




PRACE Workshop: OpenMP Programming Workshop @ LRZ

LRZ | 11-13 Feb 2020



Overview



- The course is organized as a PRACE training event by LRZ in collaboration with Appentra Solutions, Intel and RWTH Aachen.
- Guided SuperMUC-NG tour on Tuesday, 11 Feb 2020, 17:00-18:00
- Social Event on Wednesday, 12 Feb 2020, 19:00, <http://gasthof-neuwirt.org/> , tbc.
- **Lecturers:**
 - Dr. Manuel Arenaz, CEO at Appentra Solutions and professor of computer science at the University of A Coruña (Spain)
 - Dr. Michael Klemm, Principal Engineer in the Compute Ecosystem Engineering organization of the Intel Architecture, Graphics, and Software group at Intel and CEO of the OpenMP Architecture Review Board (ARB)
 - Dr. Christian Terboven, leader of the HPC group at RWTH Aachen University and co-author of the new book “Using OpenMP – The Next Step“, <https://www.openmp.org/tech/using-openmp-next-step/>
- **Assistants:**
 - Dr. Michele Martone (LRZ)
 - Dr. Volker Weinberg (LRZ)

PRACE Training



LRZ as part of the Gauss Centre for Supercomputing (GCS) and IT4Innovations belong to the 14 **PRACE Training Centres** that started in 2012-2017-2020:

- Barcelona Supercomputing Center (Spain)
- CINECA Consorzio Interuniversitario (Italy)
- CSC – IT Center for Science Ltd (Finland)
- EPCC at the University of Edinburgh (UK)
- Gauss Centre for Supercomputing (Germany)
- Maison de la Simulation (France)
- GRNET – Greek Research and Technology Network (Greece)
- ICHEC – Irish Centre for High-End Computing (Ireland)
- IT4I – National Supercomputing Center VSB Technical University of Ostrava (Czech Republic)
- SURFsara (The Netherlands)
- TU Wien – VSC Research Center (Austria)
- University ANTWERPEN – VSC & CÉCI (Belgium)
- University of Ljubljana – HPC Center Slovenia (Slovenia)
- Swedish National Infrastructure for Computing (SNIC) (Sweden)



Mission: Serve as **European hubs and key drivers of advanced high-quality training** for researchers working in the computational sciences.

<http://www.training.prace-ri.eu/>

Tentative Agenda



	Day 1	Day 2	Day 3
09:00-10:30	Introduction to the OpenMP common core	Tasking	Tools for Performance and Correctness
10:30-10:45	Coffee Break	Coffee Break	Coffee Break
10:45-12:00	Decomposing code into patterns for parallelization	Tasking	Offloading to Accelerators
12:00-13:00	Lunch Break	Lunch Break	Lunch Break
13:00-14:45	Beyond OpenMP common core with tasking and offloading	Host Performance: NUMA	Other Advanced Features of OpenMP 5.0
14:45-15:00	Coffee Break	Coffee Break	Coffee Break
15:00-17:00	Hands-on time with Parallelware Trainer	Host Performance: SIMD	Roadmap / Outlook (until 16:00)
	17:00-18:00 Guided SuperMUC-NG Tour	19:00 Social Event (tbc.)	

Training Setup



- All Material is available under:
- <https://tinyurl.com/OpenMP-Workshop>
- `lxlogin10.lrz.de:/lrz/sys/courses/openmp/Appentra`

- **General Information about the Linux-Cluster:**
- <https://doku.lrz.de/display/PUBLIC/Linux+Cluster>

- **Login to the system:**
- `ssh -X lxlogin10.lrz.de -l hpckurs??`

- **Load pwtrainer module on login node:**
- `module use -a /lrz/sys/share/modules/extfiles`
- `module load pwtrainer`

- **Run pwtrainer on login node:**
- `pwtrainer`

- **Run a program on the compute nodes:**
- `srun --reservation=homp1w19_course program`

First self-assembled Linux cluster (1999-2002)



Cluster components (2012)



SGI UltraViolet with air guides in front to improve cooling efficiency (2012)



LRZ Linux Cluster Overview




- The LRZ Linux Cluster consists of several segments with different types of interconnect and different sizes of shared memory. All systems have a (virtual) 64 bit address space:
 - **CooLMUC2** Cluster with 28-way Haswell-based nodes and FDR14 Infiniband interconnect, used for both serial and parallel processing
 - Intel Broadwell based 6 TByte shared memory server HP DL580 "**Teramem**"
 - **CooLMUC3** Cluster with 64-way KNL 7210-F many-core processors and Intel Omnipath OPA1 interconnect, for parallel/vector processing
 - **IvyMUC** Cluster with 8-way Ivy Bridge-based nodes and FDR14 Infiniband interconnect, used for parallel processing
- Based on the various node types the LRZ Linux cluster offers a wide span of capabilities:
 - mixed shared and distributed memory
 - large software portfolio
 - flexible usage due to various available memory sizes
 - parallelization by message passing (MPI)
 - shared memory parallelization with OpenMP or pthreads
 - mixed (hybrid) programming with MPI and OpenMP

CoolMUC-2 (2015): The six racks to the left




CoolMUC-3 (2017)





IvyMUC = SuperMIC - XeonPhi



Hardware	
Processor	2 Ivy-Bridge (2 x 8 cores) host processors E5-2650
Rack	1 iDataPlex rack with 32 nodes dx360 M4
Number of nodes	32
Cores per node	16
Hyperthreads per core	2
Core nominal frequency	2.6 GHz
Memory (DDR4) per node	64 GB (Bandbreite 80.8 GB/s)
	Mellanox Infiniband FDR14
Software (OS and development environment)	
Operating system	SLES12 SP2 Linux
MPI	Intel MPI 2017, alternatively OpenMPI
Compilers	Intel icc, icpc, ifort 2017
Performance libraries	MKL, TBB, IPP
Tools for performance and correctness analysis	Intel Cluster Tools

Linux Cluster Access



```
ssh -Y lxlogin5.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin6.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin7.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin8.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin10.lrz.de -l xxyyyzz
```

```
gsissh -Y lxgt2.lrz.de
```

Haswell (CoolMUC-2) login node

Haswell (CoolMUC-2) login node

Haswell (CoolMUC-2) login node

KNL Segment (CooMUC-3) login node

Ivy Bridge (IvyMUC) login node

login node for [GSI-SSH](#)

Interactive Jobs on IvyMUC



- **Submit a job:**

```
sbatch --reservation=homp1w19_course job.sh
```

- **List own jobs:**

```
squeue -M ivymuc -u hpckurs??
```

- **Cancel jobs:**

```
scancel jobid
```

- **Interactive Access (not recommended for this workshop):**

- `module load salloc_conf/ivymuc`

- `salloc --nodes=1 --time=02:00:00 --reservation=homp1w19_course`

```
srun --reservation=homp1w19_course --pty bash
```

IvyMUC SLURM Simplest Batch File



```
#!/bin/bash
#SBATCH -o /home/hpc/pr28fa/lu65fok/ivymuc.%j.%N.out
#SBATCH -D /home/hpc/pr28fa/lu65fok/
#SBATCH -J ivytest
#SBATCH --clusters=ivymuc
#SBATCH --nodes=1
#SBATCH --get-user-env
#SBATCH --reservation=homp1w19_course
#SBATCH --time=02:00:00

source /etc/profile.d/modules.sh
cd mydir
export OMP_NUM_THREADS=16
./myprog.exe
```


Further Documentation



- New LRZ Documentation server: <https://doku.lrz.de/display/PUBLIC/Linux+Cluster>
- [Access and Login to the Linux-Cluster](#)
- [Job Processing on the Linux-Cluster](#)
 - [Available SLURM clusters and features](#)
 - [Running parallel jobs on the Linux-Cluster](#)
 - [Example parallel job scripts on the Linux-Cluster](#)
 - [Resource limits for parallel jobs on Linux Cluster](#)
 - [Running serial jobs on the Linux-Cluster](#)
 - [Example serial job scripts on the Linux-Cluster](#)
 - [Resource limits for serial jobs on Linux Cluster](#)
- [File Systems and I/O on Linux-Cluster](#)
- [Linux Cluster Segments](#)
 - [Big Data Teramem](#)
 - [CoolMUC-2](#)
 - [CoolMUC-3](#)
 - [DGX-1](#)
 - [IvyMUC](#)
 - [RStudio Server Hardware](#)
- [Linux Cluster Usage Statistics](#)
- [Software for HPC](#)
- [Servicedesk for the Linux-Cluster](#)

And now ...



Enjoy the course!