



VI-HPS



40th VI-HPS Tuning Workshop

14 – 18 June 2021 | LRZ

PRACE Training Centres



LRZ as part of the Gauss Centre for Supercomputing (GCS) belongs 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/>

ZOOM Netiquette

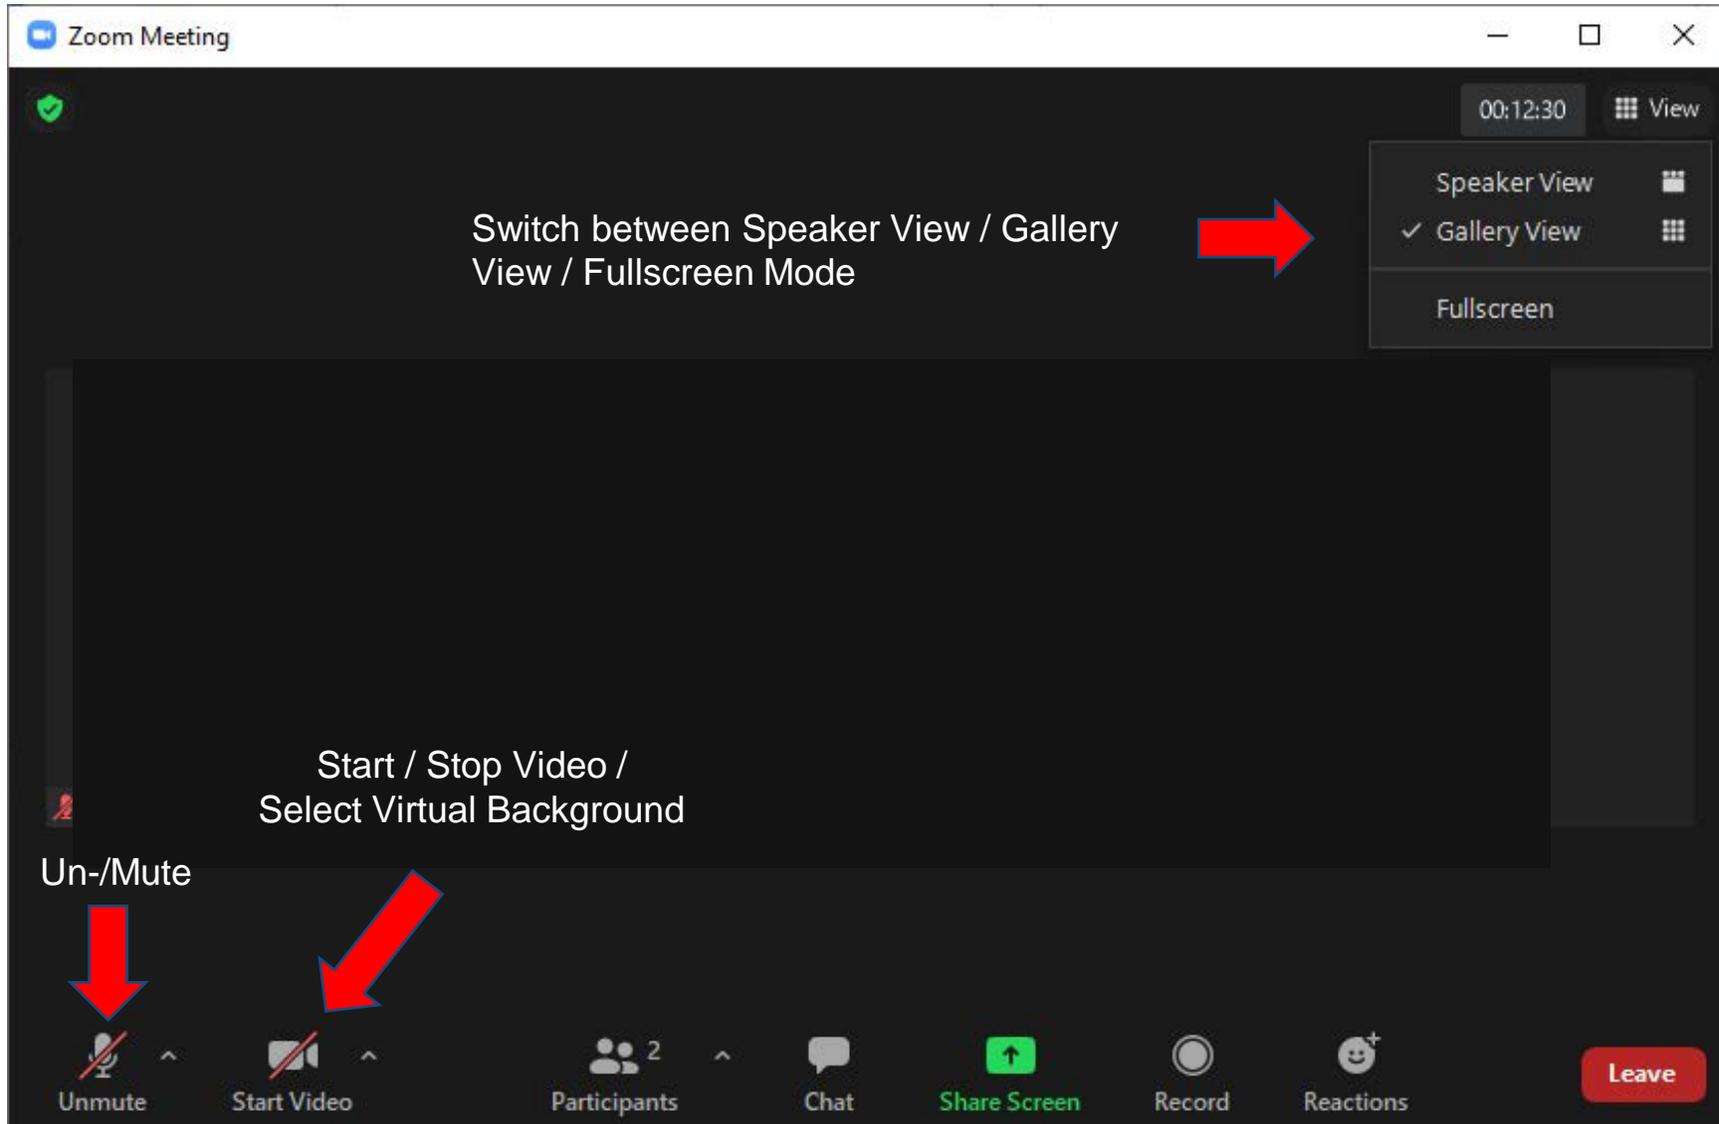
- We will use the **same link** for the complete SuperMUC-NG Status and Results Workshop:

<https://lrz-de.zoom.us/j/97159504765?pwd=MjYxWVNTd3FIZkJFcHBDdXpobnJKUT09>

Meeting ID: 971 5950 4765

Passcode: 527385

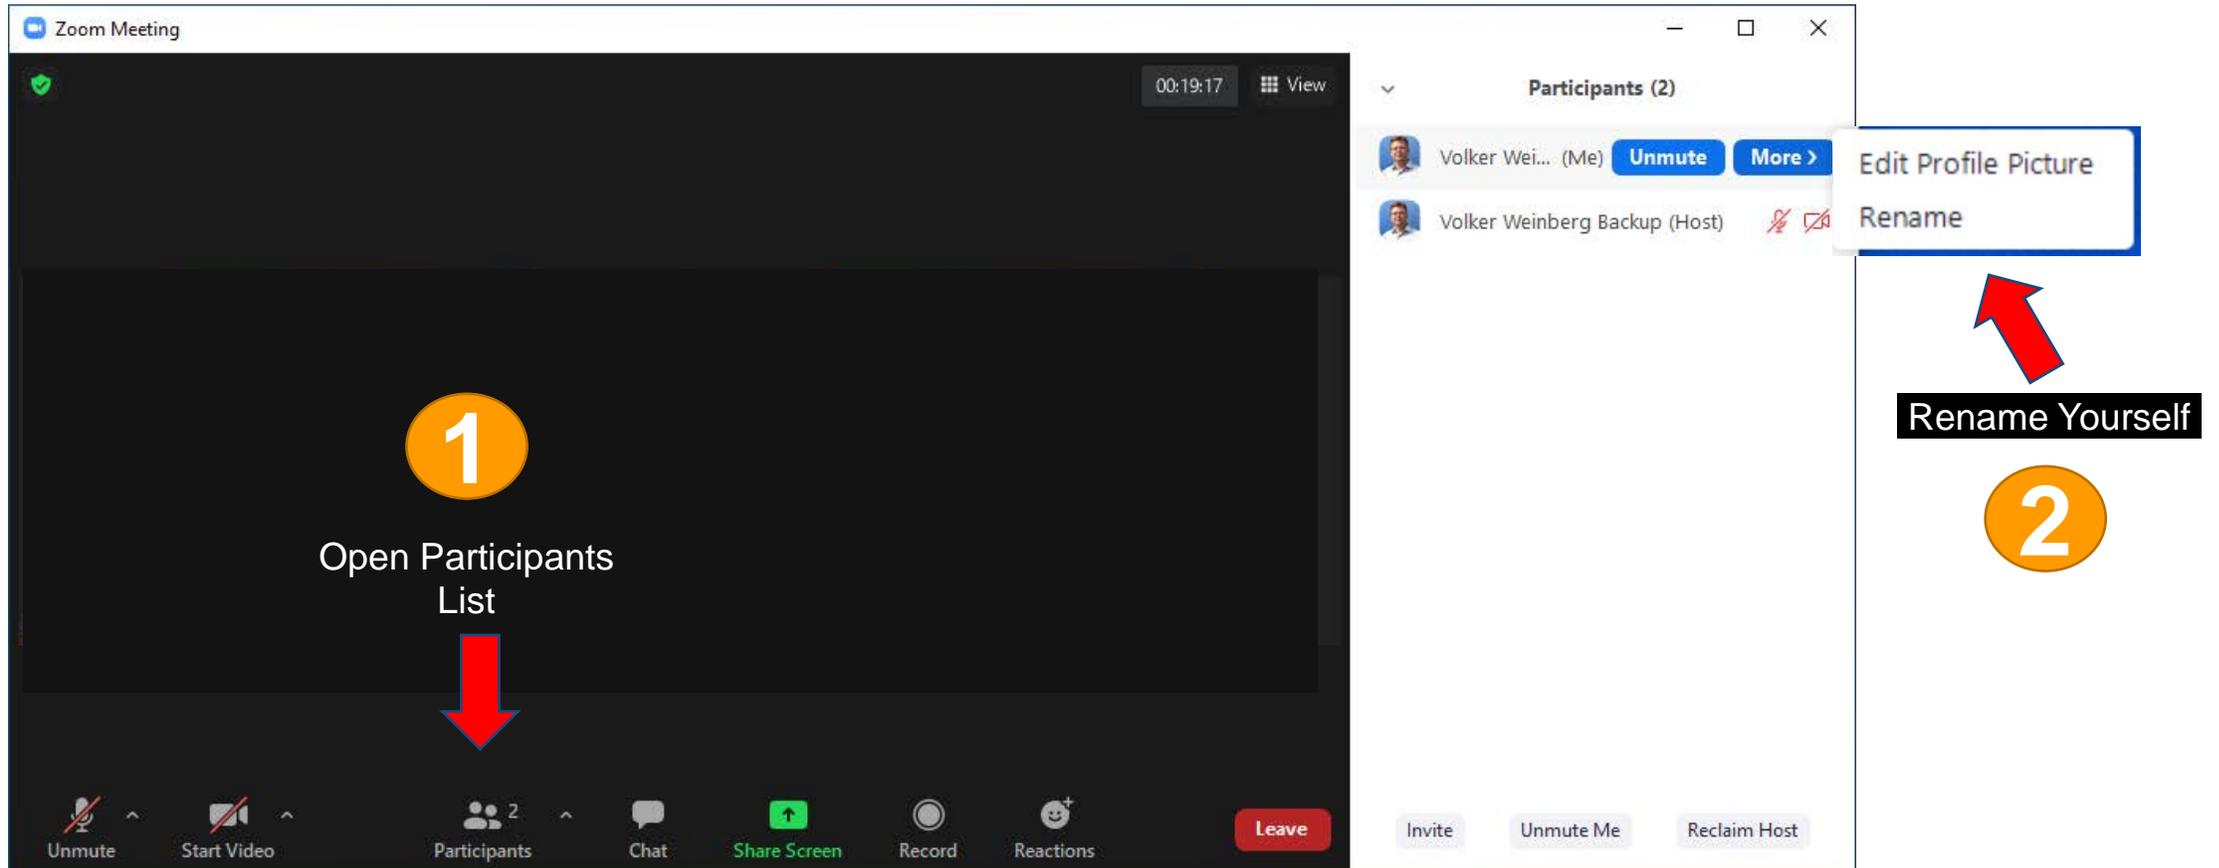
- To ensure a pleasant experience with Zoom Meeting, we encourage participants to **download and install the latest Zoom application** via <https://zoom.us/download>.
- If you have problems with your computer audio, you can also **join by phone**.
Find your local number: <https://lrz-de.zoom.us/j/97159504765?pwd=MjYxWVNTd3FIZkJFcHBDdXpobnJKUT09>



zoom

Participants List

Kindly use “first-name family-name (institute)” as your screenname.

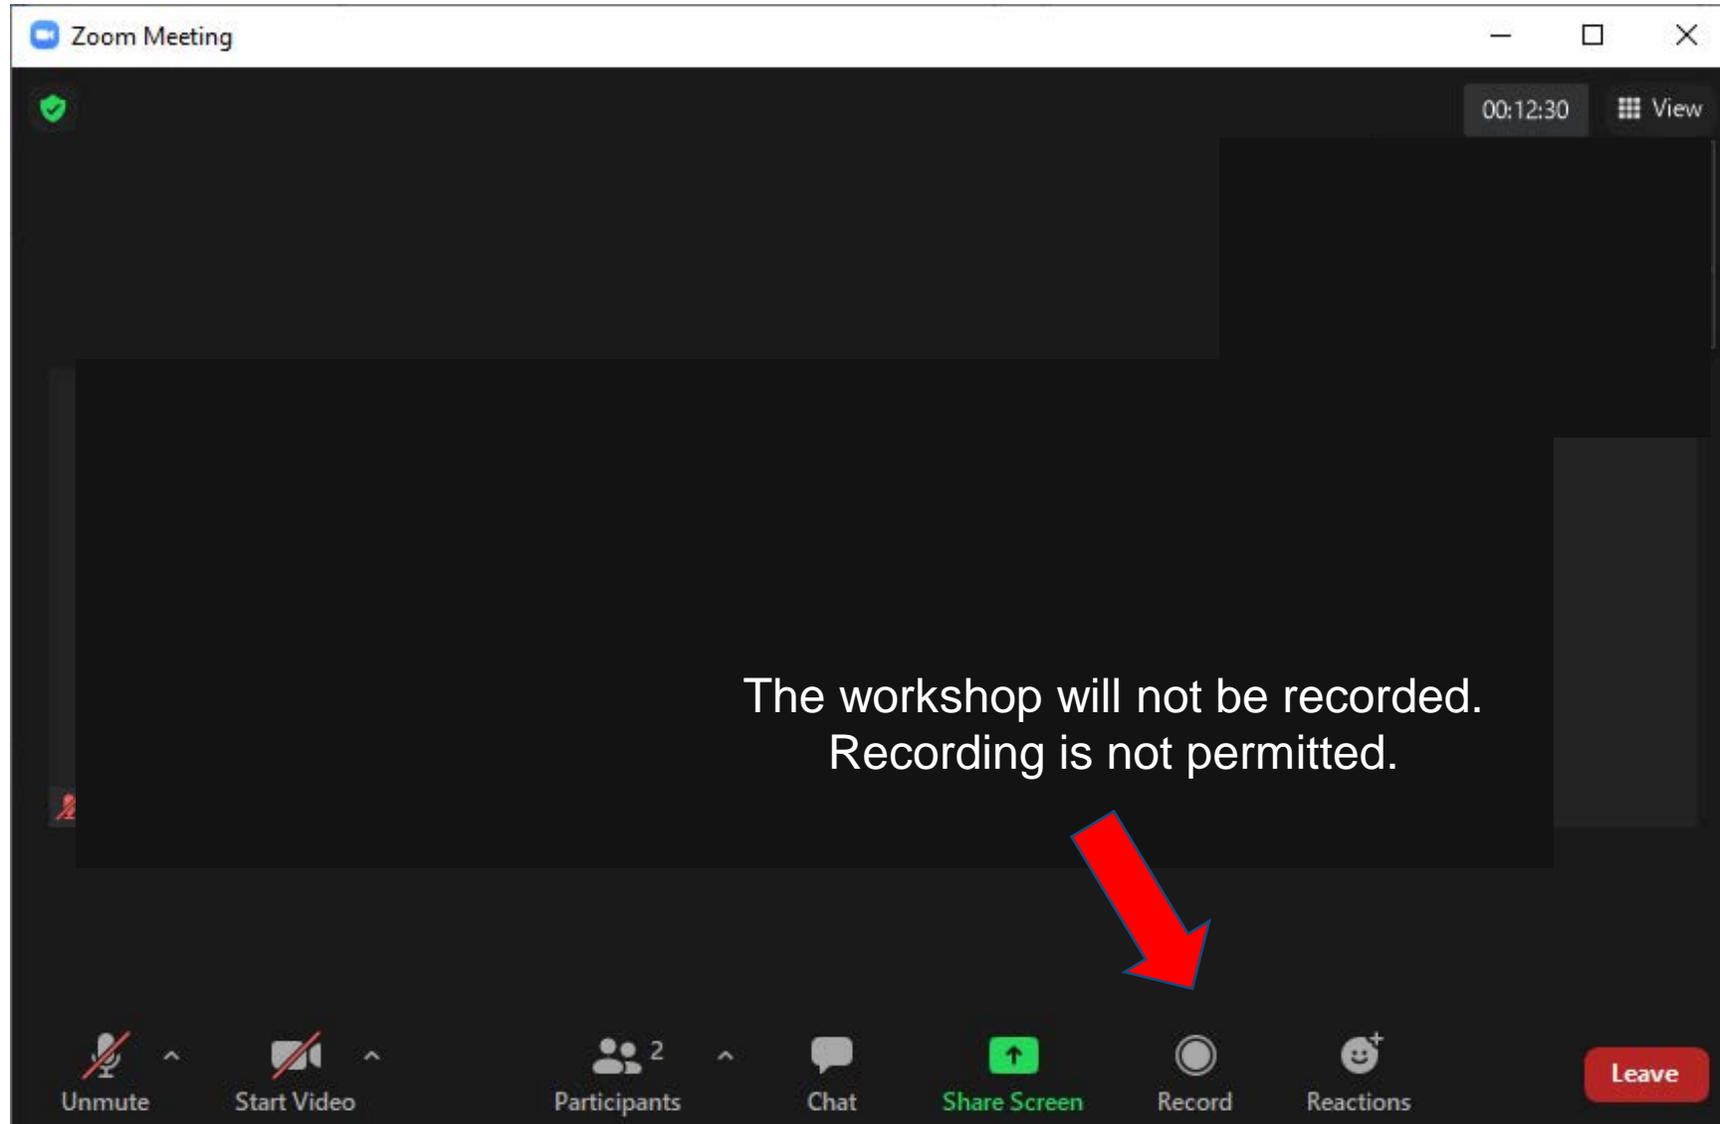


The screenshot shows a Zoom Meeting window with a 'Participants (2)' list on the right. The list contains two participants: 'Volker Wei... (Me)' and 'Volker Weinberg Backup (Host)'. A context menu is open over the 'Me' participant, showing options: 'Edit Profile Picture' and 'Rename'. A red arrow points from the 'Rename Yourself' text to the 'Rename' option. A large orange circle with the number '1' is positioned over the 'Participants' icon in the bottom toolbar, with the text 'Open Participants List' below it. A red arrow points from this text to the 'Participants' icon. A second large orange circle with the number '2' is positioned below the 'Rename Yourself' text.

1
Open Participants List

2
Rename Yourself

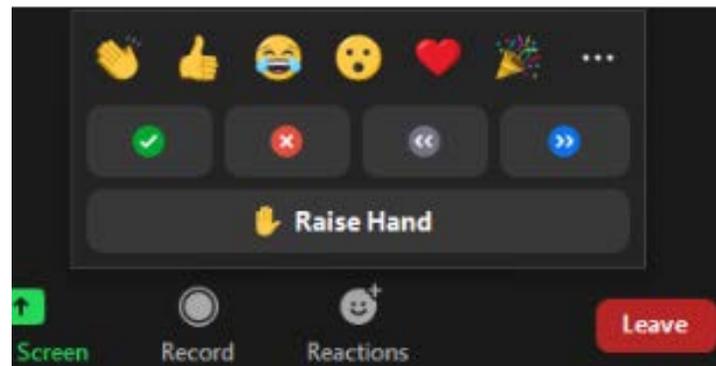
The image shows a Zoom Meeting window with a chat window open on the right. The Zoom Meeting window has a dark background and a bottom toolbar with icons for Unmute, Start Video, Participants (2), Chat, Share Screen, Record, Reactions, and Leave. A red arrow points from the Chat icon to the chat window. The chat window is titled 'Chat' and contains three messages: 'Lorem ipsum dolor sit amet.', 'From Volker Weinberg Ba... to Everyone: test', and 'From Me to Everyone: Lorem ipsum dolor sit amet,'. Below the messages is a dropdown menu for 'Who can see your messages?' set to 'Everyone'. A red arrow points from the text 'Type Messages Here' to the chat input field.

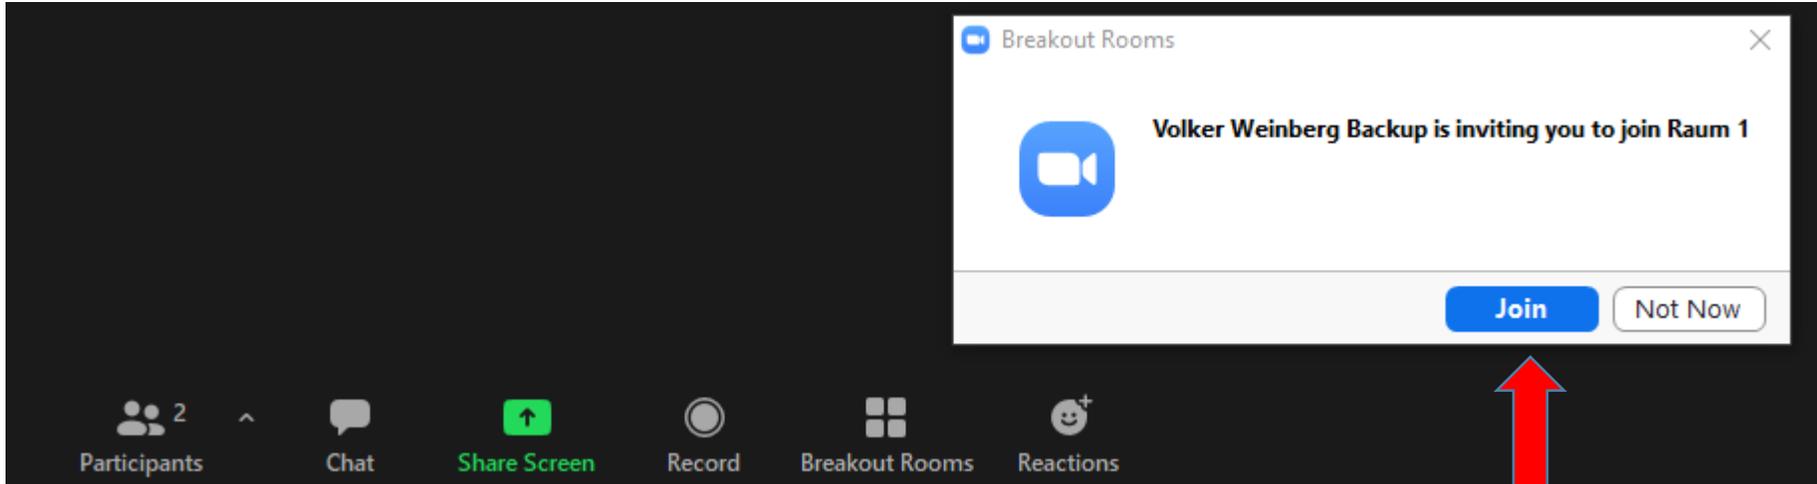


- Use **chat window** to ask questions **during the presentations**.
- You can also **raise your hand** if you have questions.
- If you do not mind, please **show your video when asking questions** to make this workshop as interactive as possible.

- **Push to Talk:** The Push to Talk feature allows you to remain muted throughout the Zoom meeting and only if you hold down the spacebar you will be unmuted.

- **Instant Feedback:**

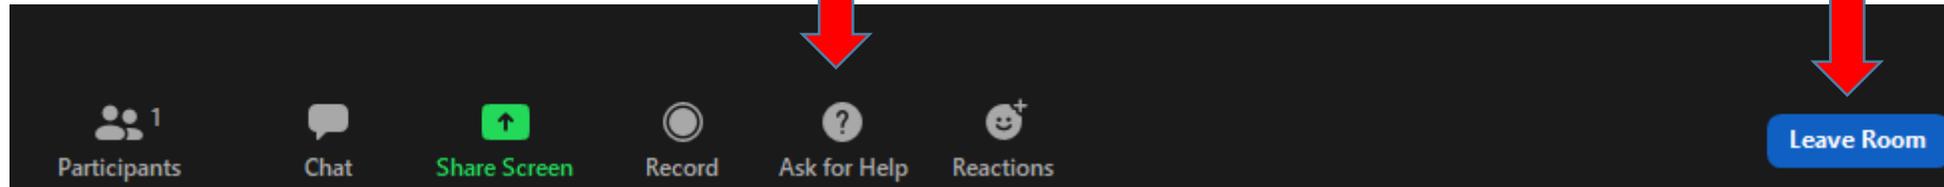


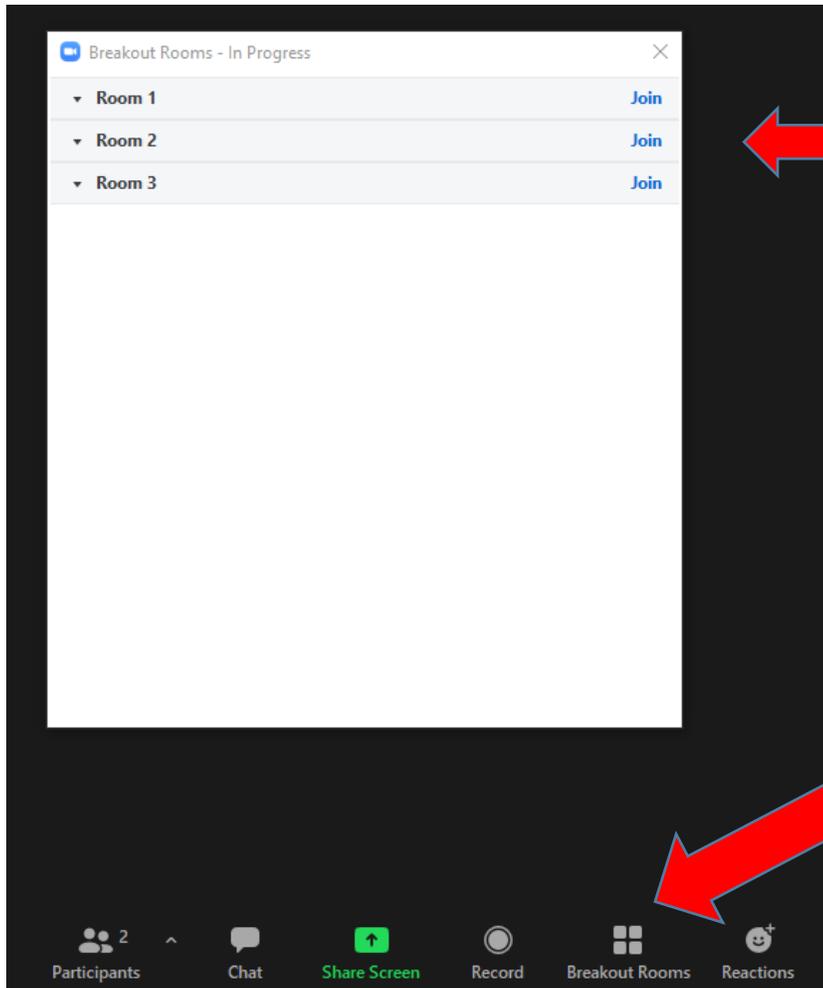


Join Breakout Room

Invite the Host to the Breakout Room

Leave Breakout Room





Join Breakout Room

2

Show Breakout Rooms

1

Using the LRZ Linux Cluster / IvyMUC

- 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
 - **secure shell** based logins and data transfer to generally accessible front end nodes
 - **development environment** with compilers, tools and libraries available on front end nodes, run time environments and applications available on batch nodes. Necessary licenses are supplied by LRZ.
 - resource assignment via **SLURM scheduler**
 - **data management:**
 - SCRATCH space for short lifetime data (removal is forced)
 - DSS/HOME area with small quota for program and configuration data
 - DSS/PROJECT area (max. 10 TByte) upon request for long lifetime data

First self-assembled Linux cluster (1999-2002)



Cluster components (2012)



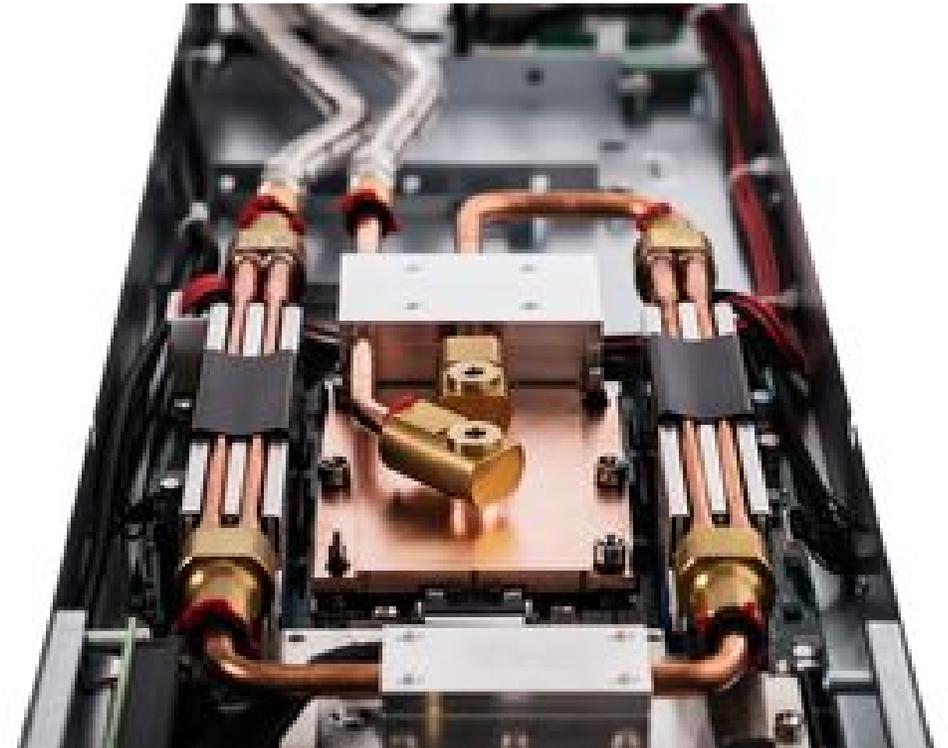
SGI UltraViolet with air guides in front to improve cooling efficiency (2012)



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 SP5 Linux
MPI	Intel MPI 2019
Compilers	Intel icc, icpc, ifort 19.0
Performance libraries	MKL, TBB, IPP
Tools for performance and correctness analysis	Intel Cluster Tools

```
ssh -Y lxlogin1.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin2.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin3.lrz.de -l xxyyyzz
```

```
ssh -Y lxlogin4.lrz.de -l xxyyyzz
```

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

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

Haswell (CoolMUC-2) login node

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

- **Submit a job:**
`sbatch --reservation=hhps1s21_workshop job.sh`
- **List own jobs:**
`squeue -M ivymuc -u hpckurs??`
- **Cancel jobs:**
`scancel jobid`

- **Interactive Access:**
 - `module load salloc_conf/ivymuc`
 - `salloc --nodes=1 --time=02:00:00 --reservation=hhps1s21_workshop --partition=ivymuc_batch`
 - **Or:** `srun --reservation=hhps1s21_workshop --pty bash`

IvyMUC SLURM OpenMP Batch File



→ /lrz/sys/courses/vihps/job-omp.sh

```
#!/bin/bash
#SBATCH -o /dss/dsshome1/0D/hpckurs10/ivymuc.%j.%N.out
#SBATCH -D /dss/dsshome1/0D/hpckurs10/
#SBATCH -J ivytest
#SBATCH --clusters=ivymuc
#SBATCH --nodes=1
#SBATCH --get-user-env
#SBATCH --reservation=hhps1s21_workshop
#SBATCH --time=02:00:00
module load slurm_setup
export OMP_NUM_THREADS=16
./myprog.exe
```

IvyMUC SLURM MPI Batch File



→ /lrz/sys/courses/vihps/job-mpi.sh

```
#!/bin/bash
```

```
#SBATCH -o /dss/dsshome1/0D/hpckurs10/ivymuc.%j.%N.out
```

```
#SBATCH -D/dss/dsshome1/0D/hpckurs10
```

```
#SBATCH -J ivytest
```

```
#SBATCH --clusters=ivymuc
```

```
#SBATCH --nodes=2
```

```
#SBATCH --ntasks-per-node=16
```

```
#SBATCH --get-user-env
```

```
#SBATCH --reservation=hhps1s21_workshop
```

```
#SBATCH --time=02:00:00
```

```
module load slurm_setup
```

```
mpiexec -n $SLURM_NTASKS ./myprog.exe
```

Further Documentation for the Workshop



<https://tinyurl.com/vihps>

And now enjoy the workshop!