDR)	

Scalable Storage Service Module

It is based on spinning disks. It is composed of 4 volume data server systems, 2 metadata servers and 2 RAID enclosures. The RAID enclosures each host 24 spinning disks with a capacity of 8 TB each. Both RAIDs expose two 16 Gb/s fibre channel connections, each connecting to one of the four file servers. There are 2 volumes per RAID setup. The volumes are driven with a RAID-6 configuration. The BeeGFS global parallel file system is used to make 292 TB of data storage capacity available.

Here are the specifications of the main hardware components more in detail:

All Flash Storage Module

It is based on PCIe3 NVMe SSD storage devices. It is composed of 6 volume data server systems and 2 metadata servers interconnected with a 100 Gbps EDR-InfiniBand fabric. The AFSM is integrated into the DEEP-EST Prototype EDR fabric topology of the CM and ESB EDR partition. The BeeGFS global parallel file system is used to make 1.3 PB of data storage capacity available.

Here are the specifications of the main hardware components more in detail:

Network overview

Currently, different types of interconnects are in use along with the Gigabit Ethernet connectivity that is available for all the nodes (used for administration and service network). The following sketch should give a rough overview. Network details will be of particular interest for the storage access. Please also refer to the description of the filesystems.

network is going to be formed to an "all IB EDR" setup soon !

Rack plan

This is a sketch of the available hardware including a short description of the hardware interesting for the system users (the nodes you can use for running your jobs and that can be used for testing).

SSSM rack

This rack hosts the master nodes, file servers and the storage as well as network components for the Gigabit Ethernet administration and service networks. Users can access the login node via **deep@fz-juelich.de** (implemented as virtual machine running on the master nodes). The rack is air-cooled.

CM rack

Contains the hardware of the DEEP-EST Cluster Module including compute nodes, management nodes, network components and liquid cooling unit.

DAM rack

This rack hosts the nodes of the Data Analytics Module of the DEEP-EST prototype and the Network Federation Gateways. The rack is air-cooled.

SDV rack

Along with the prototype systems serveral test nodes and so called software development vehicles (SDVs) have been installed in the scope of the DEEP(-ER,EST) projects. These are located in the SDV rack (07). The following components can be accessed by the users:

Prototype DAM [4 nodes]: `protodam[01-04]`
- 2 x Intel Xeon 'Skylake' (26 cores per socket)
- 192 GB RAM
- network: Gigabit Ethernet

Old DEEP-ER Cluster Module SDV [16 nodes]: `deeper-sdv[01-16]`
- 2 Intel Xeon 'Haswell' E5-v2680 v3 (2.5 GHz)
- 128 GB RAM
- 1 NVMe with 400 GB per node(accessible through BeeGFS on demand)
- network: 100 Gb/s Extoll tourmalet

KNLs [4 nodes]: `knl[01,04-06]`
- 1 Intel Xeon Phi (64-68 cores)

- 1 NVMe with 400 GB per node (accessible through BeeGFS on demand)
- 16 GB MCDRAM plus 96 GB RAM per KNL
- network: Gigabit Ethernet

GPU nodes for Machine Learning [3 nodes]: `ml-gpu[01-03]`

- 2 x Intel Xeon 'Skylake' Silver 4112 (2.6 GHz)
- 192 GB RAM
- 4 x Nvidia Tesla V100 GPU (PCIe Gen3), 16 GB HBM2
- network: 40GbE connection

Further information

- [Information about the batchsystem](#)
- [Filesystems](#)
- [Information on available software and tools](#)
- [Use the DEEP-EST Cluster Module](#)
- [Use the DEEP-EST Data Analytics Module](#)
- [Use the SDV Cluster](#)
- [Use the SDV KNLs](#)