CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

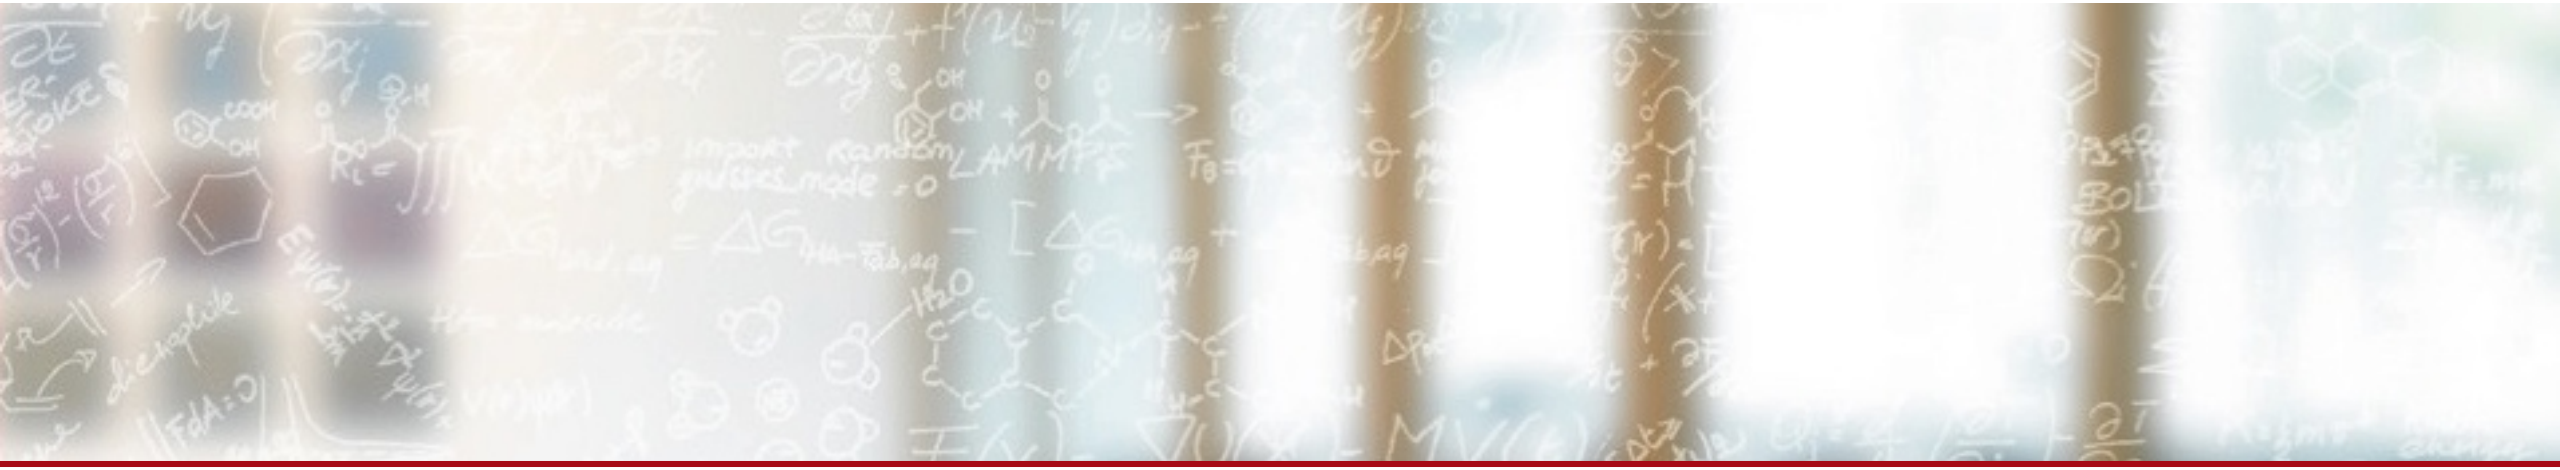

The Swiss National Supercomputing Centre (CSCS): Future developments for scientific computing

SWISS ICT AWARD

Guilherme Peretti-Pezzi

Swiss SKA Days 2019
Bern, 19.06.2019

The Swiss National Supercomputing Centre

Driving innovation in computational research in Switzerland

- Established in 1991 as a unit of ETH Zurich, Located in Lugano
- 114 highly qualified staff from 23 nations
- **Develops and operates the key supercomputing capabilities required to solve important problems to science and/or society**
- Leads the national strategy for **High-Performance Computing and Networking** (HPCN) that was initially passed by Swiss Parliament in 2009
- Has a dedicated **User Laboratory** for supercomputing since 2011 with ~1600 users and 130 projects
- Hosting Member of the Partnership for Advanced Computing in Europe

What happened at CSCS in the last 10 years

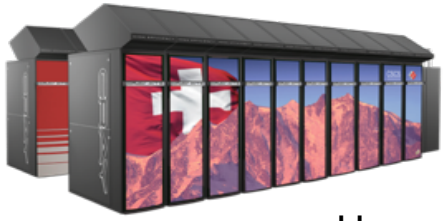

Monte Rosa
Cray XT5
14'762 cores

Hex-core upgrade
22'128 cores

Upgrade to
Cray XE6
47,200 cores

2009

2010

2011

2012

2013

2014

2015

2016

2017

Phase II:
K20X based hybrid

Phase III:
Pascal based hybrid

Upgrade

Development & procurement of petaflop/s scale supercomputer

New building complete

Begin construction of new building

Three pronged approach of the HPCN Initiative
1. New, flexible, and efficient building
2. Efficient supercomputers
3. Efficient applications

1st Chinese EF-system

1st USA EF-system

Sustained EF-system (USA)

Upgrade Piz Daint
26+8 cabinets

Commitment for PRACE 2 ends

Final extension
28+10 cabinets

Install of test system

Installation of successor of Piz Daint

2016

2017

2018

2019

2020

2021

2022

2023

2024

Flexible infrastructure, since we can only predict that technology will change

Flexible Infrastructure:

- power/cooling: 12 MW,
- upgradable to 25 MW

Free cooling with water from Lugano lake

Current Power Usage Effectiveness (PUE) 1.2

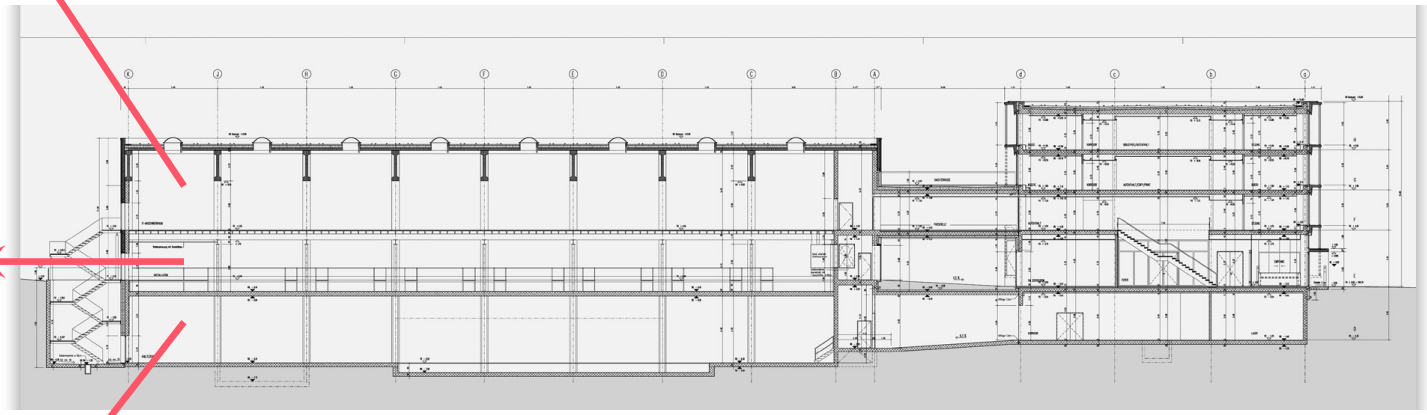

Customer List

Core Mission

peer reviewed access at no charge
about 120 projects and 1000 users

- User Lab (two calls per year)
- PRACE Tier 0 (two calls per year)

Housing

- BlueBrain for EPFL
- Euler for ETH Zurich

Hosting (dedicated services)

- MeteoSwiss
- LHC Tier-2 for CHIPP
- Paul Scherrer Institute (PetaByte archive)

Services on non dedicated systems

- Empa & Eawag
- ETH Zurich
  - C2SM (storage)
  - Prof. Petros Koumoutsakos
  - Prof. Nicola Spaldin
  - Prof. Matthias Troyer
  - Institute for theoretical physics
- Hilti
- HTW Chur
- NCCR MARVEL
- PartnerRe
- Paul Scherrer Institute
- Swiss Data Science Centre
- Università della Svizzera italiana
- University of Geneva / CADMOS
- University of Zurich

Service Catalog – Published on www.cscs.ch/services

COMPUTE RESOURCES

MultiCore (@Piz Daint)

Computing node-hours on the multicore partition (Cray XC40) of the Piz Daint supercomputer. This service includes access to the high-performance Lustre scratch file system.

→ **Info**

GPU Hybrid (@Piz Daint)

Computing node-hours on a CPU/GPU hybrid partition (Cray XC50) of the Piz Daint supercomputer. This service includes access to the high-performance Lustre scratch file system.

→ **Info**

Manycore (@Grand Tavè)

STORAGE RESOURCES

Home Directory

Every user is assigned their own home directory with sufficient storage space to store codes and documents (/home)

→ **Info**

Online Storage

A shared parallel file system to store project data (/store or /project)

→ **Info**

OTHER ACCESS TOOLS

Data Mover

Bidirectional data transfer services between

INTERACTION MODELS / SERVICE LAYERS

Job Queue

Access to compute nodes is provided with the use of an integrated workload manager (Slurm). This provides traditional batch processing services through a pre-defined queue structure

→ **Info**

Containers

The user can "containerize" the Operating System of choice (must be a Linux-based OS) plus all the libraries and applications that are needed on his/her own, and then launch it inside a job. This gives the user the flexibility of "almost" an Infrastructure Service.

→ **Info**

cscs2go – Compute and Data Services

The cscs2go platform

- powered by the Swiss National Supercomputing Centre
- gives access to world-class compute and storage HPC resources
- lowers barriers and simplifies access avoiding going through peer-review process,
- subscription through simple, web-based purchasing form,
- competitive prices approved by SNF

http://2go.cscs.ch

cscs2go offering :

- a state-of-the-art environment to store and analyze scientific data (from small to huge),
- applications for data science and machine learning (Spark, TensorFlow, Theano),
- variegated programming environments,
- 50% first year discount on subscription fees.

Flagship Supercomputer “Piz Daint”

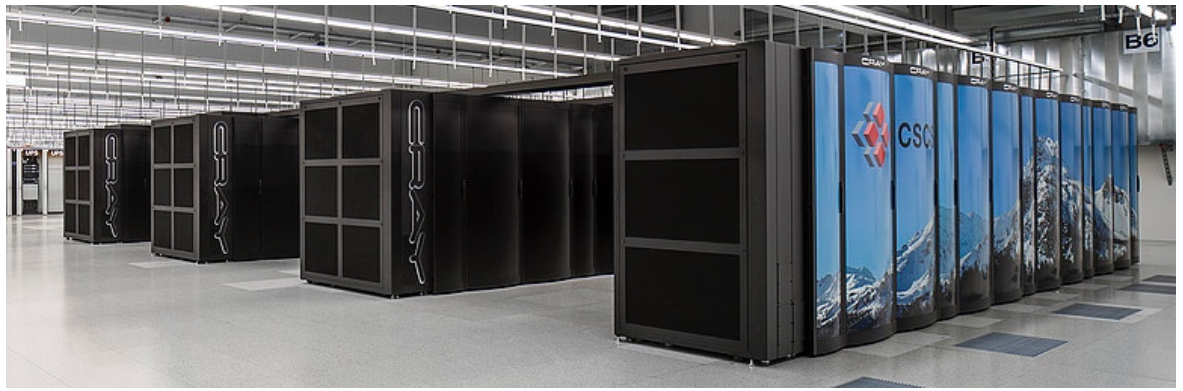

- Cray XC40 / Cray XC50
- Operational since April 2013
  - Extension + upgrade to hybrid in late 2013
  - Upgrade to new GPU in 2016
- Compute nodes
  - 5‘704 dual-socket nodes with Intel Xeon CPU and NVIDIA Tesla P100 GPU
  - 1‘810 dual-socked nodes with Intel Xeon CPUs
- Total system memory 521 TB RAM
- Peak Performance
  - Hybrid partition ~27 Petaflops
  - Multicore partition 1.7 Petaflops
- Ranked No. 6 among the World‘s TOP500 Supercomputers with ~21 PFlop/s Linpack performance and one of the greenest Petaflop supercomputers according to Green500

Infrastructure for Data Intensive Applications

- Internet connection
  - 100 + 100 Gbit connection to SWITCH
  - CSCS is the largest user of SWITCH in terms of transmitted data
- File systems for the storage of user data
  - ”Project” for User Lab projects: 6.0 PB
  - “Store” for paying customers: 4.4 PB
  - “Scratch” for the intermediate storage of computational results: 8.7 PB
  - Cray’s DataWarp (SSD layer integrated into fabric on Piz Daint)
  - Data Mover services supporting bi-directional parallel data transfer services between Scratch and Store and between Store and any external location
- Offline storage and Archival storage
  - Last generation tape library: 120 PB
  - Archival services allowing to store and access data on long term (>10 years) under development with NCCR project MARVEL

All infrastructure elements listed here are modular and can easily (and quickly) be extended by orders of magnitude.

Paul Scherrer Institute – SELVEDAS project

- PSI Mission
  - Study the internal structure of a wide range of different materials
  - Research facilities: the Swiss Light Source (SLS), the free-electron X-ray laser SwissFEL, the SINQ neutron source and the SμS muon source

- PSI facility users reserve a scientific device for a period of time
  - Compute power should also be available
  - Storage and archive availability during the experiment
  - Data retrievable after experiment by the users of PSI facilities (not PSI)

- Proposal to interface Piz Daint with their workflow
  - Use an API to access compute and data services (job scheduler, data mover)
  - Create a reservation service to reserve computation nodes
  - Provide a portal running on OpenStack to let PSI users access archived data at CSCS

SELVEDAS: Workflow proposal

- Web portal for creating a reservation service to reserve compute nodes
- Portal for providing PSI users access archived data at CSCS

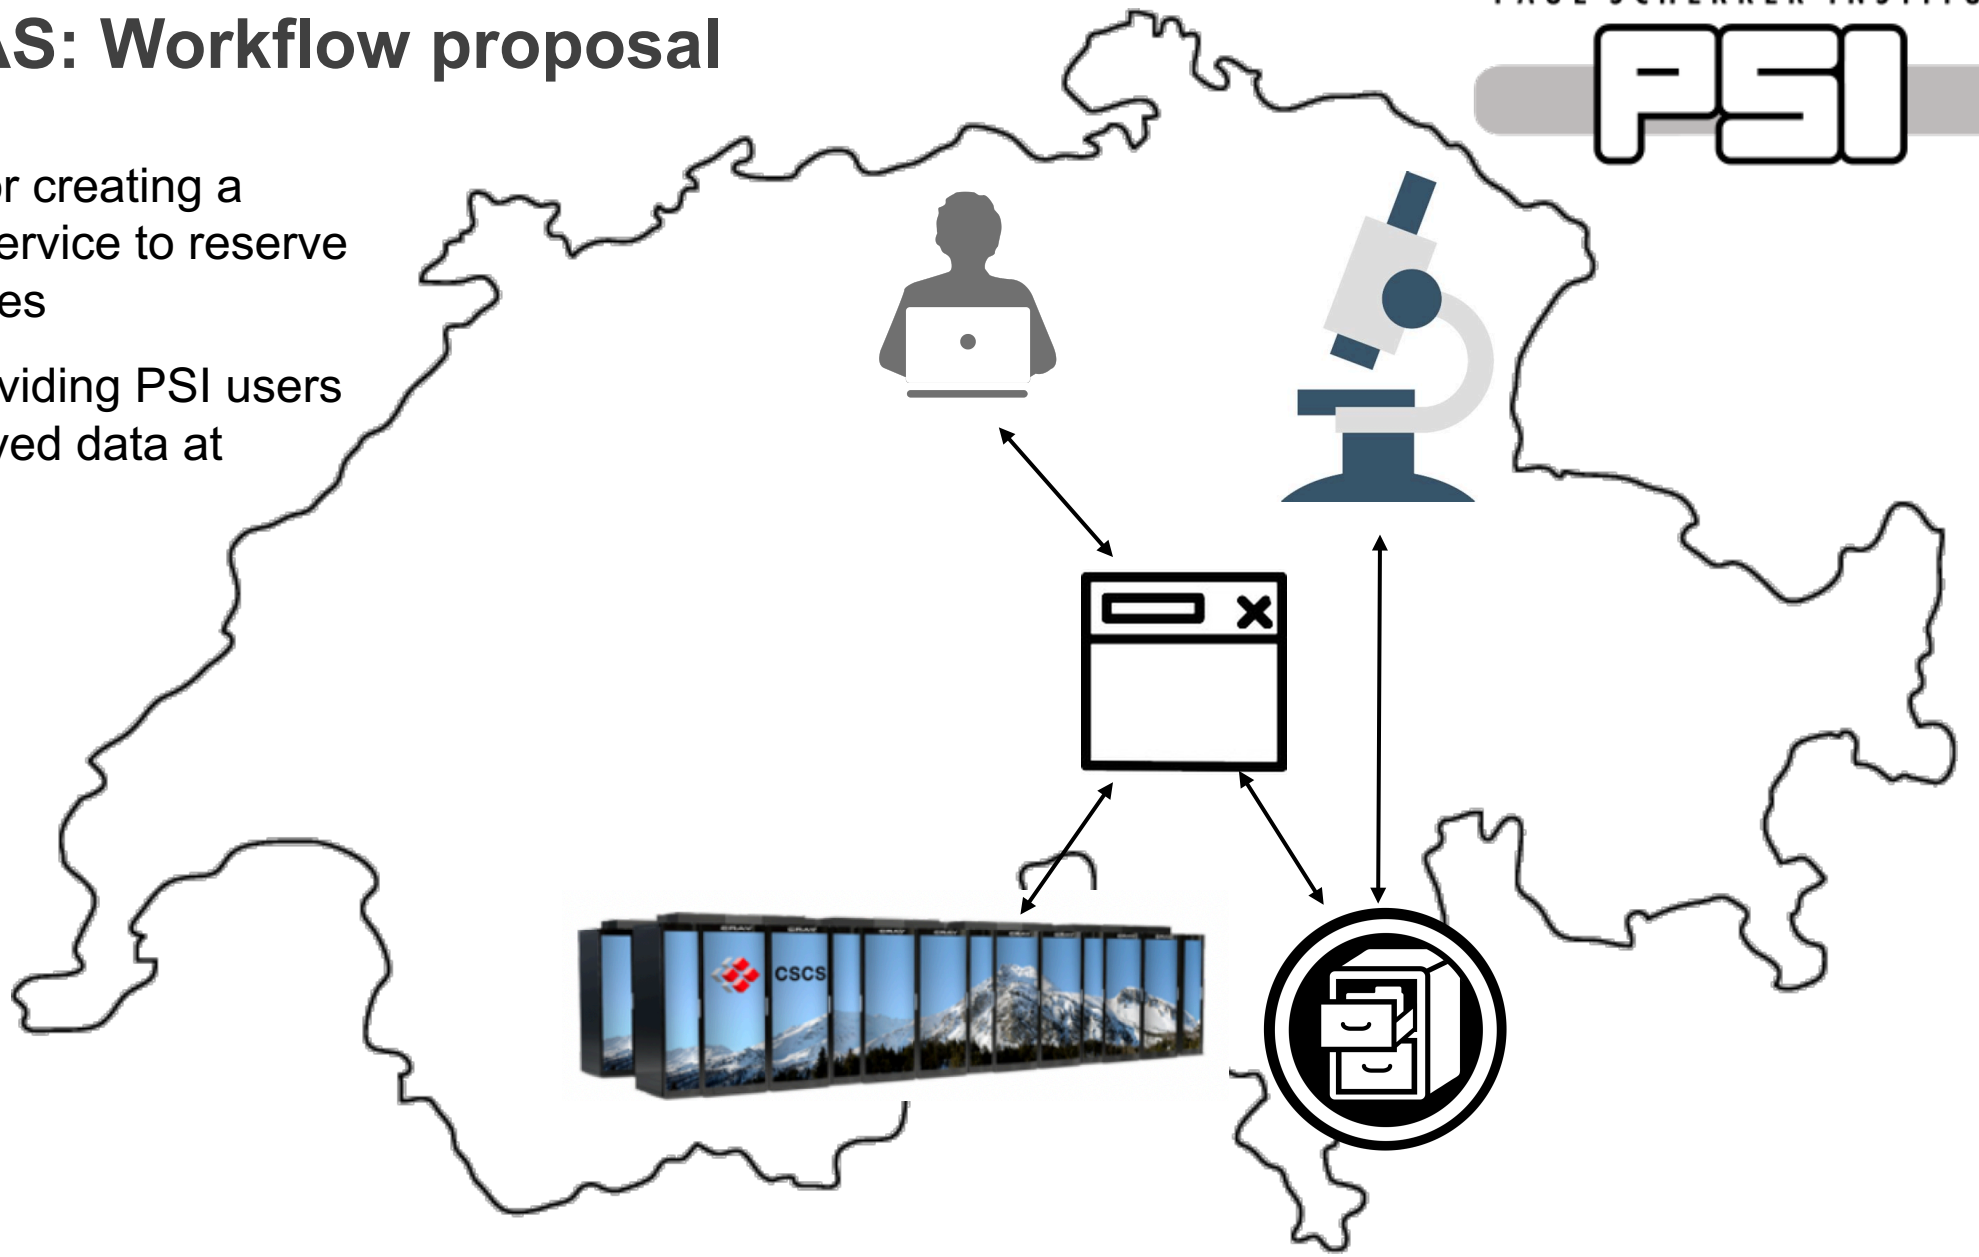

Users by scientific field - 2018

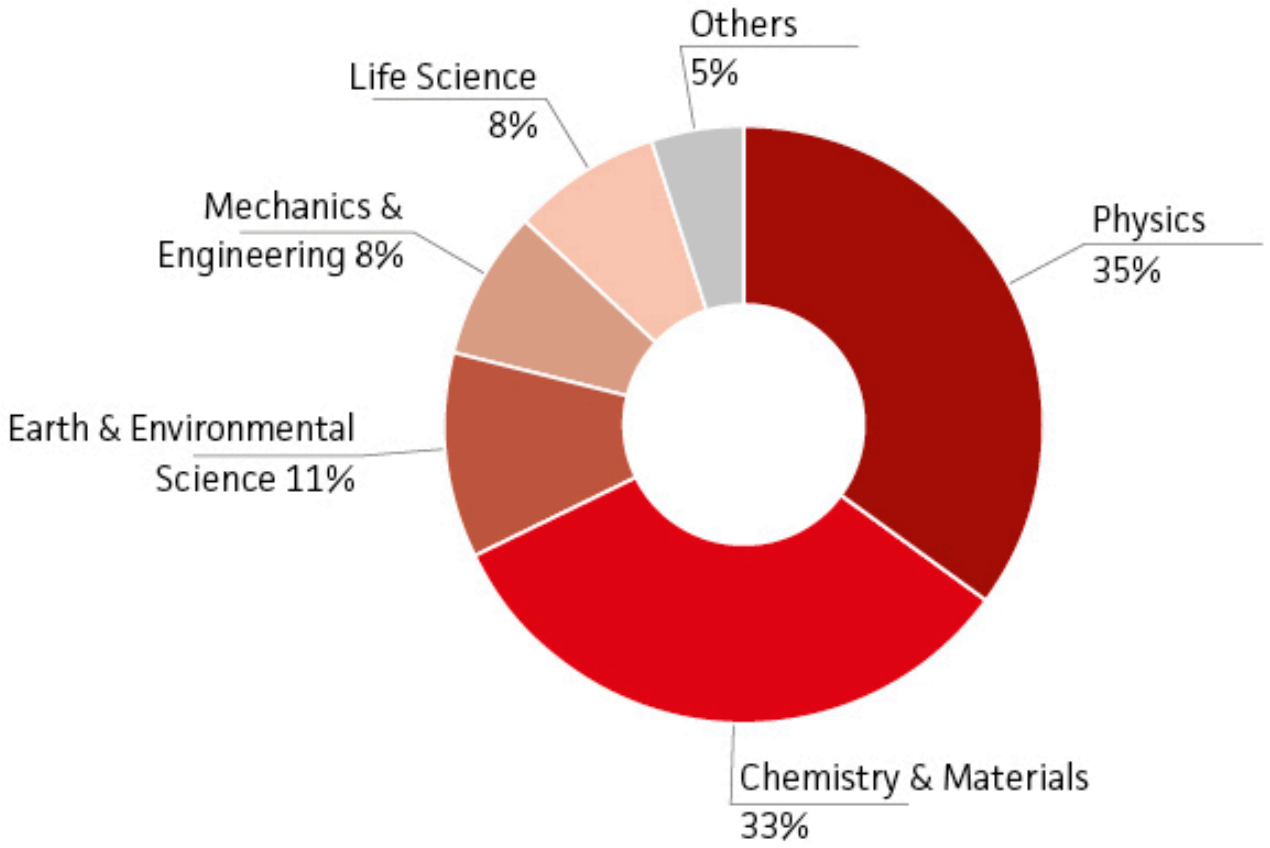

Users by organisation - 2018

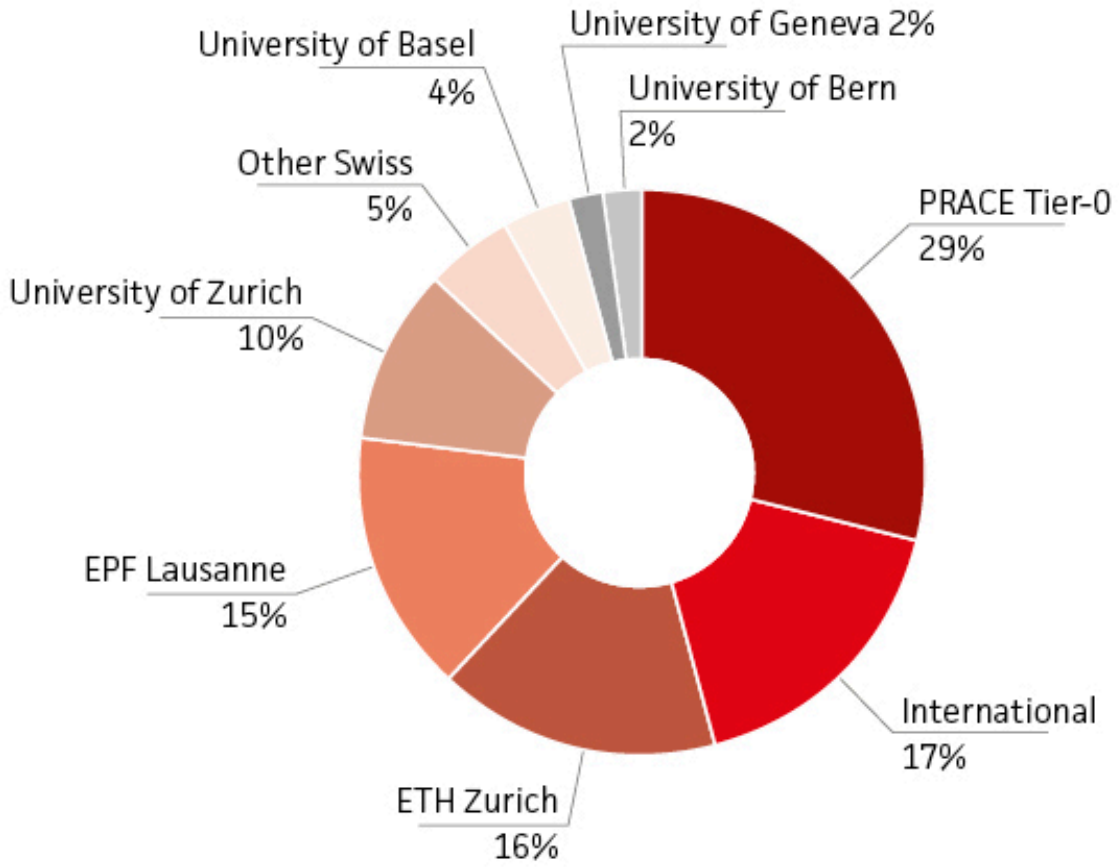

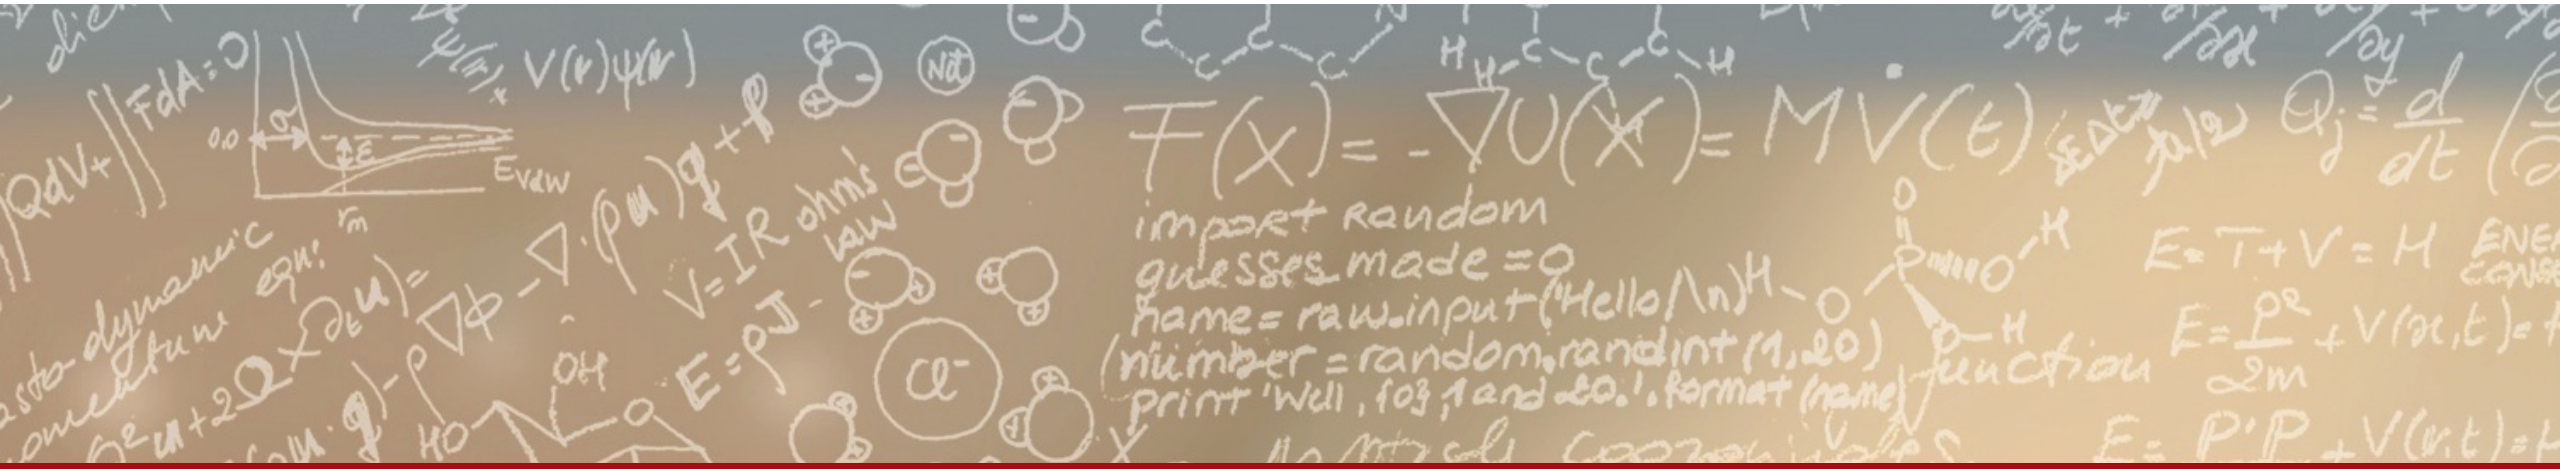

Thank you for your kind attention