ViCE – Virtualized Research Environments Cloud, HPC and Classrooms

Bernd Wiebelt, Michael Janczyk, Dirk von Suchodoletz

Rechenzentrum der Universität Freiburg (eScience)

UNI FREIBURG

Albert-Ludwigs-Universität Freiburg

Karlsruhe, 01.09.2017

GridKa School 2017

Structure of the talk

UNI FREIBURG

- Challenges in Scientific Computing
- How to make everybody unhappy
- Virtualization – a game changer
- bwLehrpool
  - Virtualized Teaching Environments for the Classroom
- bwCloud
  - Virtualized Research Environments in a scientific cloud service
- bwHPC (bwForCluster NEMO in Freiburg)
  - Virtualized Compute Environments
- Outlook

Challenges in Scientific Computing

UNI FREIBURG

- Very diverse scientific communities and broad set of software, tool demands
- Different, contradicting demands regarding software environments
- Short notice demands for hardware to be used at least for five years (amortization of equipment in economic terms)
- Human resources to operate all the (small, diverse) hardware servers expensive and in limited supply
- Save on money and hardware resources, most resources under-untilized most of the time
- Save on rackspace and energy – the computer center has significant energy costs each year

Challenges in Scientific Computing

UNI FREIBURG

- Freiburg: Work started a couple of years ago
- Optimize operation of computer pools
- Very diverse user base requiring very different software environments for different courses: Lecturers, students
- No real common denominator
- Windows OS for „standard“ software packages like text processing, spreadsheets or interactive statistics, web publishing, …
- Linux OS for software development, many open source packages like R, ...
- Very different ideas on how even a common software base should be configured (modules, examples installed or not, preconfiguration of packages, …)
- Demand to offer e-exam environments

Pre-Virt: Make everyone unhappy

UNI FREIBURG

- Scientists (*sigh*). Can't live with them, can't live without them...
- Lecturers expect from the computer center to comply to their expectations on installed software
  - Should be available in 20 minutes time (when the course starts) on 20+ machines
  - Works for me, should work in the PC pools too (where is the problem??)
  - Install some evaluation software which is only valid for 20 days (even a month would be to short)
  - Why didn't you install the exhaustive example collection?
  - Why did you changed/updated the software base, was fine in my last course!?
- Traditional software deployment doesn't work that way

Not all software and versions easily
live together in a single installation

Pre-Virt: Make everyone unhappy

UNI FREIBURG

- Admins (*sigh*). Can't live with them, can't live without them...
- Pool admins are annoyed by ever changing expectations of different lectureres
  - Often not available in the right time
  - PC pools never free during normal working hours or at least not long enough for proper software roll-out and testing
- Utilization of pools suboptimal
  - Software environment for the morning course might significantly differ from the one for afternoon and evening again and again for the week to follow
  - Difficult to schedule courses to pools
  - Software installation in one pool is completely different to the one in the other (no option to change the room if a particular one not available)

Pre-Virt: Make everyone unhappy

UNI FREIBURG

- Students (*sigh*). Can't live with them, can't live without them...
- Students are unhappy not to find the software environment of the course they are attending
  - Might be only available in the pool booked by some other course
  - Might just got removed because of conflicting demands of some other course
- Tight hardware – software coupling introduces inflexibility in time and space
  - Inflexibility increases operational costs

Virtualization: A game changer

UNI FREIBURG

- 180° turn:
  - Provide a virtual machine to the computer (instead of using it bare metal)
  - With all necessary drivers, connectors, etc before the environment is used for research or teaching
  - Then it becomes easy to archive the complete environment
- Only approved images will be used
  - Higher reliability in teaching and research
  - Reduced complexity – all computers run the same environment; same artefacts
  - Coexistence with local environments and environments of other users
  - Saves time – no need to install software
  - Talk about science rather than setup + configuration problems

Optimize operation of PC pools

- Virtualization comes into play: Break the tight link between software and hardware
- Developed a new form of Desktop virtualization (presented at 8th DFN-Forum)
- Only PXE boot a Linux system on the PC (without the need of locally installed software at all)
- Even allows to maintain local installation

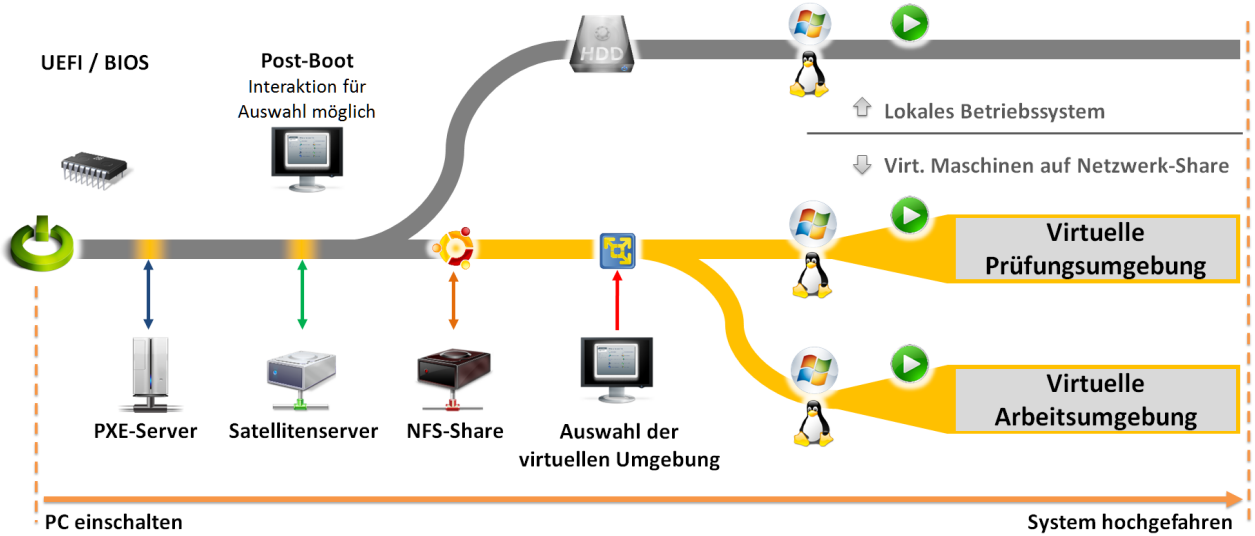

Optimize operation of PC pools

UNI FREIBURG

- Let the user choose from a wide selection of different environments (which are actually made available as images from a network share)
- Then configure a hypervisor to run selected environment
- Linux base system tries to handle all locally relevant stuff
  - Authentication of users
  - Mounting home directory and further shares if desired
  - Provide printing services

Avoid any site dependencies within the virtual teaching and learning environment

Selection of teaching environment

- Hypervisor runs locally on the PC with selected virtual environment
- Excellent for class rooms
- Change rapidly from Programming C in Linux to Desktop Publishing in Windows without „leftovers“

Loginmaske

Login

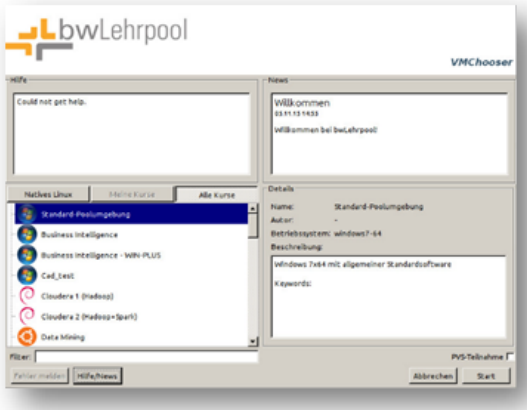

VMChooser

Auswahl

Virtuelle Maschine

bwLehrpool: Separation of tasks

- Administration of hardware is independent of the administration of the netbooted base Linux
- Software environment and configuration is absolutely up to the lecturer who wants to teach a certain course
- No time and physical dependencies for installations
- Lecturer receives just a base image and extends it to his/her needs
- All booted systems are exactly the same as using the same image to boot from

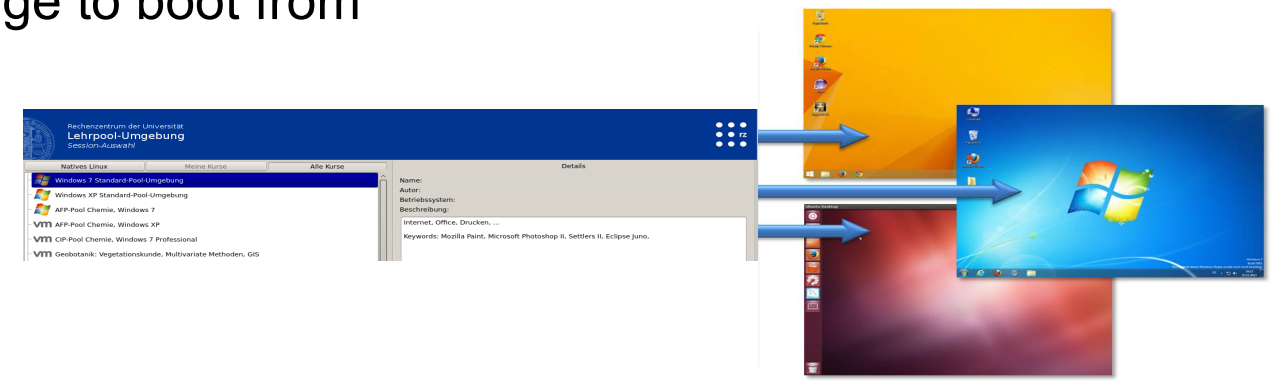

New ways of cooperation

- Next possible step:
  - Courses / images of virtual teaching and learning environments could get exchanged offering cooperation between different entities
  - Community provided environments e.g. created by students of a certain semester or faculty
- Widen it‘s application to further domains

Successful model of task separation
between users and computer center operation

Virtual science – how far we got?

UNI FREIBURG

- Challenge: convince science that virtualization approach is necessary
  - If only for convenience and reliability of results
- Virtual(ized) research environments (VRE) for various scientific communities
  - Build virtual machines which suit a whole discipline → at least for their research interest
  - Long term usage and reproducibility
  - Long term archive (possibly create the electronic lab book as a side effect)
  - Right from the outset
- CERN already uses „standard virtual machines“
  - Some disciplines are on the right track

The need for electronic „lab books“

UNI FREIBURG

- Digital science often a „good“ example for fire&forget principle
  - Results created/generated by computer program(s) which are heavily dependent on a certain environment that exists during a short time in a single location
- Status Quo in digital science
  - Data and results become worthless without proper context
  - „inherit“ the data and do extra research on how it was produced
  - Not everything is known, although promised
  - Some inter-dependencies forgotten
  - Exact state of the original computer environment is not recoverable
  - Updates, modification to the system, twists imposed by manufacturers of research hardware
  - Malicious modifications of the computing environment?
  - By accident – or intentionally?

Make the mistake and worry later (or let others worry)

Research infrastructure: bwCloud

- OpenStack based state wide self service cloud to cater for scientific, computer center services and student use (science, operations, education)
- State wide cooperation of university computer centers in Karlsruhe, Ulm, Mannheim and Freiburg

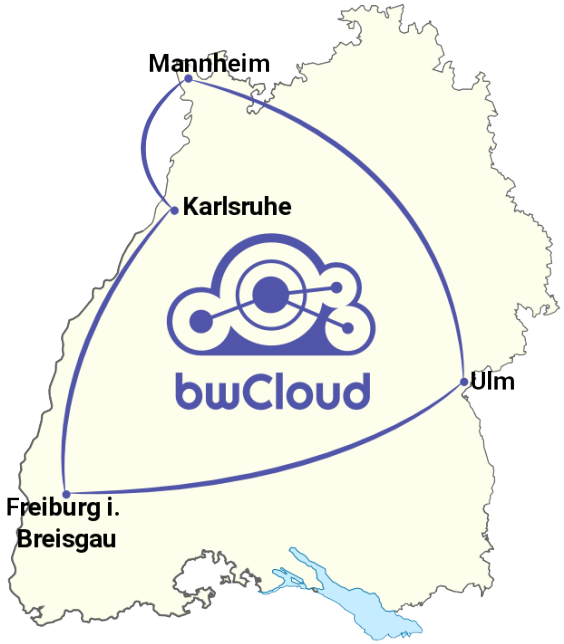

bwCloud

openstack
CLOUD SOFTWARE

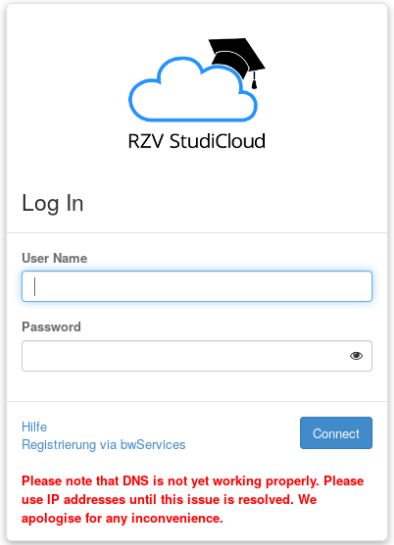

Bridging worlds

- ViCE – state sponsored cooperation project of multiple disciplines and computer centers
- Separate infrastructure / provider from core scientific tasks

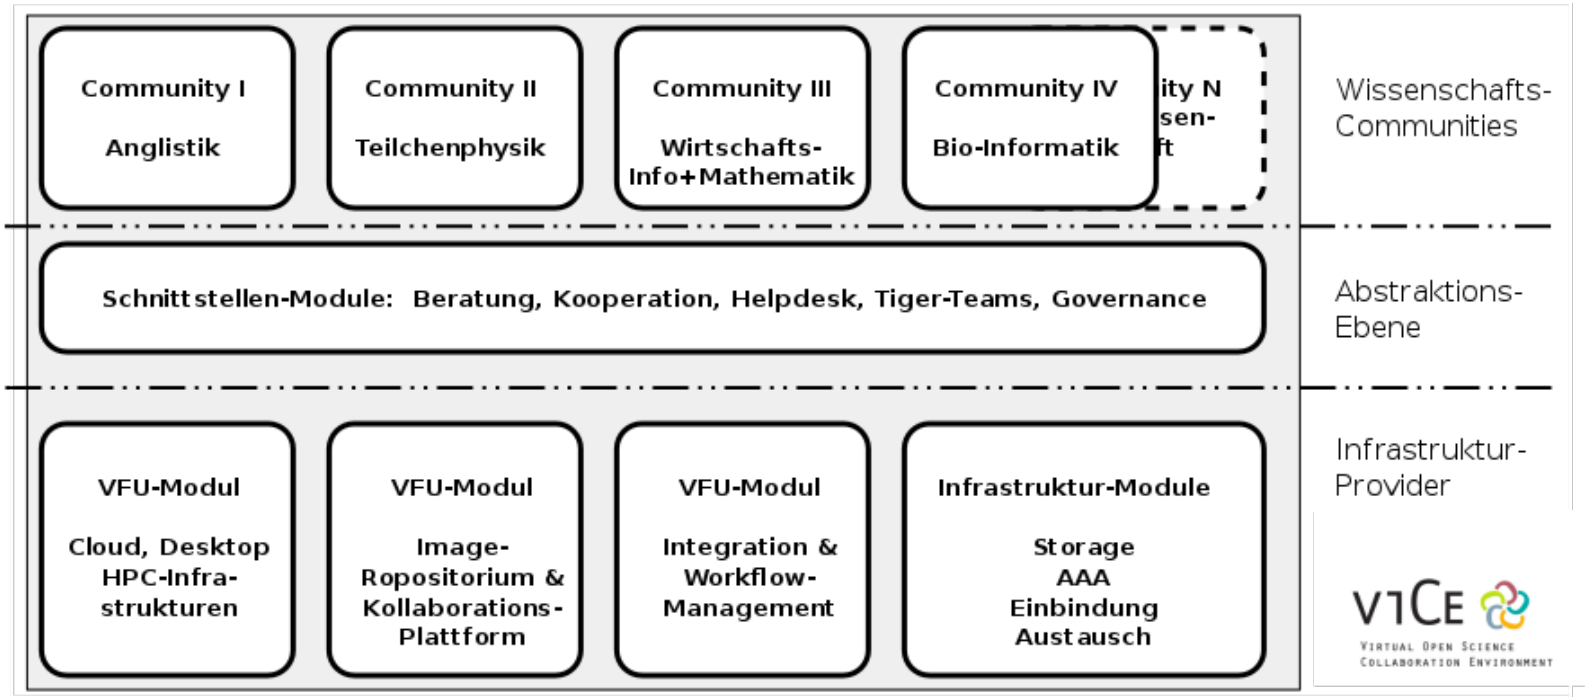

Virtualized Research Environments

- Tailored to the scientific application
- Customizable by the users:
  - VRE per scientist
  - VRE per scientific workgroup
  - VRE per scientific field
- Enables reproducibility of results
- Abstraction from underlying hardware
- Enables versioning of research environments
- Allows citation and referencing of software methods
- Requires Open Data
- Requires Workflow Management
- Requires Data Management Plan

ViCE
Virtual Open Science Collaboration Environment

CiTAR
Citing & Archiving Research

Virtualized Research Environments

- Optimally: Allow flexibility over the different domains
- Start to create and test a scientific workflow interactively on your desktop
- Move it into the cloud in long running and not dependent on massive resources
- Move it to the cluster if larger resources required or massive parallelism possible

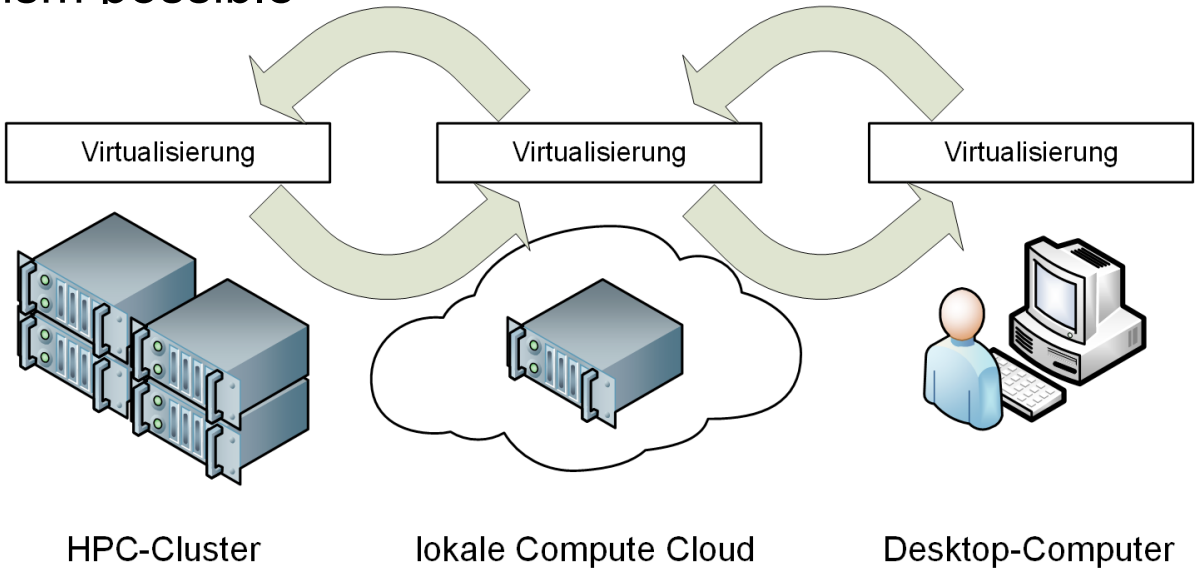

HPC for science: NEMO

- Computer center of University of Freiburg one of the operation sites of Tier 3 High Performance Computing (super computer)
- Roughly 800 (1000 in near future) dual CPU nodes of Intel Broadwell platform architecture

Introducing bwHPC

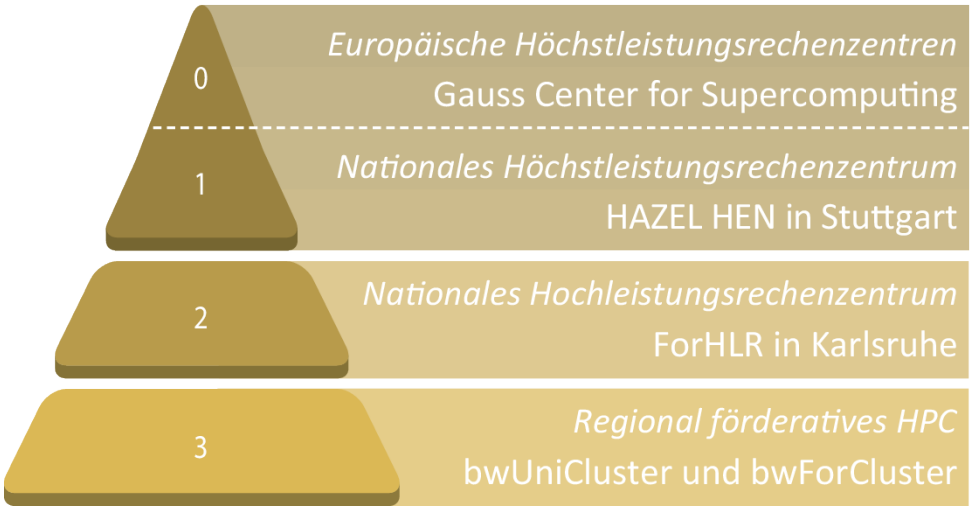

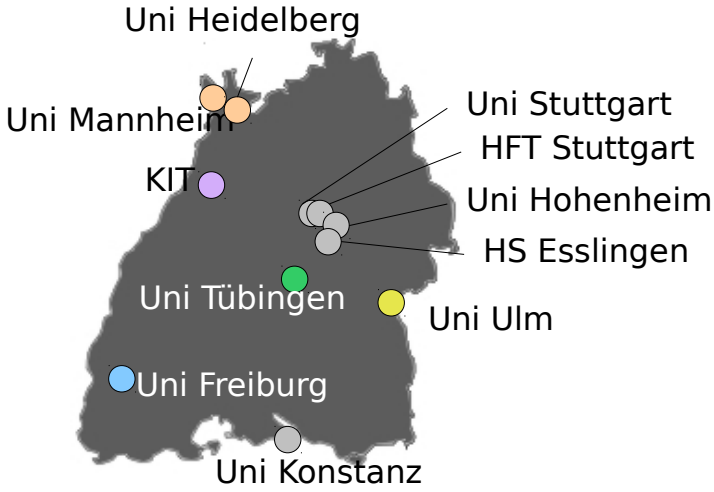

User Support

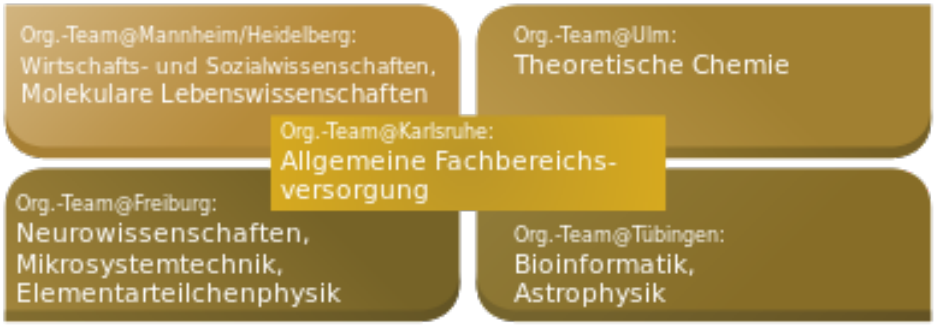

(*) https://mwk.baden-wuerttemberg.de/fileadmin/redaktion/m-mwk/intern/dateien/pdf/Forschung/Umsetzungskonzept_bwHPC.pdf

Usecase 1: HEP CMS

UNI FREIBURG

- CMS is a large scale, world-wide run experiment based on CERN data
- Environment has to be exactly the same everywhere in the world
- Software updates are slow because of exhaustive regression tests
- Software environment can not easily be reproduced directly on the NEMO cluster

Motivation

Goals we want to achieve:

- Render shared HPC resources accessible → virtualization
- Dynamic allocation of resources (no static VMs)
- Integration of new resources transparent to HEP user

Our complete "virtualized HEP node" tool set:

1. Hybrid HPC Cluster : **OpenStack** (IaaS) — @ Uni Freiburg

+

2. Flexible batch system: **HTCondor**

+

3. On-demand cloud manager: **ROCED** — @ KIT

Hybrid HPC setup

- Same hardware, different usecases: Create flexibility by abstracting hardware and software (again)

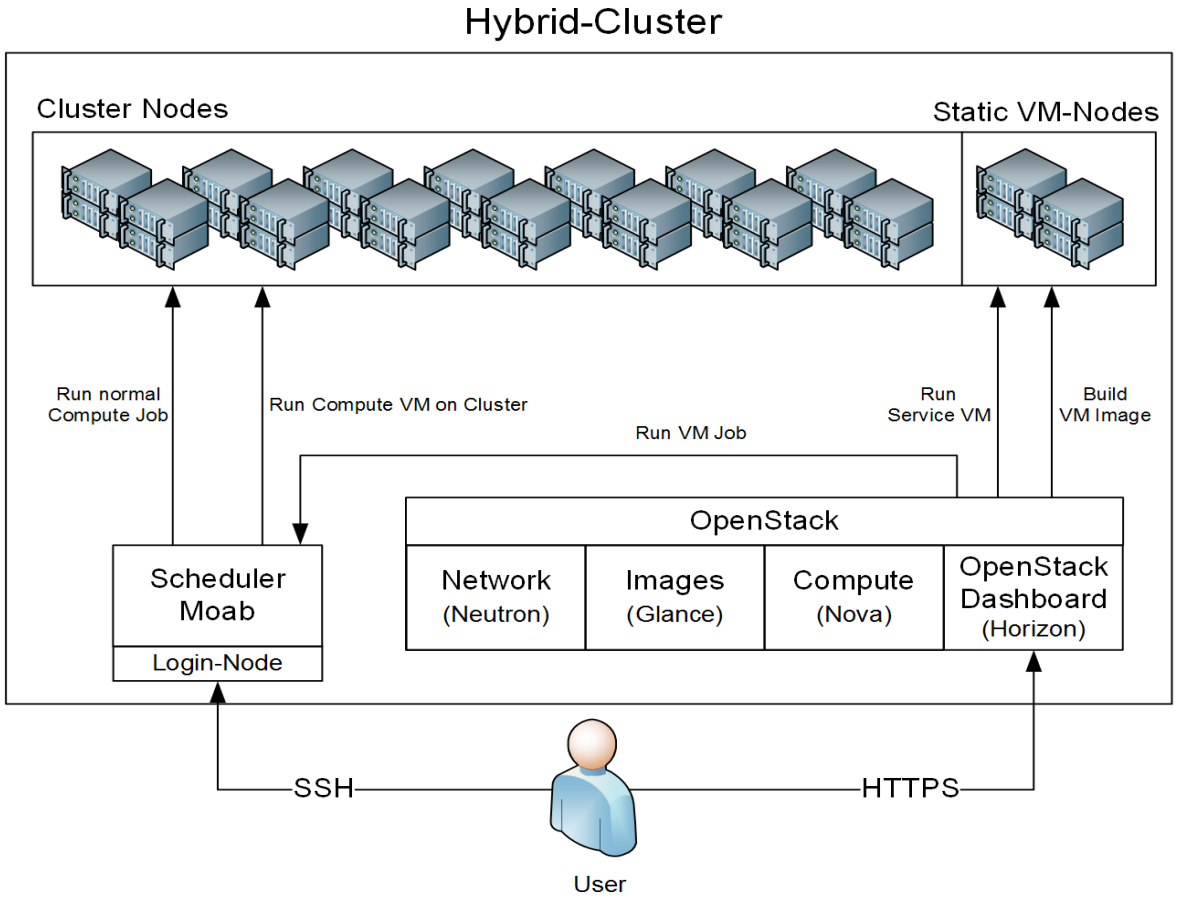

Dynamic Virtualization @ NEMO

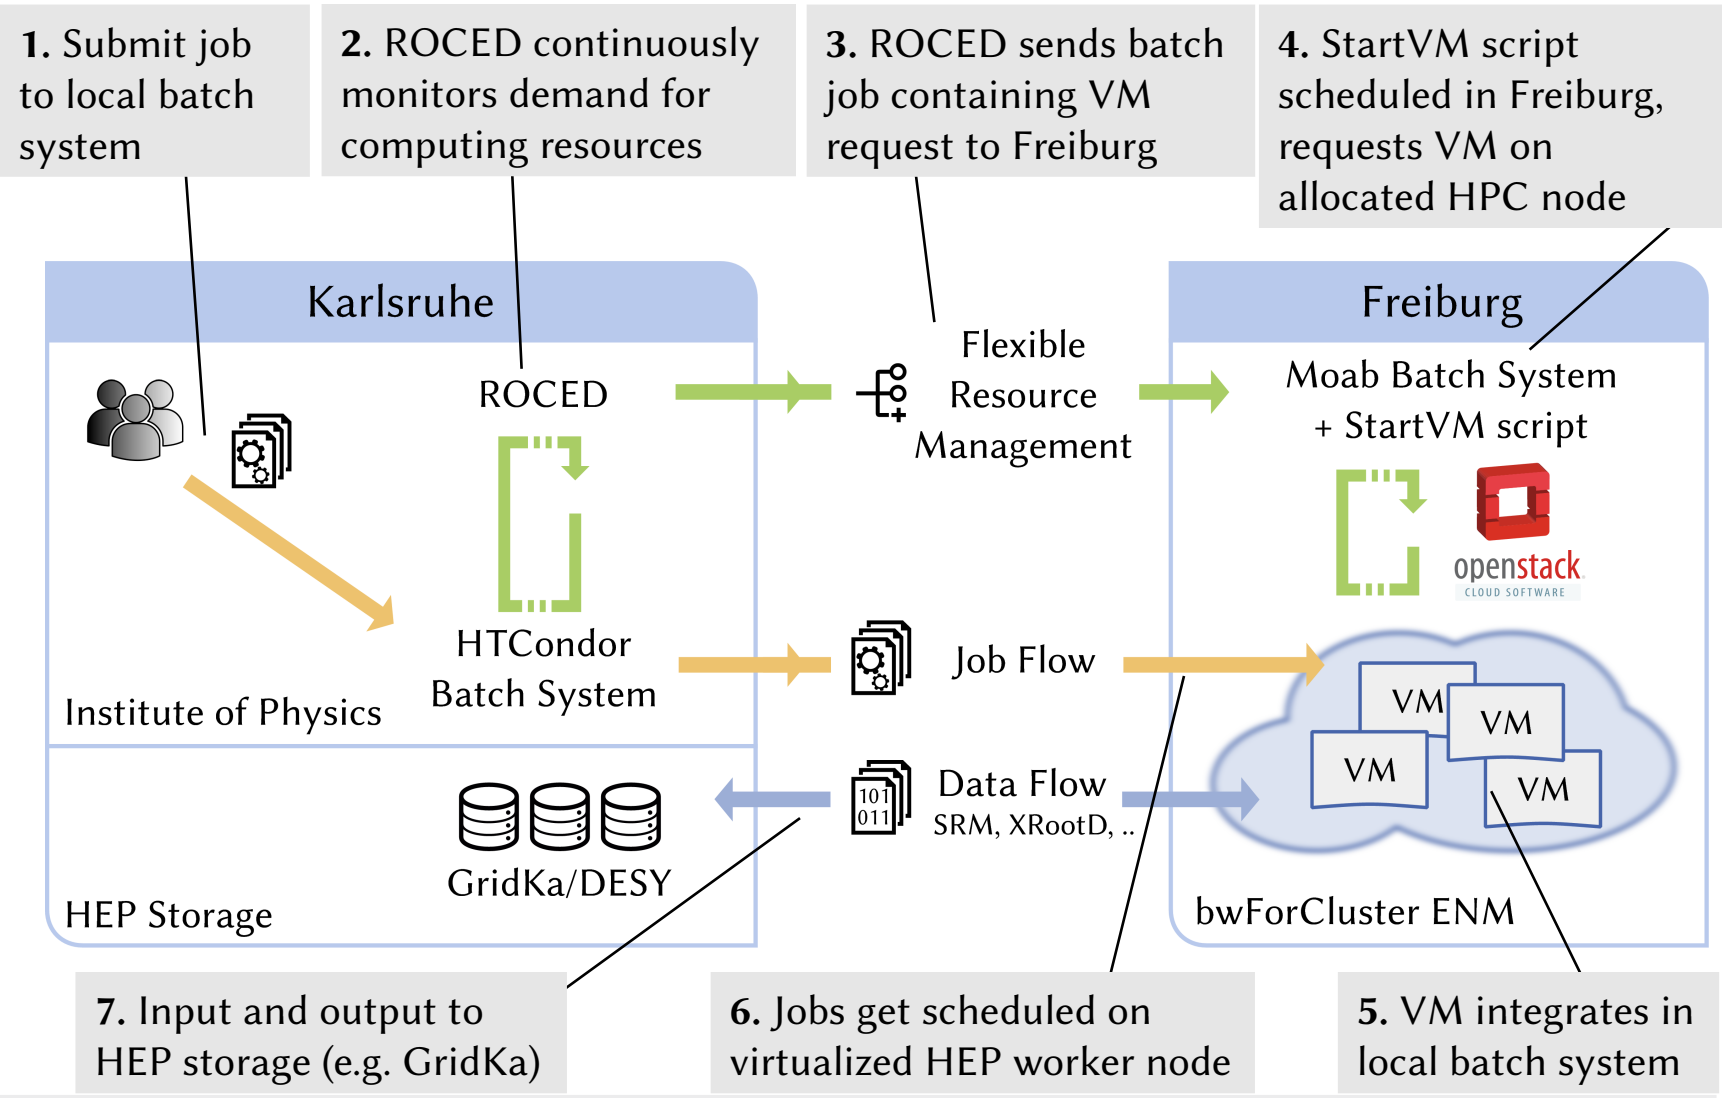

Usecase 2: Bioinformatics

- Offering (resouce demanding) Galaxy services ontop of an HPC system / mixed cloud

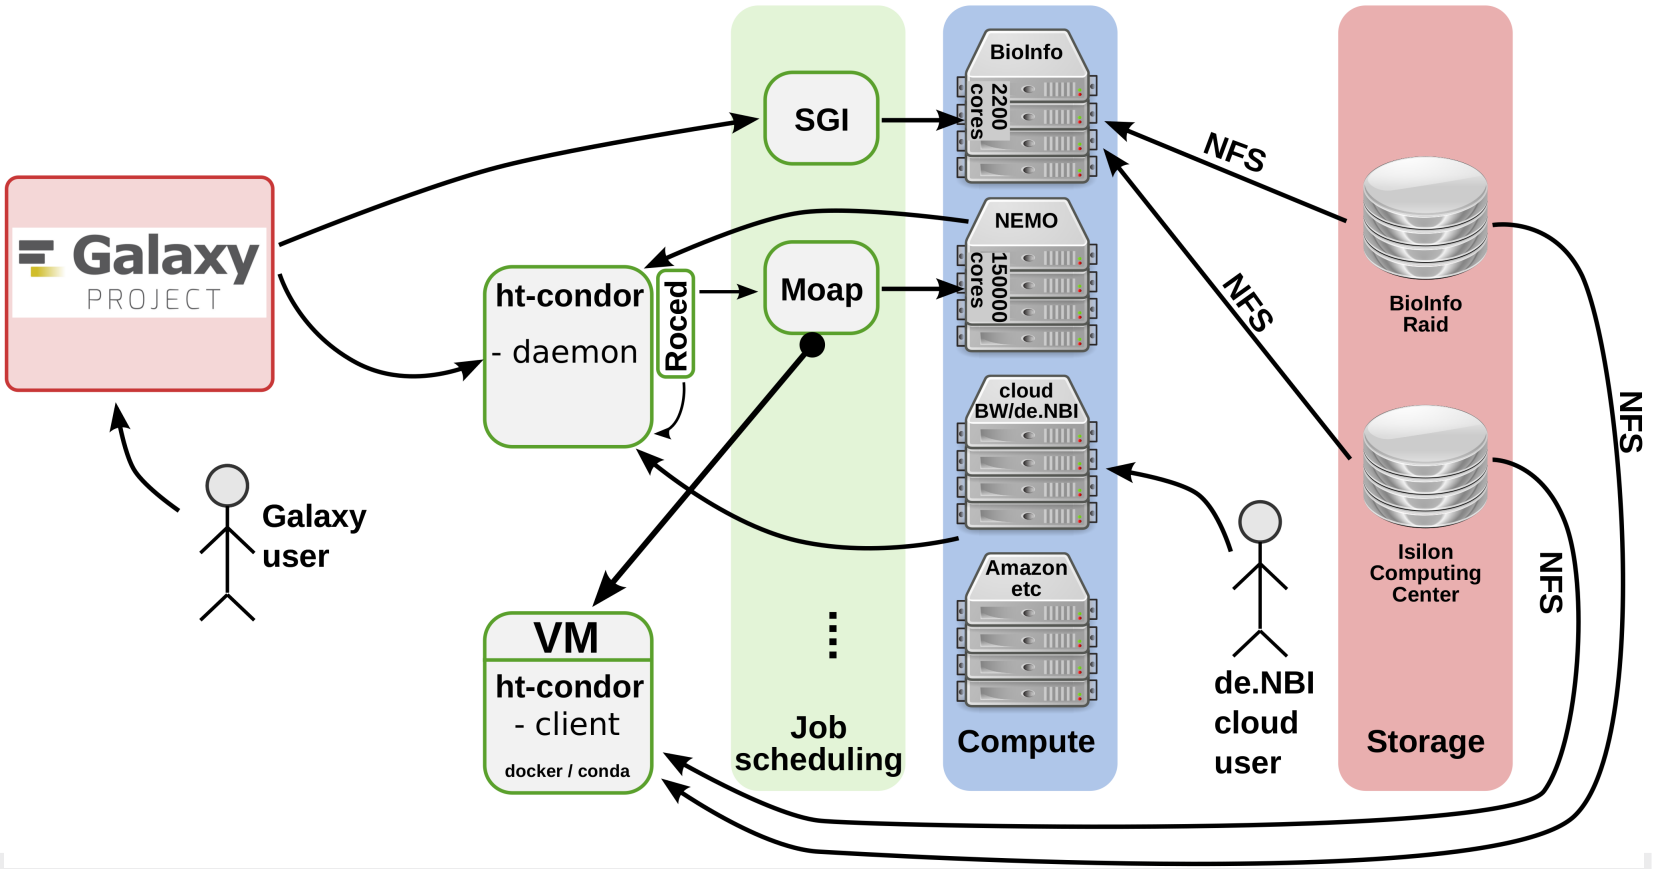

Usecase 3: HEP - ATLAS

UNI FREIBURG

- ATLAS is a large scale, world-wide run experiment based on CERN data
- Environment has to be exactly the same everywhere in the world
- Software updates are slow because of exhaustive regression tests
- Software environment can not easily be reproduced directly on the NEMO cluster

Usecase 3: HEP - ATLAS

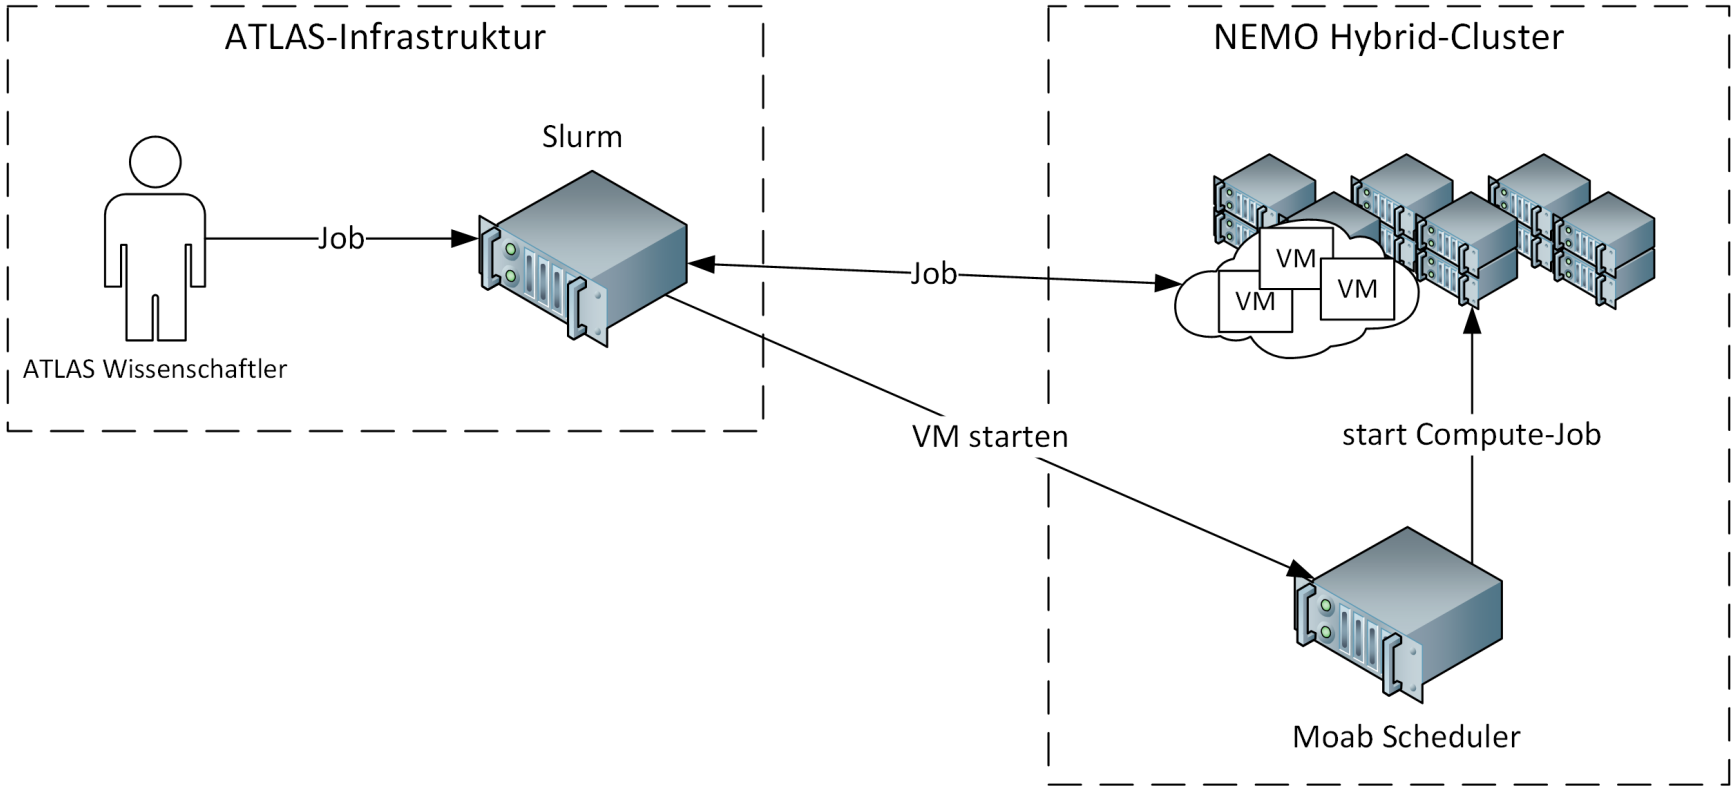

HPC and virtualization

UNI FREIBURG

- Additional benefits: VREs could get suspended and resumed
- Plenty of new options for scheduling / to improve scheduling
  - Offer long running jobs on a cluster with a standard 4 days walltime
  - Suspend (expensive, because of long running) jobs before cluster maintanance
  - Create a „fast lane“ and let certain jobs overtaking long running ones (which otherwise would clog the cluster)

Flexible HPC: Virtualization

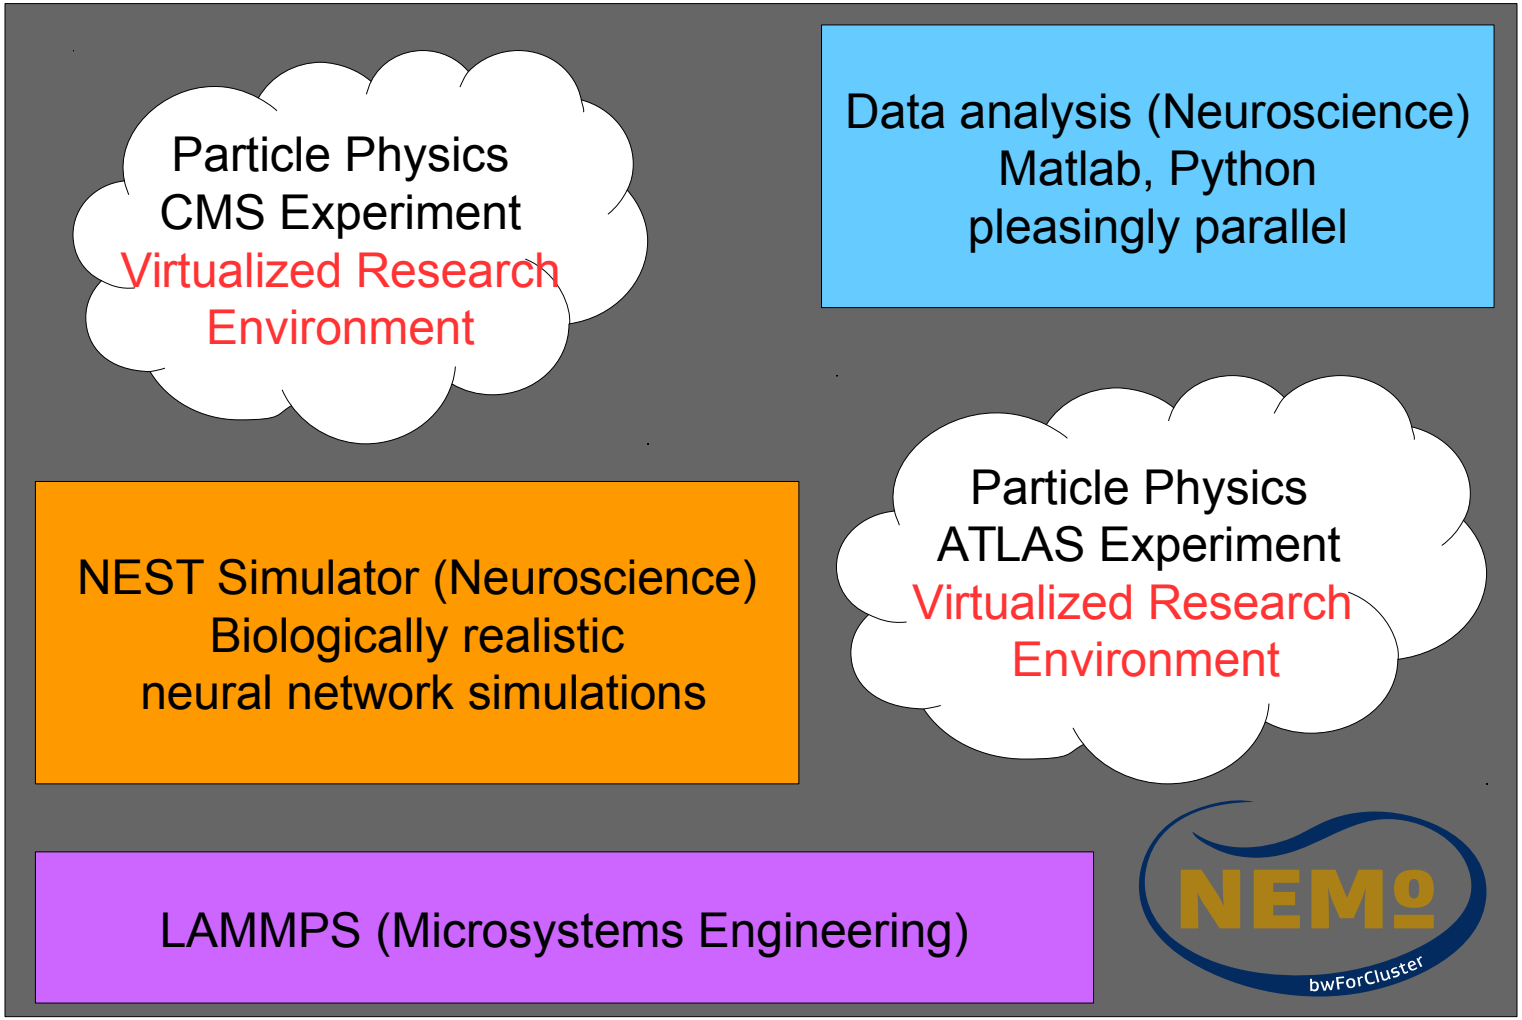

Virtual science – how far did we get?

UNI FREIBURG

- Challenge: convince science that virtualization approach is necessary
  - If only for convenience and reliability of results
- Virtual(ized) research environments (VRE) for various scientific communities
  - Build virtual machines which suit a whole discipline → at least for their research interest
  - Long term usage and reproducibility
  - Long term archive (possibly create the electronic lab book as a side effect)
  - Right from the outset
- HEP already uses „standard virtual machines“
  - Some disciplines are on the right track

Challenges / Research Questions

UNI FREIBURG

- How to properly describe scientific software and/or complete VREs?
  - Which meta data, schema to use?
  - E.g. suggestion of a very abstract schema definition (dataset), http://schema.org/Dataset or discussion on data publishing: http://blog.wolfram.com/2017/04/20/launching-the-wolfram-data-repository-data-publishing-that-really-works/#the-data-curation-hierarchy
- Scheduling challenge
  - Hybrid clusters generate nice „scheduling nightmares“
  - Running non-interactive batch jobs for a rather wide selection of different computations and simulations
  - Batch scheduler tries to allocate appropriate resources and optimally fill the cluster

Challenges / Research Questions

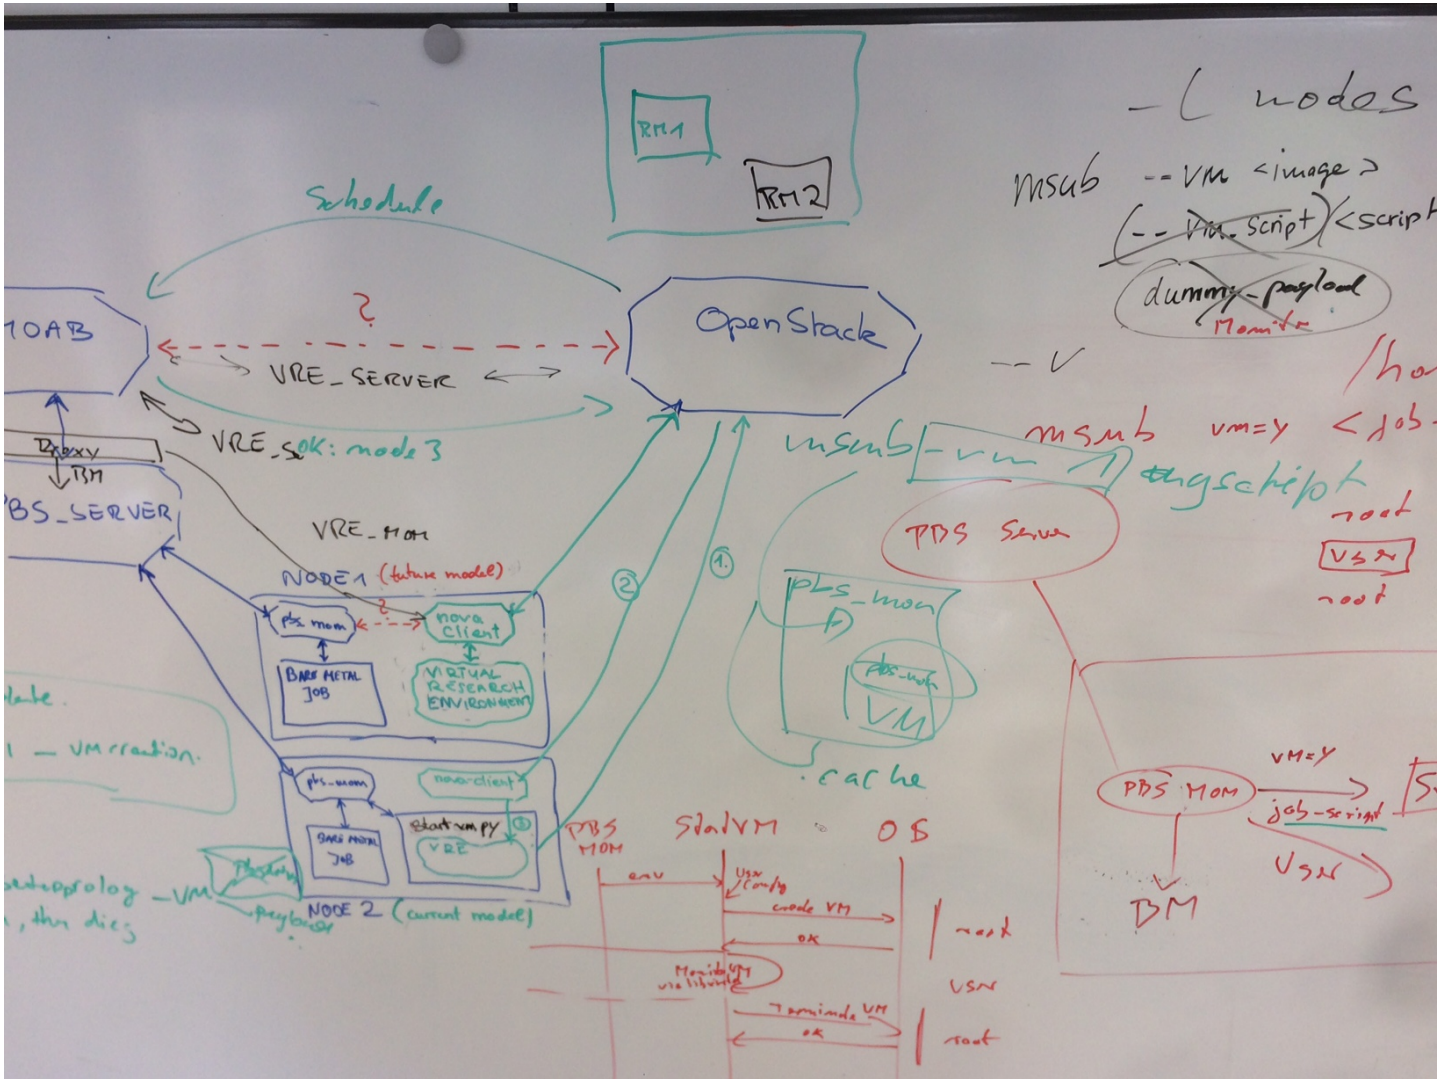

Thank you / Questions!?

- Further information: ViCE people
  - CMS / Physics: Thomas Hauth, Günter Quast
  - ATLAS / Physics: K. Meier, U. Schnoor, A. Gamel
  - Bio informatics: B. Grüning, C. Blank
  - English language studies, computer linguistics:
  - Computer center FR, UL, MA: J. Bauer, M. Janczyk, B. Wiebelt, J. Vollmer, D. v. Suchodoletz, Ch. Hauser, J. Schulz
- Project description:
  - https://www.alwr-bw.de/kooperationen/vice