Mitglied der Helmholtz-Gemeinschaft

JÜLICH
FORSCHUNGSZENTRUM

Clouds for HPC

Potential?
Challenges?

Session: Cloud Computing and HPC – Synergy or Competition?

ISC09, June 24, 2009

Thomas Lippert
Institute for Advanced Simulation
Jülich Supercomputing Centre

HPC

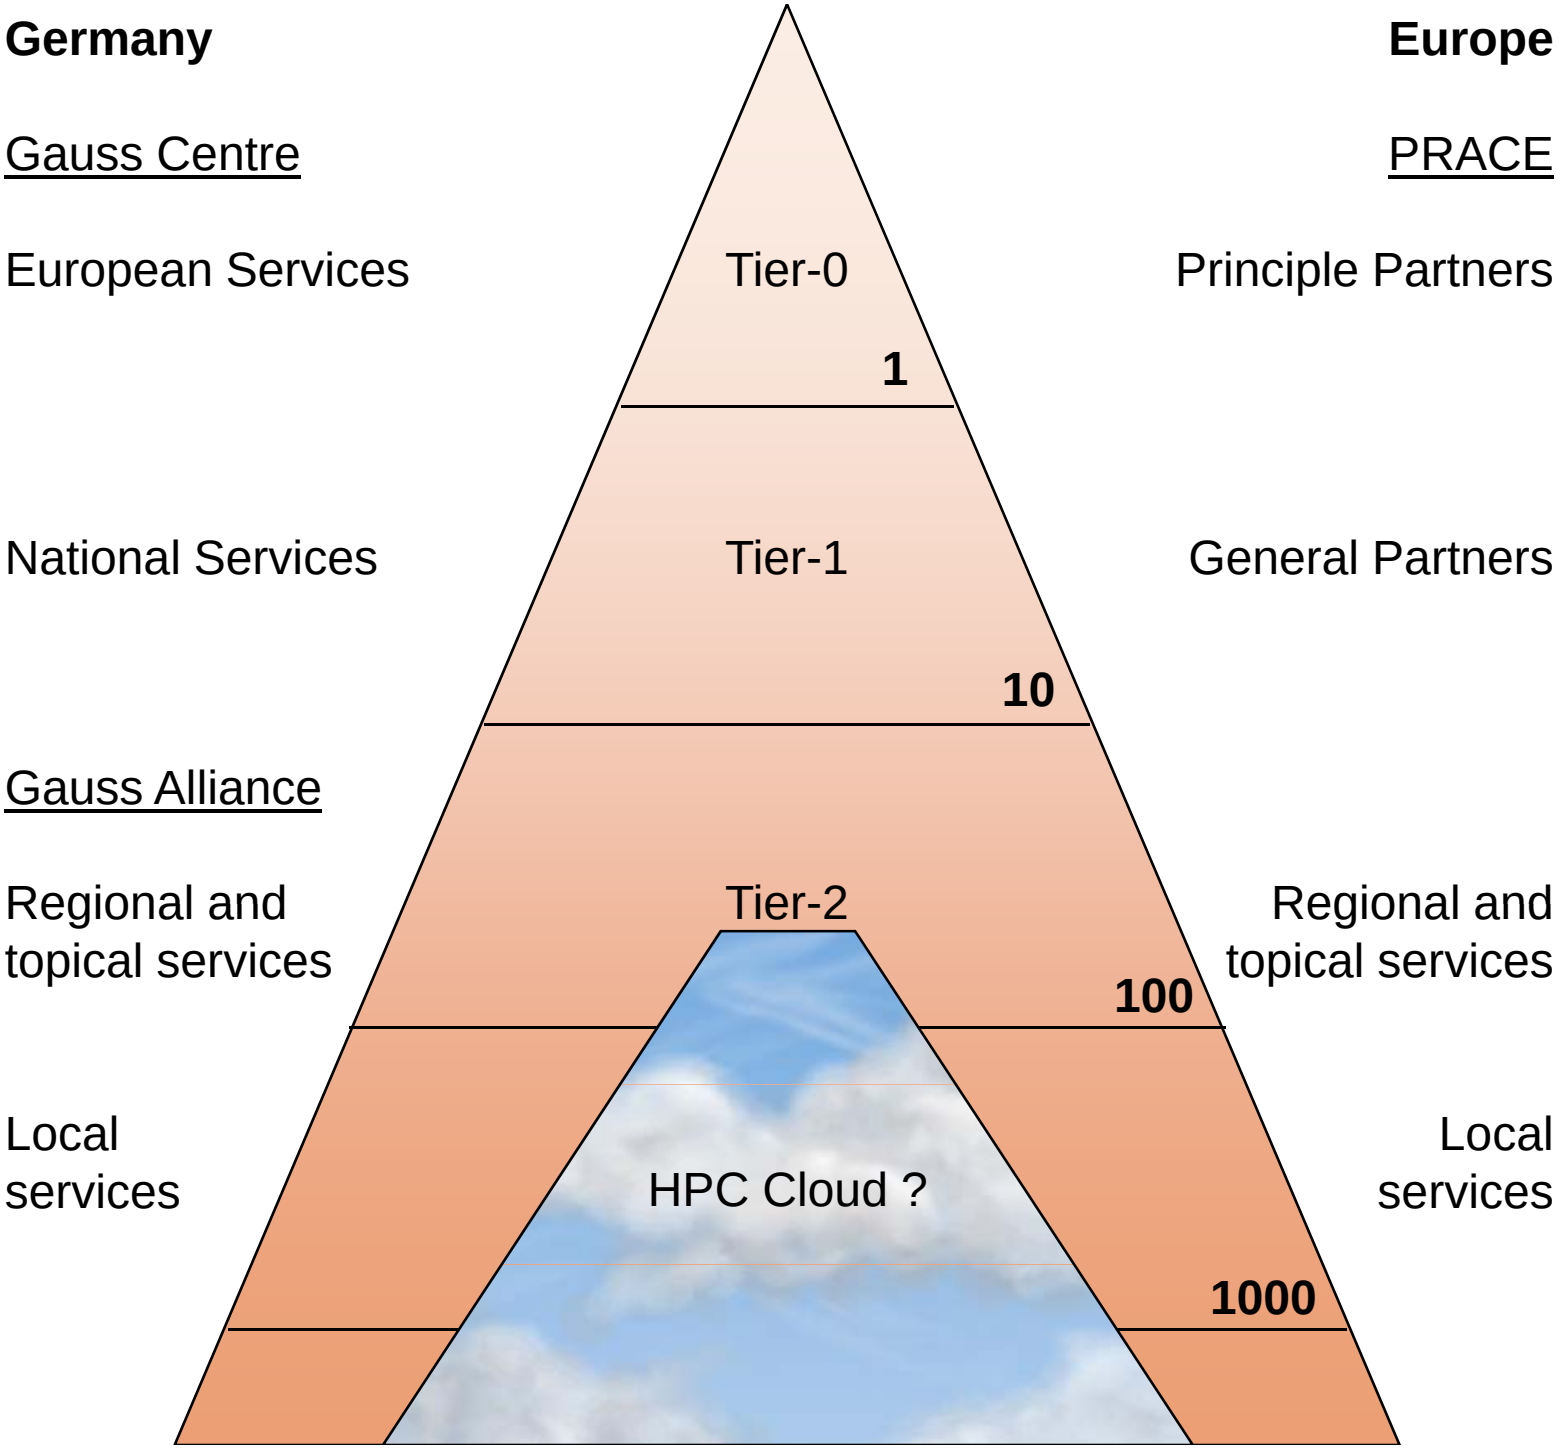
Germany
Europe
Gauss Centre
PRACE
European Services
Tier-0
Principle Partners
1
National Services
Tier-1
General Partners
10
Gauss Alliance
Regional and topical services
Tier-2
Regional and topical services
100
Local services
HPC Cloud ?
Local services
1000

Hardware Aspects

Leadership HPC systems

- Similar to large experimental projects
- Machine life cycle of 3 to 5 years
- Time scale of know-how: 15 to 30 years
- Usage: 24x7 h / week
- Most industries are more than 6 years behind leadership HPC

Tendency

- „Full transparency“ of machine becomes more and more utopia
- User needs to know her machine like physicist need to know math
- Assembler, SSE, MPI, parallelization strategy, scalability

JUGENE, JuRoPA + HPC-FF @ Jülich

IBM Blue Gene/P

72 racks, 294912 cores

1 Petaflop/s peak

144 Tbyte memory

6 Pbyte disks

25 PByte tape capacity

Highest scalability

Cluster computer

SUN-blades

2208 nodes, 17664 cores

207 TF peak, Intel Nehalem

48 GB memory

Infiniband QDR (SUN M9)

ParaStation Cluster-OS

General Purpose HPC

Cluster computer

Bull NovaScale R422-E2

1080 nodes, 8640 cores

101 TF peak, Intel Nehalem

24 GB memory

Infiniband QDR (Mellanox)

ParaStation Cluster-OS

HPC for Fusion

Need of HPC Users

For effective usage of tier-0, tier-1 and tier-2-systems

- → **high-level** support structures

Jülich: more than 50 % of staff works as domain scientists, mathematician and computer scientist in **simulation labs**

- Support
- Research
- Community oriented
- Integrated in community
- Parallelization has come closer to theory and model

Simulation Labs

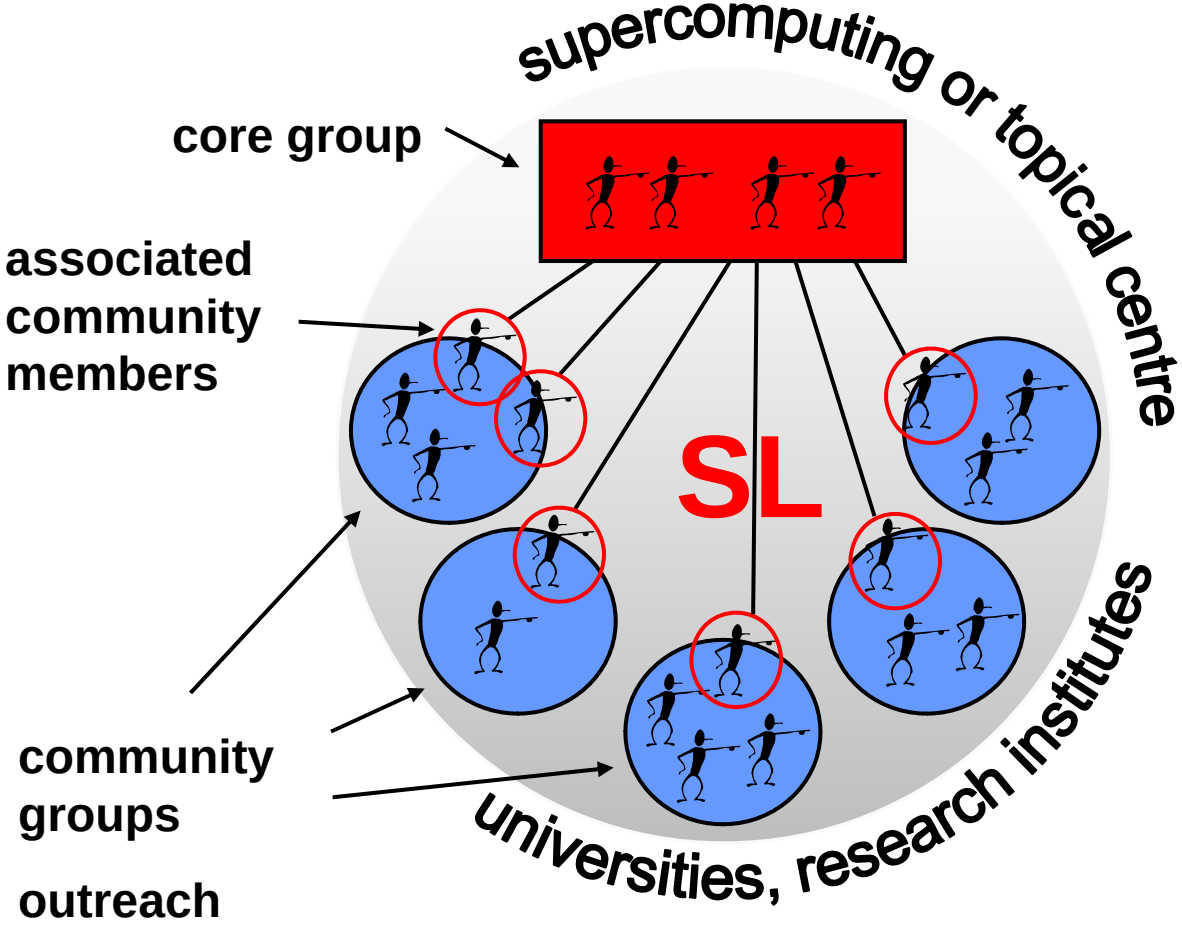

Core group tasks:

- disciplinary research
- supportive tasks

Steering Committee

Expansion towards distributed European SLs

Simulation Labs @ FZJ and KIT

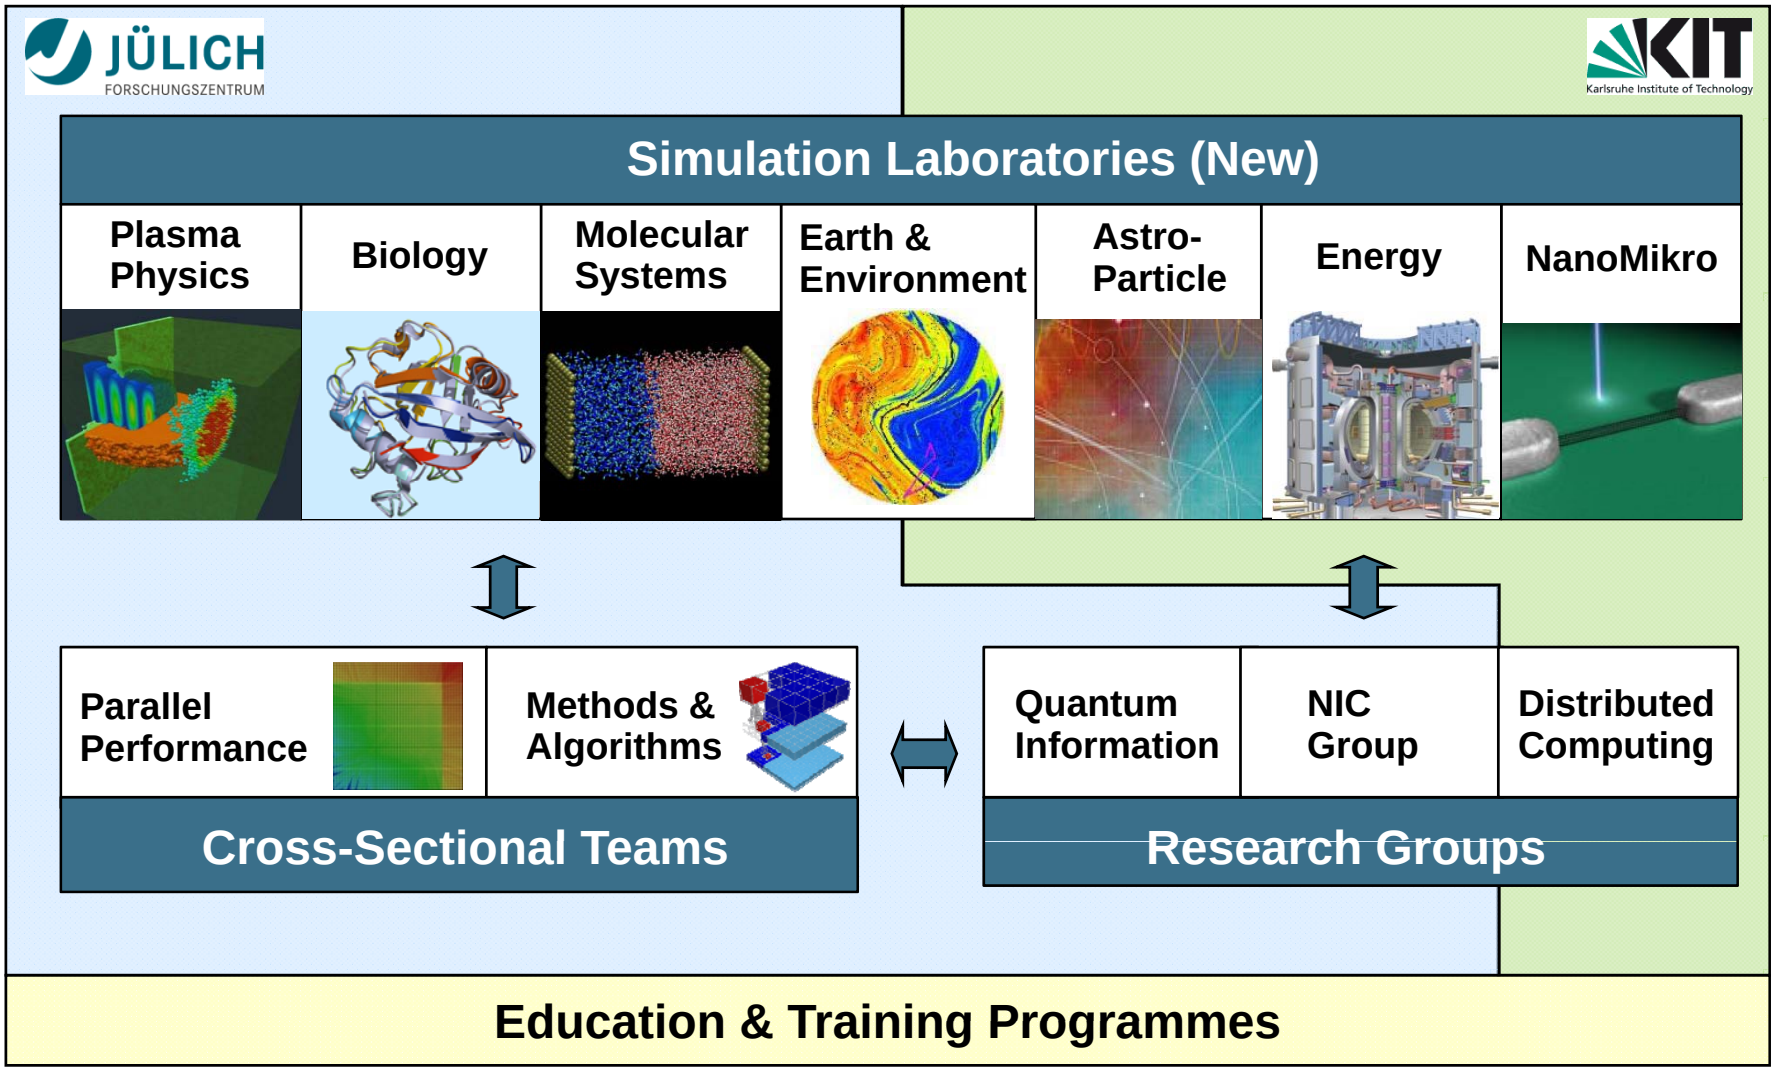

Example: Simulation Lab Biology

Research
- Protein folding & interaction
- Structure prediction
- Systems biology

Support
- Libraries, Bio databases
- Benchmarking
- Monte Carlo, FFT docking, Machine learning

Codes
- *PROFASI*, *SMMP*

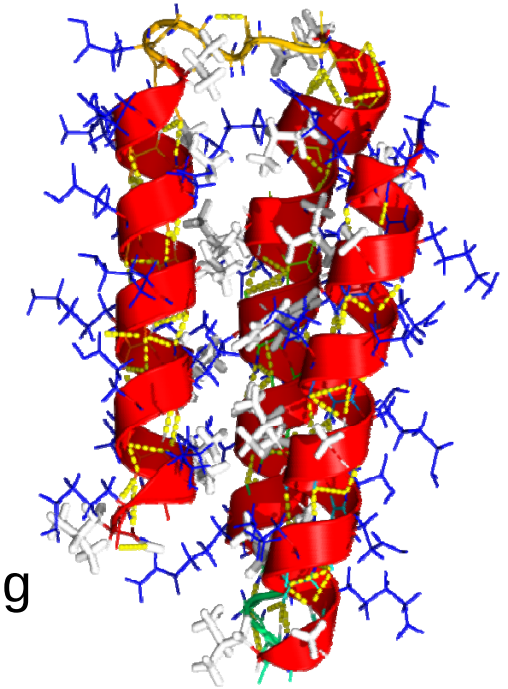

Protein 1LQ7

Outreach
- FZJ: Biological institutes (ISB, INM), Helmholtz Groups
- Regional: ABC of Life Science Informatics
- International: UC Berkeley, Michigan Tech

National and European User Group

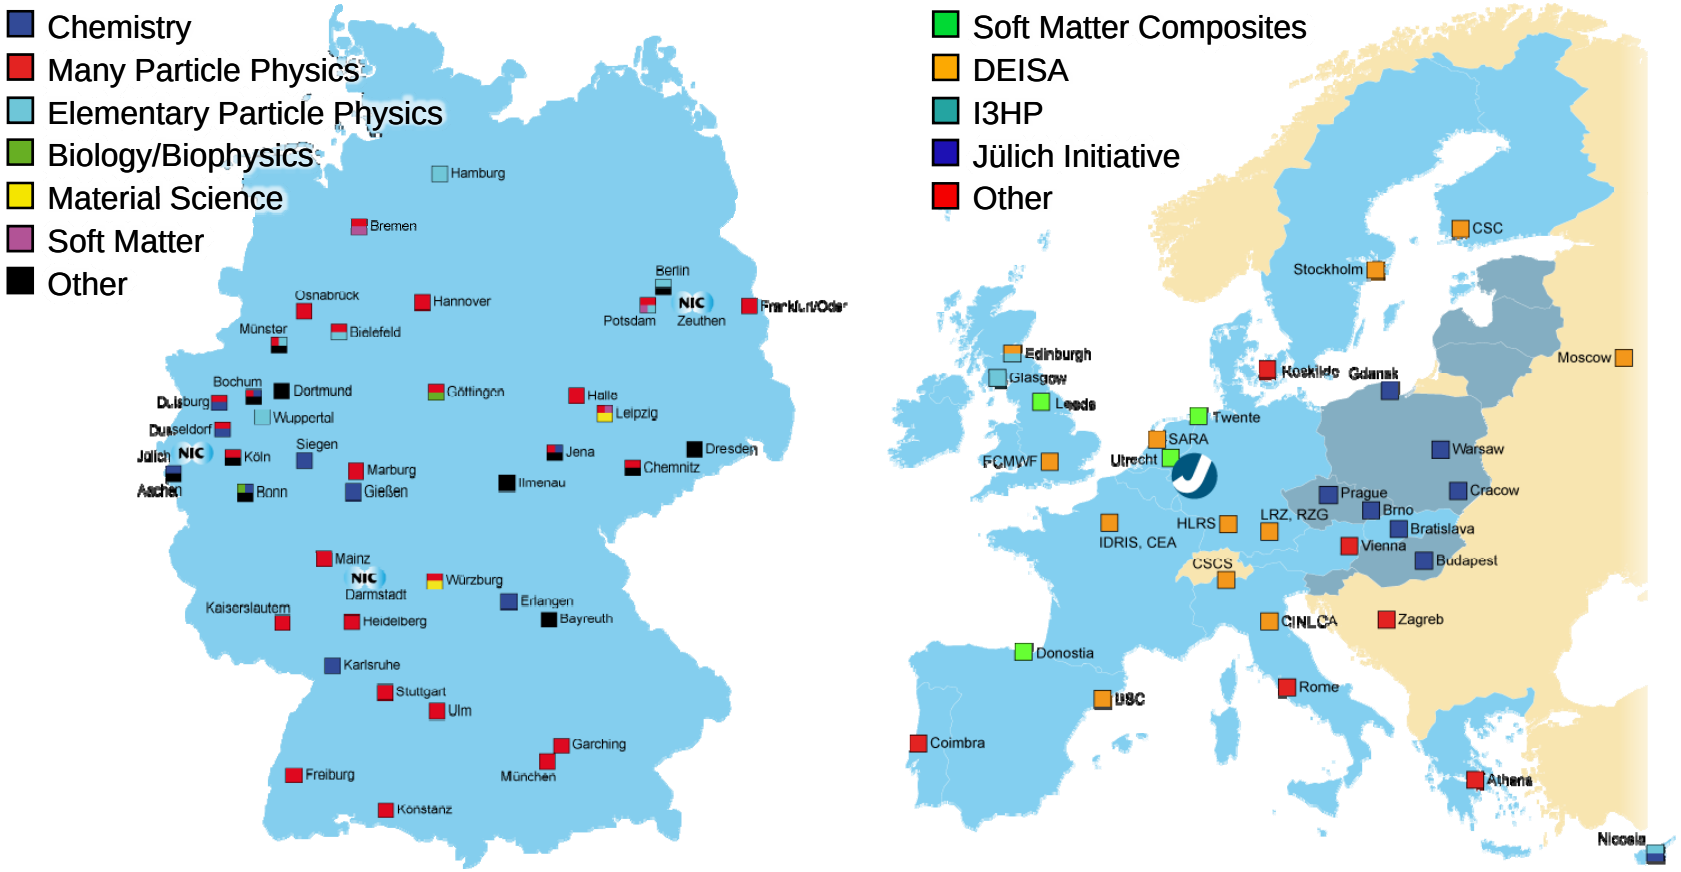

- Proposals for computer time accepted from Germany and Europe
- Peer review by international referees

Cloud @ Jülich

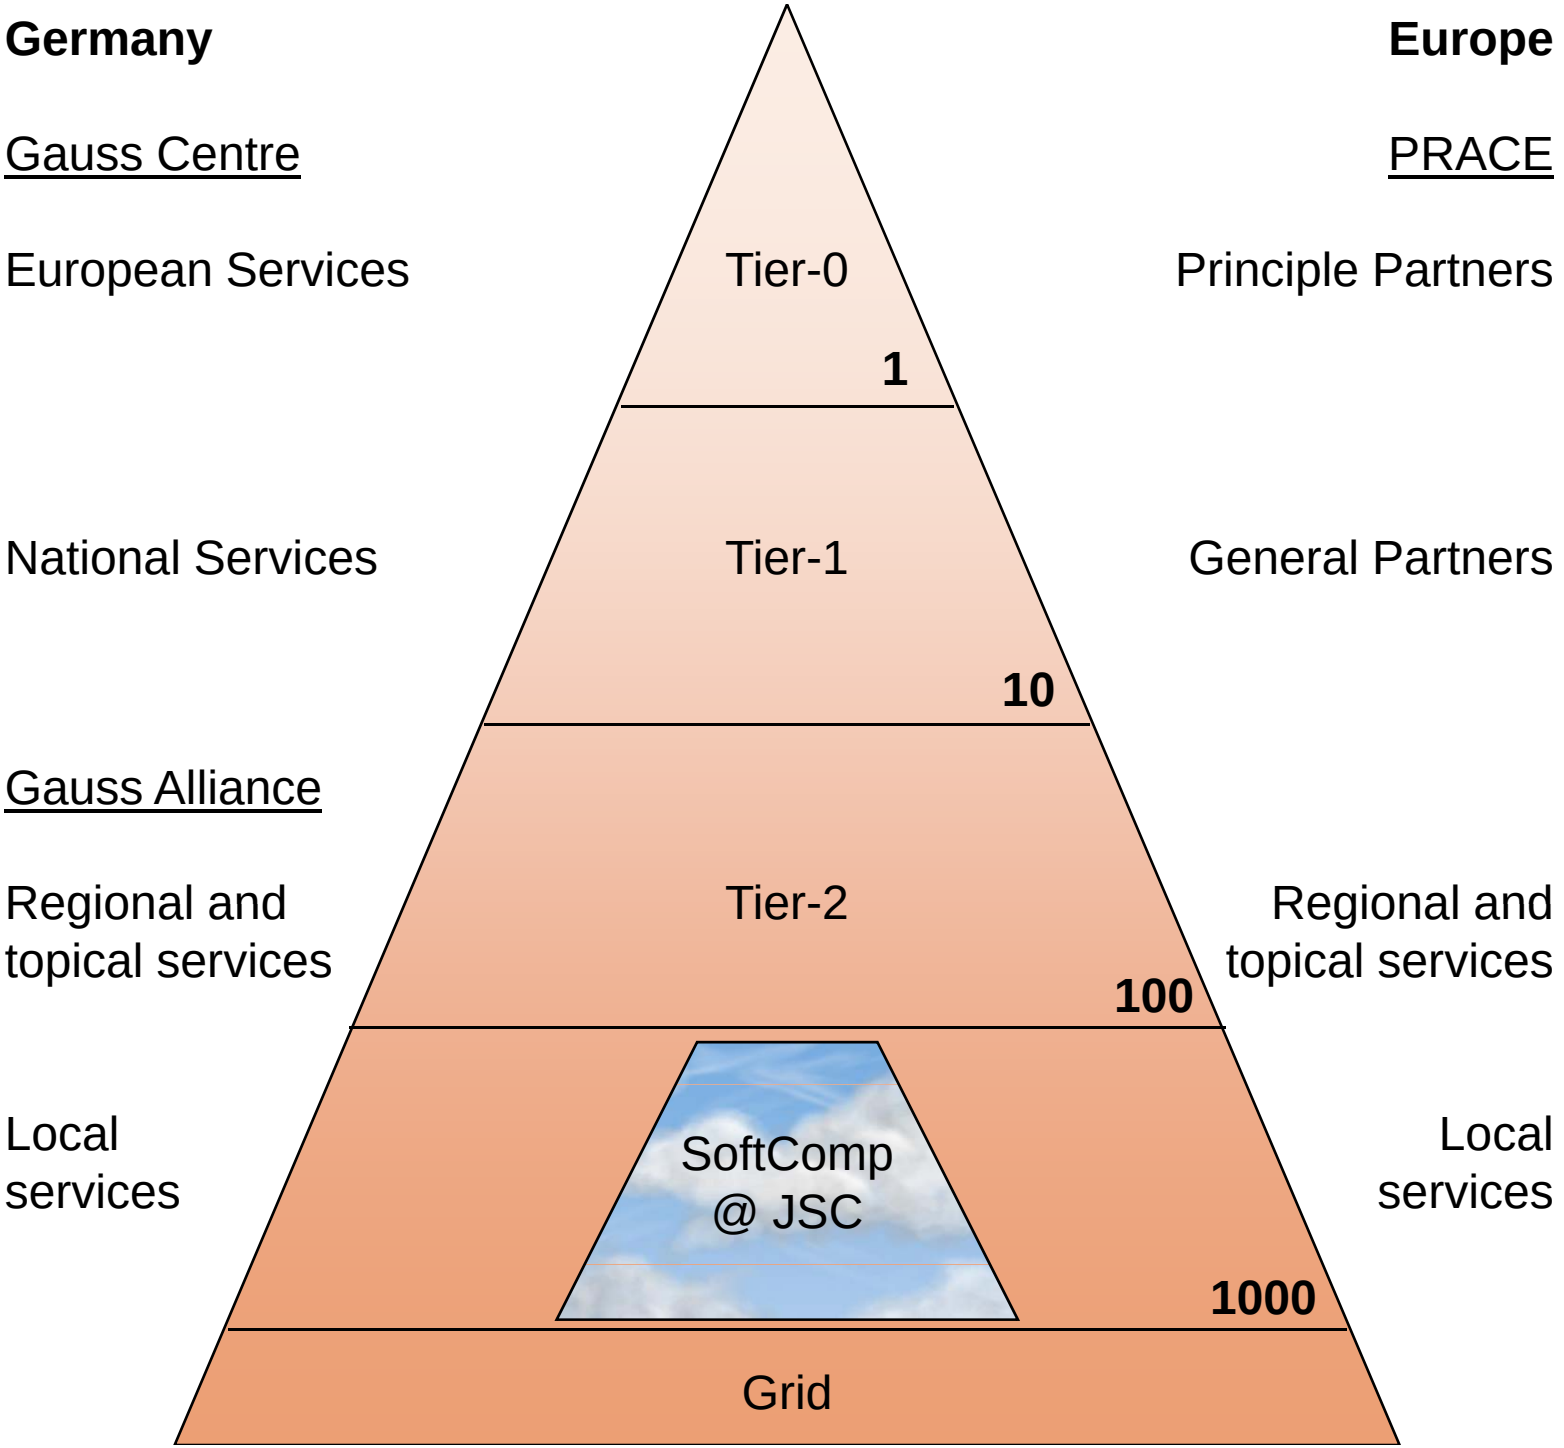
Germany
Europe
Gauss Centre
PRACE
European Services
Tier-0
Principle Partners
1
National Services
Tier-1
General Partners
10
Gauss Alliance
Regional and topical services
Tier-2
Regional and topical services
100
Local services
SoftComp @ JSC
Local services
1000
Grid

Users and Members of SoftComp

Member groups of the NoE SoftComp

- fzjg: Forschungszentrum Jülich
- jogu: Johannes Gutenberg Universität Mainz
- scr: Schlumberger Cambridge Research Limited
- unid: Heinrich-Heine Universität Düsseldorf
- upv: Universidad del Pais Vasco, Euskal Herriko Unibertsitatea
- utcdr: University of Twente
- uutr: Utrecht University
- Ulcrl: Unilever UK Central Resources Limited

German Federal Ministry for Education and Research (BMBF)

- Promotion of applications in economy and science by grid infrastructures

Our Cloud Computer: SoftComp Linux Cluster

125 compute nodes (500 cores)

Heterogeneous, AMD Opteron

IB and GigE, ParaStation, Unicore

2.5 TF

UNICORE

JÜLICH FORSCHUNGSZENTRUM

Access

- Open, extensible, interoperable
- Strong security, workflow support, powerful clients, application integration
- Widely established in academia (D-Grid, DEISA, EGI, SKIF-GRID, SoftComp) and industry (T-Systems, Philips, 52° North)
- Average downloads per month ~ 1200
- Management of scientific data with metadata and scalable storage
- User-defined execution environments with Clouds and virtualisation

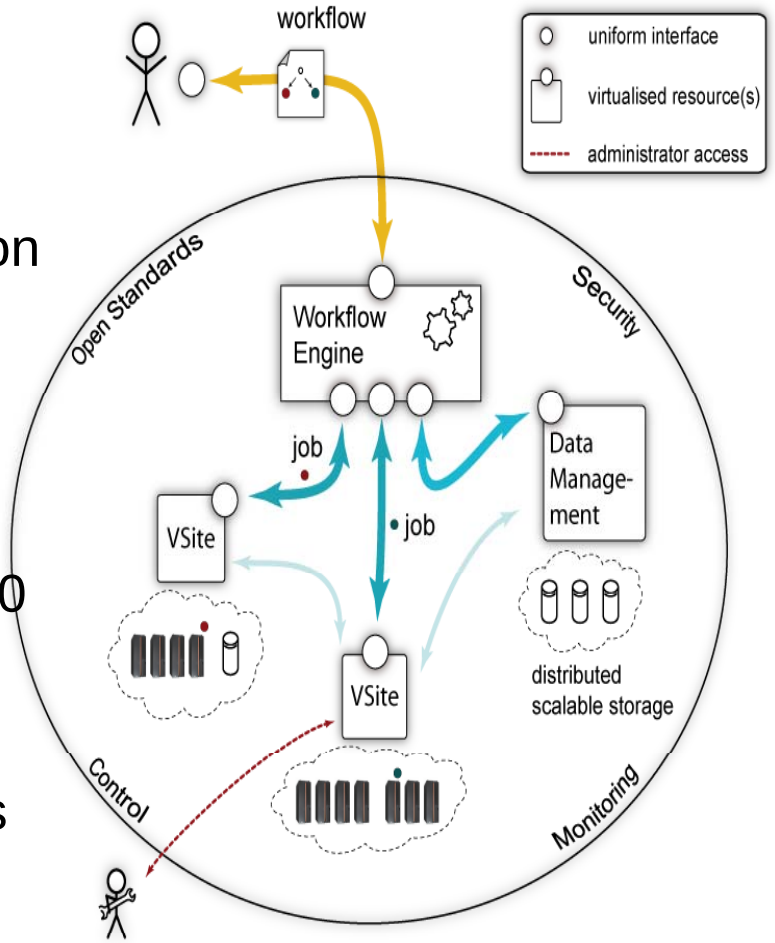

SoftComp WallCpuTime

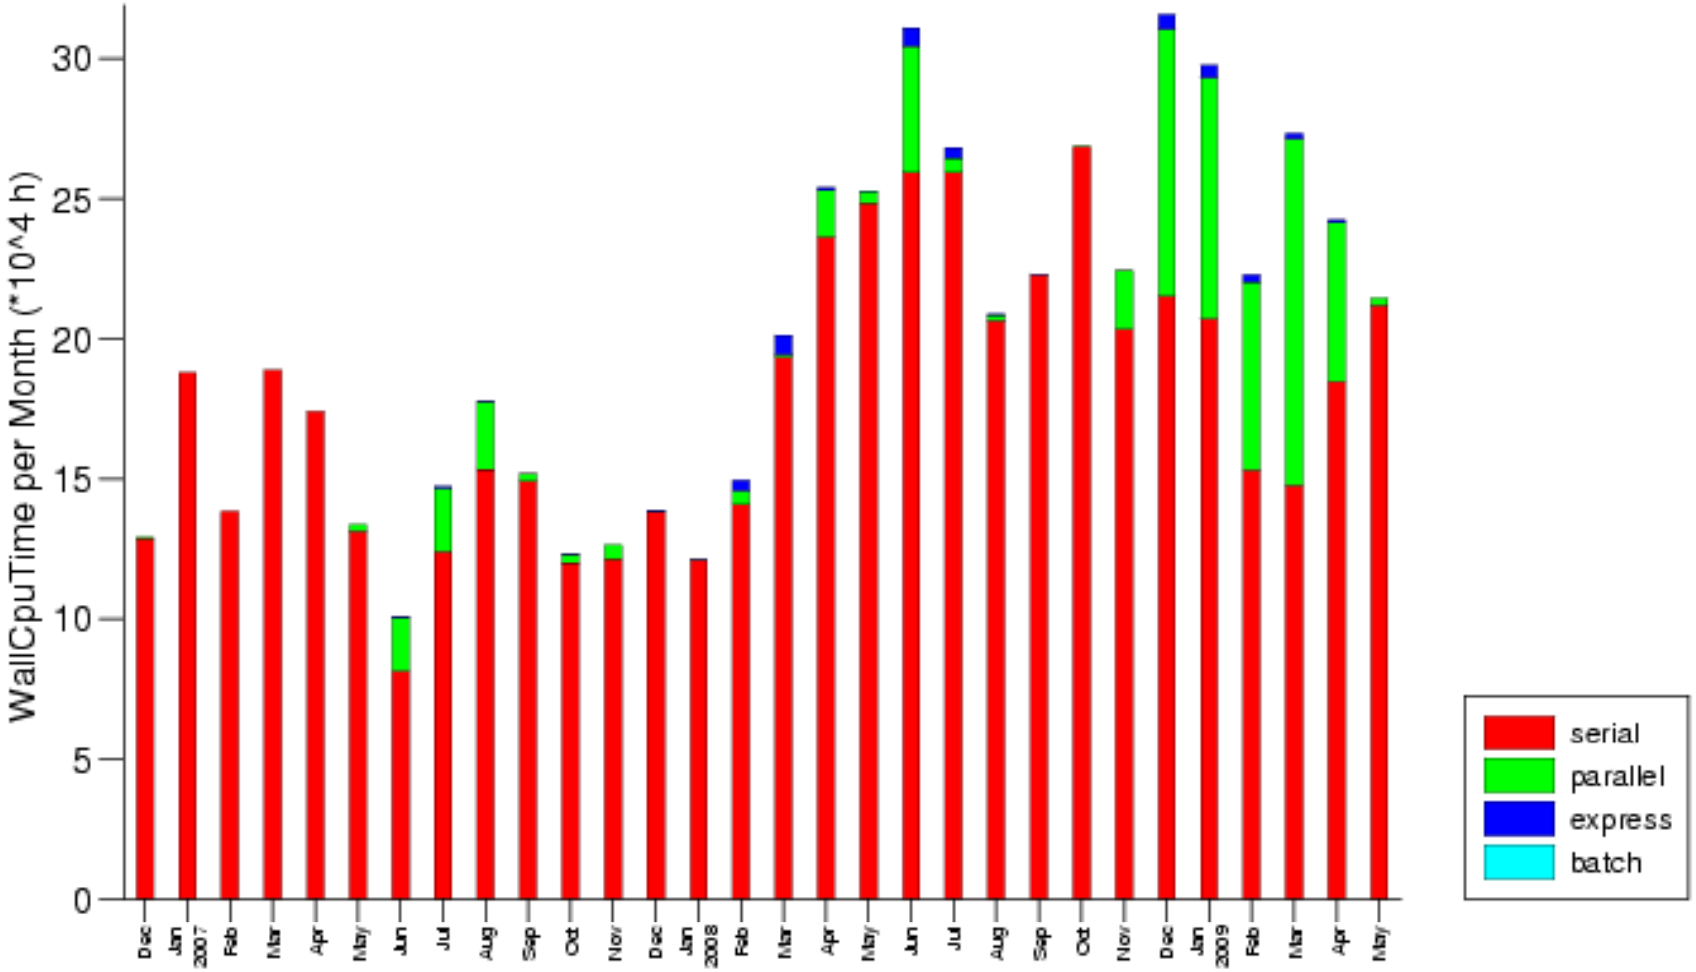

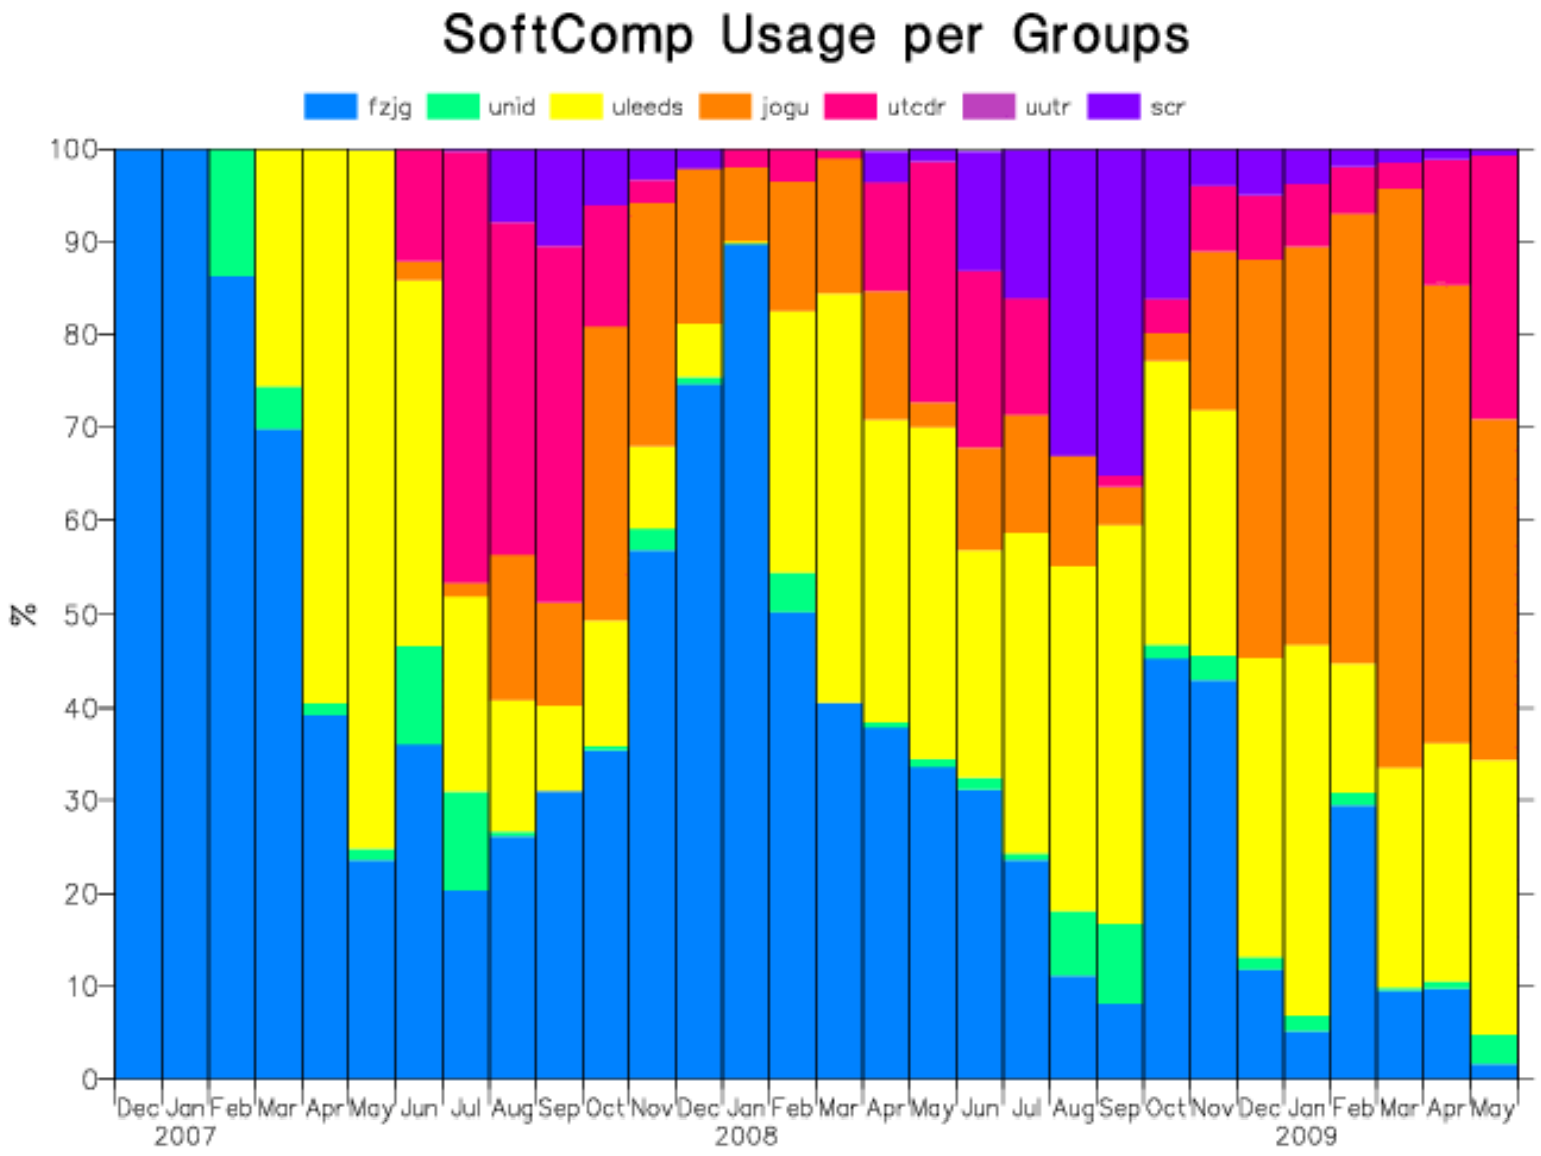
SoftComp Usage per Groups
fzjg
unid
uleeds
jogu
utcdr
uutr
scr
%
100
90
80
70
60
50
40
30
20
10
0
Dec Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan Feb Mar Apr May
2007
2008
2009

Parallel Codes

Why parallel?

- Results in a shorter time
- Results more precise
- Local system memory not sufficient
- But: Programs have to be adapted to run in parallel mode

Required

- Support of SimLab and Cross Cutting Groups
  - *Mathematics*
  - *Performance Analysis*

Lessons Learnt:

Challenges for the HPC Cloud Provider

To serve beyond tier-3 and desktop applications a cloud provider must

- offer leading edge tier-3,2,1,0 high performance systems
- guarantee absolute security
- guarantee absolute privacy
- care for long-term data storage and curation
- guarantee uninterrupted service for critical applications
- actively offer highest level support and research for science communities and industry
- provide SimLab-like support structure for HPC