Leibniz-Rechenzentrum
der Bayerischen Akademie der Wissenschaften

Using virtualisation for tailor-made HPC solutions

Ferdinand JAMITZKY,
Deputy Leader Application Group at Leibniz Supercomputing Centre, Garching

TERATEC 2016 Forum
Workshop 1 - Wednesday, June 29, 2016

Pilot Project: Bioinformatics Workflows at LRZ

Take existing bioinformatics projects at LRZ and implement them using workflow tools

- Metagenomics in Childhood Asthma research

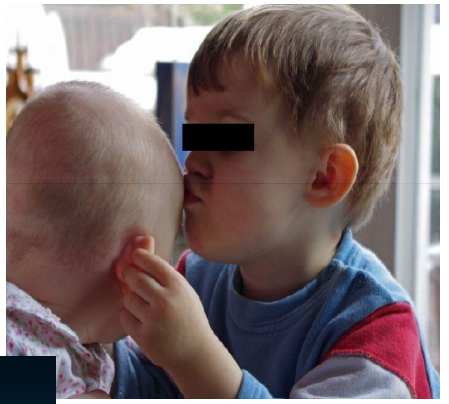

- Systems Biology: Predict effects of small changes in amino-sequence on protein function

Co-Design: Infrastructure/Platform-Provider, Workflow-Developers, End-Users

Bioinformatics Workflows

Proteomics

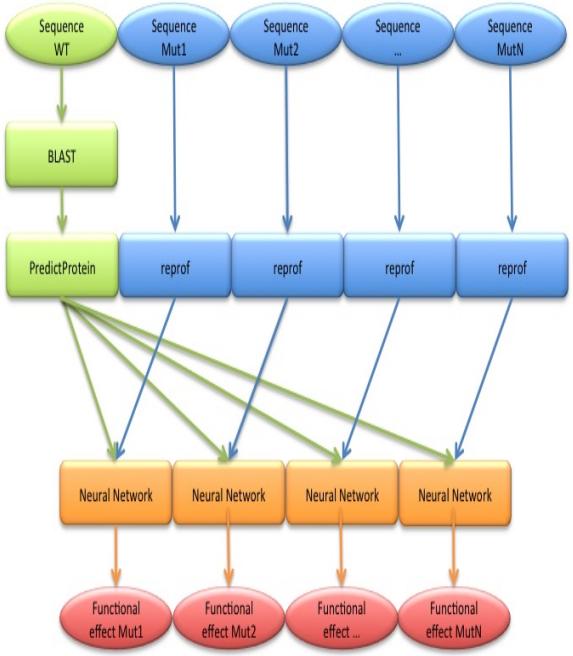

NGS

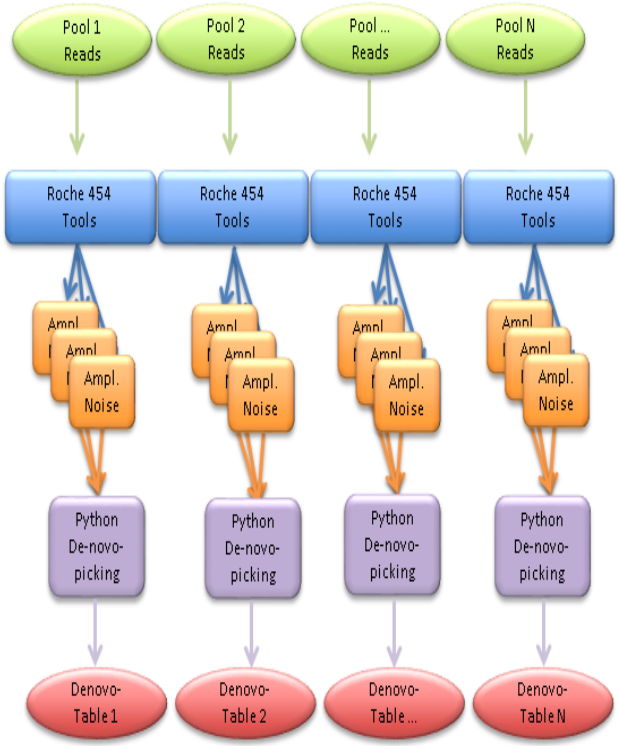

Systems Biology

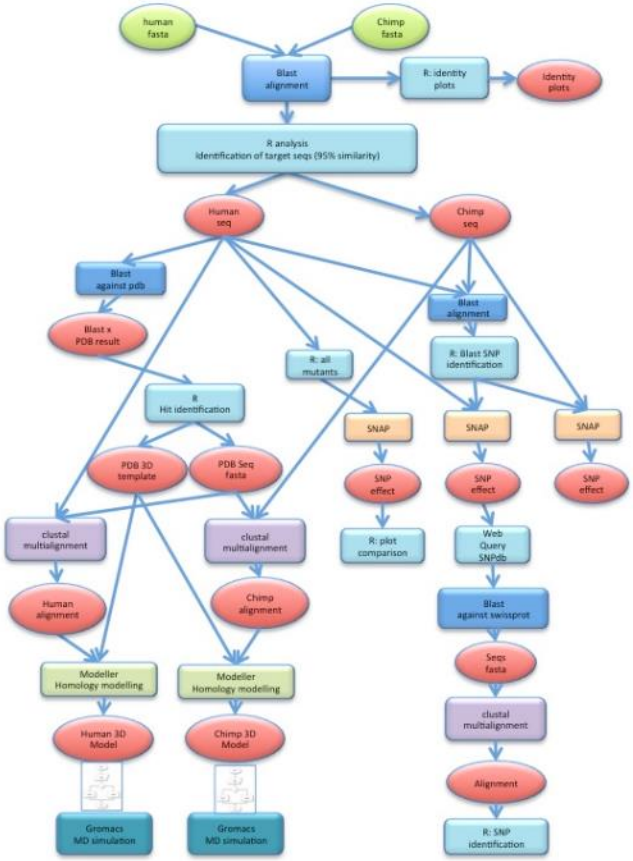

Mapping on LRZ Hardware

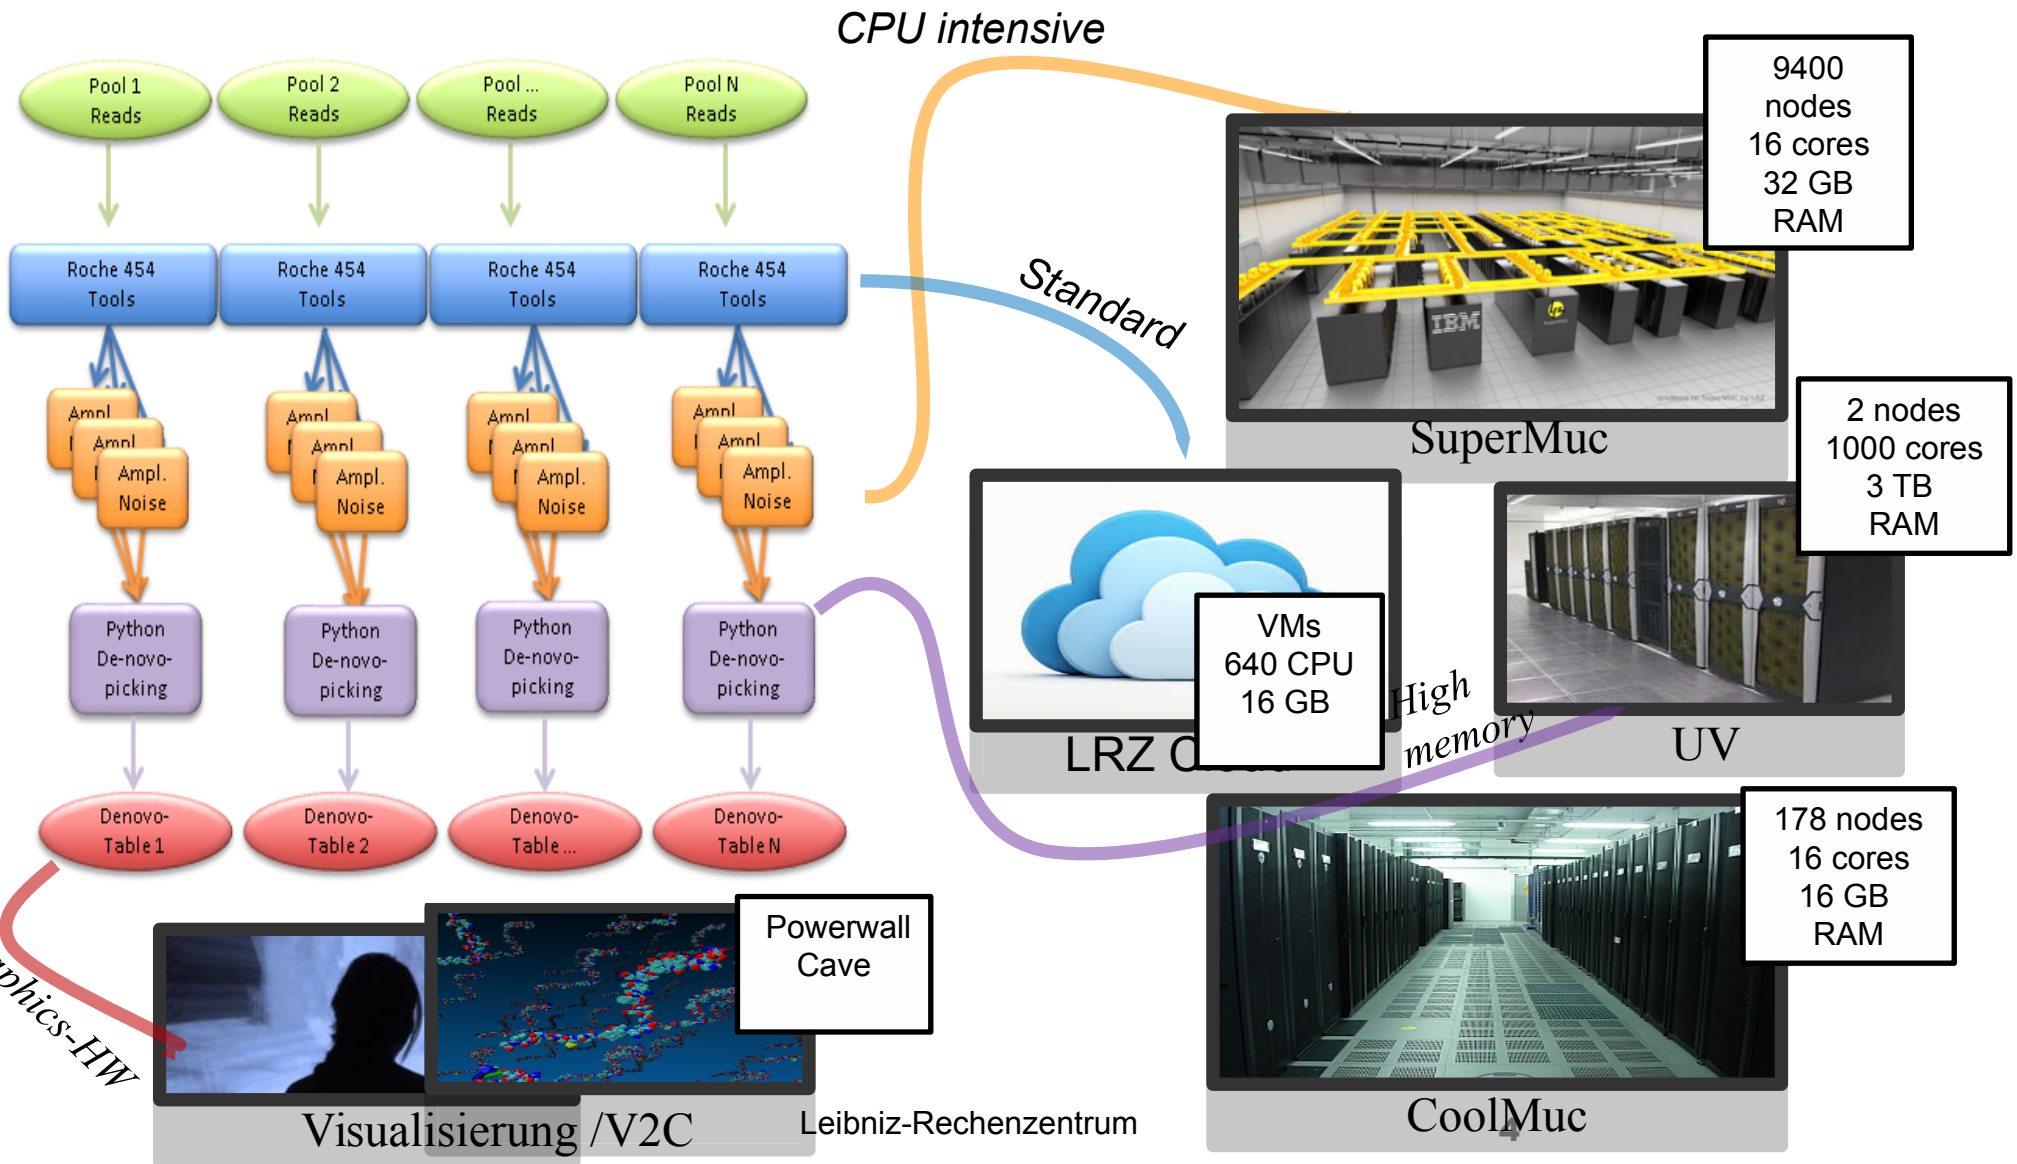

Overview: Before vs After

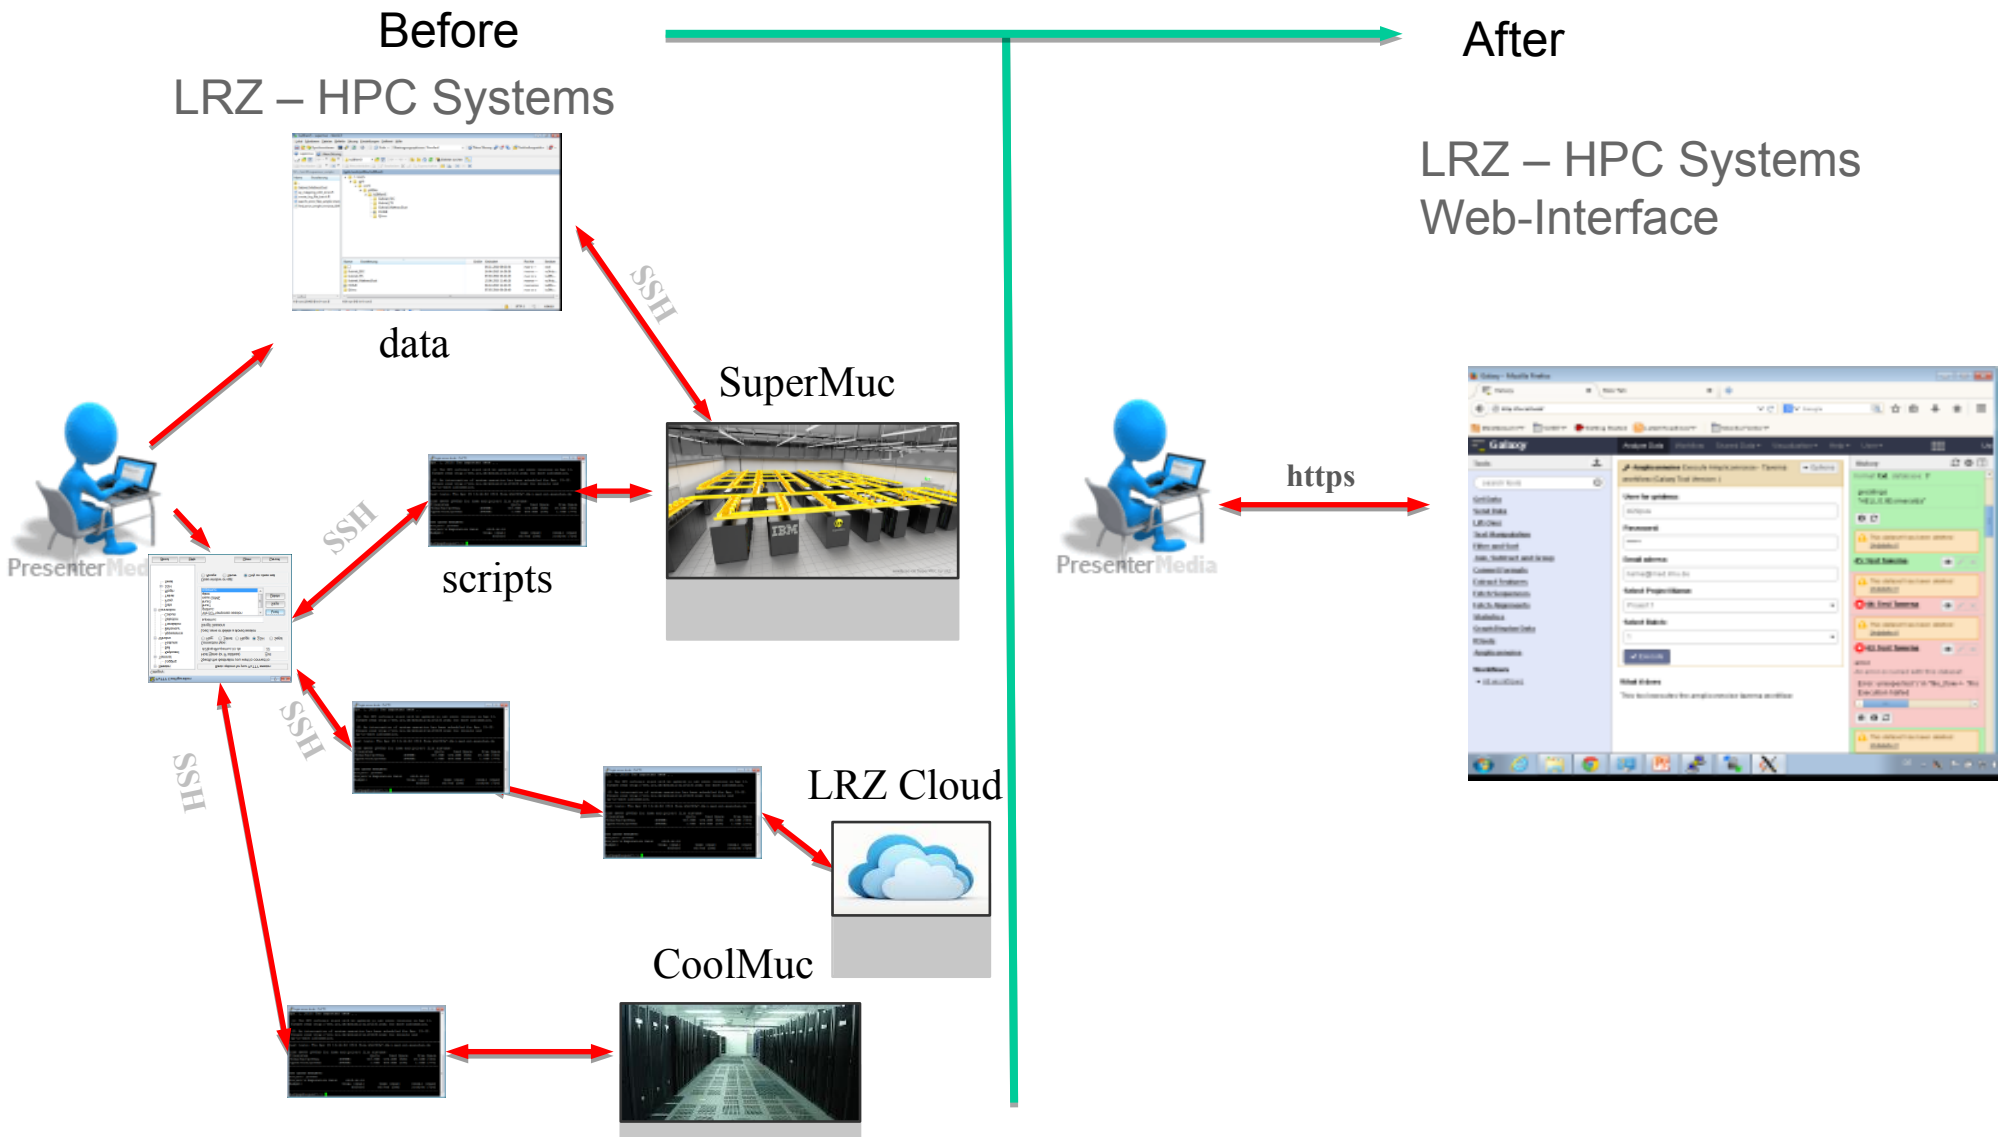

Obstacle: Separated HPC Systems

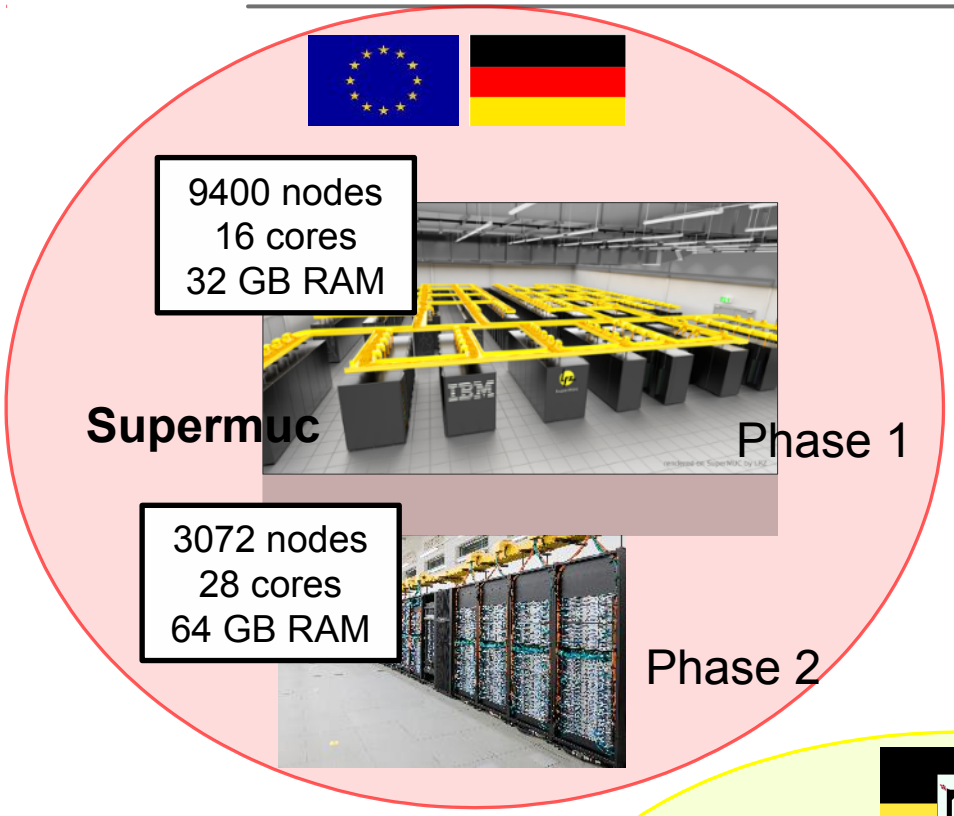

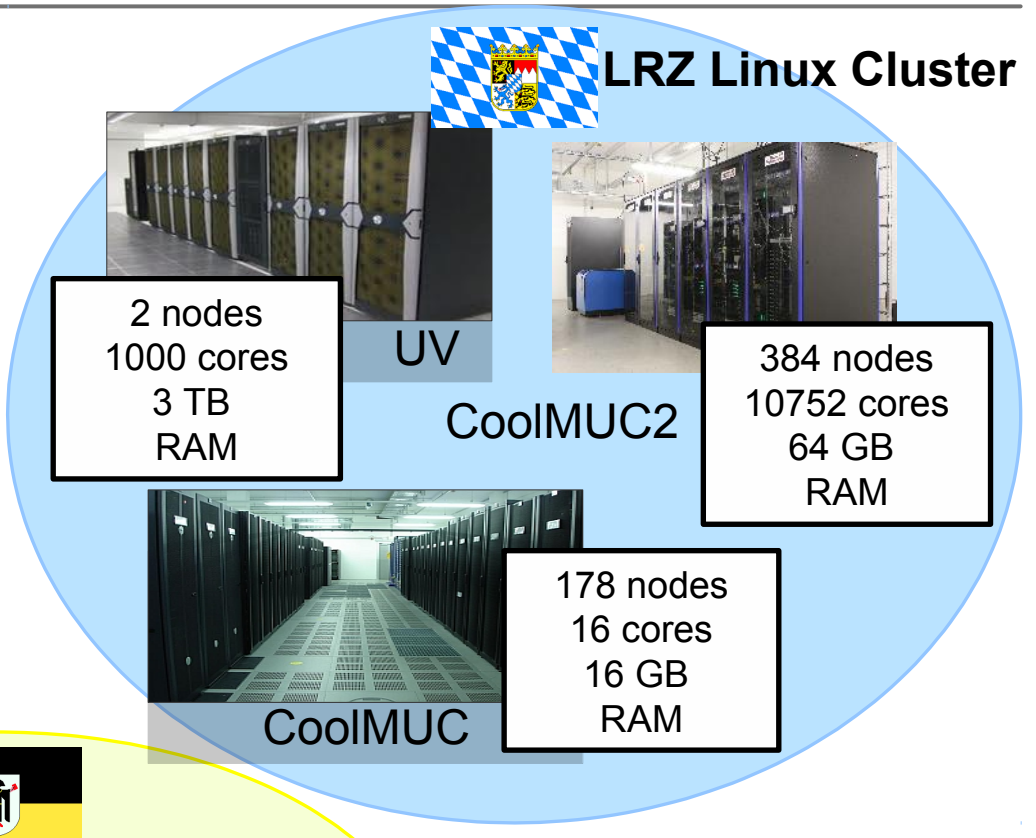

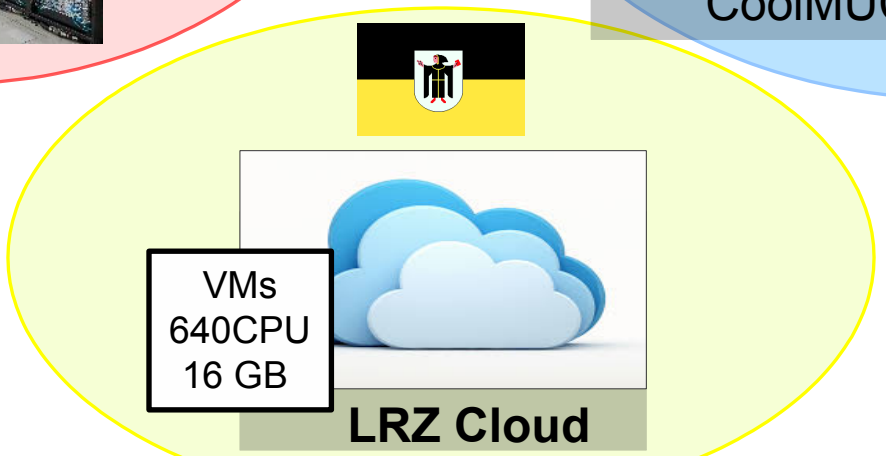

Solution: Connect via Globus Grid Services

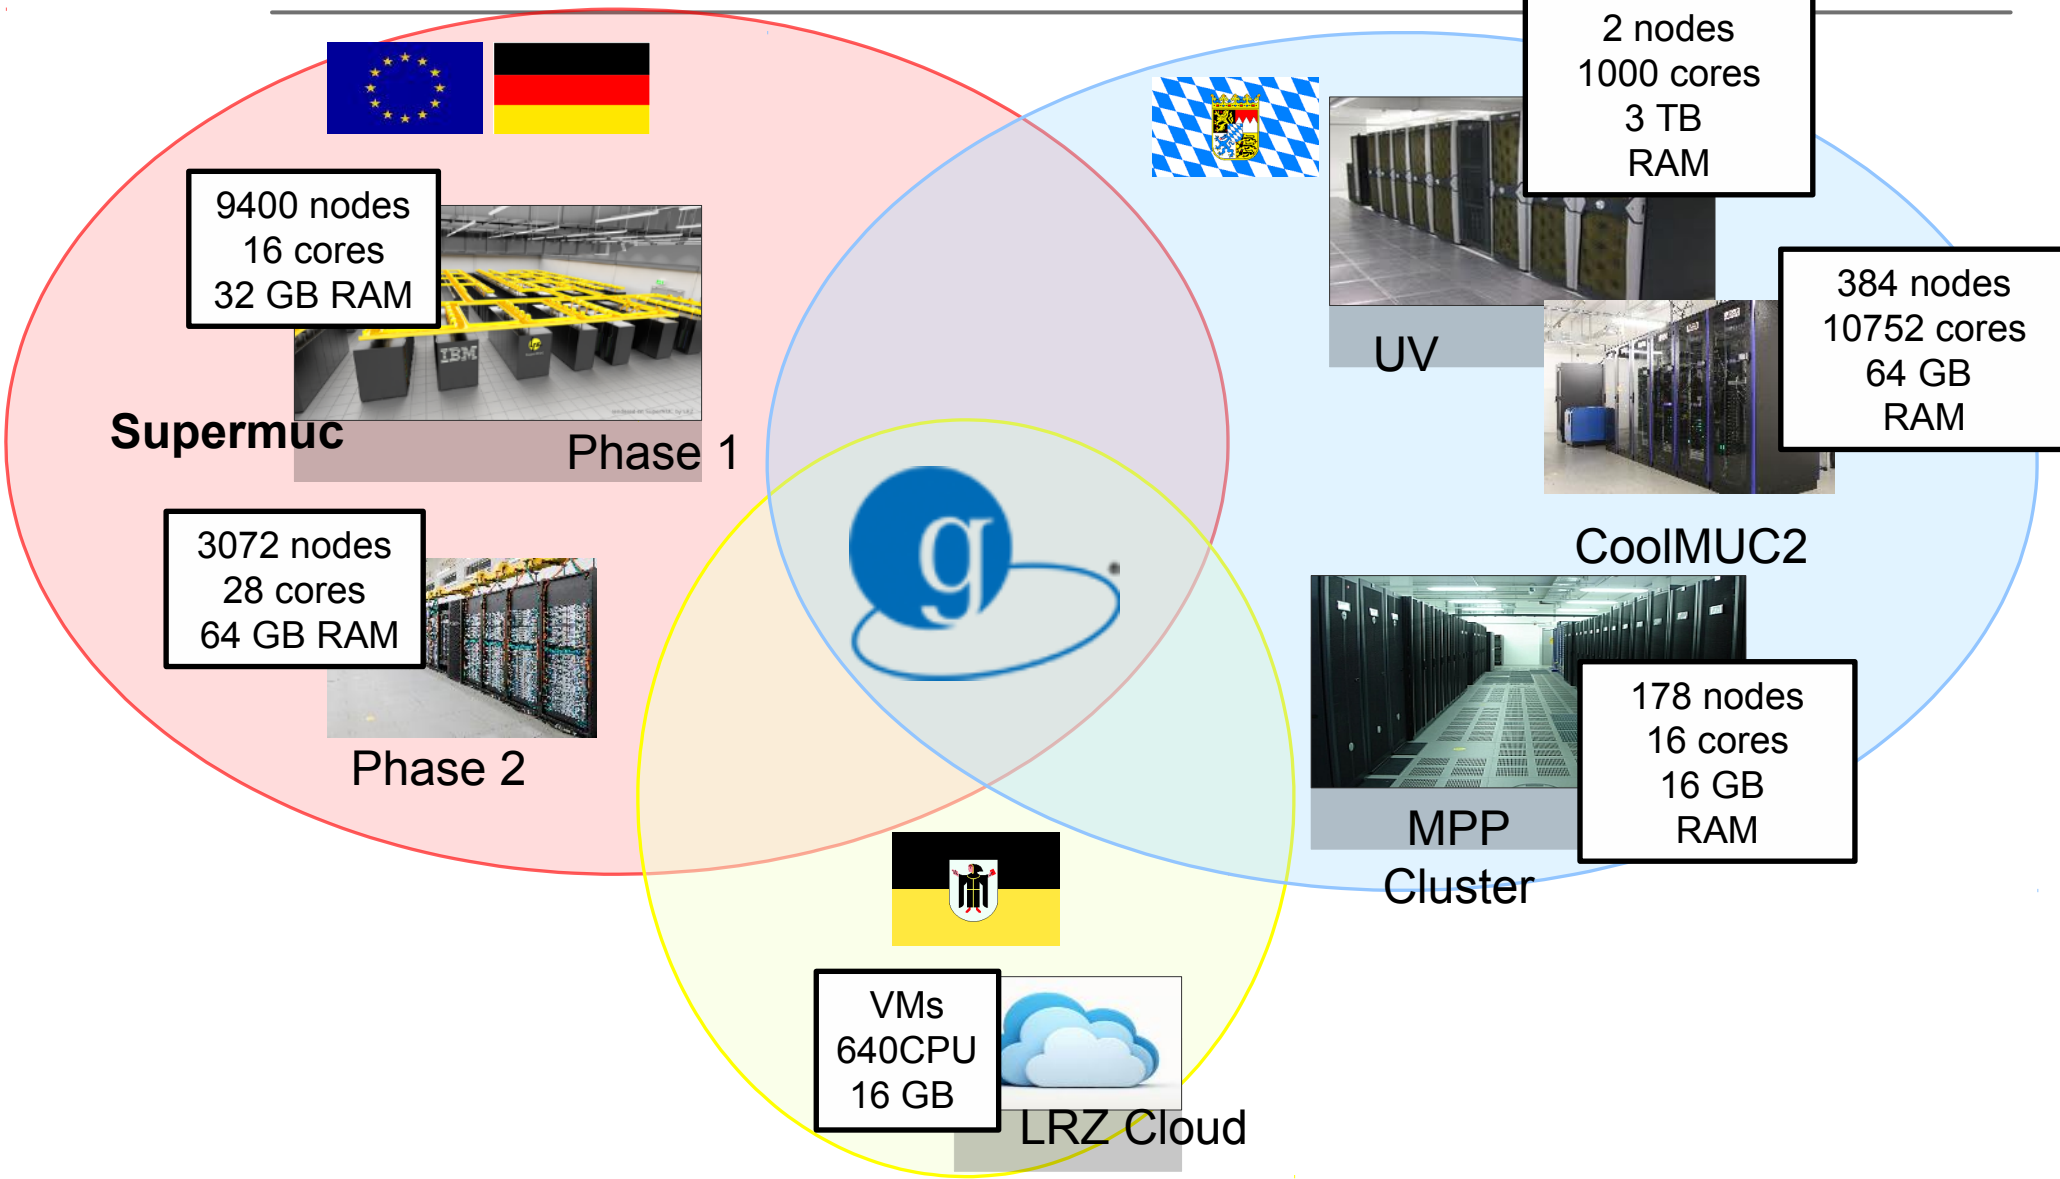

Workflows: Taverna Workbench using R

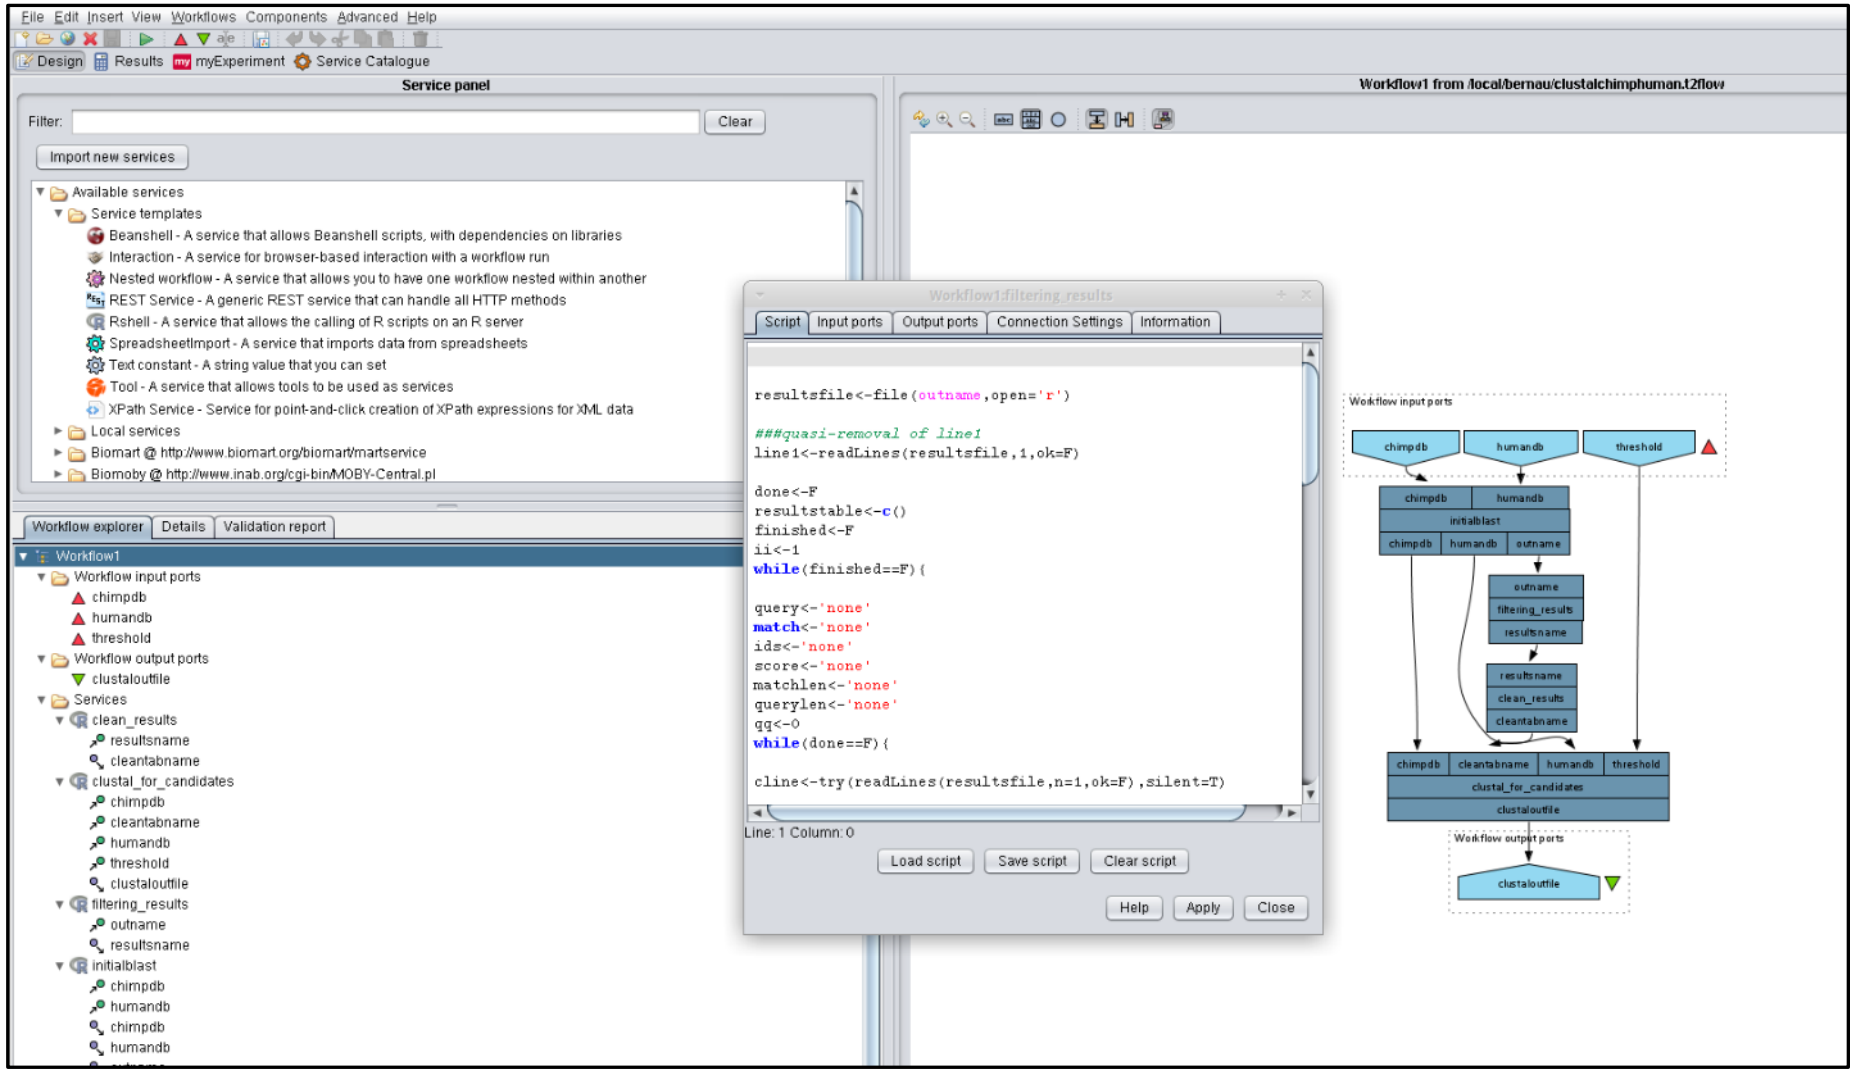

Example: Taverna workbench

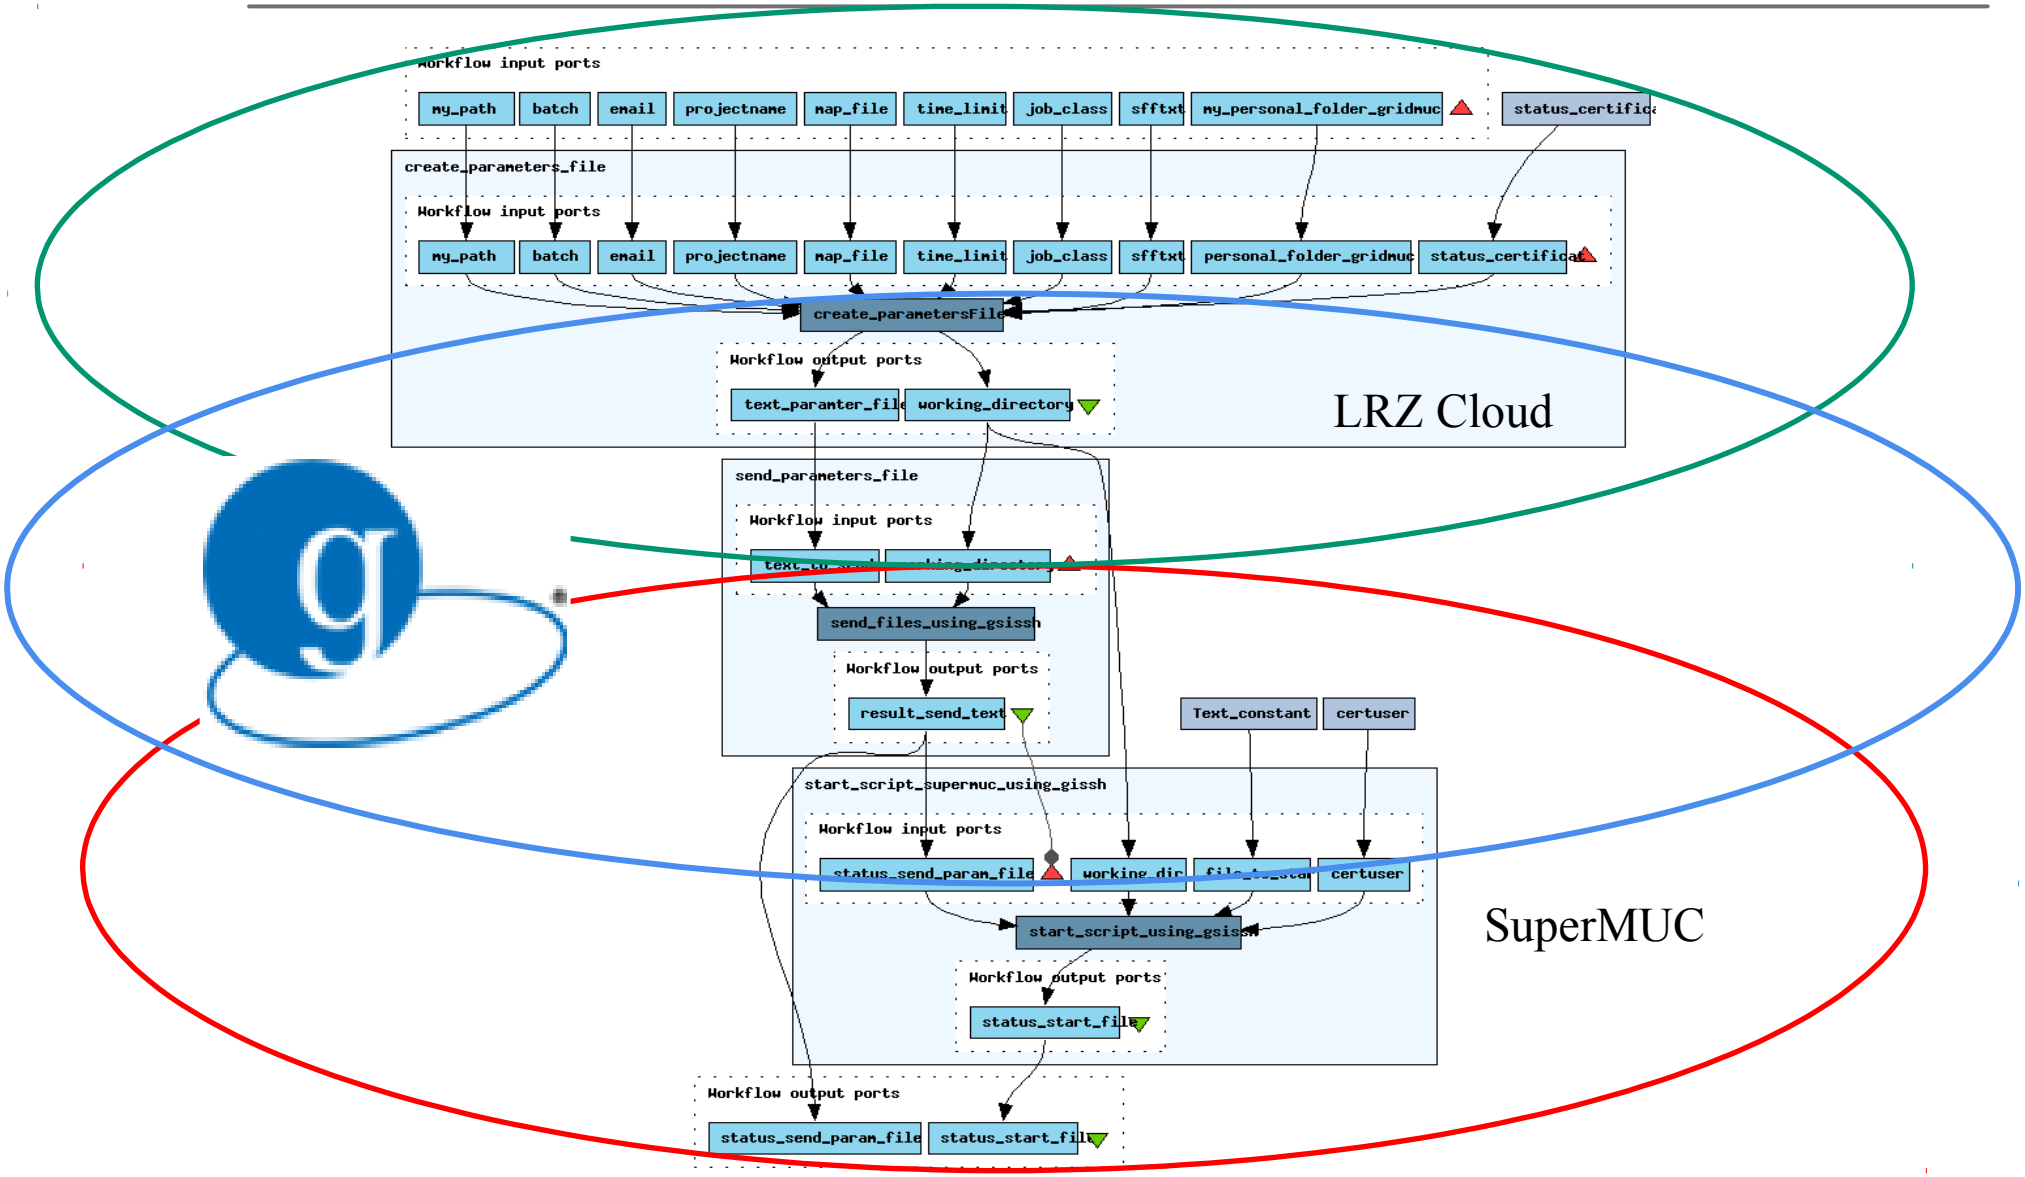

User Web Interface: Galaxy

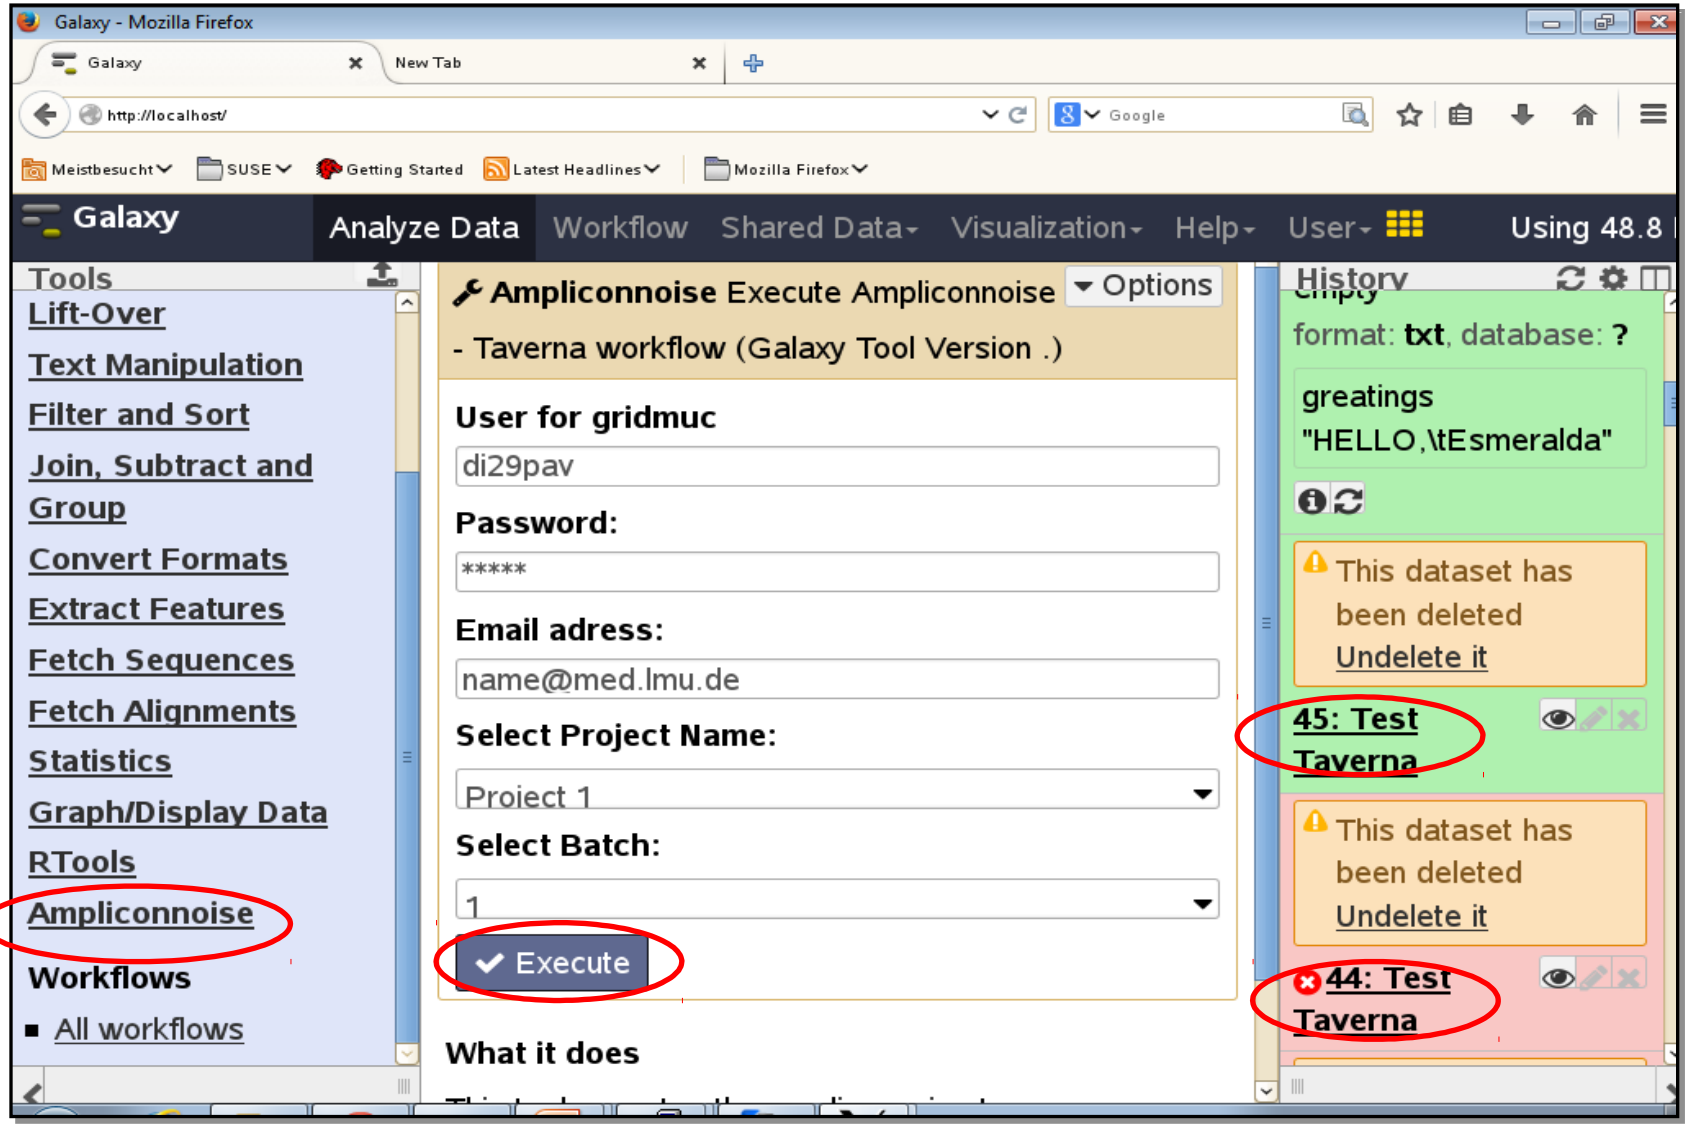

Even more obstacles:

- Bioinformatic software is distributed as debian packages and cannot be easily installed on SuperMUC (SLES11)
- Solution: Linux Containers
  - docker: security problems, has to run as root
  - proot: linux containers in userspace
- re-mount your own $HOME/var to /var, re-mount your own root folder and fake root account

```
$ proot -b $HOME/var:/var -r $HOME/debian8 -0
```

Advantages/Disadvantages

- You can trick software packages into believing they are root and they can write to system folders
- Everything is fake. Programs may break if they really need root access.
- You can fake e.g. a debian distribution running on suse. Debian can access /proc, /var, /dev and /etc like it was its own.
- Download at http://proot.me
- proot can not modify kernel and drivers!
- Who uses it?
  - Sony, STMicroelectronics, Ericsson, Cisco

Putting it all together

openNebula + taverna + globus + proot + virtualGL + galaxy + R/python

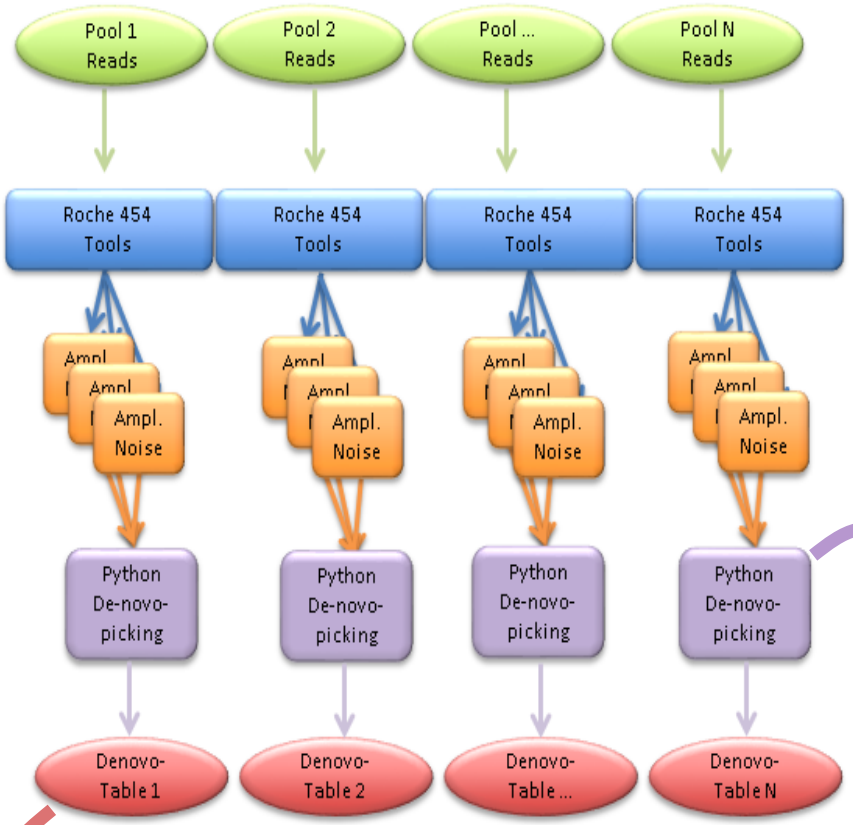

9400 nodes 16 cores 32 GB RAM

SuperMuc

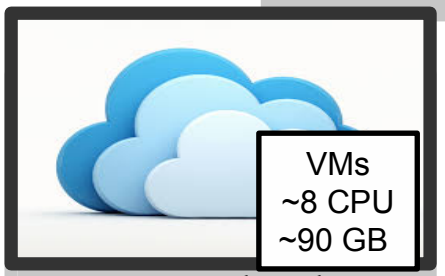

VMs ~8 CPU ~90 GB

LRZ Cloud

2 nodes 1000 cores 3 TB RAM

UV

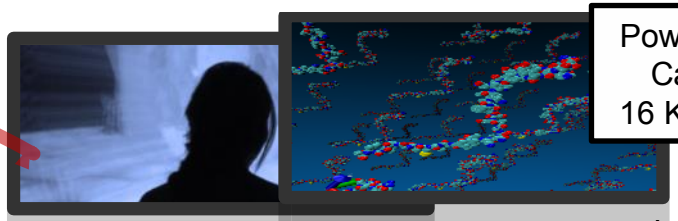

Powerwall Cave 16 Kepler

Visualisierung /V2C

178 nodes 16 cores 16 GB RAM

CoolMuc

Summary and Conclusions: Virtualization

Taverna/R

- Taverna workbench is a flexible, powerful tool for designing workflows
- tools are designed as web services and can thus easily be moved to other resources
- Workflow developers in our test projects clearly favoured it over Galaxy
- Dislike: Taverna does not provide a working web interface (no multi-user interface)

Galaxy/python

- Nice web interface to make workflows available to end users
- Inconvenient for workflow developers (server needs to be restarted for changes to take effect)
- Dislike: runs everything under same account, no user file protection

Globus

- Available and working on all LRZ systems (almost an unique characteristic)
- Provides certificates and user mapping
- Provides a convenient way to automatically request certificates at login or workflow runtime via LRZMyProxy

Proot

- Run software in lightweight container
- Good security model
- Dislike: Some programs may not work correctly (system files, low level drivers)

Thanks for your attention!

Project Members (LRZ Application Labs):

- Ferdinand Jamitzky
- Christoph Bernau
- Shaila Rössle-Blank
- Yu Wang
- Helmut Satzger
- Momme Allalen
- Gerald Mathias

Collaborators:

- LMU Munich (Esmeralda Vicedo, Markus Ege)
- TU Munich (Andrea Schafferhans, Timothy Karl, Burkhard Rost)
- Max-von-Pettenkofer-Institute (Debora Garzetti)

Outlook 2020: BG.DAF
Bavarian Genome Data and Analysis Facility

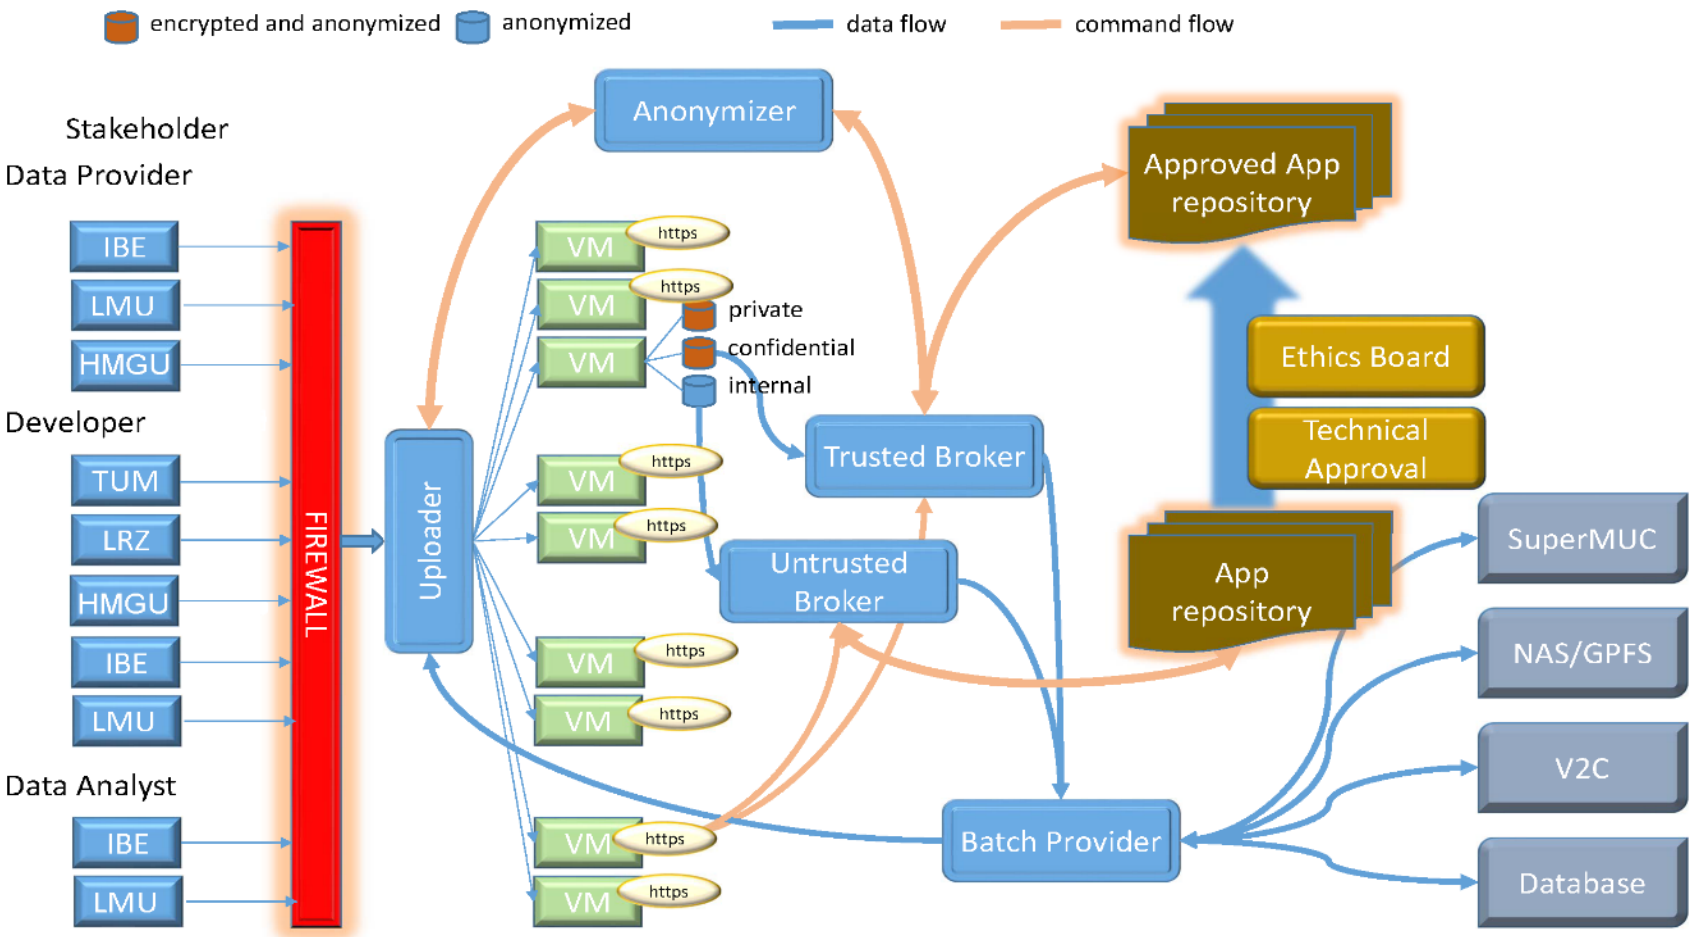

Platform overview: current status

Registers distinguished name of his certificate

login

Password or certificate

APACHE

Forwards username

validation

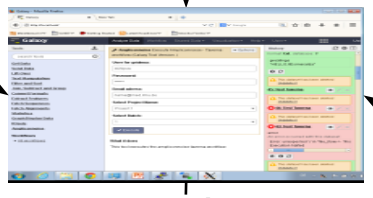

Requests short-living certificate

Sends certificate

Embeds username and certificate in workflow

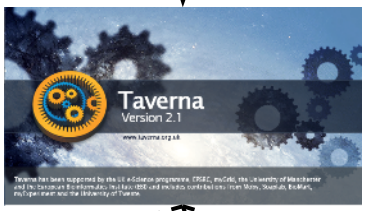

ID-Portal

LDAP

validation

MyProxy

validation

Job submission + gridftp

gsissh + gridftp

Job submission + gridftp

gridmuc.lrz.de

forwarding

Linux Containers: Why?

- Software package xxx needs to write to a system directory during installation (e.g. write to /var)
- You have a rpm and want to install it in another place (e.g. /lrz/sys)
- You want to install debian software packages on SUSE
- Your program requires root user to be run, but does not need root privileges. (e.g. apt-get)

How is it done?

ptrace is a system call found in several Unix and Unix-like operating systems. By using ptrace (the name is an abbreviation of "process trace") one process can control another, enabling the controller to inspect and manipulate the internal state of its target. ptrace is used by debuggers and other code-analysis tools, mostly as aids to software development.

Solution: Linux Container using proot

- re-mount your own $HOME/var to /var:

```
$ proot -b $HOME/var:/var
```

- re-mount your own root folder

```
$ proot -r $HOME/debian8
```

- fake root account

```
$ proot -0
```

-