Leibniz Supercomputing Centre

of the Bavarian Academy of Sciences and Humanities

Introduction to the usage of the Linux Cluster and the Compute Cloud@LRZ

Ferdinand.Jamitzky@lrz.de

The Leibniz Supercomputing Centre

Academy Institute of the Bavarian Academy of Science and Humanities

- ✓ IT Service Provider for the **Munich** Universities
- ✓ Regional Computing Centre for Research Institutions in **Bavaria**
- ✓ **German** National Supercomputing Centre
- ✓ **European** Supercomputing Centre

LRZ Overview by regions

IT Service Provider for **Munich** Universities

Email, Web,
Multimedia,
IT Security,
HelpDesk,
Virtual Reality,
Trainings, etc.

Regional Computing Centre for **Bavarian** Universities and Research institutions

~ 150 PByte Storage/Archive
Digital Archive of the
Bavarian State Library

Munich Scientific Network

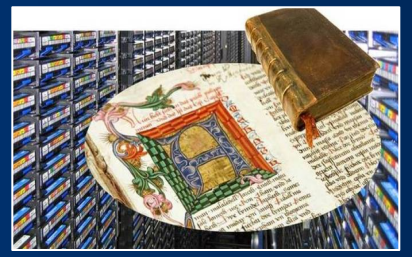

German National Supercomputing Centre

SuperMUC-NG

25 Pflop/s peak
300k compute cores
0.7 PB main memory
50 PB HDD

European Supercomputing Centre

Participating in large European e-Infrastructures

High Performance Computing

High Speed Networks

Grid/Cloud Computing

LRZ as IT Competence Centre:
Providing Comprehensive IT Services for Science

LRZ by Hardware
Operating Cutting-edge IT Infrastructure

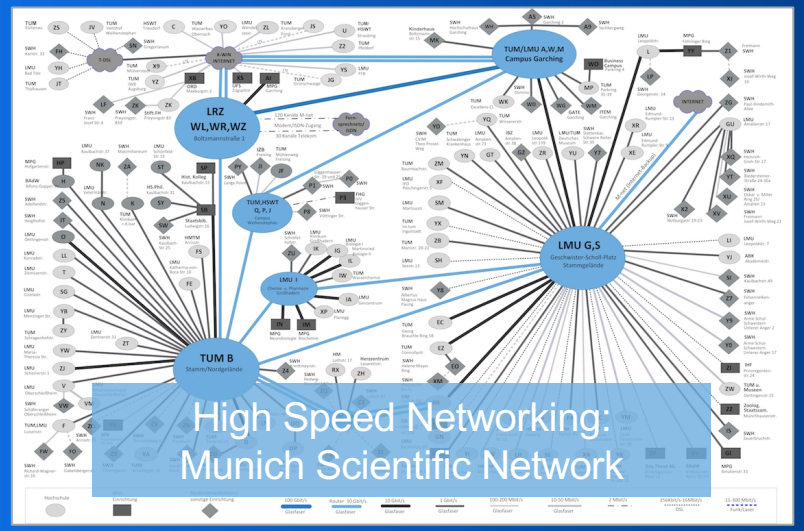

HPC@LRZ

Linux Cluster
- Massively parallel Cluster (CoolMUC-2, CoolMUC-3, IvyMUC)
- Big shared memory system (Teramem)
- Serial Cluster
- Remote Visualisation

SuperMUC-NG
- Massively parallel Cluster (Worker Nodes)
- Big shared memory nodes (Big Nodes)

Cloud Systems (Compute and Storage)
- Compute Cloud (openNebula and OpenStack)
- Long running instances (vmware)
- Data Science Storage DSS
- Machine Learning Systems (https://datalab.srv.lrz.de) Single GPU DGX-1, DGX-1v

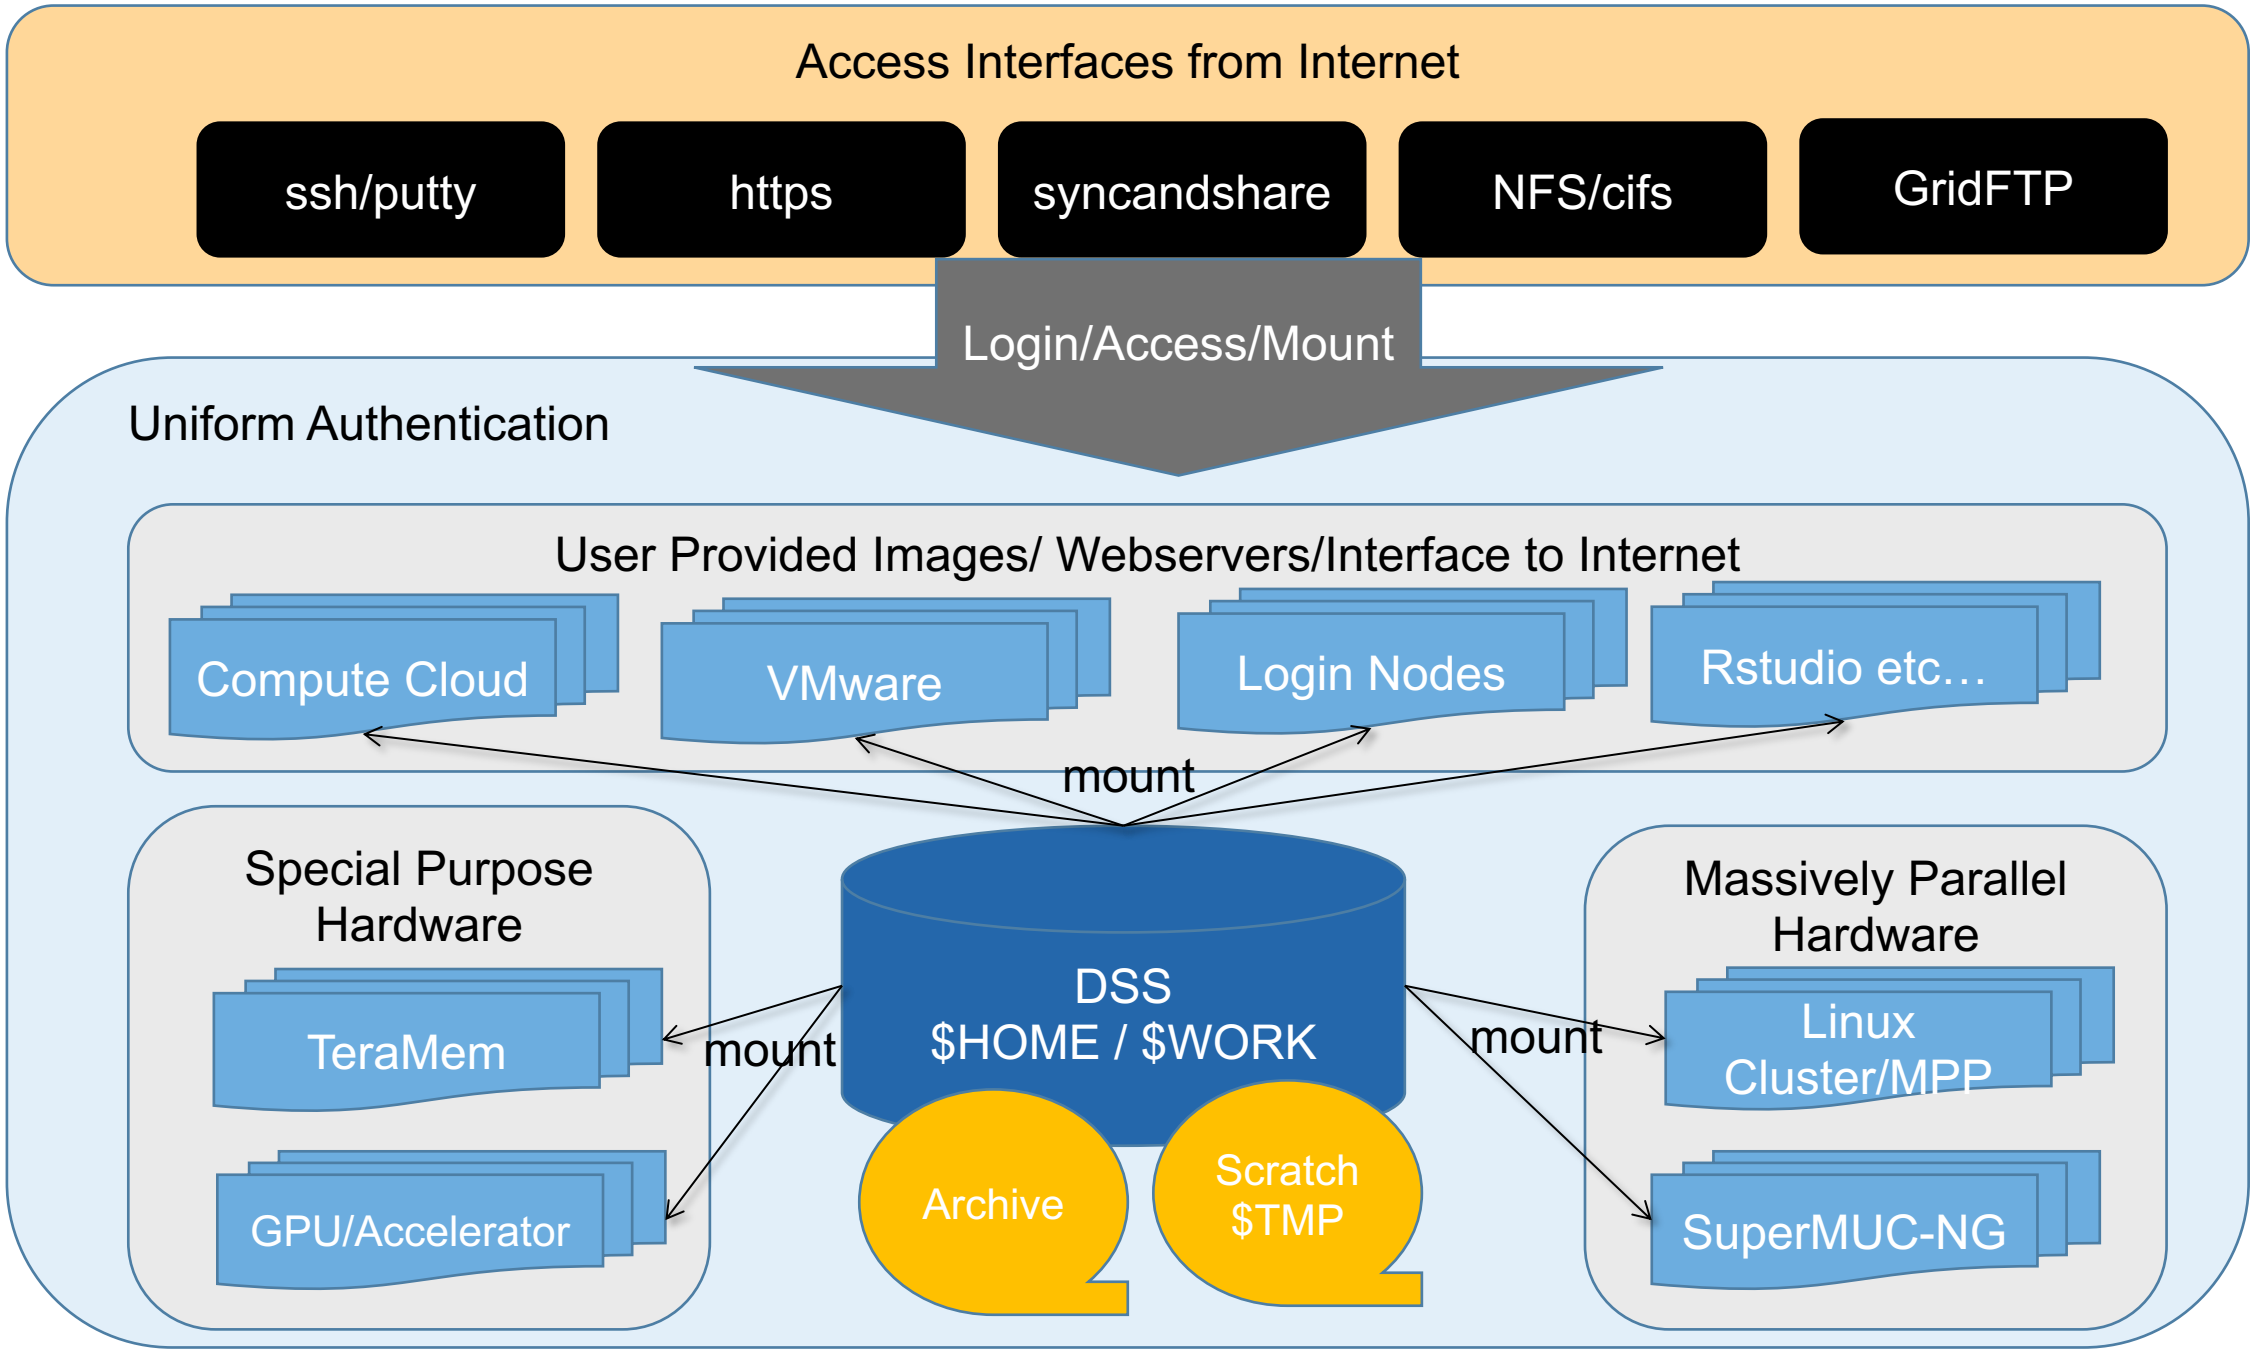
Access Interfaces from Internet
ssh/putty
https
syncandshare
NFS/cifs
GridFTP
Login/Access/Mount
Uniform Authentication
User Provided Images/ Webservers/Interface to Internet
Compute Cloud
VMware
Login Nodes
Rstudio etc…
mount
Special Purpose Hardware
TeraMem
GPU/Accelerator
mount
DSS
$HOME / $WORK
Archive
Scratch
$TMP
mount
Massively Parallel Hardware
Linux Cluster/MPP
SuperMUC-NG

LRZ "Linux Cluster" und "Cloud"

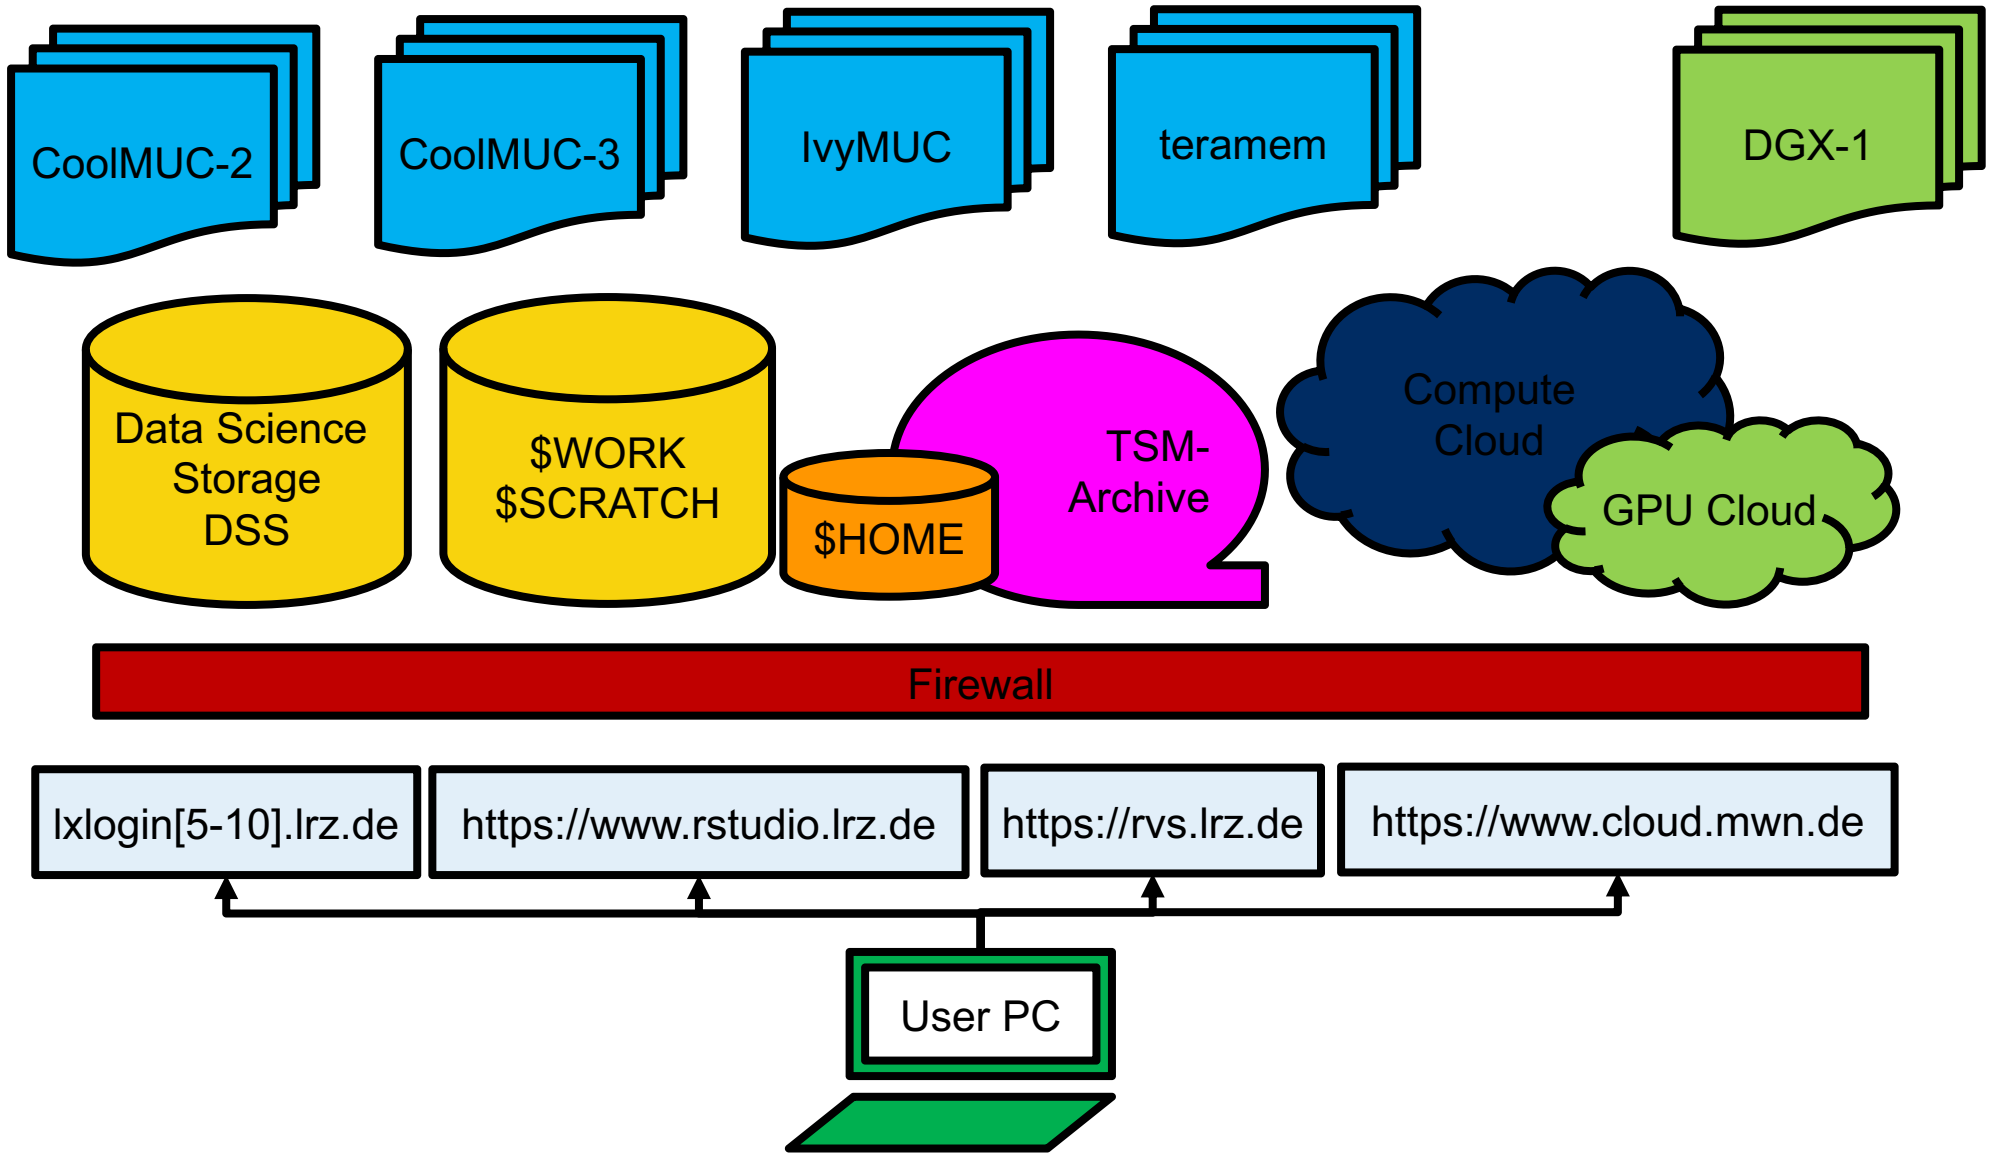

Linux Cluster Hardware

Architecture				Total		Max Jobs Limits				Access	
System	CPU	#core	RAM GB	#nodes	#cores	#nodes	#cores	#time	RAM	queue	login node
Linux-Cluster CoolMUC-2	Intel Xeon E5-2690 v3 ("Haswell")	28	64	384	10,752	60	1680	48h	3.8 TB	mpp2	lxlogin5-7 .lrz.de
Linux-Cluster Serial	Intel Xeon E5-2690 v3 ("Haswell")	28	64	1	28	1	28	96h	64 GB	serial	lxlogin5-7 .lrz.de
Linux Cluster CoolMUC-3	Intel Xeon Phi (Knights Landing)	64	96	148	9,472	148	9472	48h	8,9 TB	mpp3	lxlogin8.lrz.de
Linux-Cluster IvyMUC	Intel Xeon E5-2660 v2 ("Sandy Bridge")	16	64	31	496	12	192	72h	768 GB	ivymuc	lxlogin10.lrz.de
Linux-Cluster Teramem	Intel Xeon E7-8890 v4	96	6,144	1	96	1	96	48h	6.1 TB	teramem_inter	any cluster login node

Cloud Ressources

Architecture				Total		Access
System	CPU	#core	RAM GB	#nodes	#cores	login node
Machine Learning System	Nvidia Pascal P100	8 GPU	128	1	n.a.	contact servicedesk
OpenNebula Compute Cloud	Intel Xeon E5540 X5650 E5-2660v2	1-20	1-512	95	852	www.cloud.mwn.de
OpenStack Compute Cloud	Intel Xeon („Skylake“)	1-48	96	64	3072	cc.lrz.de
LRZ Virtual Machines	Intel Xeon E5-2660 v2 ("Sandy Bridge")	1-8	1-32	90	1800	http://www.lrz.de/services/serverbetrieb/

Moore’s Law (1965-2015)

Number of transistors doubles every 2 years

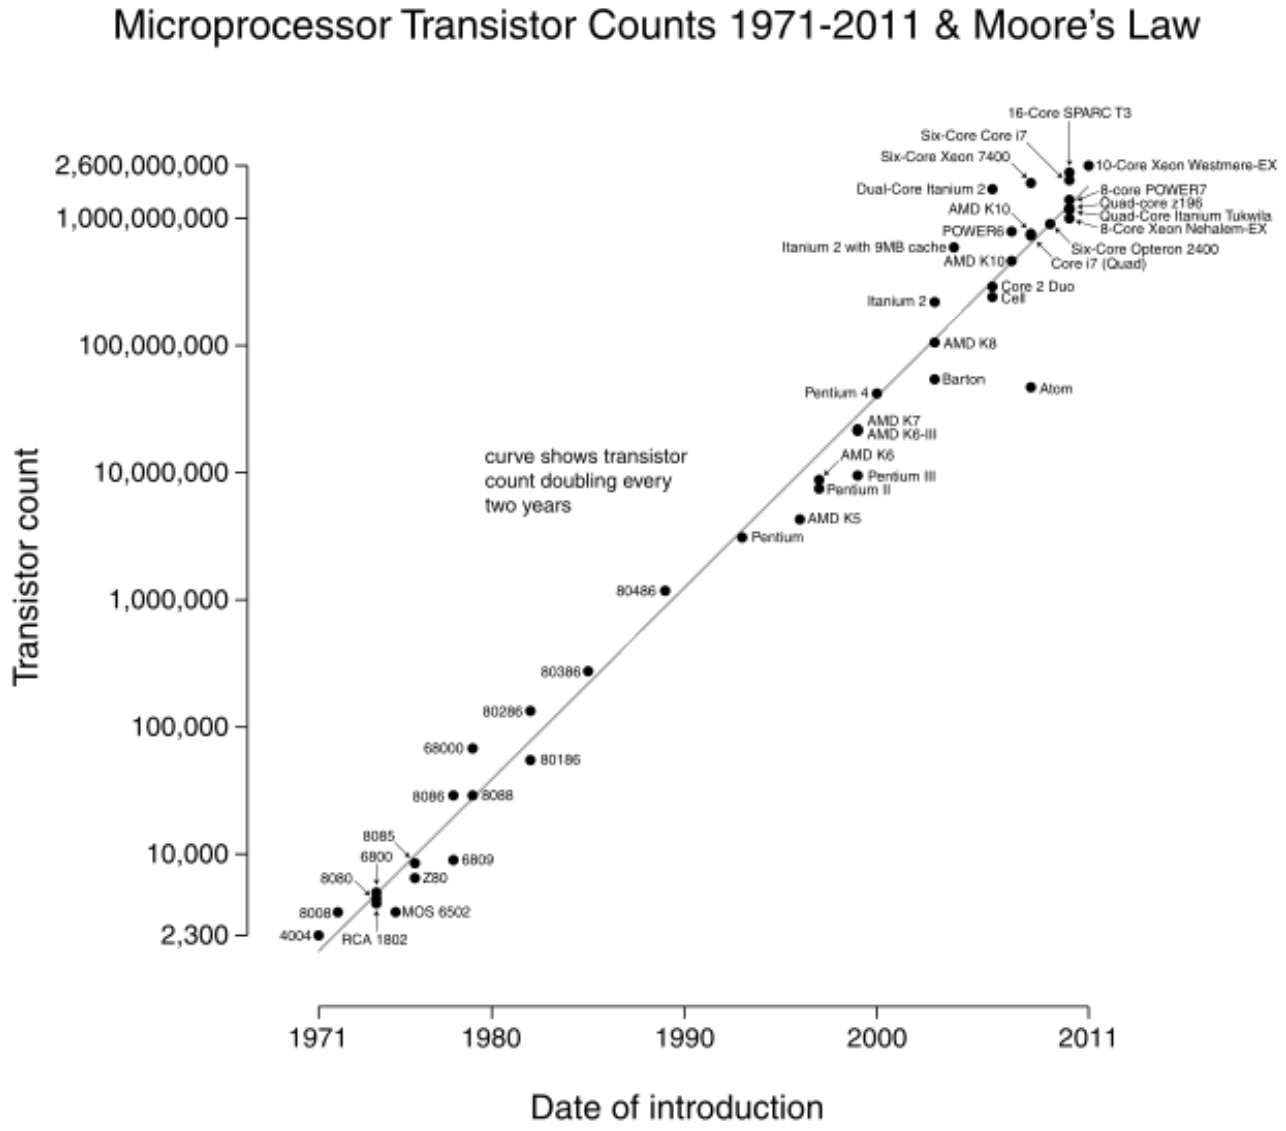

Why parallel programming?

End of the free lunch
in 2000 (heat death)

Moore's law means
not faster processors,
only more of them.

But!
2 x 3 GHz < 6 GHz

(cache consistency,
multi-threading, etc)

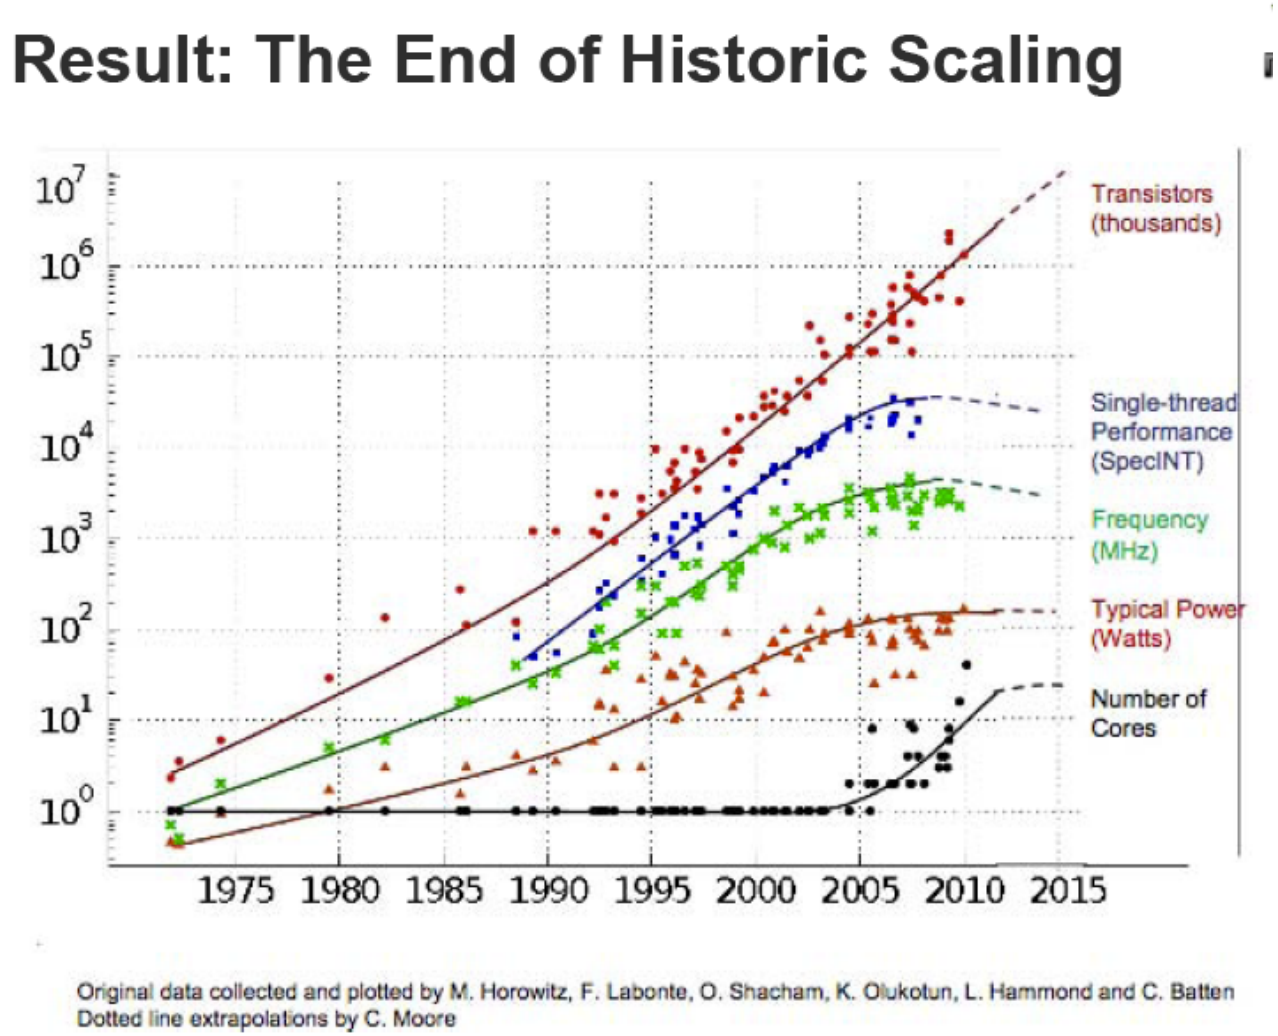

C Moore, *Data Processing in ExaScale-ClassComputer Systems,* Salishan, April 2011

Supercomputer scaling

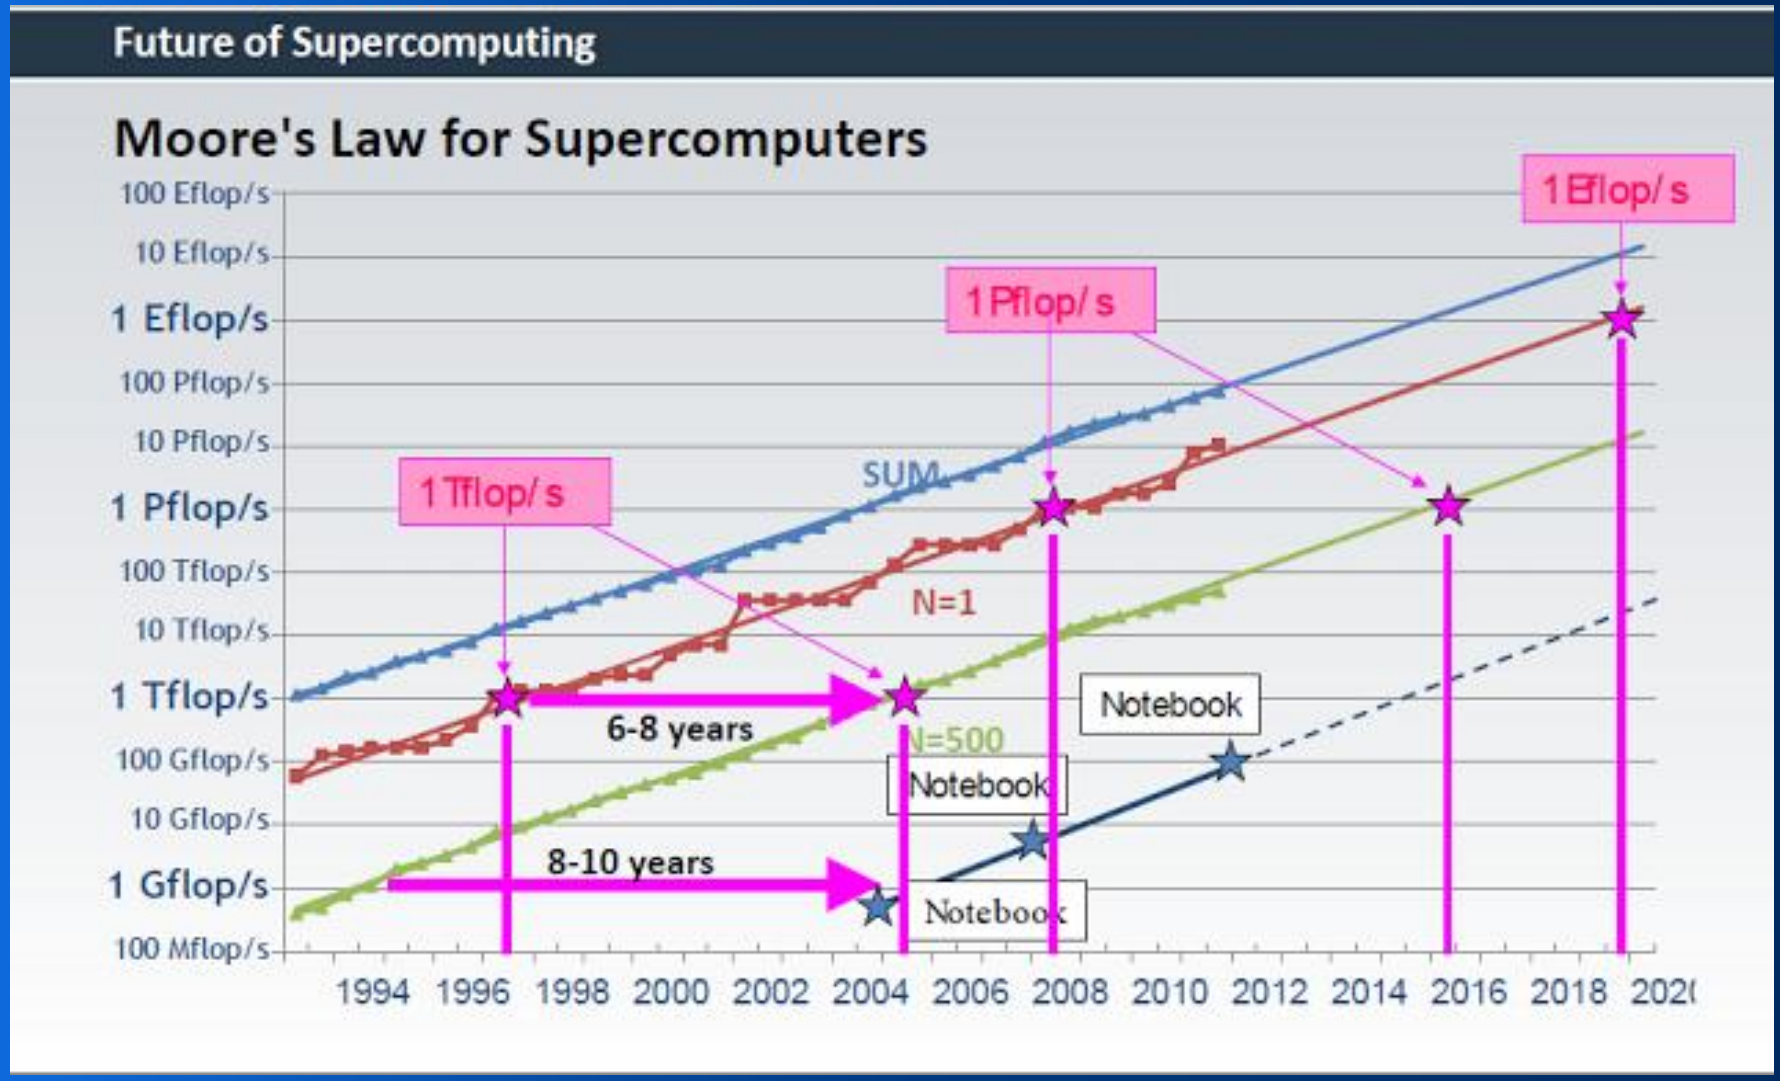

Evolution of Peak Performance and Memory (Sum over all LRZ systems)

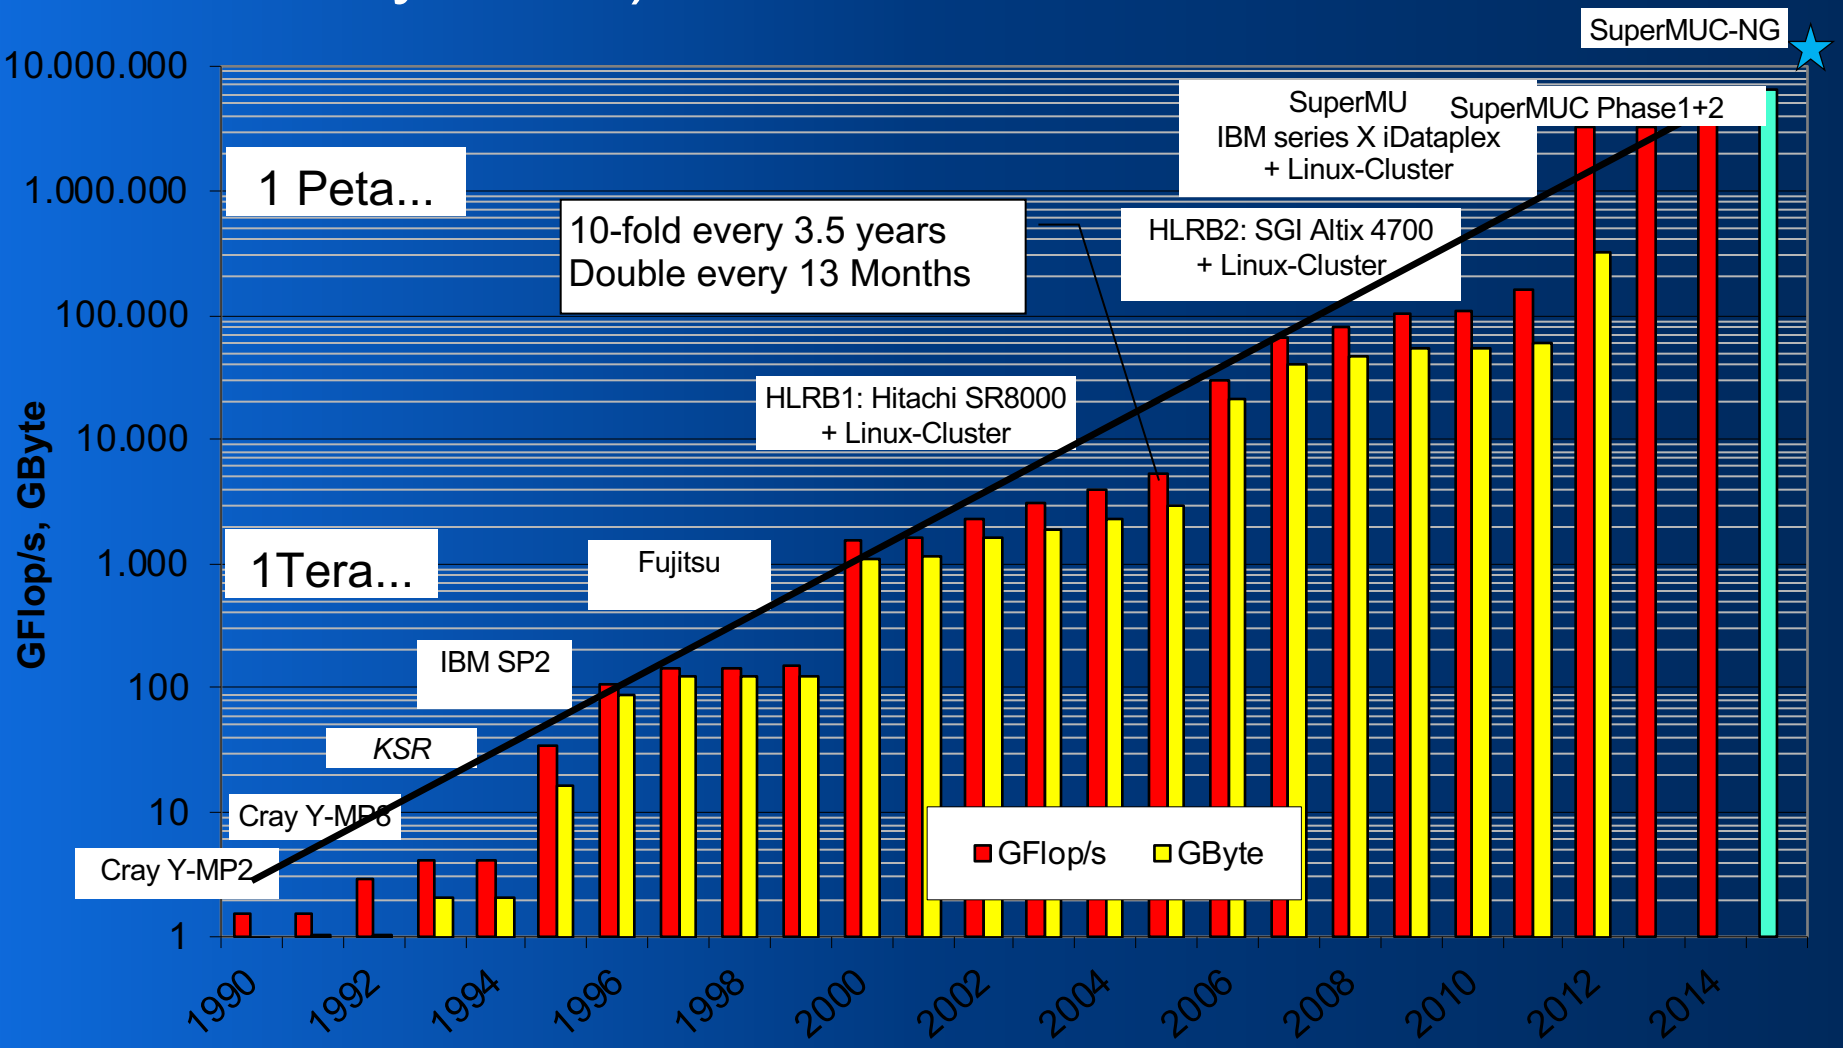

SuperMUC Utilization 2017

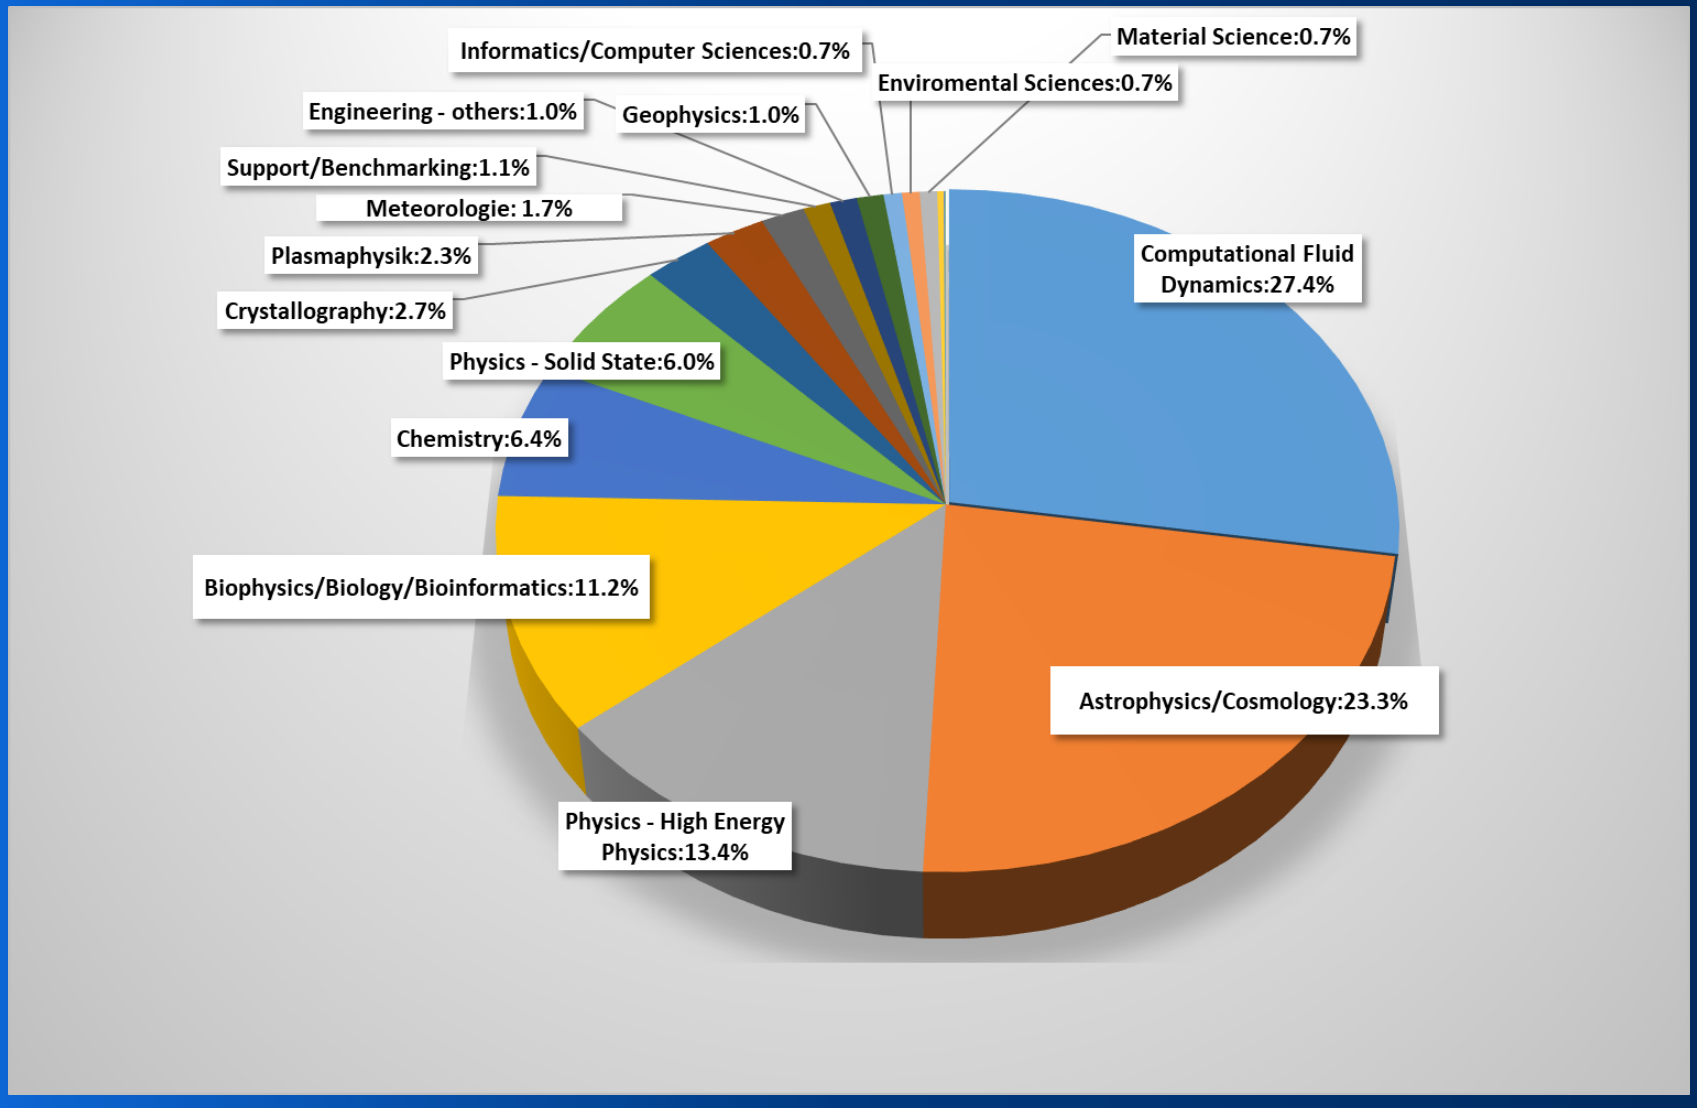

SuperMUC Jobsizes 2012-2017

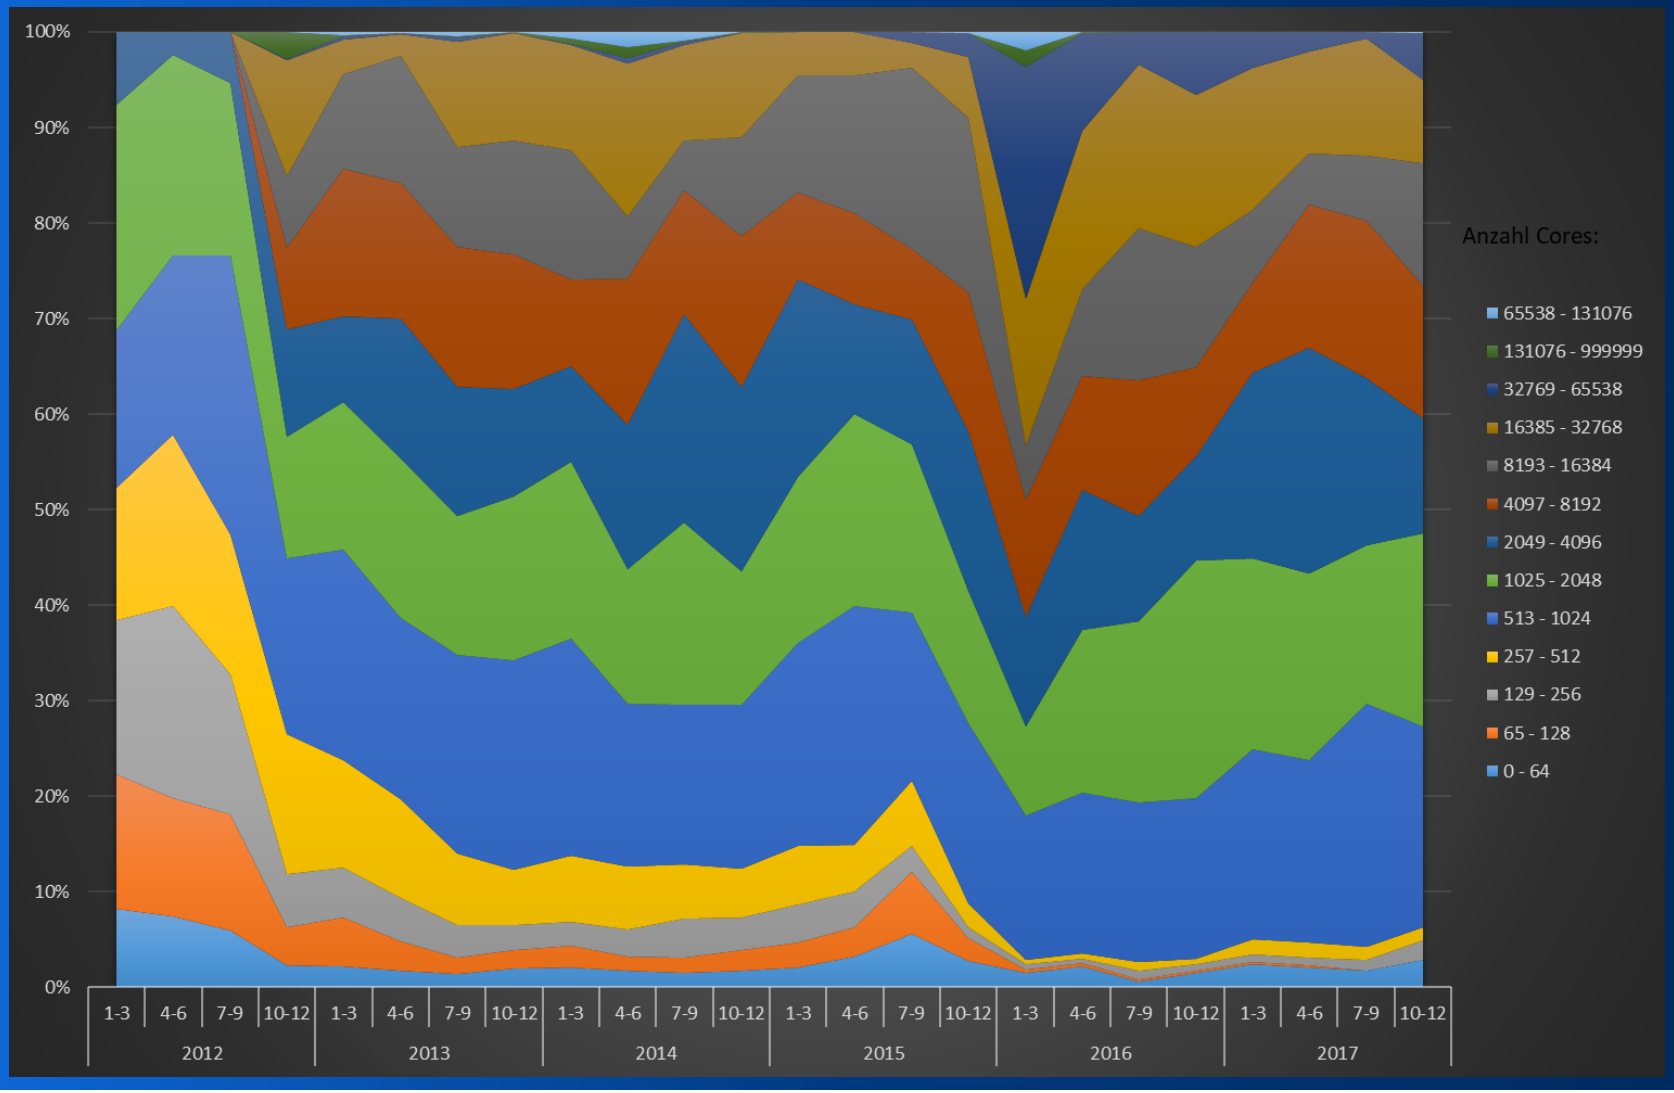

Supercomputer Layout

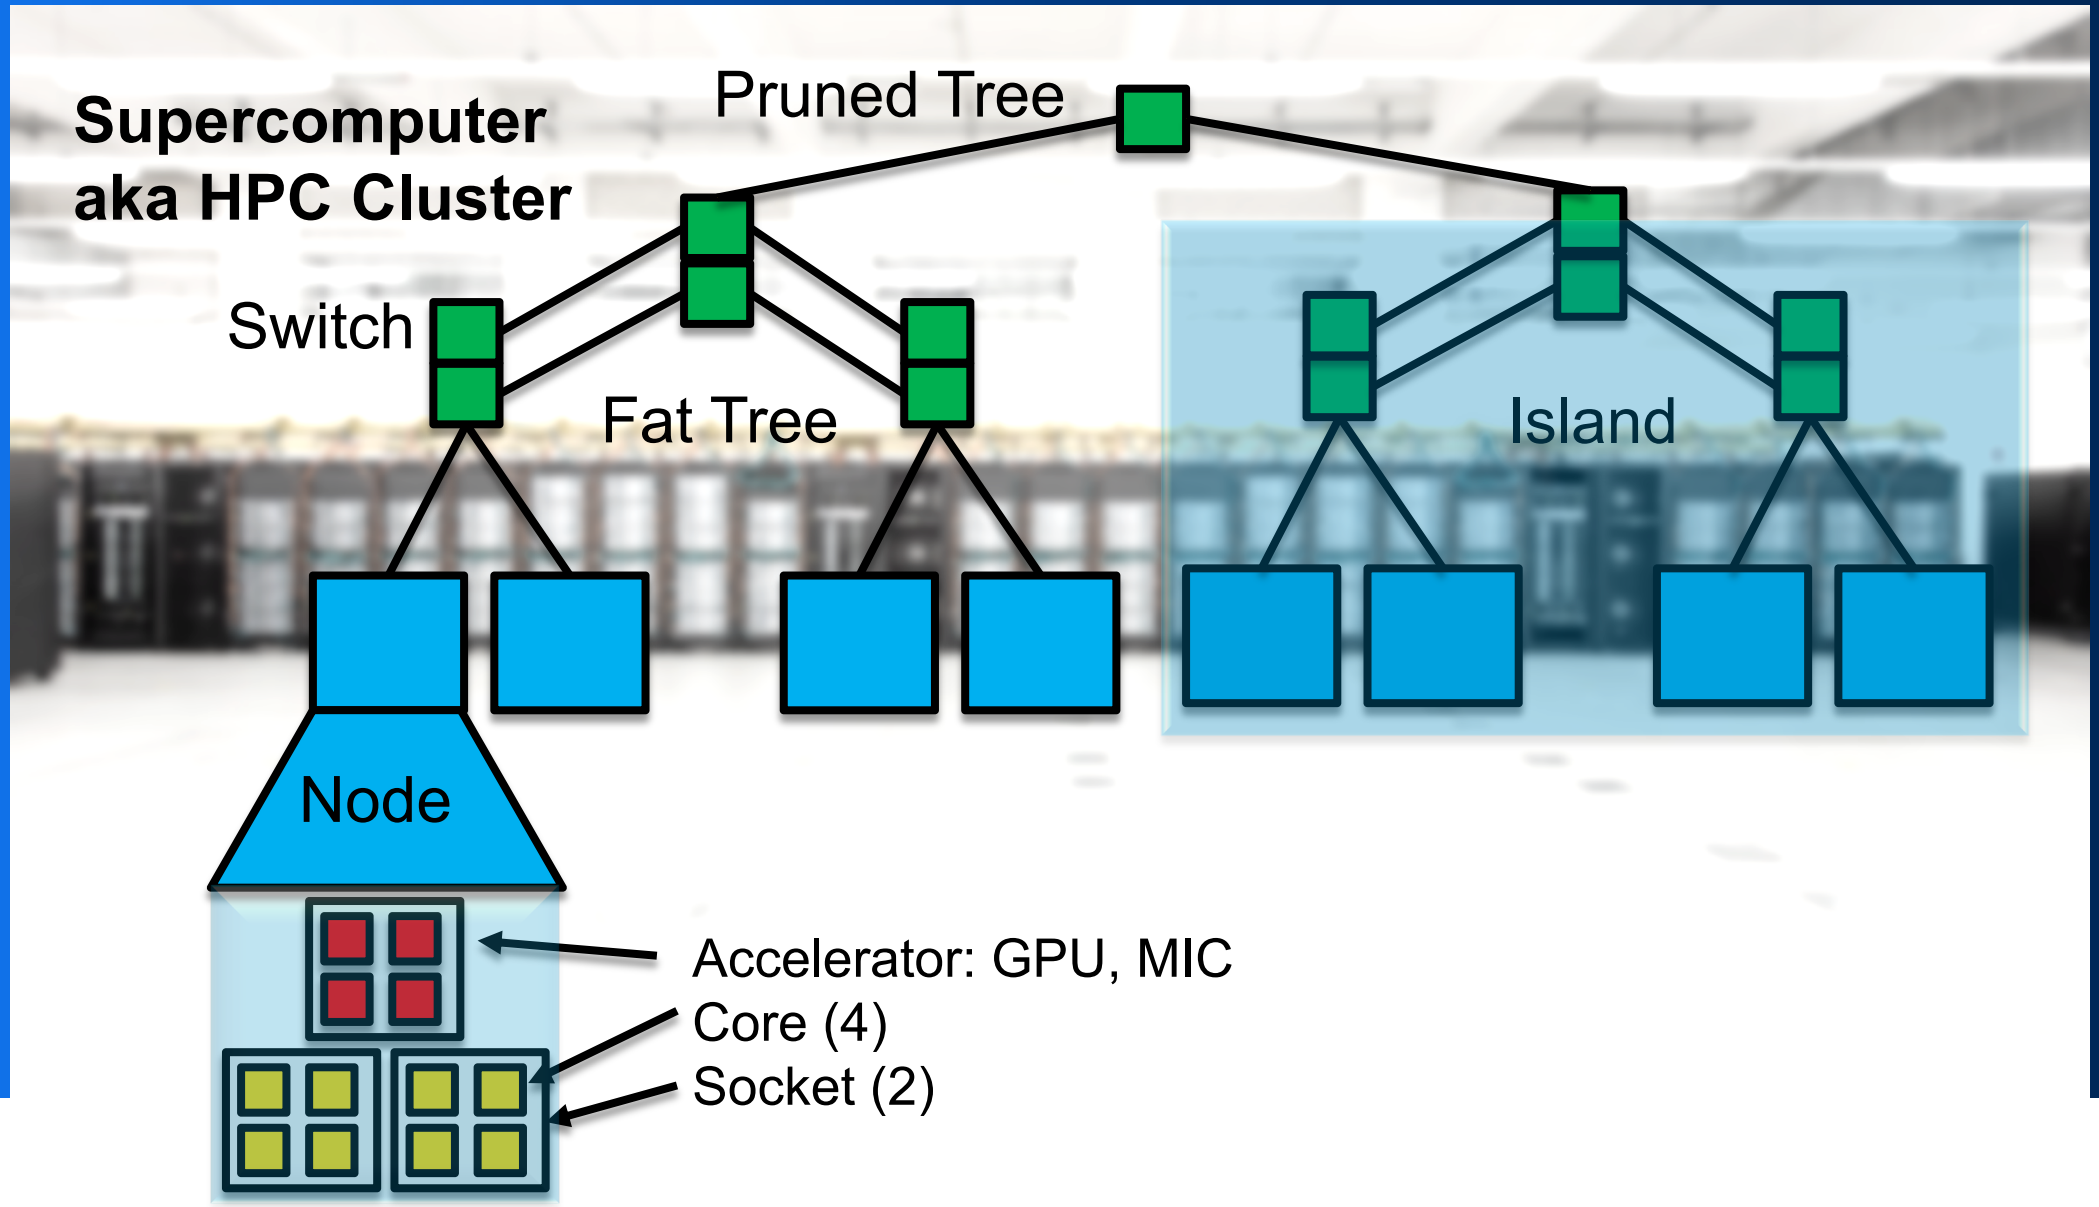

So... what is a Supercomputer?

- It has many off-the-shelf CPUs with vector instructions (AVX, AVX2, AVX512)
- The diskless nodes are connected by a high-speed internal network (Infiniband, OmniPath)
- The compute nodes (no ssh access) have to be programmed using Message Passing (MPI, GPI, ibverbs)
- All nodes are connected to a parallel file system (GPFS, Lustre) which needs special libraries for full speed (MPI-I/O)
- Programs cannot be run interactively, but have to be submitted to the batch scheduler (LoadLeveler, SLURM)
- The Operating System is a version of UNIX(e.g. Linux)

What is a Supercomputer (not)?

- It has overclocked high-speed processors? **NO**
- It has a big internal RAM? **NO (maybe)**
- It runs MS Windows? **NO**
- The CPU runs faster than a desktop PC? **NO**
- It will run my software without changes? **NO (maybe)**
- It will run my software with millions of threads? **NO**
- It will run my old trusted executable? **NO (maybe)**
- It will run my Excel spreadsheets? **NO**
- You can use it interactively? **NO (maybe)**

Levels of Parallelism

- Node Level (e.g. SuperMUC-NG has 6336 nodes)
- Accelerator Level (e.g. DGX-1 has 2 CPUs and 8 GPUs)
- Socket Level (e.g. teramem has 4 CPUs with 24 cores)
- Core Level (e.g. CoolMUC-3 has 64 cores with AVX512)
- Vector Level (e.g. AVX512 has 32 vector registers)

SuperMUC-NG Peak Performance: **25.3 PFlop/s** =

6336 Nodes x **2** Sockets x **24** Cores x **32** Vectors x **2.6** GHz

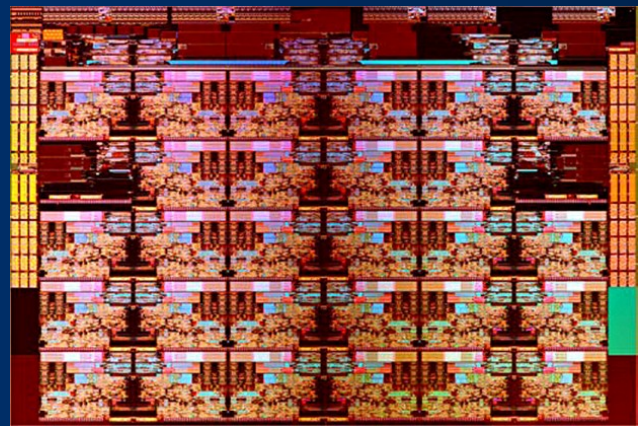

RStudio

Web GUI for R and Linux Shell

- read and edit code
- enter commands
- open data files
- plot data
- save plots to pdf/jpg

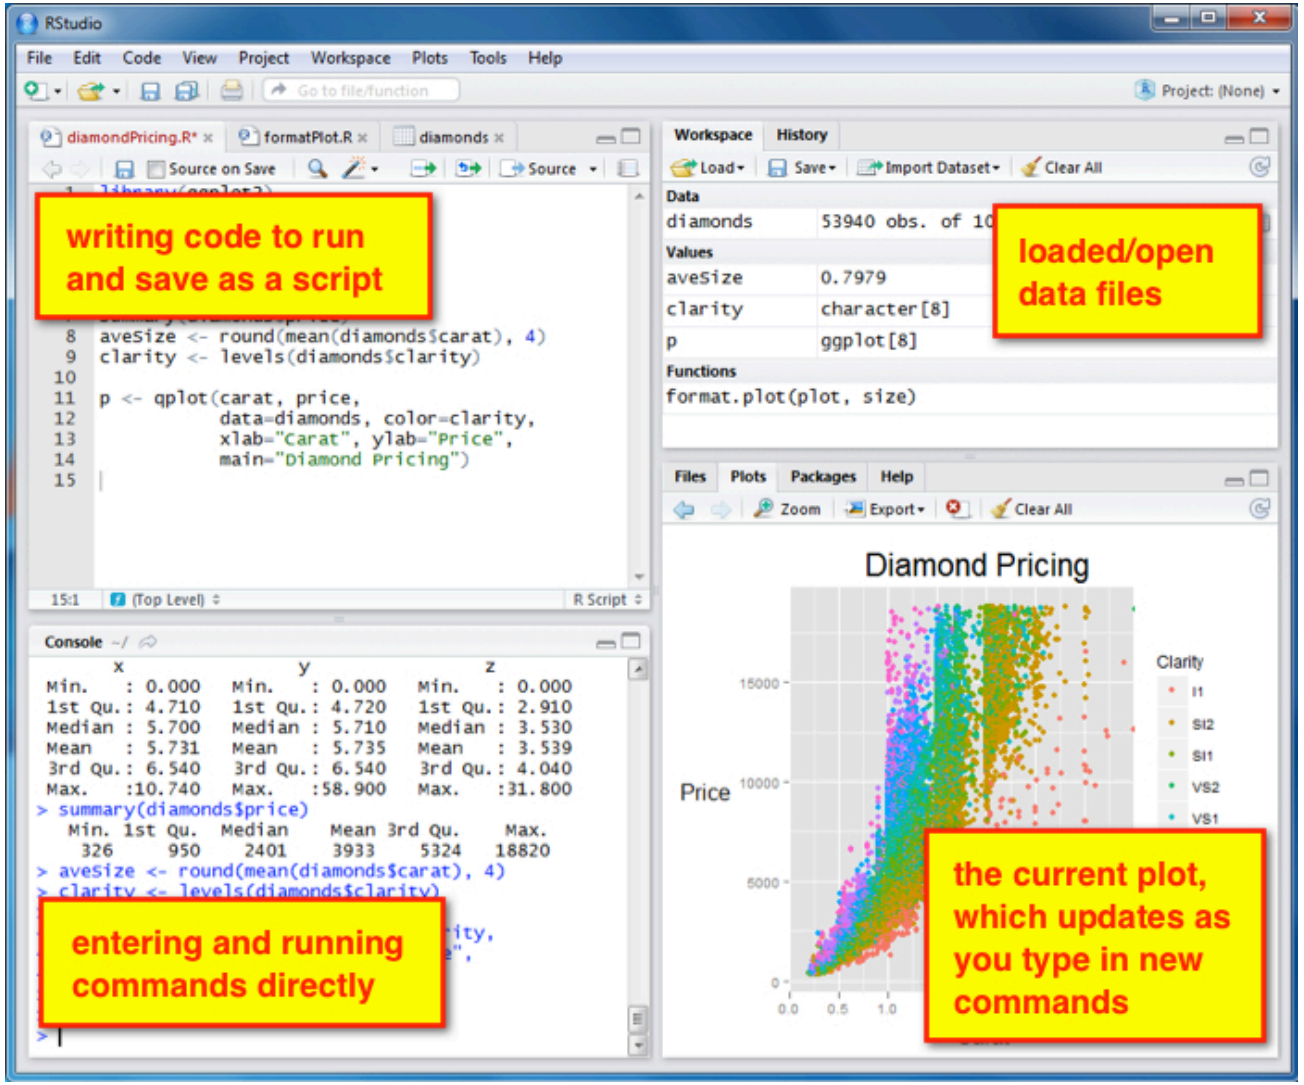

putty

Hostname:

lxlogin5.lrz.de
or
lxlogin6.lrz.de
or
lxlogin7.lrz.de

Enter userid and password when asked for

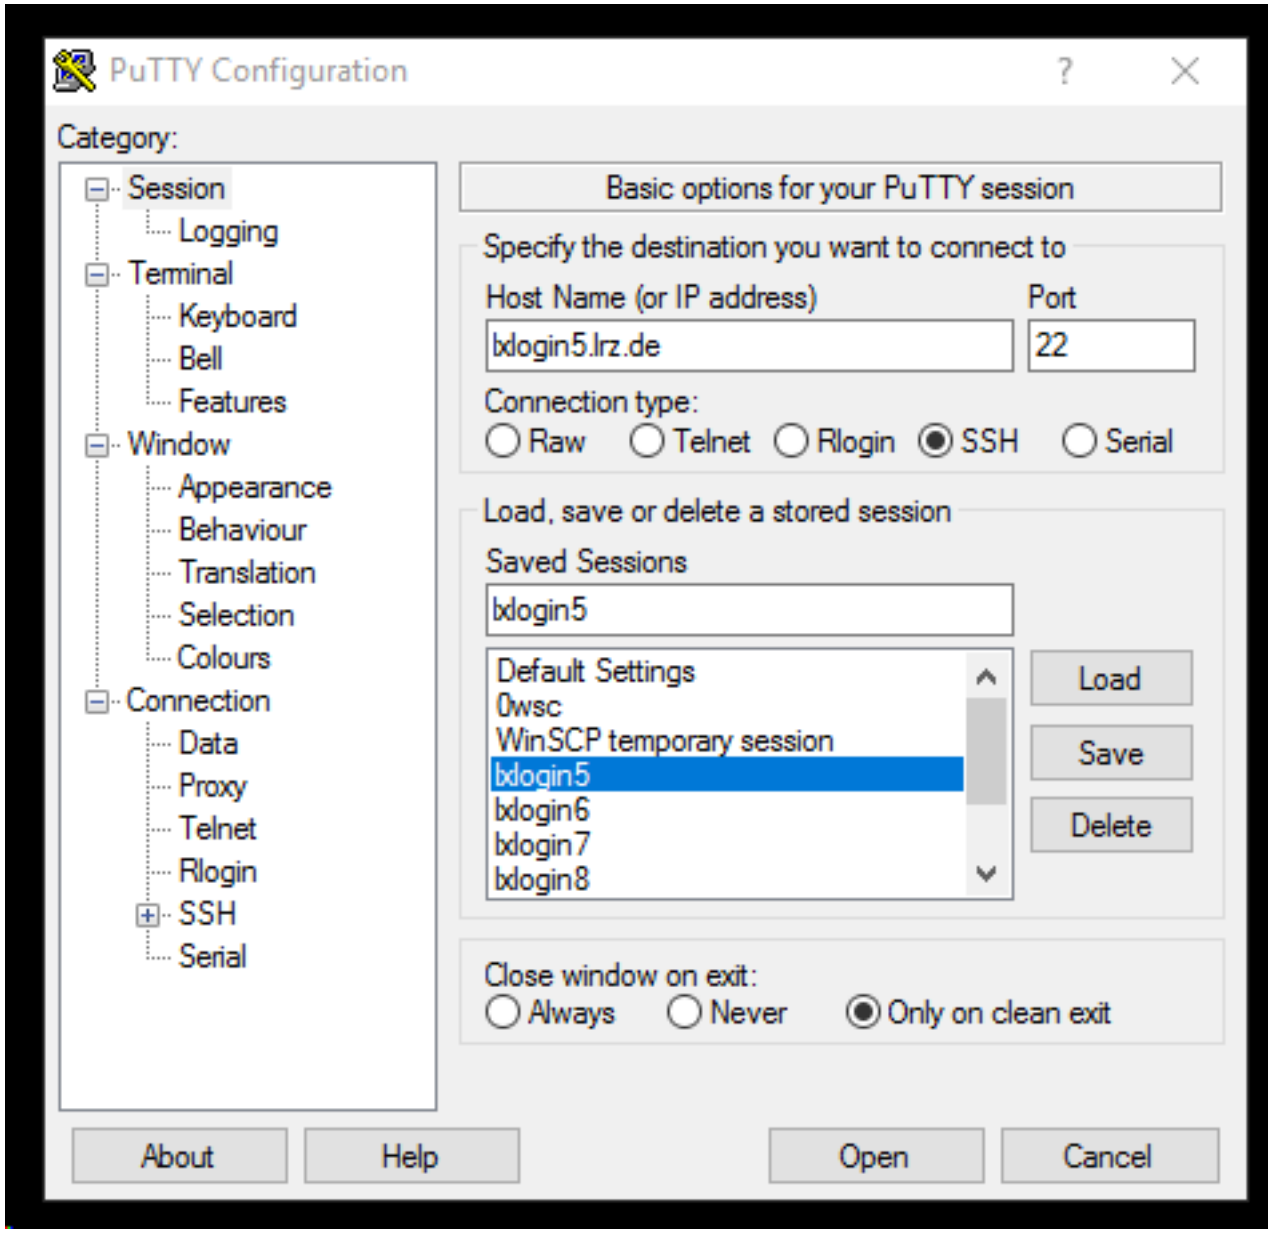

WinSCP

File Manager

Explorer like GUI

Copy files from/to Linux Cluster

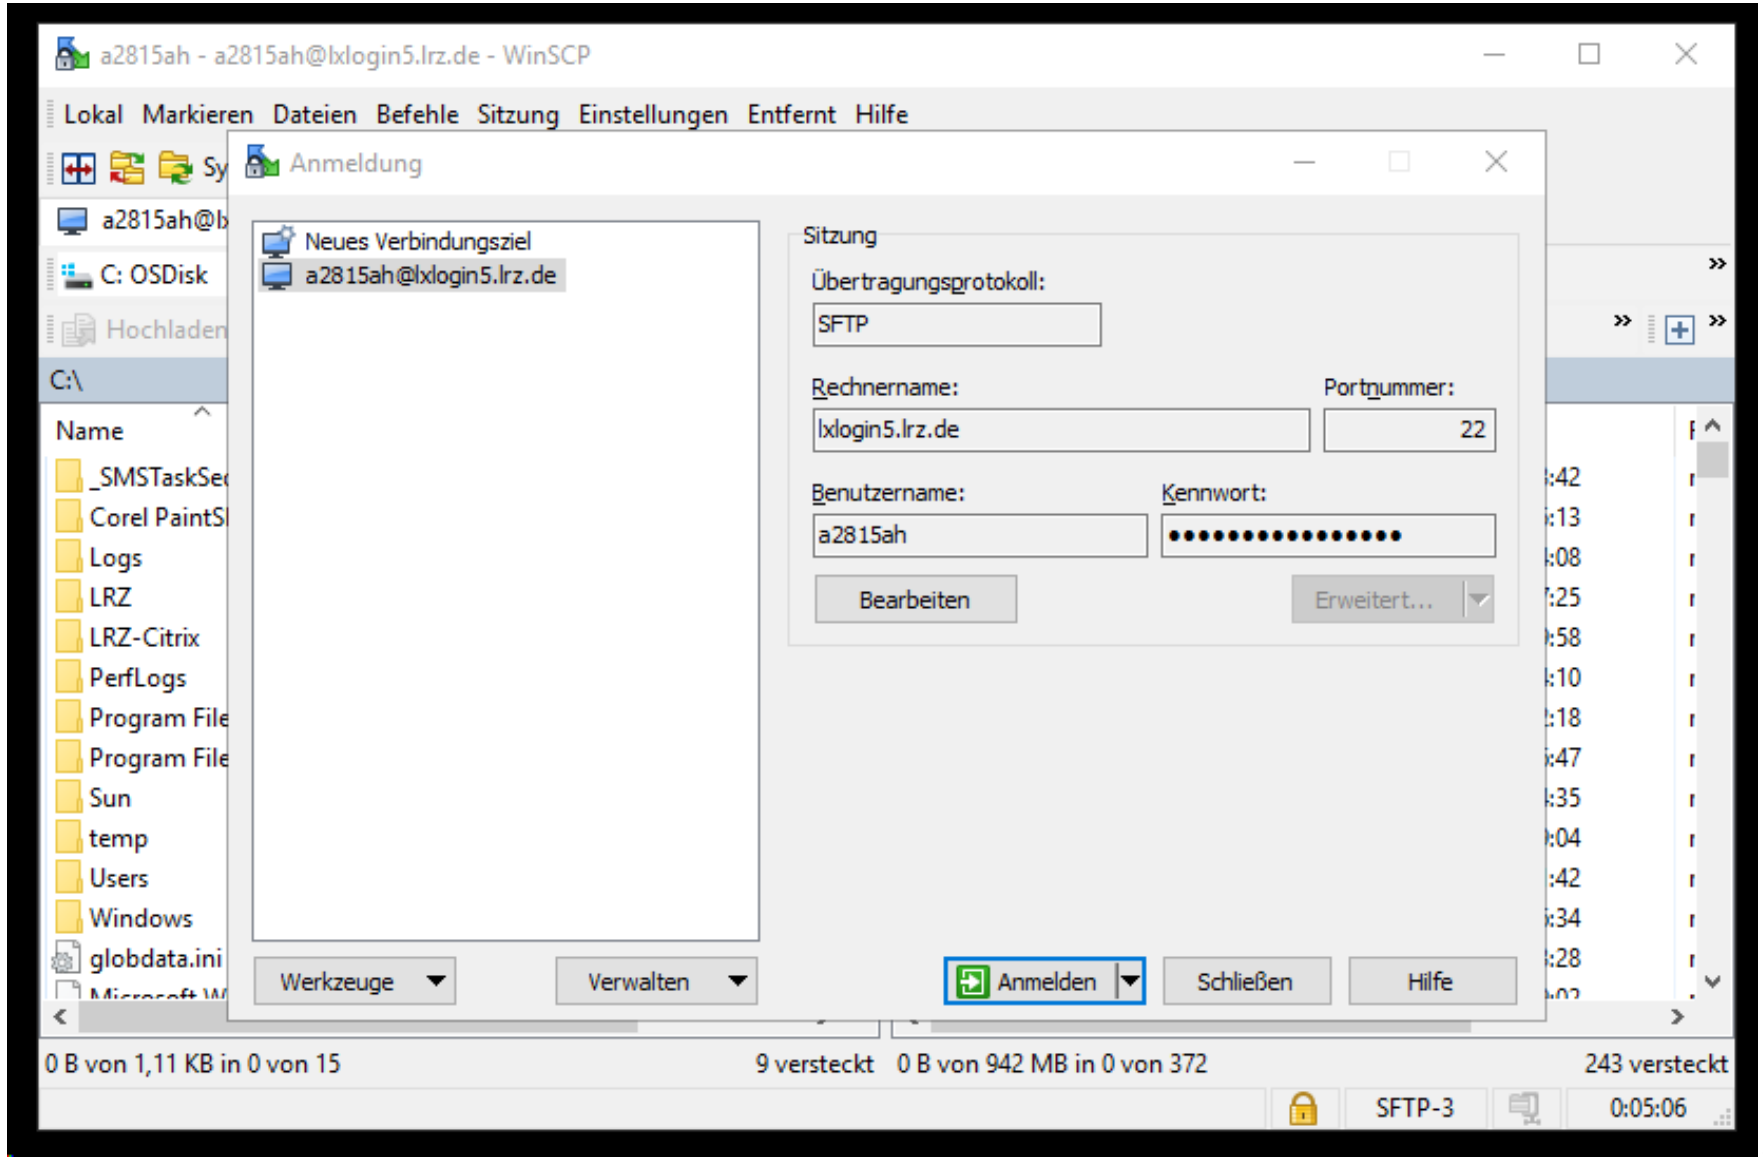

The LRZ Compute Cloud

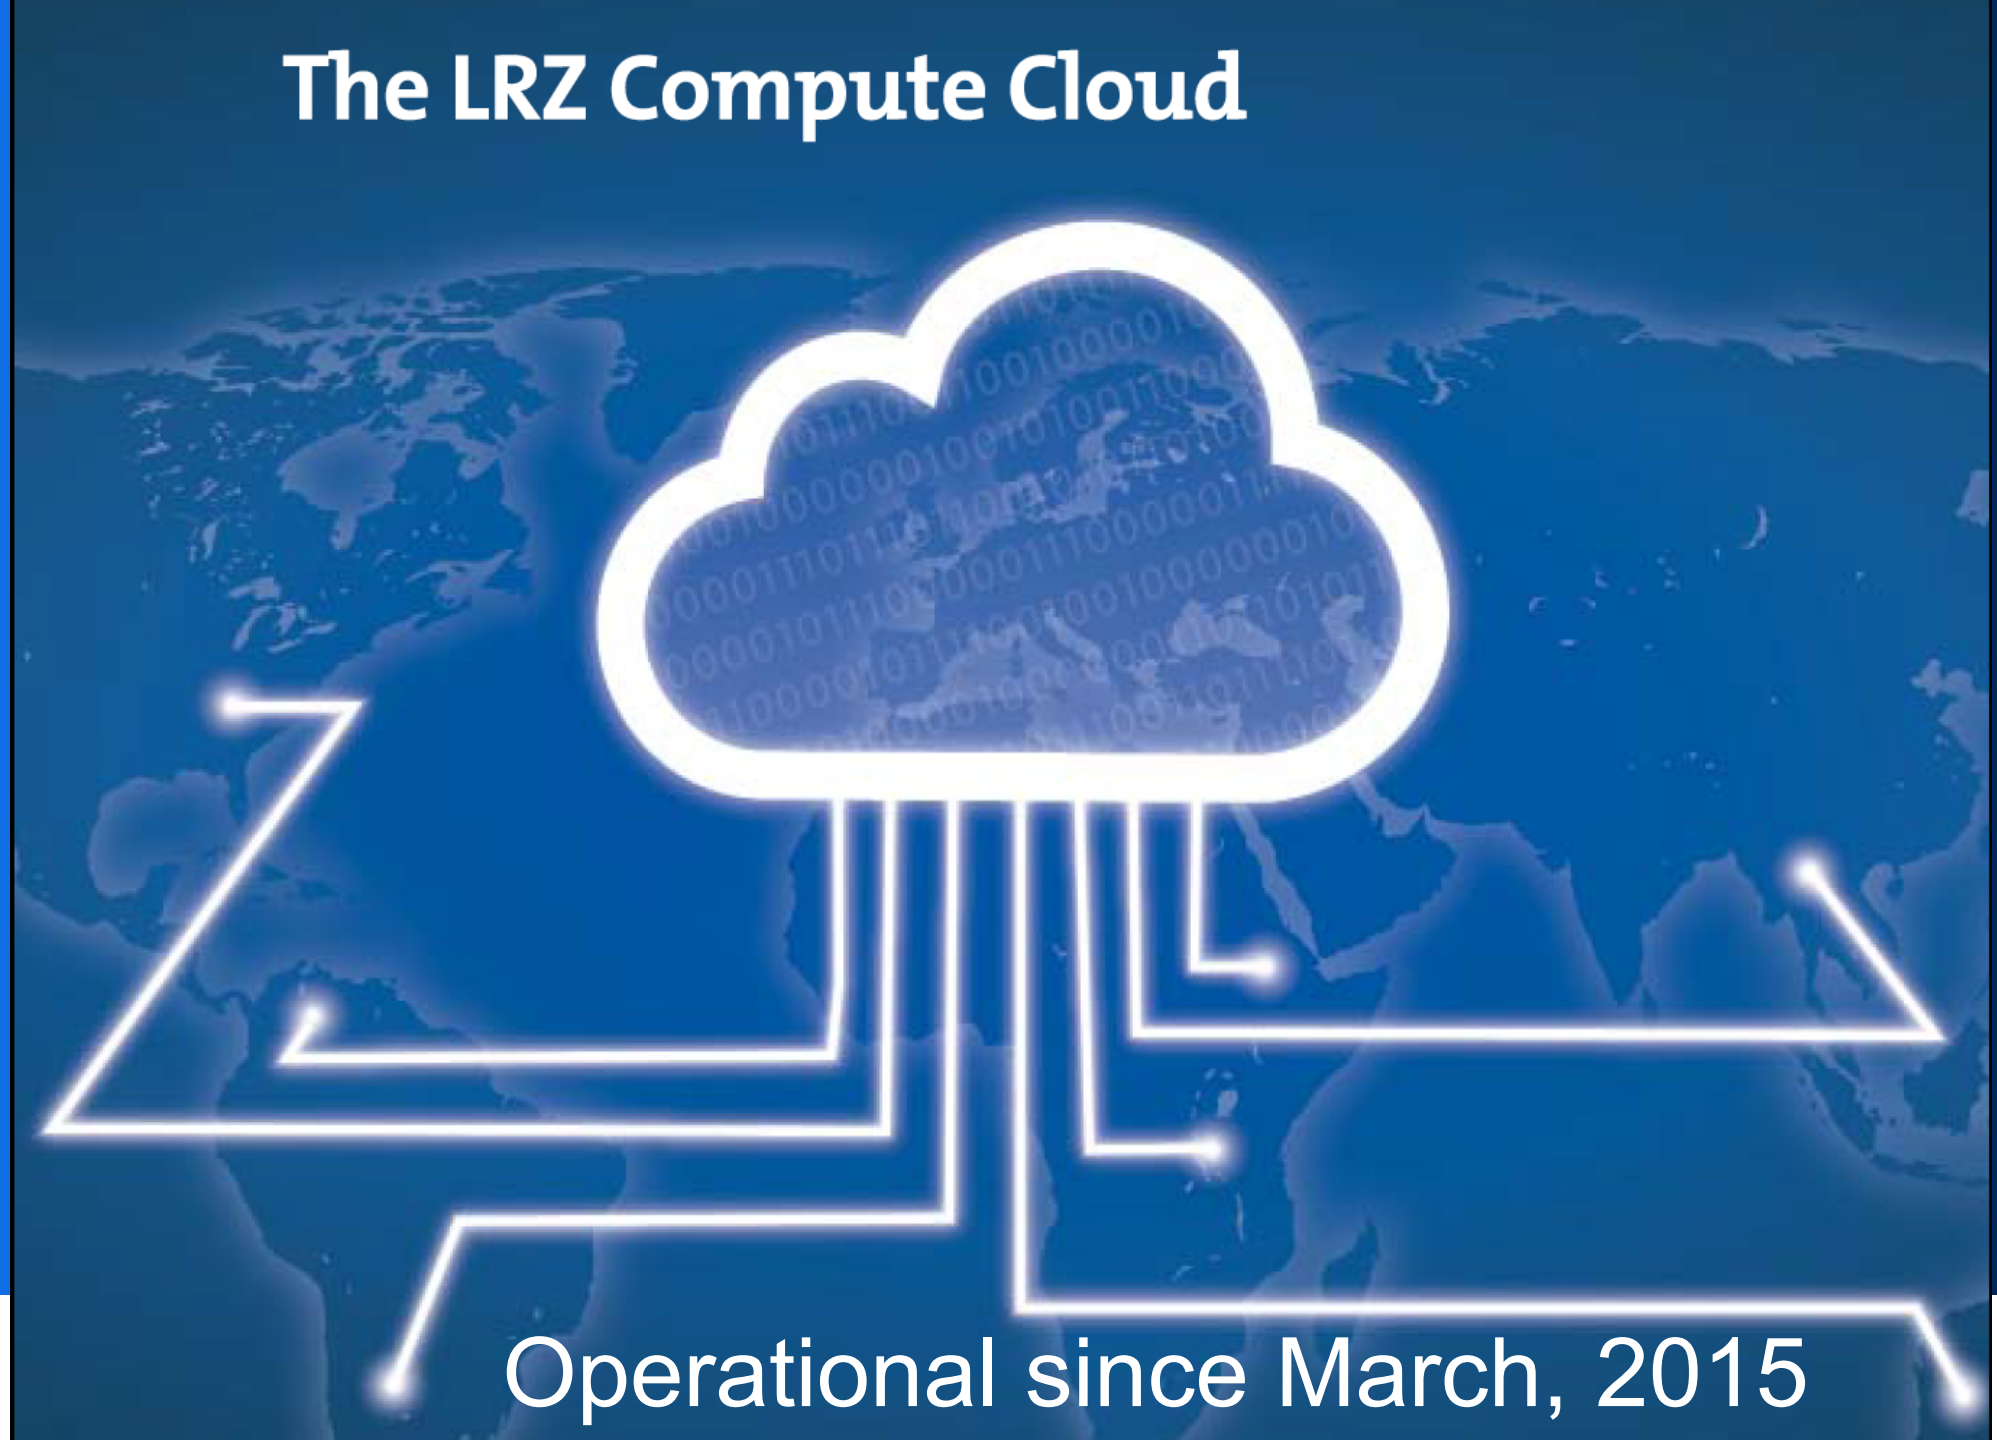

Operational since March, 2015

LRZ Compute Cloud: OpenNebula/OpenStack

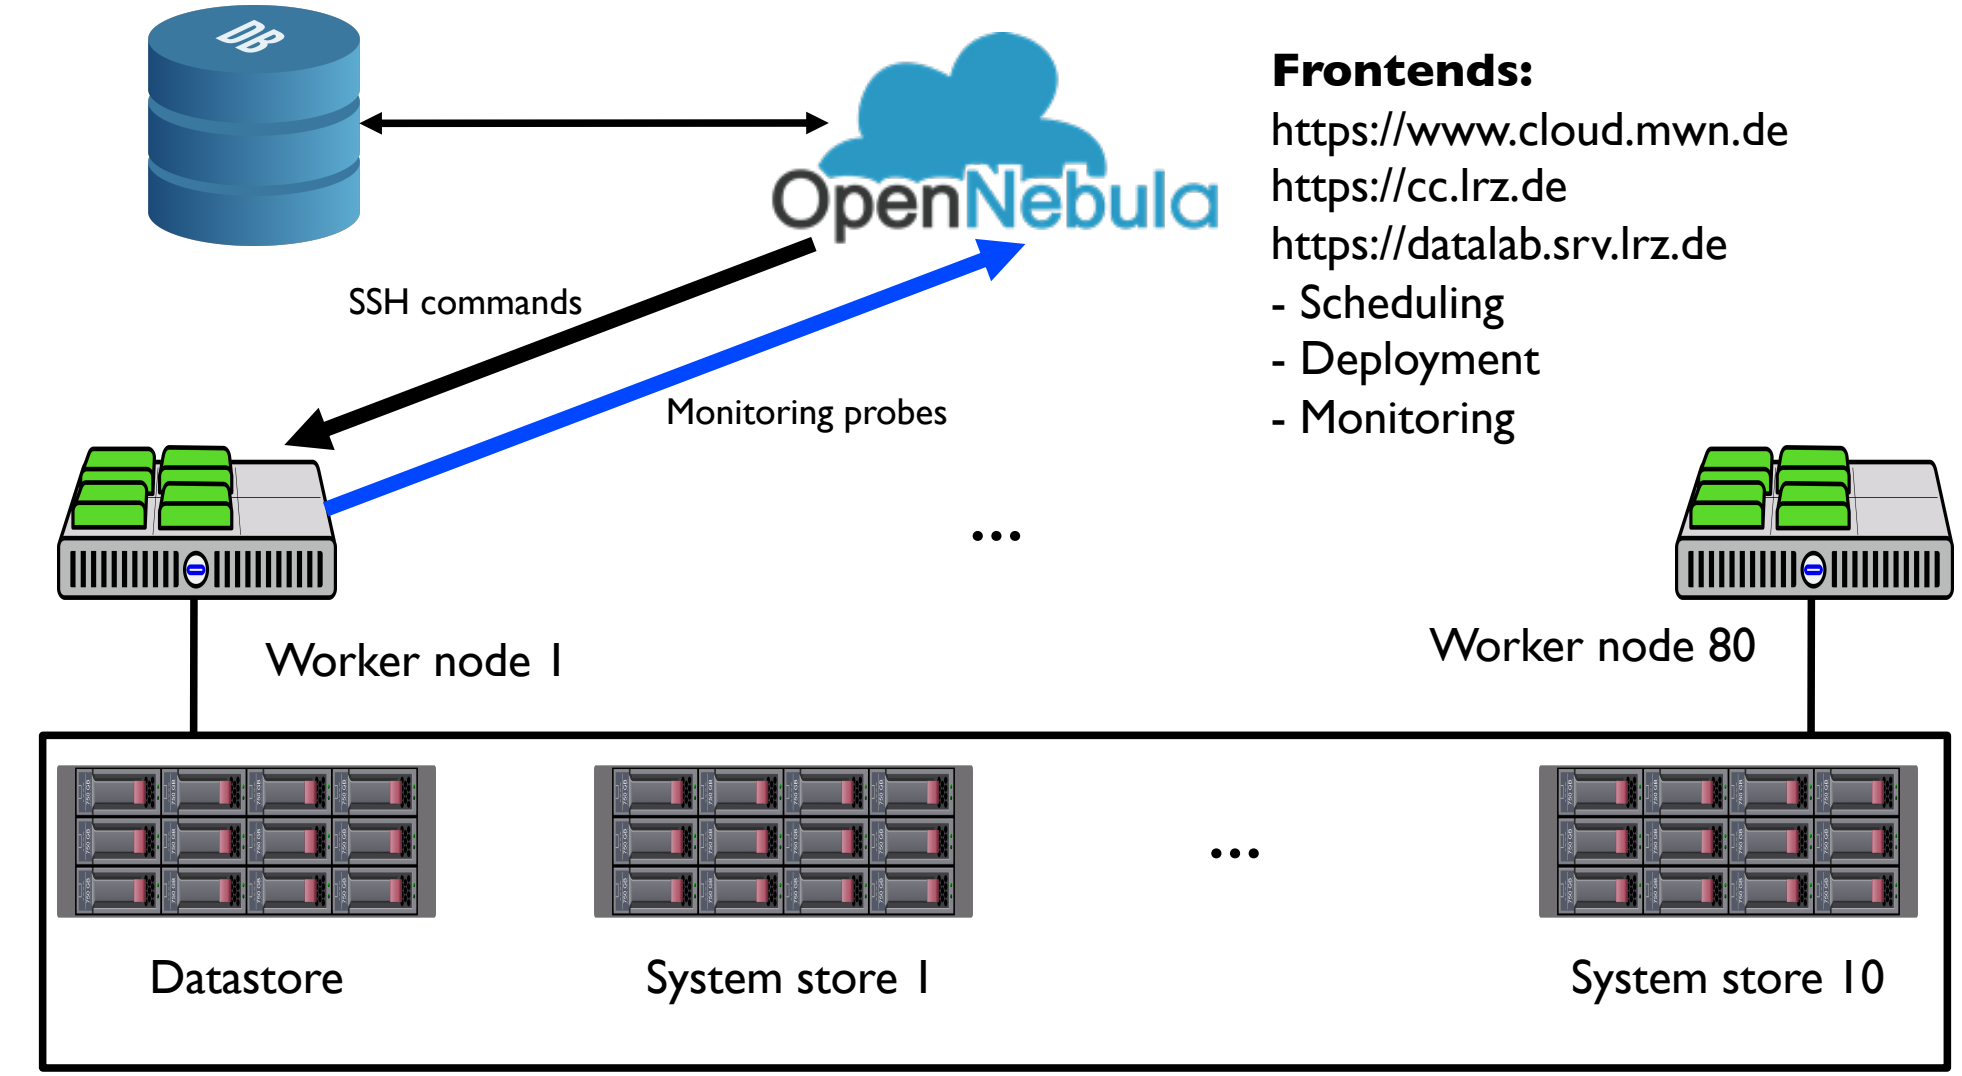

“Supercomputer in a box“: DGX-1 and Teramem1

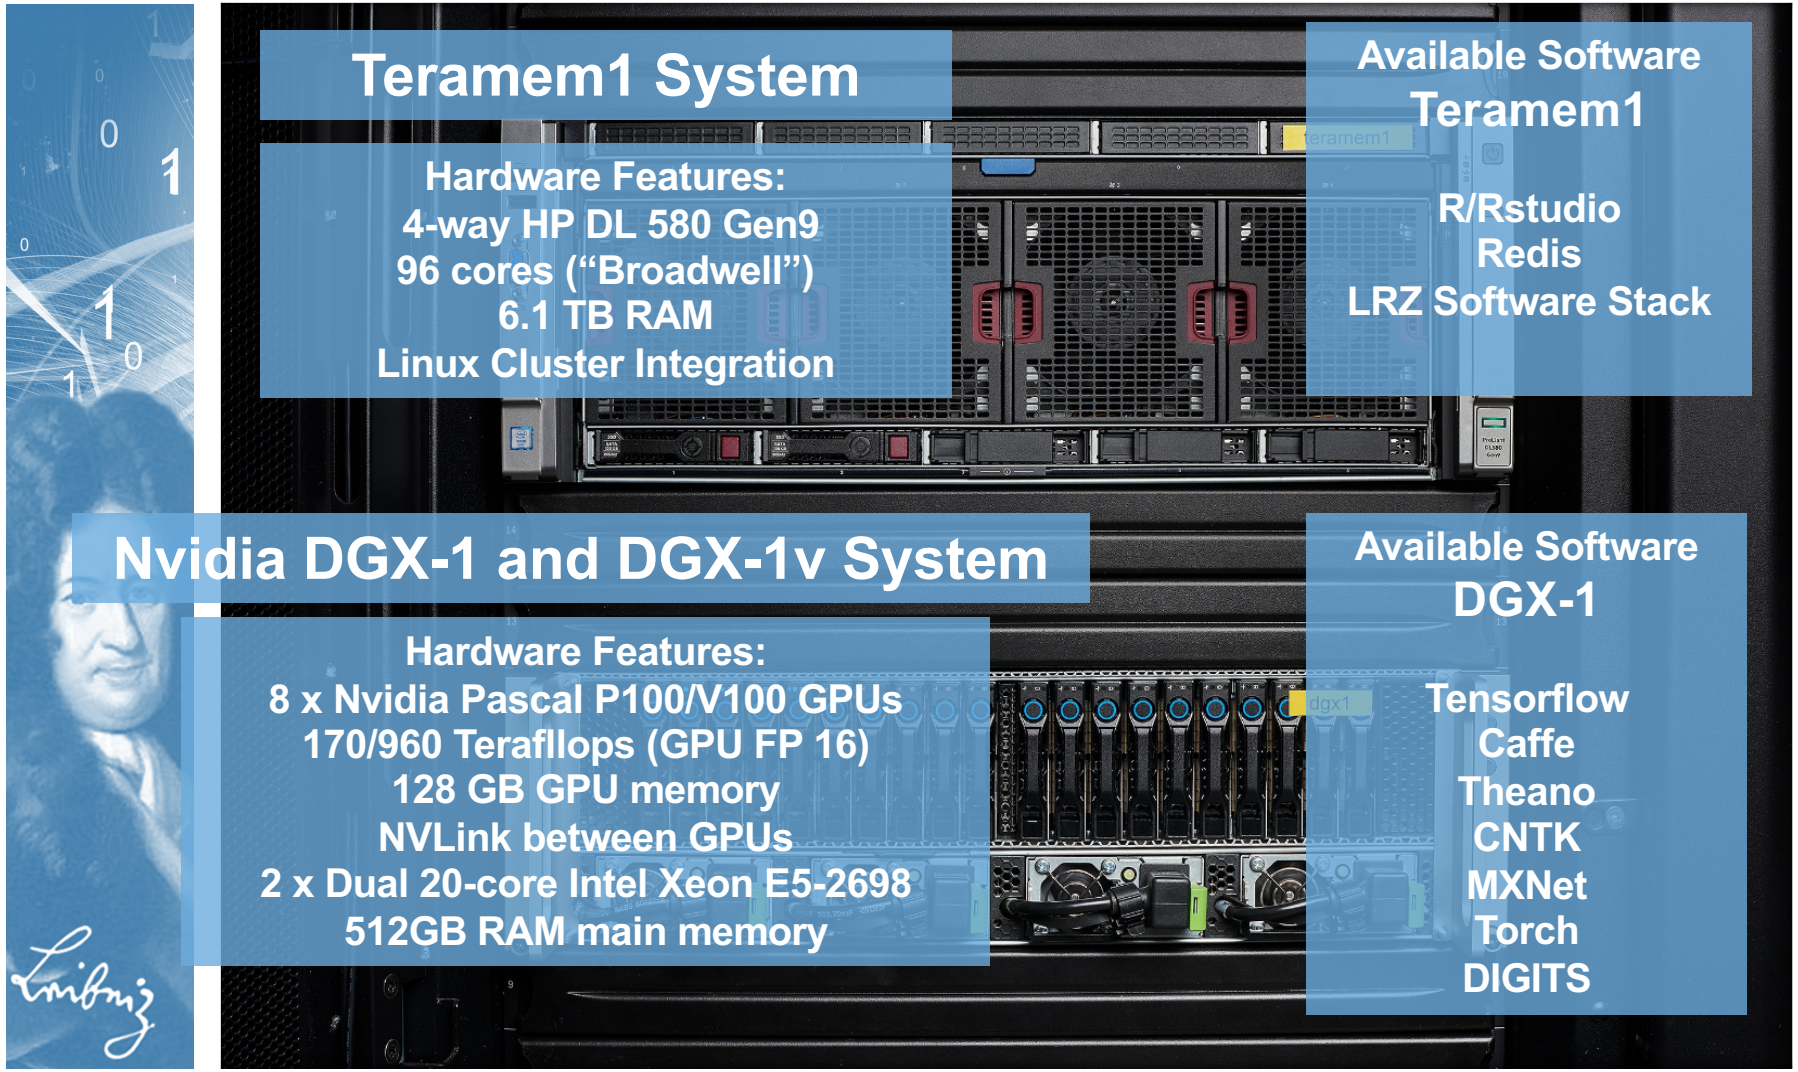

Jupyter notebooks

Run your interactive program in the browser

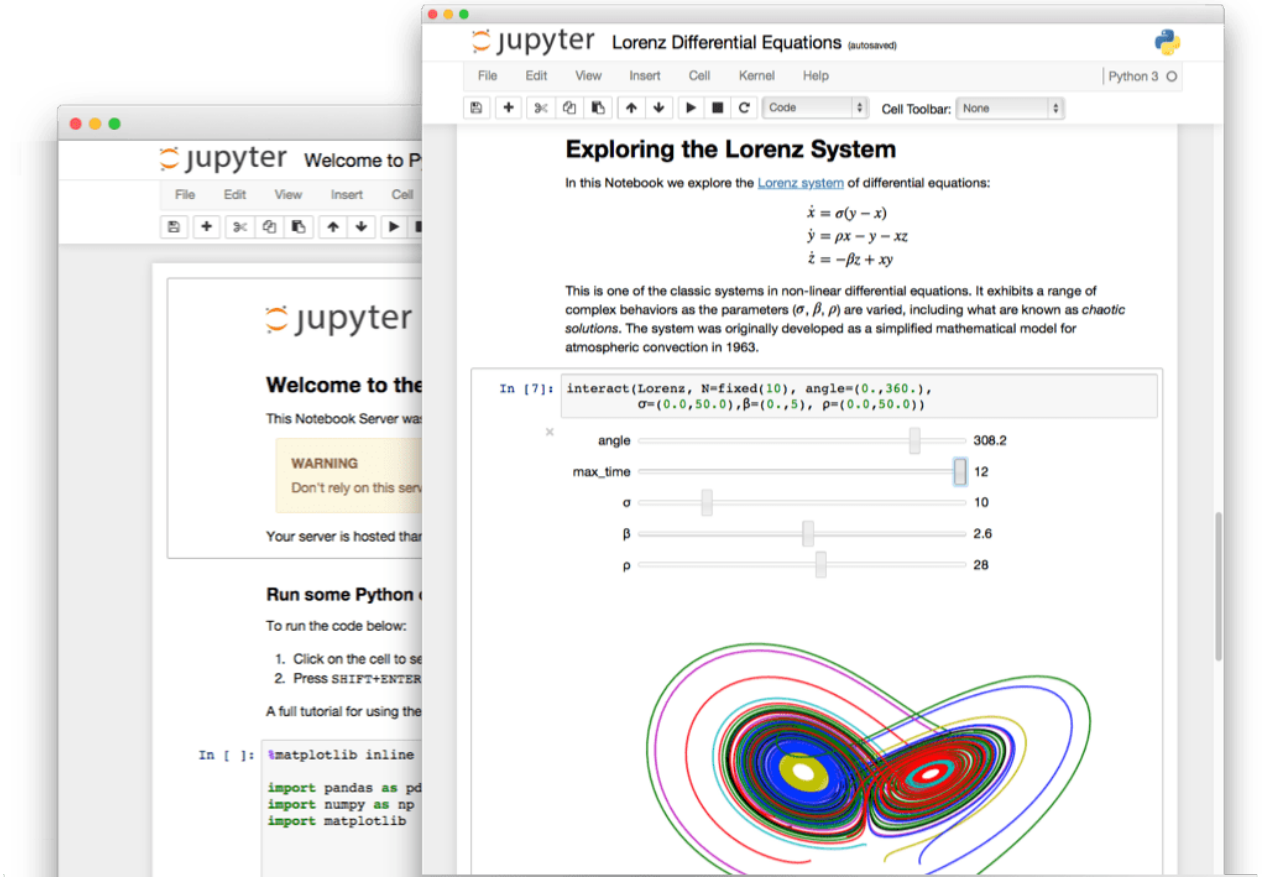

Docker

```
$ docker run hello-world
Hello from Docker!
```

Advantages:

- reproducible
- portable
- lightweight
- Docker hub repository
- Jupyter out of the box

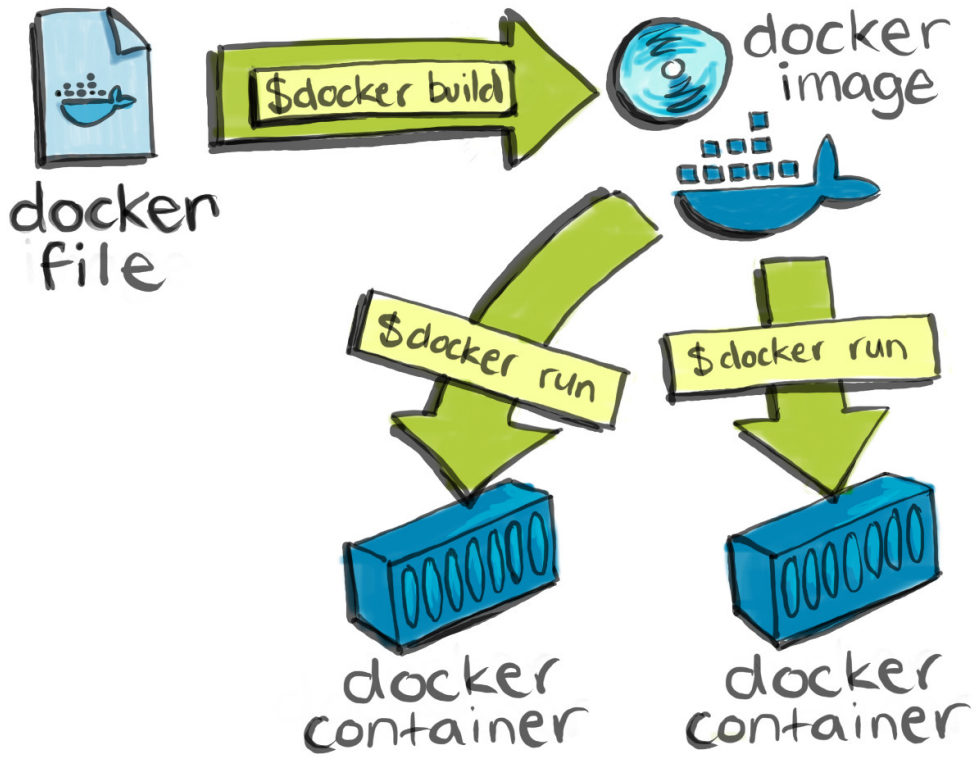

https://www.cloud.mwn.de/

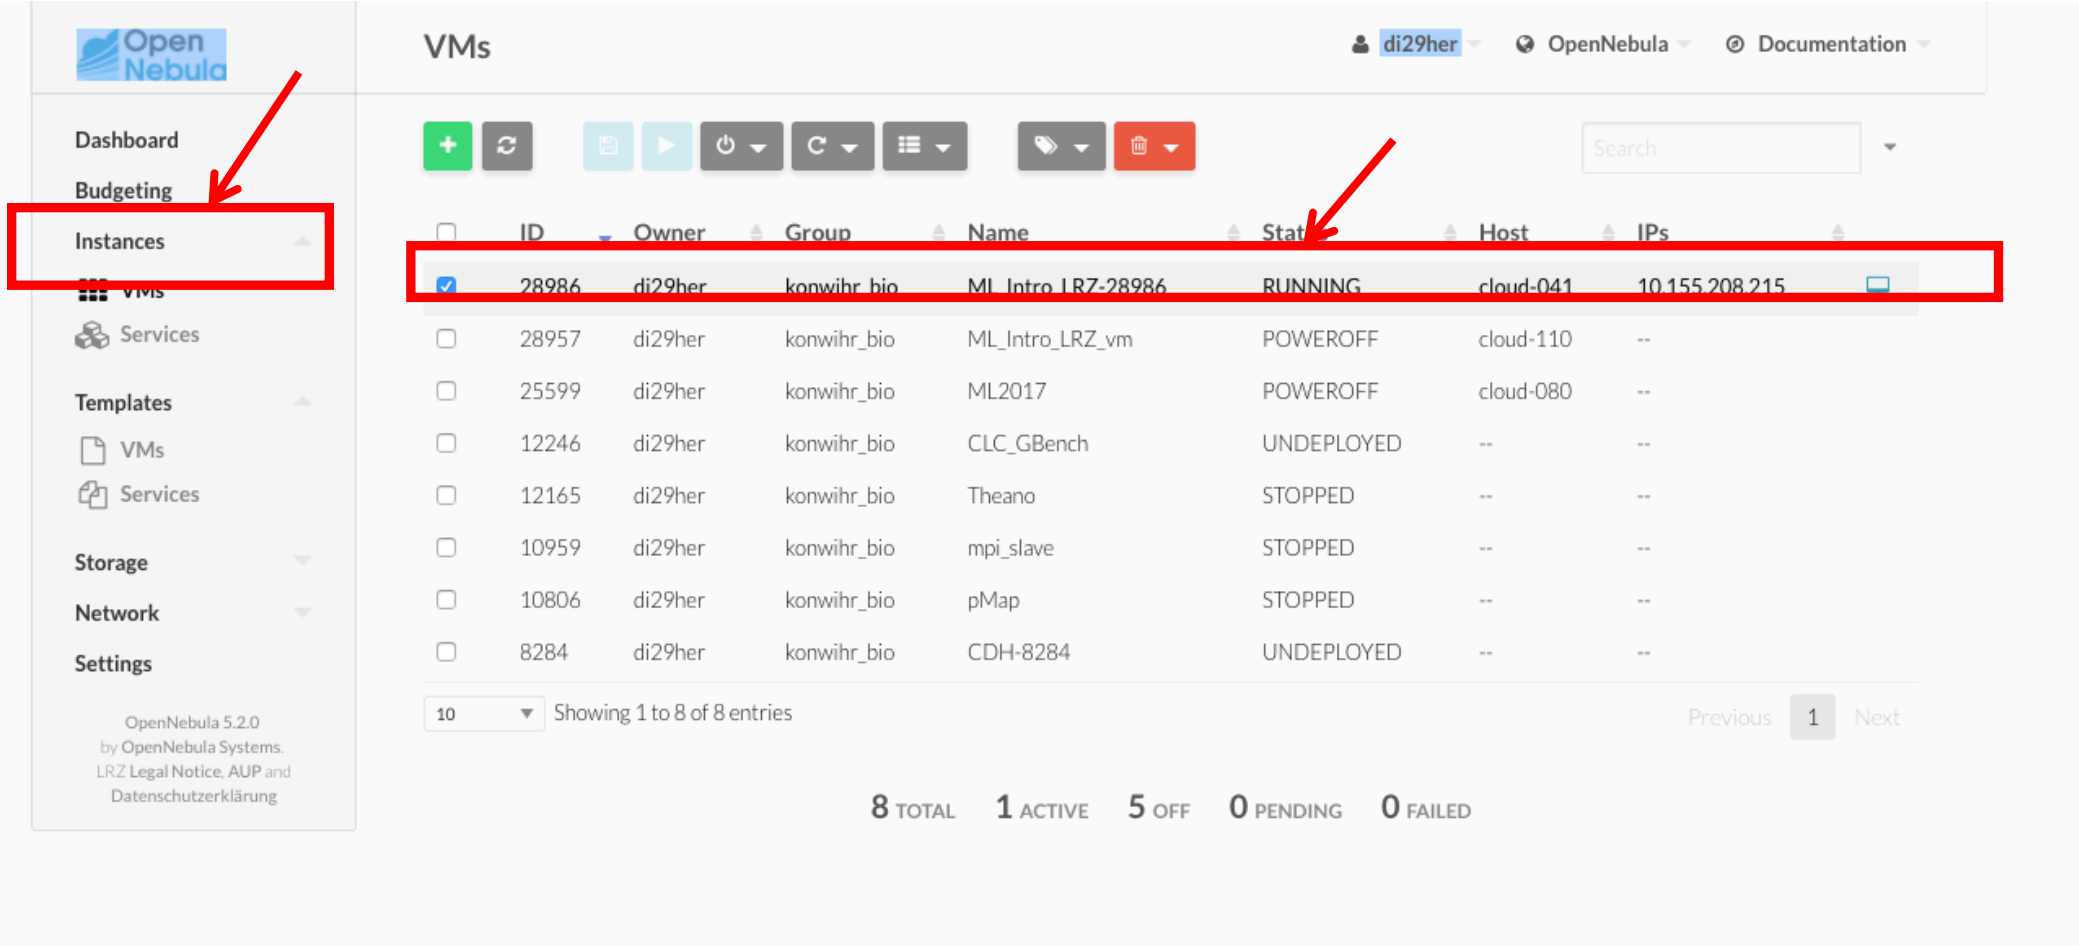

Terminal in the browser