

Overview of our systems

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

[Insert Figure?](#)

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

- Cluster Module
- Extreme Scale Booster
- Data Analytics Module

In addition to the previous compute modules, the Scalable Storage Service Module provides the shared storage infrastructure for the DEEP-EST prototype.

The modules are connected together by the Network Federation, composed by different types of interconnects and briefly described below.

Cluster Module

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

<p>Cluster [50 nodes]: <code>dp-cn[01-50]</code></p> <ul style="list-style-type: none"> • 2 Intel Xeon 'Skylake' Gold 6146 (12 cores (24 threads), 3.2GHz) • 192 GB RAM • 1 x 400GB NVMe SSD (for OS only, not exposed to users) • network: InfiniBand EDR (100 Gb/s) 	<p>Insert figure?</p>
---	---------------------------------------

Extreme Scale Booster

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

<p>Extreme Scale Booster [75 nodes]: <code>dp-esb[01-75]</code></p> <ul style="list-style-type: none"> • 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 ?? GB SSD (for boot and OS) • network: EXTOLL 100 (Gb/s) 	<p>Insert figure?</p>
---	---------------------------------------

Attention: the Extreme Scale Booster will become available in January 2020.

Data Analytics Module

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

<p>Data Analytics Module [16 nodes]: <code>dp-dam[01-16]</code></p> <ul style="list-style-type: none"> • 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 + 2 or 3 TB non-volatile memory (14 nodes with 2, 2 nodes with 3) • 2 x 1.5 TB Intel Optane SSD • 1 x 240 GB SSD (for boot and OS) • network: EXTOLL (100 Gb/s) + 40 Gb Ethernet 	<p>Insert figure?</p>
--	---------------------------------------

Network overview

Different types of interconnects are in use along with the Gigabit Ethernet connectivity (used for administration and service network) that is available for all the nodes. 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](#).

Attention: performance measurements for the Network Federation will be provided in the future.

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, files 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).

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.

SDV rack

Along with the prototype systems several test nodes and so called software development vehicles have been installed in the scope of the DEEP(-ER,EST) projects. These are located in the SDV rack. 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]: `kn1[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
- Old DEEP-ER NAM SDV:
 - size: 2 GB
 - network: Extoll
 - details: <https://www.deep-projects.eu/hardware/memory-hierarchies/49-nam>

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](#)