

# High Performance Computing



Figure 1: Ethernet and Omni-Path connections at the backside of node 38 and 40 inside rack 8.

## Introduction

High Performance Computing (HPC) is the discipline of running scientific simulations and calculations on large-scale computer clusters. Ranging from particle science to genome sequencing and economics, there is a wide variety of disciplines using HPC systems. All have in common an extensive demand for computational power and memory that vastly exceeds the resources of conventional computers and workstations. For this purpose the ZIV is hosting the PALMA High Performance Compute Cluster. In its second generation PALMA-II went online in 2018 and with a computational performance of around 800 TFlop/s made it into the *Top500* of the world's fastest supercomputers.

## Hardware



|                      |  |
|----------------------|--|
| Manufacturer:        | MEGWARE                                    |
| Cores:               | 15,120                                     |
| Memory:              | 72,384 GB                                  |
| Nodes:               | 412  |
| Processor:           | Intel Xeon Gold 6140 18C@2.30GHz (Skylake) |
| Interconnect:        | 100 Gbit/s Intel Omni-Path                 |
| Storage:             | 1 PB (GPFS)                                |
| Linpack Performance: | Rmax 800 TFlop/s                           |

### GPUs Legacy

|               |                                  |
|---------------|----------------------------------|
| 8x RTX 2080Ti | 44x Broadwell Nodes (E5-2683 v4) |
| 4x Tesla V100 | 4x Xeon Phi KNL Nodes            |
| 8x TitanXP    |                                  |
| 12 Tesla K20  |                                  |



Figure 2: Open backside of rack 8.

## SLURM Job Scheduler

In order to coordinate thousands of jobs on hundreds of nodes, a capable job scheduling system is necessary. On PALMA, we use Slurm for this task. Users cannot start computations directly on the nodes, but will have to write so-called *job scripts* to communicate their requirements.

An example for a batch script:

```
#!/bin/bash

#SBATCH --export=NONE           # Start with a clean environment
#SBATCH --nodes=1              # the number of reserved nodes
#SBATCH --ntasks-per-node=72   # the number of CPU cores per node
#SBATCH --mem=64G              # how much memory is needed per node (units can be: K, M, G, T)
#SBATCH --partition=normal     # which partition to use
#SBATCH --time=24:00:00        # the max wallclock time (time limit your job will run)
#SBATCH --job-name=MyJob123    # the name of your job
#SBATCH --output=output.dat    # the file where output is written to (stdout & stderr)
#SBATCH --mail-type=ALL        # receive an email when your job starts, finishes normally or is aborted

#SBATCH --mail-user=your_account@uni-muenster.de # your mail address

# START THE APPLICATION
srun /path/to/my/application
```

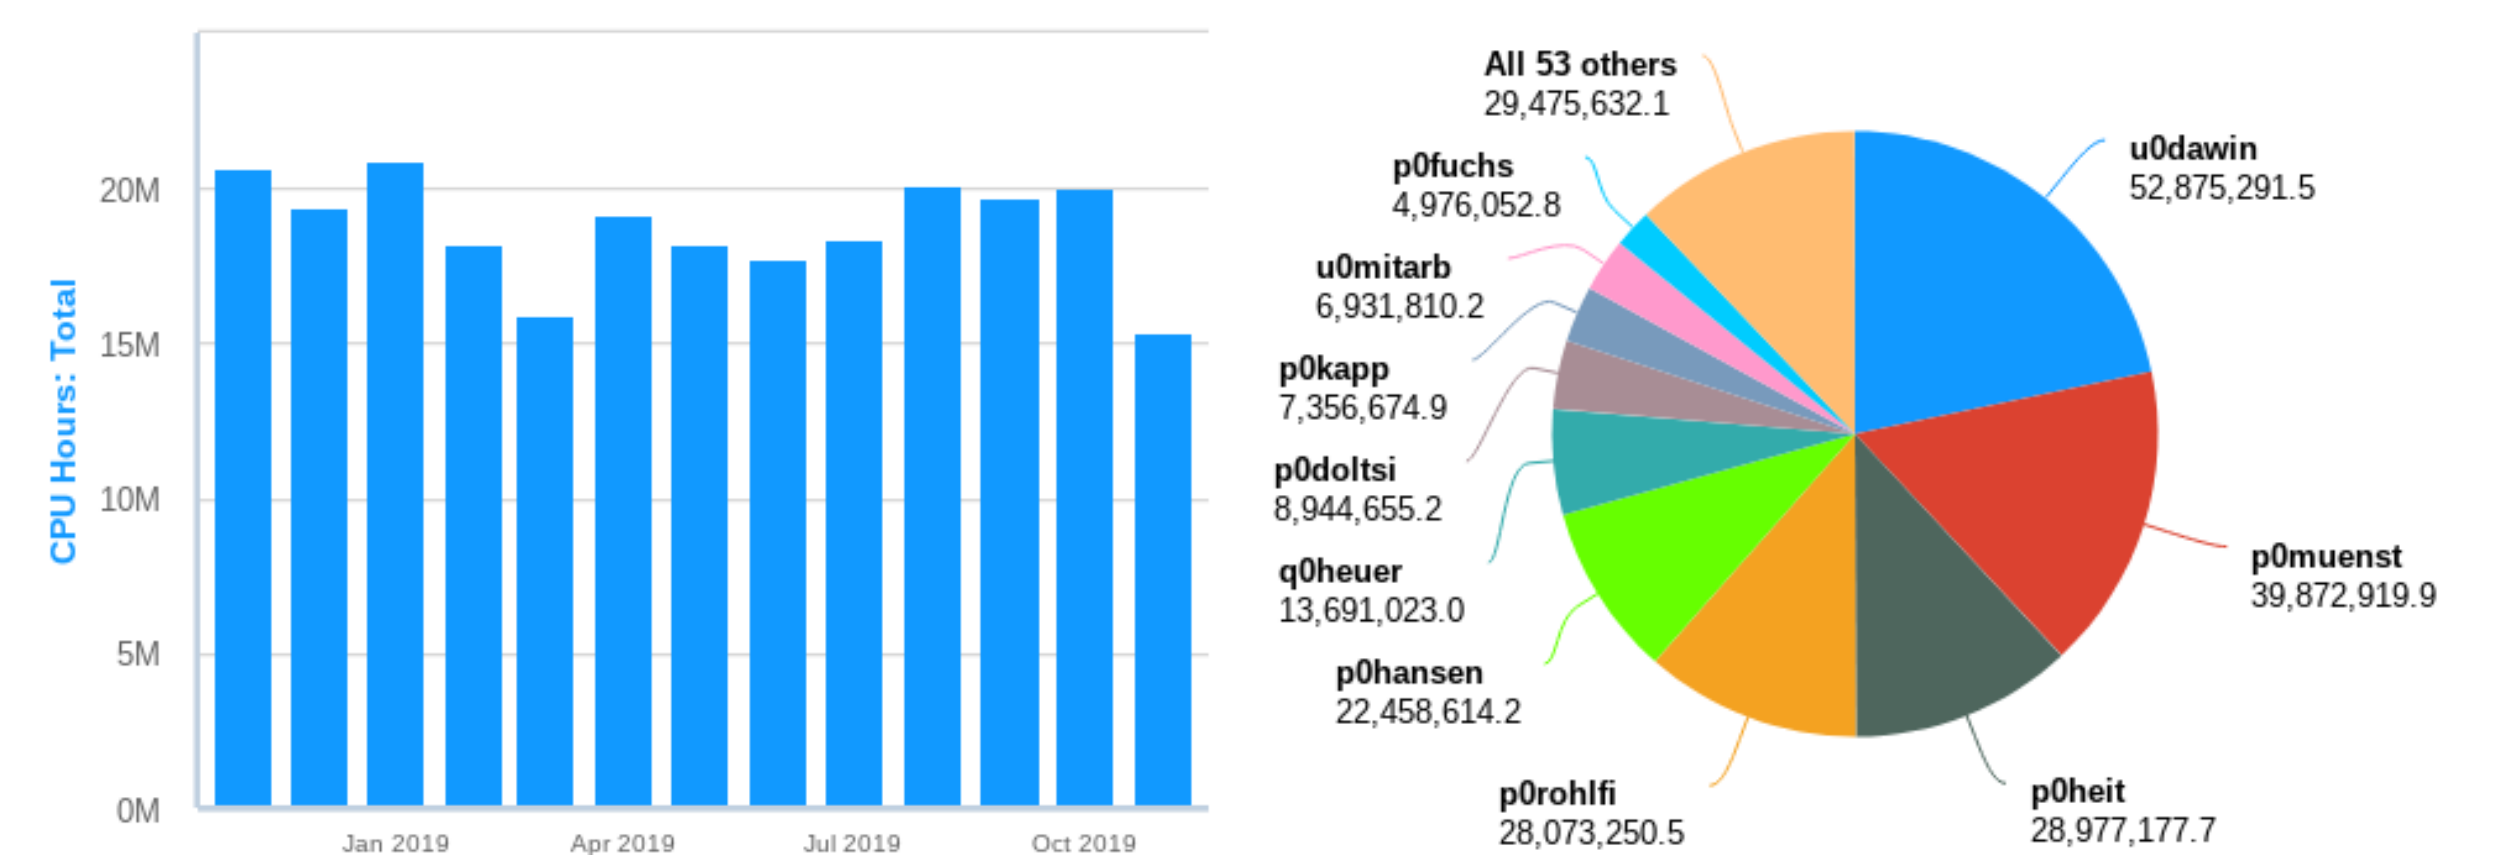
## Providing Scientific Software

Software on PALMA can be accessed via *environment modules*. These are small scripts that set environment variables (like PATH and LD\_LIBRARY\_PATH) pointing to the locations where the software is installed. The tool to manage environment modules is called **LMOD**. On PALMA a so called *hierarchical module naming scheme* is employed in combination with the **Easybuild** framework to deploy scientific software. In doing so, we can offer up to date applications for the users and at the same time ensure reproducibility of results.

At this point we are providing software inside of two different *toolchains*: The *free and open source software (foss)* toolchain consists of the GCC compiler, the OpenMPI MPI library as well as OpenBLAS, ScaLAPACK and FFTW for numerical tasks. On the other hand we provide the proprietary *intel* toolchain including among others the intel compiler, MPI library and math kernel library (MKL). Most available libraries and applications are being built within these toolchains.



## Usage Statistics



PALMA can deliver about 26 million CPU core hours per month. Most of these resources are being filled via the batch system. There are around 140 active users with more than 50% coming from the departments of physics and chemistry. Almost 200,000 jobs are started on the cluster each month.

## Gaining Access

All scientific researcher as well as students who are member of a registered project group are eligible to use the HPC resources. To gain access, you have to perform the following steps:

- Apply for a project at PALMA with your research group
- Register for the group **u0clstr** at MeinZIV
- Change your main group according to your work group at MeinZIV
- Register at our HPC mailing list: <https://listserv.uni-muenster.de/mailman/listinfo/hpc-l>
- Study the online documentation!




Figure 3: Closeup of one of the skylake nodes of PALMA-II.

## References

Find the documentation at:  
<https://zivconfluence.uni-muenster.de/display/HPC>  
Monitoring:  
<https://palma.uni-muenster.de/ganglia>  
Contact: [hpc@uni-muenster.de](mailto:hpc@uni-muenster.de)

