CURTA

A General-purpose High-Performance Computer

at

ZEDAT, Freie Universität Berlin

16th March 2020

Loris Bennett
Boris Proppe
Bernd Melchers

Curta

A General-purpose High-Performance Computer

at ZEDAT, Freie Universität Berlin

Loris Bennett[1], Bernd Melchers[1], and Boris Proppe[1]

[1]Scientific Computing Service, ZEDAT, Freie Universität Berlin, Germany

16th March 2020

Abstract

Curta is a general-purpose high-performance computer operated by the Zentral Einrichtung Datenverabeitung (ZEDAT) of Freie Universität Berlin and designed to provide enhanced computing resources for a wide range of disciplines within the university.

1 Introduction

Freie Universität Berlin has provided its researchers with high-performance computing (HPC) resources via its central computing facility *Zentral Einrichtung Datenverarbeitung* (ZEDAT) for many years. Funding for the current HPC system, *Curta* [Wik20], was provided by *Deutsche Forschungsgemeinschaft* through the program 'Forschungsgroßgeräte' (Project reference: INST 130/1063-1 FUGG). Following a European tender process, the system was purchased from *ClusterVision BV* and went into full operation on 6th March 2018.

2 System Details

2.1 General Description

The system comprises 170 standard compute nodes with varying mounts of memory, 12 GPGPU nodes, an administration node, and a parallel scratch file system all housed in 7 air-cooled racks. The various nodes and the scratch file system communicate via a fast, low-latency fabric.

2.2 Hardware

2.2.1 Compute Nodes

The compute nodes comprise 170× Intel Server System R1208WFTYS barebone systems[Int20a], each with the following specifications:

- 2× Intel Xeon Gold 6130 2.10 GHz, AVX-512 intruction set extension [Int20c]
- 8× Intel Omni-Path Host Fabric Adapters 100 Series [Int20b]
- 96–786 GB Samsung DDR4 2666 MHz ECC RAM

There are four classes of node with varying amounts of RAM, as shown in Table 1.

Table 1: Node Classes

Node Class	RAM	# of nodes
CPU-Bound	96 GB	102
CPU-Bound Mem+	192 GB	50
Memory-Bound	384 GB	14
Memory-Bound Mem+	768 GB	4

2.2.2 GPU Nodes

The GPU nodes are configured as follows:

- 12× ASUS ESCS400-G4 2U 4-GPU server barebone [ASU20]

each with the following specifications:

- 2× Intel Xeon Gold 6130 2.10 GHz, AVX-512 intruction set extension[Int20c]
- 12× 8 GB Samsung DDR4 2666 MHz ECC RAM
- 2× NVIDIA GeForce GTX 1080 Ti GPU [nvi20]

2.2.3 Network

The high-performance network comprises an Omni-Path fabric. The topology, shown in Figure 1, is such that each of two core-switches is connected to six edge-switches, with each node being connected to a single edge-switch.

2.2.4 Storage

Storage for temporary files is provide in the form of a 1.3 PB Lustre[Lus20] file system running over ZFS. The system comprises 2 metadata servers, 7 object data servers (ODS) with 6 object data targets per ODS.

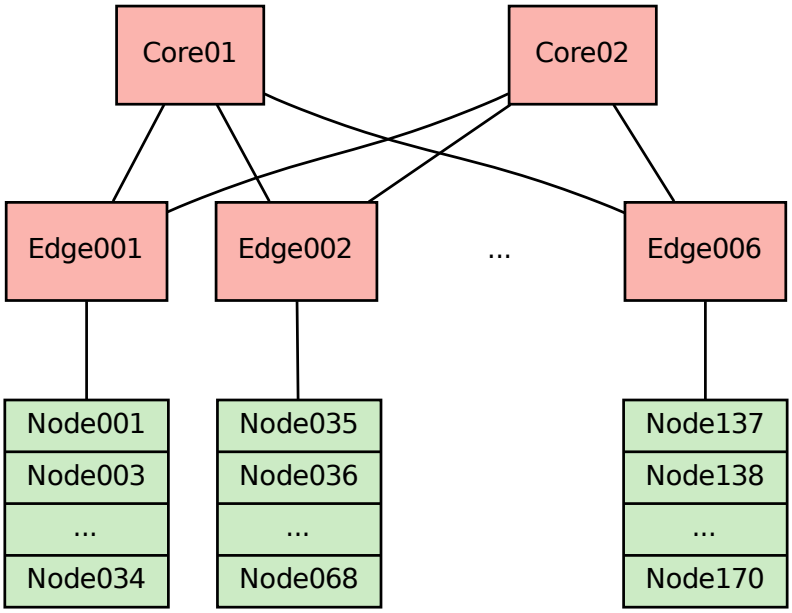

Figure 1: Cluster Topology (switchs in red, nodes in green)

2.3 Software

2.3.1 System Software

The following software stack is installed on Curta:

- CentOS 7 Linux operating system [Cen20]
- Slurm resource manager [Slu20]
- Zabbix monitoring tool [Zab20]
- TrinityX cluster manager [Tri20]

2.3.2 Scientific Software

Scientific software is provided for a wide range of disciplines. To facilitate installation and enable packages to be built in a reproducible way, the framework *EasyBuild* [Eas20] is used. In addition this enables users of the system to be able to refer not only to the version of a piece of software, but also to the versions of the compiler and any libraries used to build the software.

3 Access

All members of the university who are members of an active research group can apply for access to the system. Accounts are created by the central Identity Management service of the university but sorted into various groups by the HPC team using the tool Grouper [Gro20]. The groups are the following:

- primary groups
  - the main research group a user is affiliated to
- secondary groups
  - other research groups a user is affiliated to
  - software packages with restricted access

From this information the Identity Management service generates a dedicated LDAP branch for containing the HPC users and groups. Users then log in against a local copy of this branch is this branch.

References

[ASU20] ASUS. *Product Specification ASUS ESC4000 G4 Server*. 2020. URL: `https://www.asus.com/Commercial-Servers-Workstations/ESC4000-G4/` (visited on 2020-02-10).

[Cen20] CentOS. *The CentOS Project*. 2020. URL: `https://www.centos.org/` (visited on 2020-02-12).

[Eas20] EasyBuild. *Software Build Framework*. 2020. URL: https://easybuilders.github.io/easybuild/ (visited on 2020-02-12).

[Gro20] Grouper. *Access Management System*. 2020. URL: https://www.internet2.edu/products-services/trust-identity/grouper/ (visited on 2020-02-10).

[Int20a] Intel. *Product Specification Intel® Server System R1208WFTYS*. 2020. URL: https://ark.intel.com/products/89011/Intel-Server-System-R1208WFTYS (visited on 2020-02-10).

[Int20b] Intel. *Technical Specification Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x8*. 2020. URL: https://www.intel.com/content/www/us/en/products/network-io/high-performance-fabrics/omni-path-host-fabric-interface-adapters/1-port-pcie-x8-100-series.html (visited on 2020-01-08).

[Int20c] Intel. *Technical Specification Intel® Xeon Gold 6130 Processor*. 2020. URL: https://www.intel.com/content/www/us/en/products/processors/xeon/scalable/gold-processors/gold-6130.html (visited on 2020-01-07).

[Lus20] Lustre. *Open Source Parallel File System*. 2020. URL: http://lustre.org/ (visited on 2020-03-04).

[nvi20] nvidia. *Product Specification GeForce GTX 1080 Ti*. 2020. URL: https://www.geforce.com/hardware/desktop-gpus/geforce-gtx-1080-ti/specifications (visited on 2020-01-07).

[Slu20] Slurm. *Slurm Workload Manager*. 2020. URL: https://slurm.schedmd.com/ (visited on 2020-02-12).

[Tri20] TrinityX. *HPC Management System*. 2020. URL: https://trinityx.eu/ (visited on 2020-02-12).

[Wik20] Wikipedia. *Curta*. 2020. URL: https://en.wikipedia.org/wiki/Curta (visited on 2020-02-12).

[Zab20] Zabbix. *Zabbix Monitoring Solution*. 2020. URL: https://www.zabbix.com/ (visited on 2020-02-12).