Leibniz Supercomputing Centre

of the Bavarian Academy of Sciences and Humanities

Convergence of HPC and Data Analytics from a Supercomputing Centre Perspective

14.01.2020 | 711. WE-Heraeus-Seminar |

L. Iapichino, S. Hachinger (LRZ) with AstroLab/APP and RDM/FOR teams

Leibniz Supercomputing Centre

Infrastructure & Service for Science

Bavarian Academy of Sciences and Humanities

Leibniz Supercomputing Centre (Leibniz-Rechenzentrum / LRZ)

We are the computing backbone for advanced research science in Bavaria

250
employees
approx.

58
years of
IT support

Computer Centre
for all Munich Universities

Regional Computer Centre
for all Bavarian Universities

National Supercomputing Centre
(GCS)

European Supercomputing Centre
(PRACE)

Operating Cutting-Edge IT Infrastructure

- Email
- Network
- Storage
- Cloud Computing
- Cluster
- HPC
- Training
- Consultancy

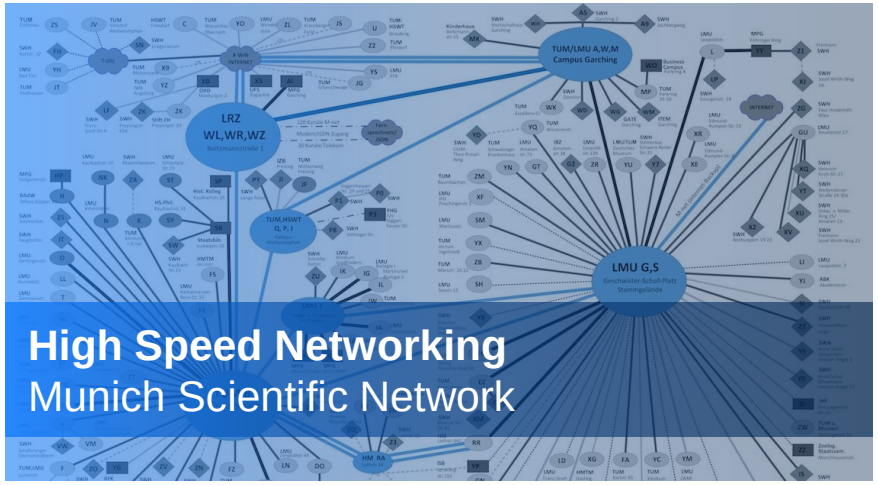

High Speed Networking
Munich Scientific Network

High Performance Computing
SuperMUC-NG, Linux cluster

Big Data
Bavarian Competence Centre

Virtual Reality and Visualisation
V2C (CAVE, Powerwall)

… providing Services on National and European level

Partnership for Advanced Computing in Europe (PRACE)

Federated, pan-European Tier-0 supercomputing infrastructure

25 Countries
Hosting Members:

- GCS (LRZ, HLRS, JSC)
- BSC (Spain)
- CSCS (Switzerland)
- CINECA (Italy)
- GENCI (France)

PRACE 2: 2017 – 2020

SuperMUC-NG (Next Generation) – HPC Top 500 #9

Specs

- Peak Performance: 26.7 Pflop/s
- 719 Tbyte main memory and
- 70 Pbyte disk storage
- 6,480 Lenovo ThinkSystem nodes with Intel Xeon processors (Skylake)
- 311,040 compute cores
- Intel Omni-Path interconnects
- Direct hot water cooled + Adsorption coolers (47 C)

HPC + Cloud

- Usage of own and individual virtual machines (integrated cloud)
- Pre- and post-processing with user's individual software
- Integrated development, ability to use familiar software and tools
- Remote visualization and integration to V2C

Data Storage & Sharing Facilities at LRZ

(a simplified overview)

Large-volume Data | Data for/from HPC

Systems based on IBM Spectrum Scale (GPFS)

- **HPC cluster file systems (SMUC-NG: 50PB)**
  ($HOME, $WORK, $SCRATCH)
- **Data Science Storage (DSS, >20PB)**
  - NFS (Network File System) export
  - access from all LRZ systems
  - configurable via Website and API/CLI
  - data transfer and sharing: GLOBUS

Large-volume Data | Backup/Archive

- **Tape+Disk Archive (→ 100PB)**
  (IBM Spectrum Protect)

Small-/medium-volume Data

- **“Cloud Storage” Personal/Institutional**
  CIFS, NFS (up to 10TB per client)
- **LRZ Sync & Share**
  with access via app or web GUI
  “Dropbox”-like (<50GB per client)
- **LRZ GitLab**
  Git Repository Platform (some GB per client)

How we witness Convergence

… of systems: Supercomputers, Cloud, GPU systems …

… and communities/techniques: HPC, HTC, HPDA, AI

A New World is Emerging: High Performance AI (HPAI)

New User Communities with New Workflows

Ability and Expertise to Target Large Scale Problems

High Performance Data Computing at LRZ

Data-Analytics Communities

Increasing
computing demands

HPC User Communities

Increasing
analytics demands

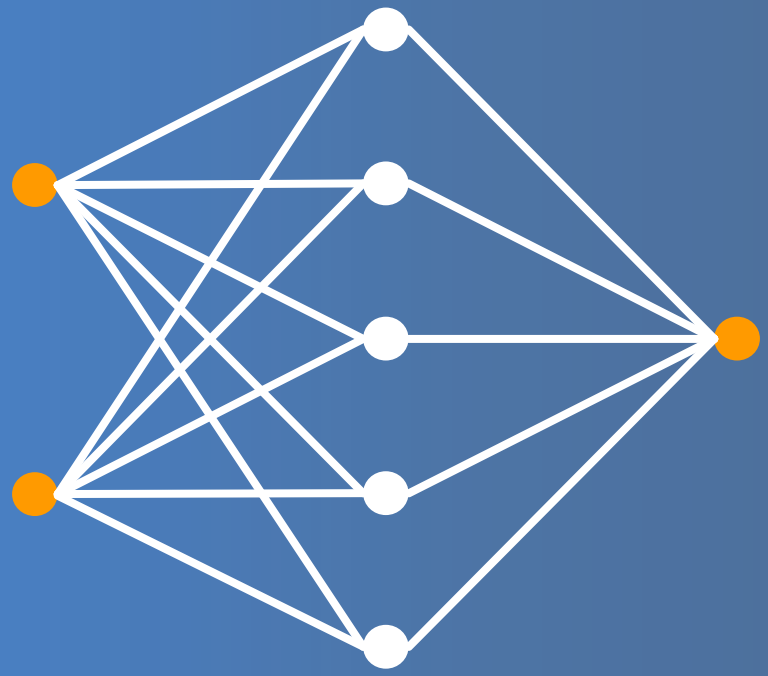

Data Intensive Computing, Data Analytics & AI

Wait, but I've always done simulations with preprocessing and data analysis !?

Yes, but you've (probably) not leveraged the full spectrum of data analytics! Just look at some big companies…

Wait, but I've always used computers for my AI !?

Yes, but you've (probably) not used a supercomputer, but a workstation with a graphics card.

High Performance Data Computing at LRZ

Data-Analytics Communities

Increasing
computing demands

HPC User Communities

Increasing
analytics demands

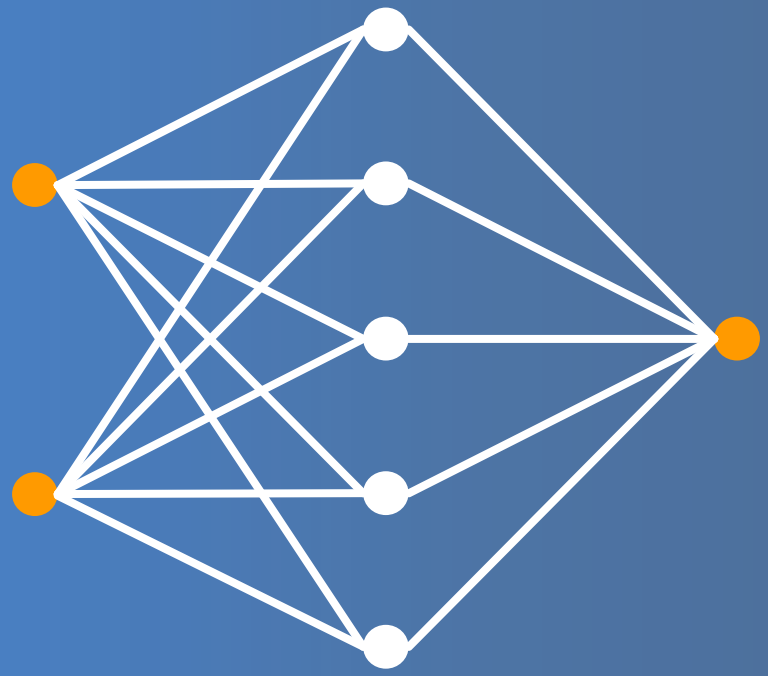

Data Intensive Computing, Data Analytics & AI

What does it mean for LRZ?

- More Data-Centric field of work than before
- New Systems (e.g. Compute Cloud)
- New work for the Application Labs (AstroLab, …)
- Completely new fields of work (Big Data, Research Data Management, …)

LEXIS: A systematic approach to HPC-Cloud-Big-Data Workflows

Large-scale EXecution for Industry and Society (LEXIS)

LEXIS quick facts:

Lead: IT4I (CZ), Jan Martinovic

Funding: H2020

Sum: 12.2M€

2019/01 – 2021/06

Co-funded by the Horizon 2020 Framework Programme of the EU Grant Agreement Number 825532

- Platform for **Distributed Computing / Data Analytics** (Science & Engineering)
  - YORC (Ystia Orchestrator) → **workflow modelling & MGMT**
  - HEAppE → **execution middleware** (HPC job MGMT)
  - EUDAT-B2SAFE / iRODS → **distributed data store**
  - **AAI** with OpenID support

  … and other features: **Billing, Burst-Buffer** hardware, …
- **Pilot Use Cases: HPC-Cloud** workflows **(data analysis + computing)**

Weather & flood/forest-fire prediction workflow based on WRF model

LEXIS: Weather & Climate Use Case

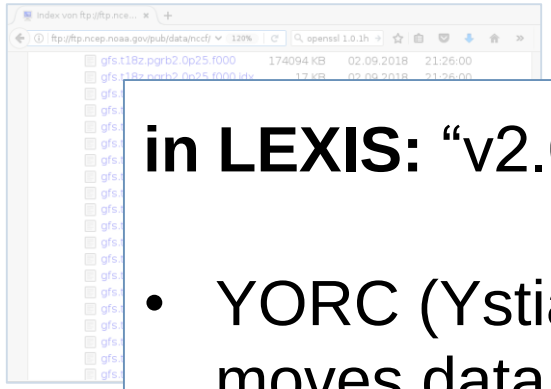

Weather Data Archive

loud VMs

ed data

Forest Fire prediction
Flood prediction

(running on LRZ *VMWare cluster*)

in LEXIS: “v2.0 of this workflow” (and others)

- YORC (Ystia Orchestrator, Bull-Atos)
  moves data, launches computation
- EUDAT-B2SAFE- / iRODS-based Distributed Data Infrastructure (DDI)
  for accessing inputs & outputs from LRZ and IT4I (CZ)
- Weather and Climate Data API from ECMWF
  API endpoints + Database for storing weather model data

LEXIS: Cloud-HPC Workflow Pilots to be implemented

- **WEATHER & CLIMATE**
  - Aim: rich workflow with massive data assimilation (model forced to follow sensor-network data) and post-processing
  - Orchestrate data collection (including Edge Computing) & simulation jobs
- **AERONAUTICS**
  - Aim: simulations as part of “real time” component-design process
  - Code & workflow optimization, visualization / analytics on result
  - Utilisation of novel hardware (GPUs/FPGAs, Burst Buffers)
- **EARTHQUAKE & TSUNAMI:**
  - Aim: Near-Real-Time decision support after earthquake
  - Get measurements, predict tsunamis and other damages
    → Simulations and data analytics under strong time constraint

Technology behind LEXIS

Computing (Cloud/HPC/Burst Buffers)

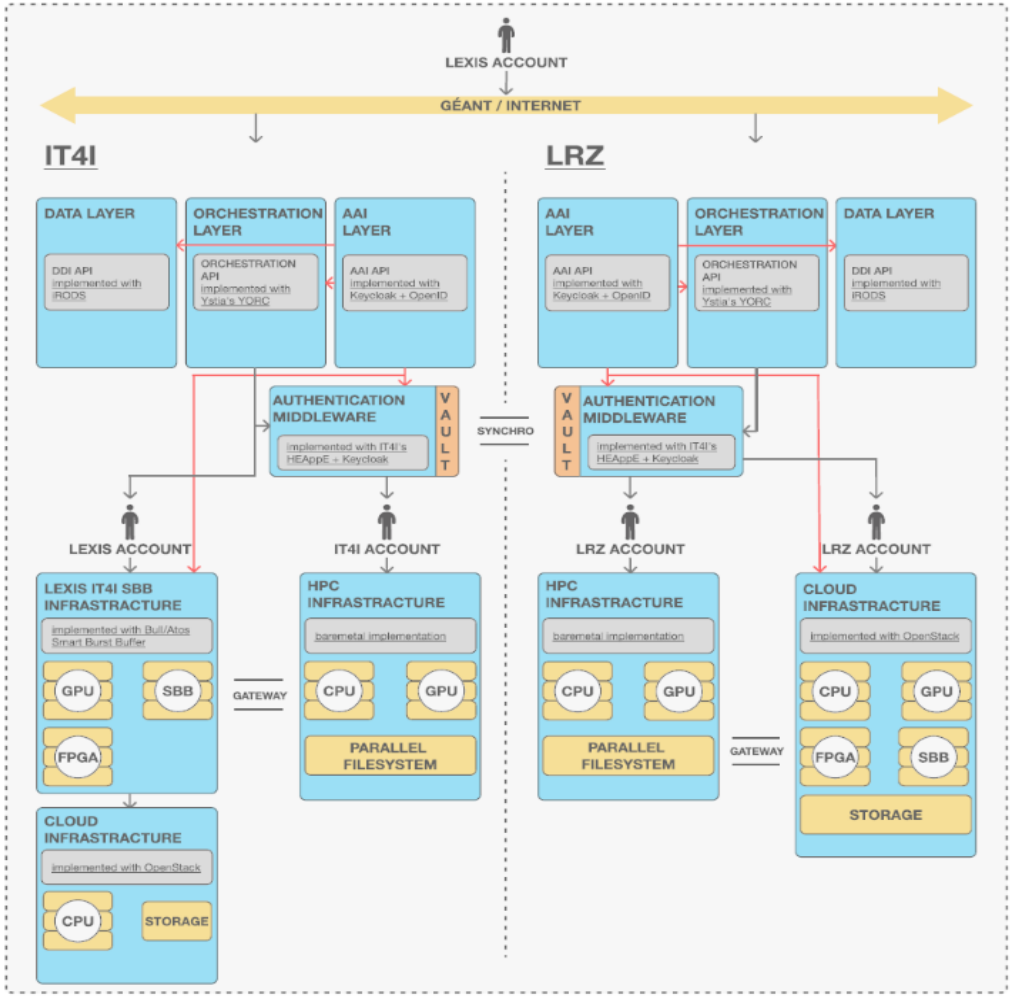

Data (iRODS, EUDAT-B2SAFE)

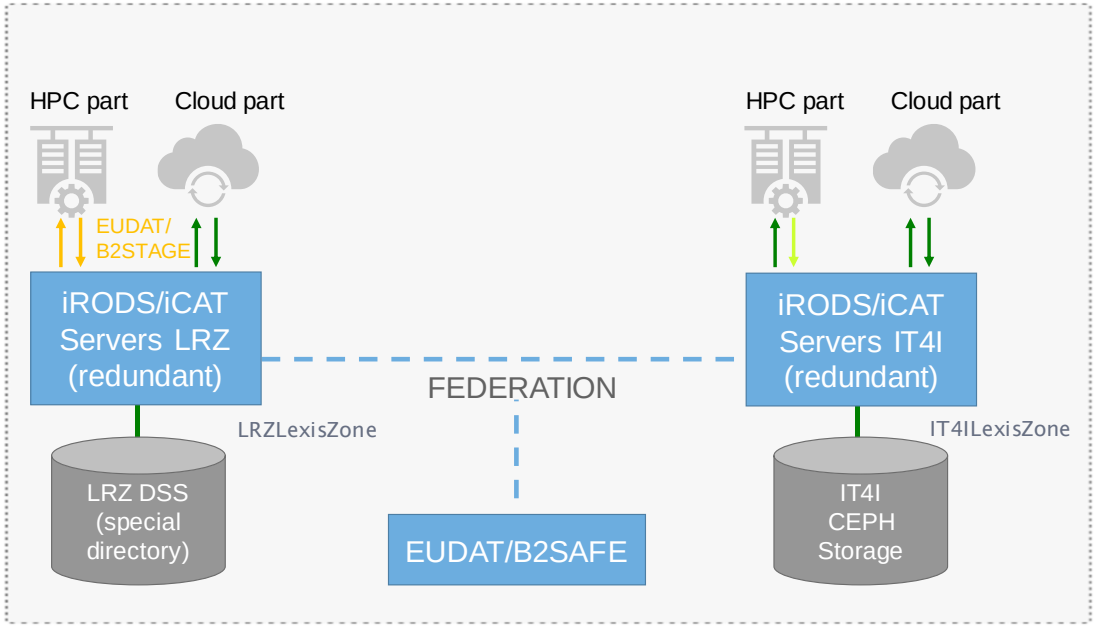

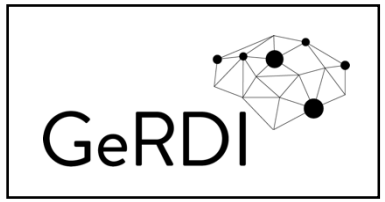

Search Engines
→ FAIR data

Simulation- and data-intensive Use Cases from LRZ's AstroLab

Supporting a forward-looking research community

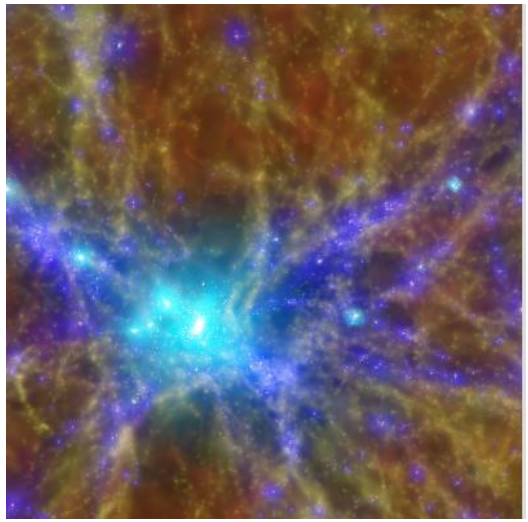

Dolag et al. (LMU)

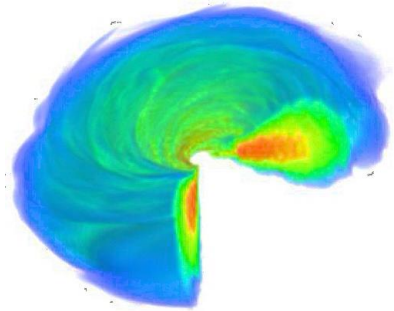

Bugli et al. (CEA Saclay)

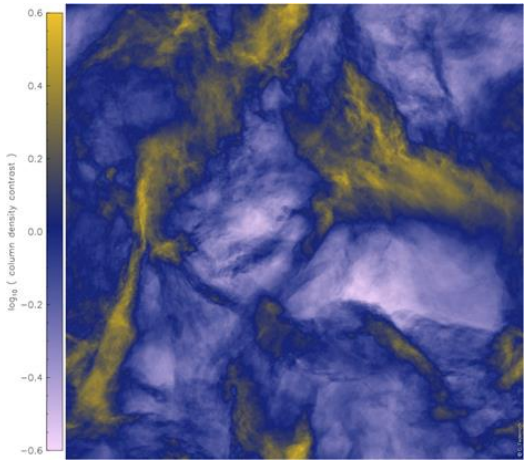

Federrath et al. (ANU / ITA)

High-level support (also GCS, PRACE): code modernisation on SuperMUC and SuperMUC-NG.

Highlights:

- Intel Parallel Computing Center @ LRZ/TUM
- Magneticum simulation suite (Dolag)
- The world's largest turbulence simulation (Federrath & Klessen)
- Numerical simulations for the Event Horizon Telescope (Bugli; Porth)

AstroLab: new steps as HPC evolves

Supporting cutting-edge research

- Runs at machine scale (Gadget, Arepo, GReX, FLASH)

Machine learning and HPC

- AI at scale on SuperMUC-NG

Data Science

- Collaboration with Astro@NFDI (National Research Data Infrastructure)
- Engagement in the SKA and LSST communities

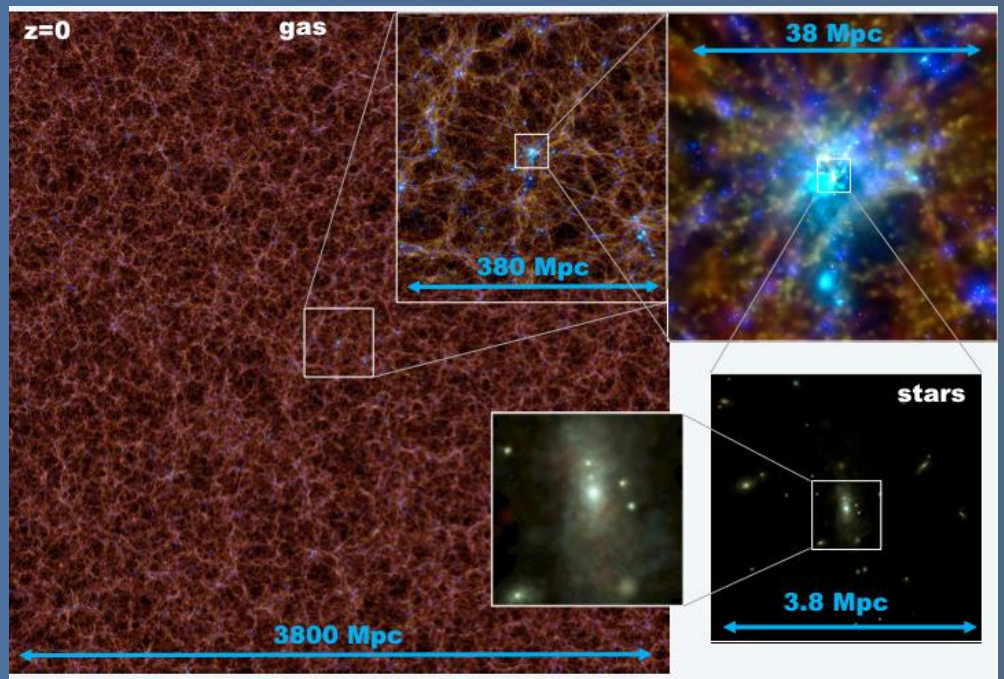

Dolag et al. (LMU)

What do the *first image of a black hole* and HPC have in common?

Image reduction:
PBytes of data

Theoretical modeling:
comparison with
simulations is necessary

ECHO-3DHPC has been part of the comparison project of general relativity simulation codes (Porth et al. 2019, arXiv: 1904.04923) to validate the observational results

Both ECHO and BHAC (Black Hole Accretion Code; Porth et al.) have been optimized in collaboration with LRZ

The use of both SuperMUC and Hazel Hen (HLRS Stuttgart) is acknowledged in the EHT papers

(image credits: EHT Collaboration 2019)

Visualizing the **world's largest** turbulent simulation

S. Cielo | Leibniz Supercomputing Centre (LRZ)
L. Iapichino | Leibniz Supercomputing Centre (LRZ)
J. Günther | Intel Germany
C. Federrath | Australian National University
E. Mayer | Leibniz Supercomputing Centre (LRZ)
M. Wiedemann | Leibniz Supercomputing Centre (LRZ)

How are stars born?

Physics

State of
interstellar gas

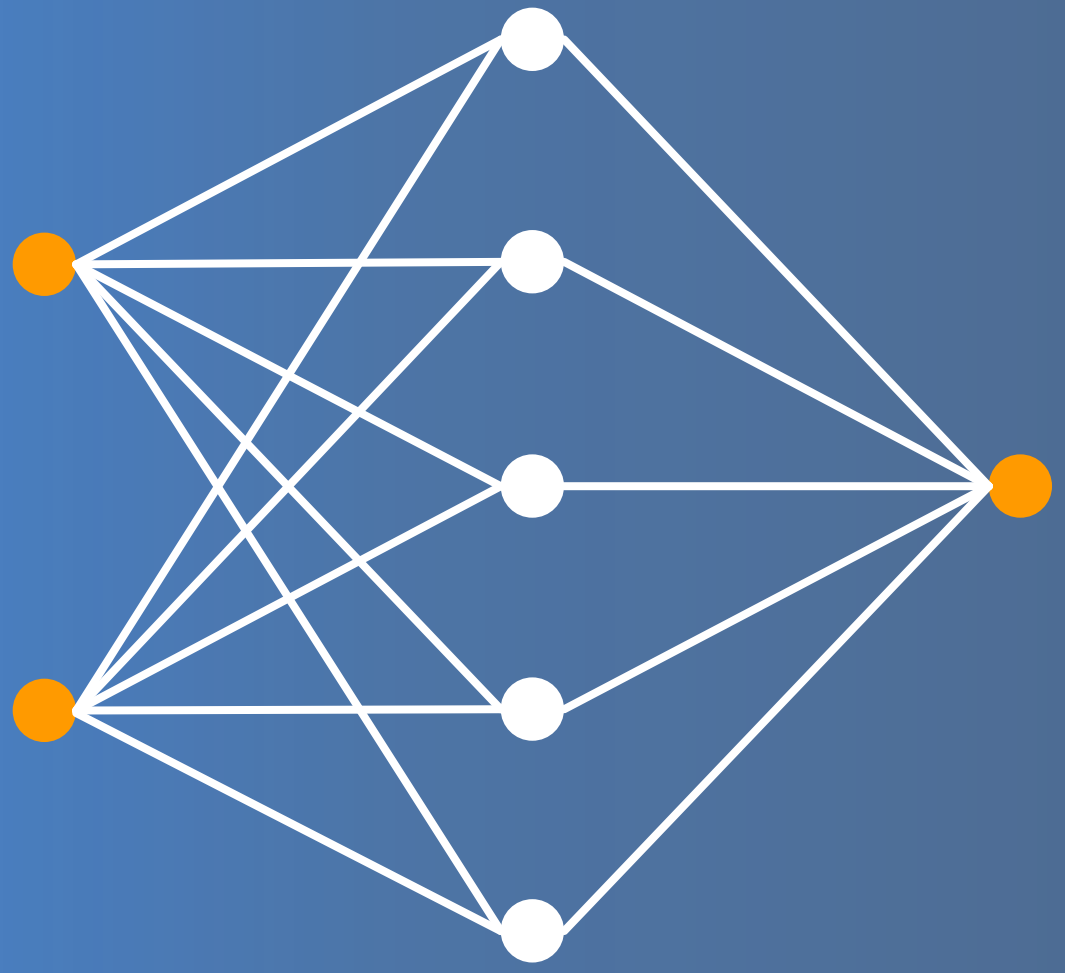

Star Formation
models

Physics → State of the gas → Simulating turbulence

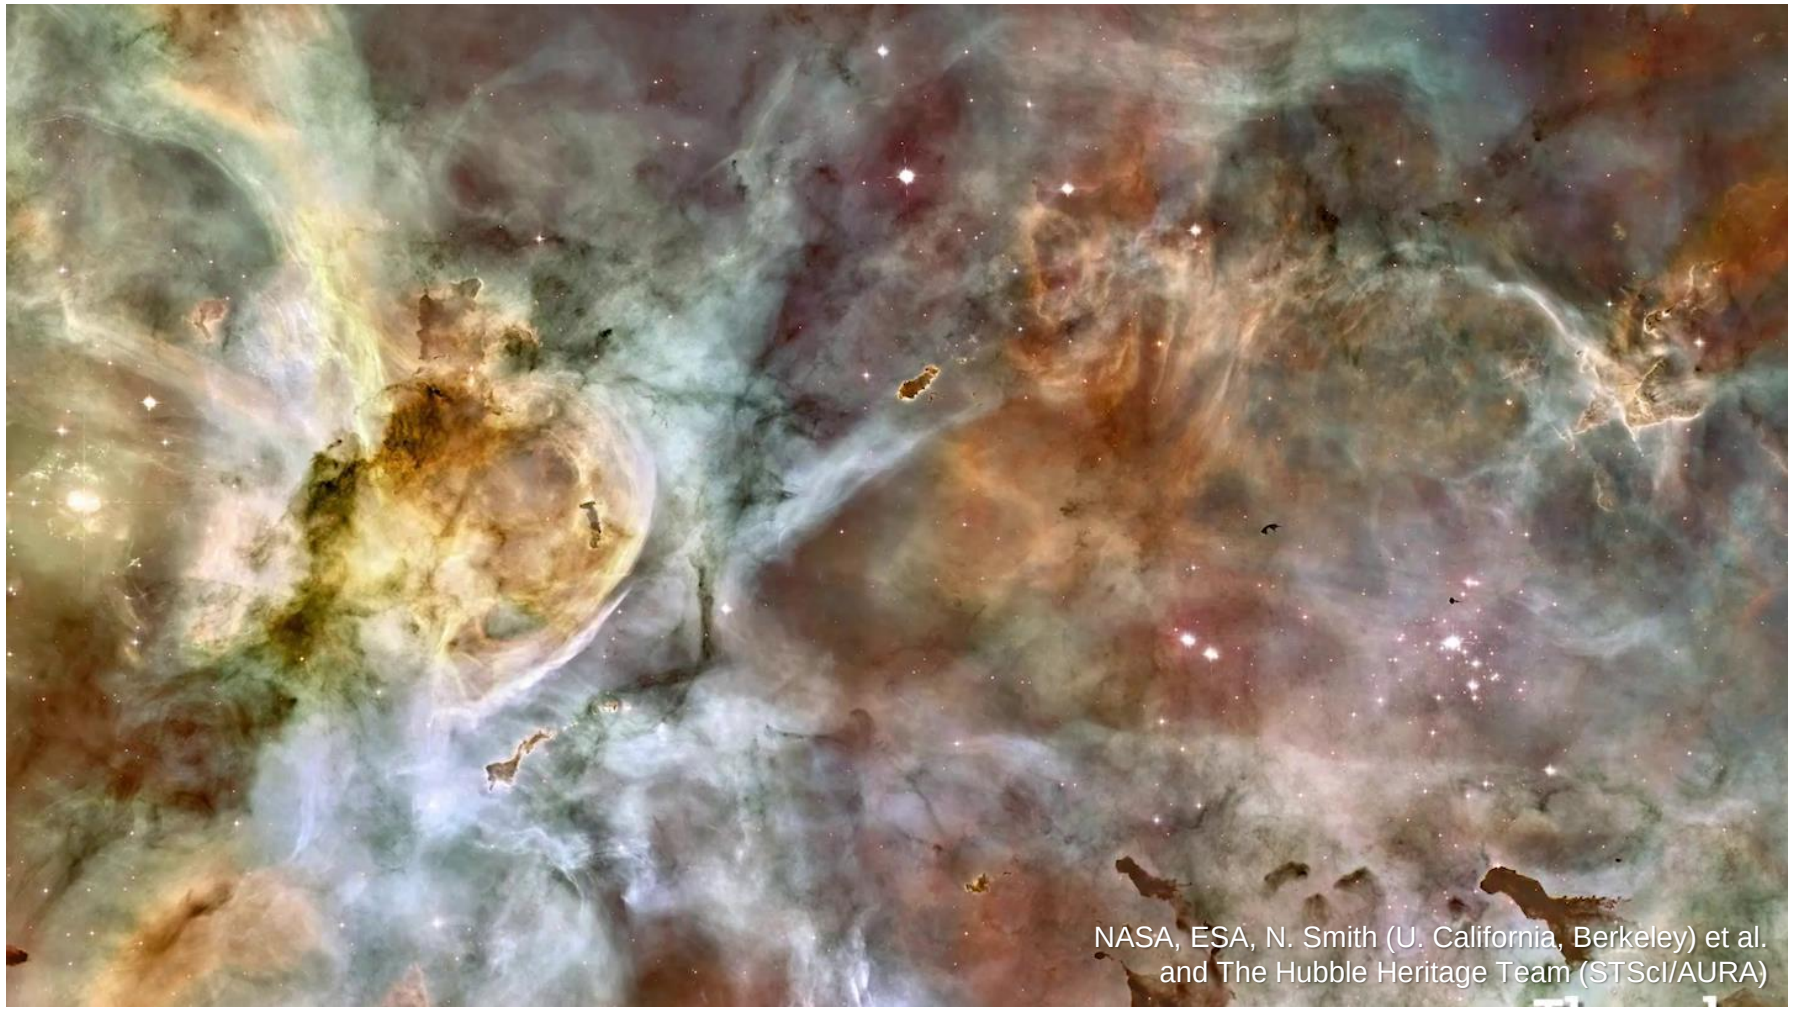

- **Star Formation**: an inefficient process
- **Shocks → Turbulence „seeds“ the stars**
- **Simulations: stirred gas in a box**

The largest **turbulence simulations**

FLASH: MHD code, optimized

- Used in (thermo)nuclear (astro)physics
- Grid 10 048^3. Data dumps: 23TB each
- Optimized with LRZ for hybrid precision
- Very good weak scaling up to full machine

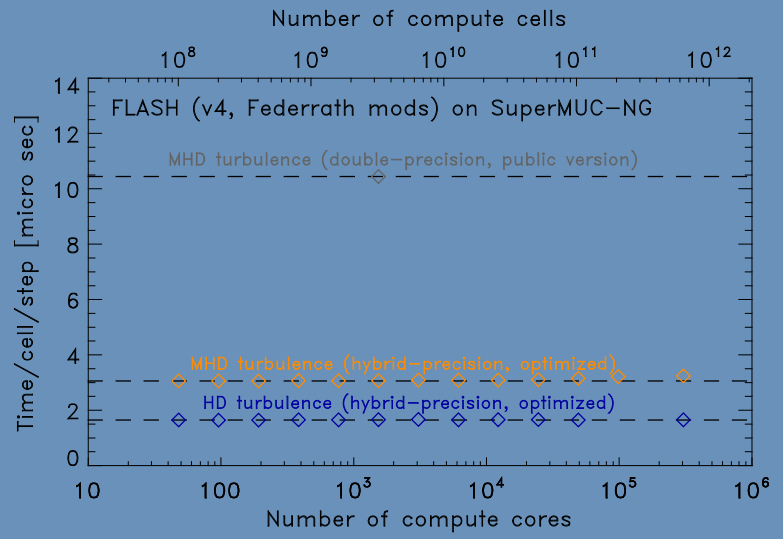

Simulations and Structure Function

- **Astrophysics:** Identified sonic scale for the first time!
- **HPC challenge:** resolving both super- and sub- sonic

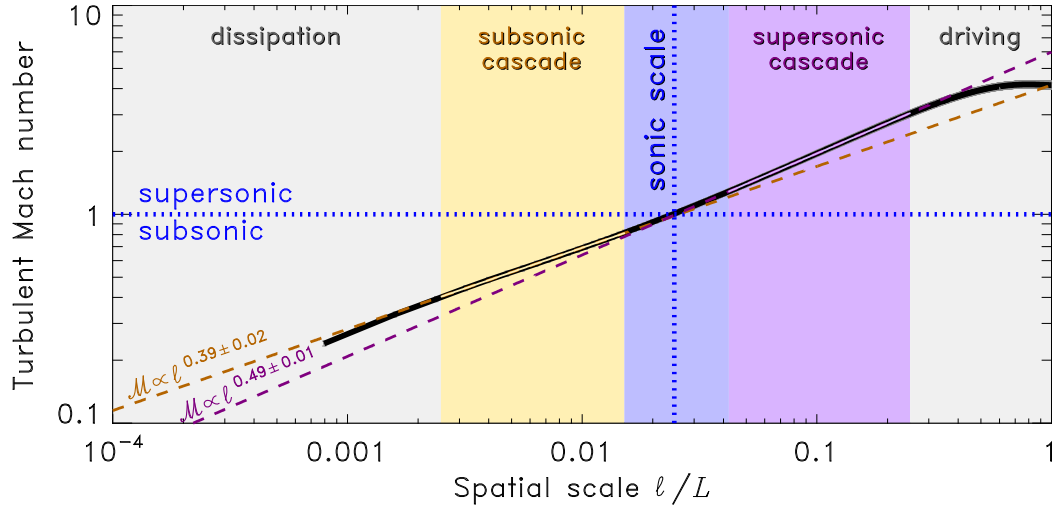

Visualization Infrastructure

OSPRay + VisIt…

- Now also public, since VisIt 3.0
- Scalable on nonaccelerated HPC systems
- TBB workers + multi-node (via MPI)
- Interactive GUI or batch, cross-code

OSPRay: **O**pen, **S**calable, **P**ortable, **Ray** Tracing Engine for High-Fidelity Visualization.

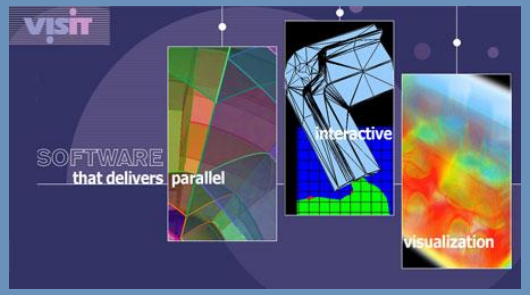

VisIt is an **open source**, interactive, scalable, multiplatform visualization, animation and analysis tool.

… on SuperMUC-NG

- Custom installation: shades with MPI
- In the software stack, for all users
- Strong-scaling to half machine:

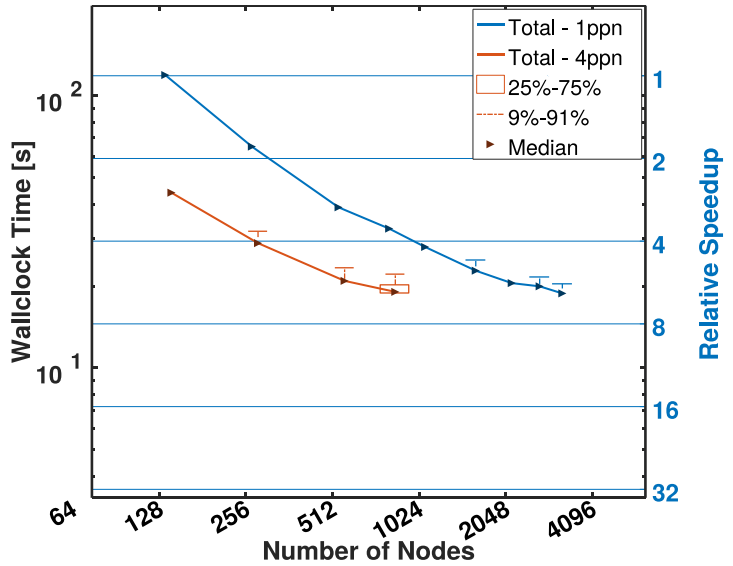

Visualizing the world's largest turbulence simulation

- Finalist at the SC19 Scientific Visualization & Data Analytics Showcase
- https://www.youtube.com/watch?v=EPe1Ho5qRuM
- Now there is time only for a short trailer…

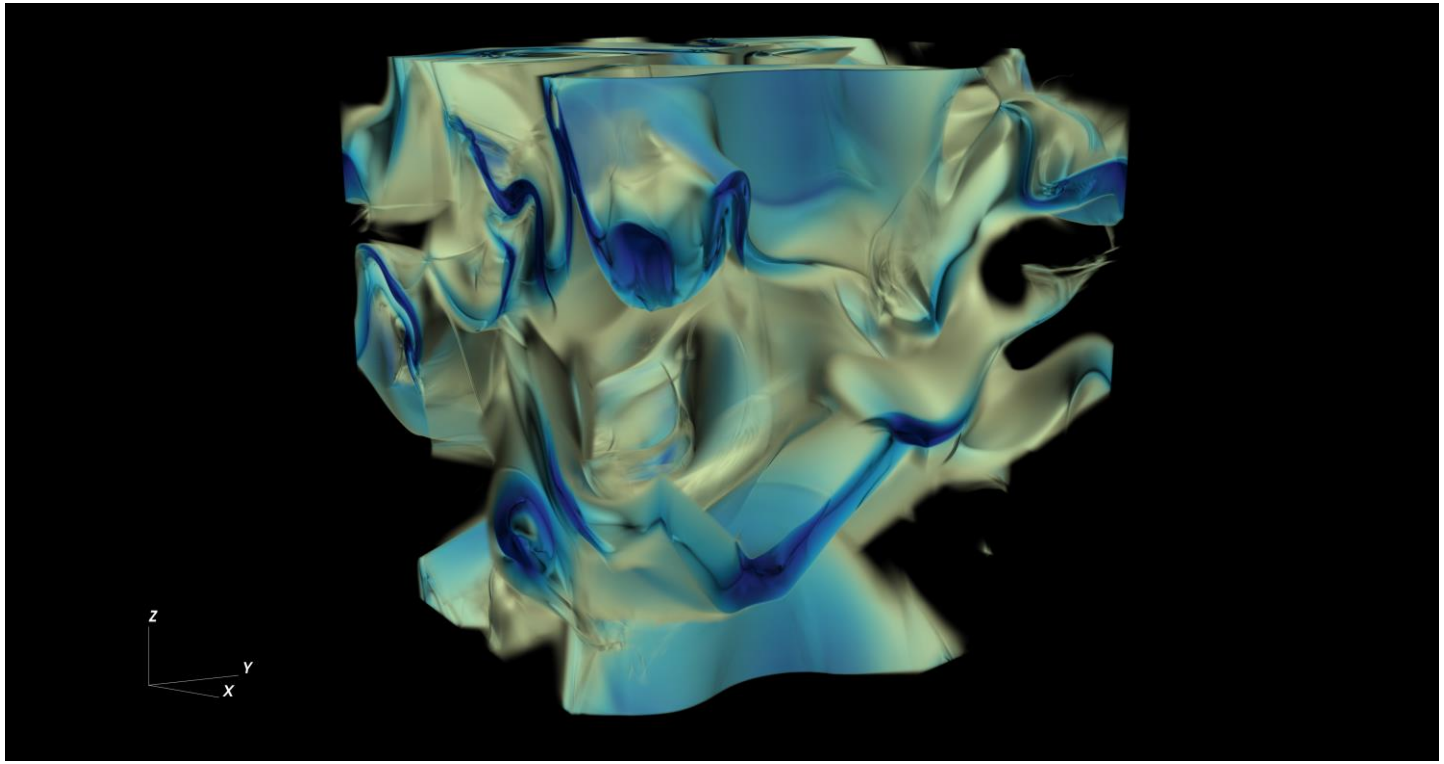

MAGNETICUM – A simulation suite of cosmological structure formation

Cosmological structure formation in the Universe (10^{11} Particles / P-GADGET3)

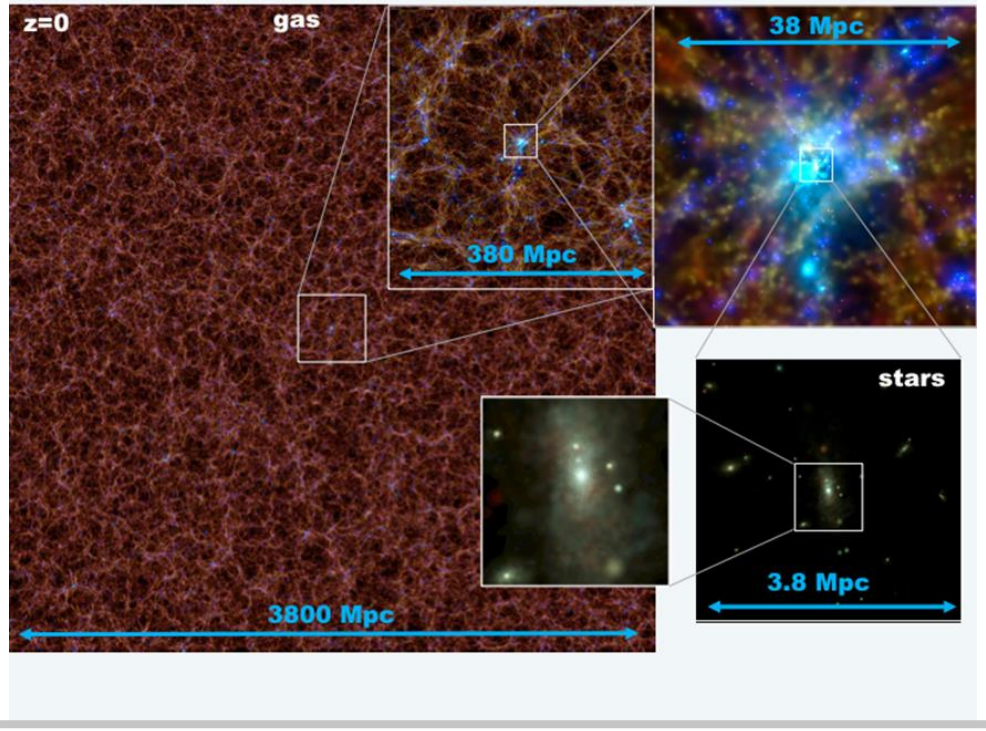

Full machine-scale production runs

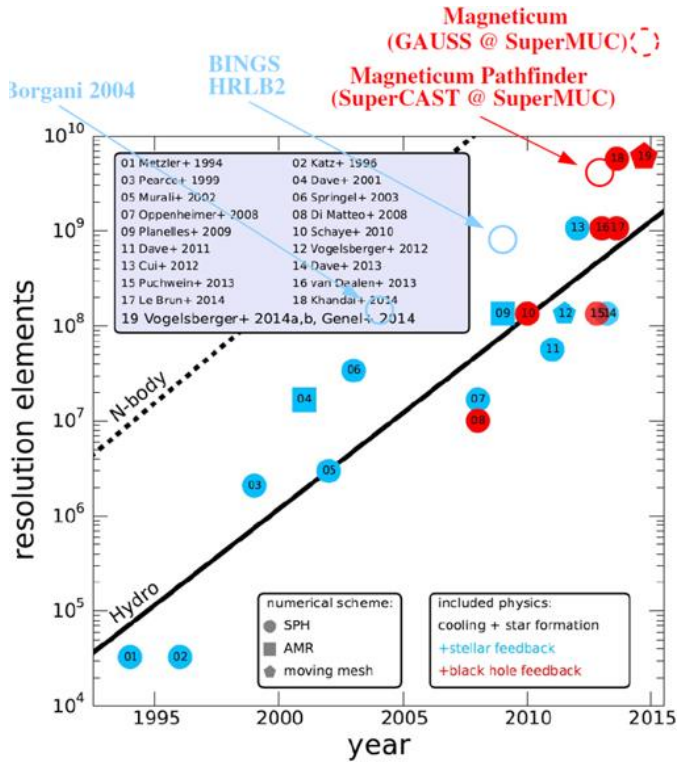

Collaboration partners: USM (Klaus Dolag), C²PAP, LRZ, IAP, INAF

Ragagnin et al. 2017, Astronomy & Computing

A Web Portal for hydrodynamical, cosmological simulations

- Web Interface: LRZ
- Computing cluster: C²PAP
- Storage: DSS (LRZ)

Browsing simulations throught the Web Interface

Analysis performed on a dedicated cluster

Results retrieved through FTP

More details in the movie (originally for ISC'17)

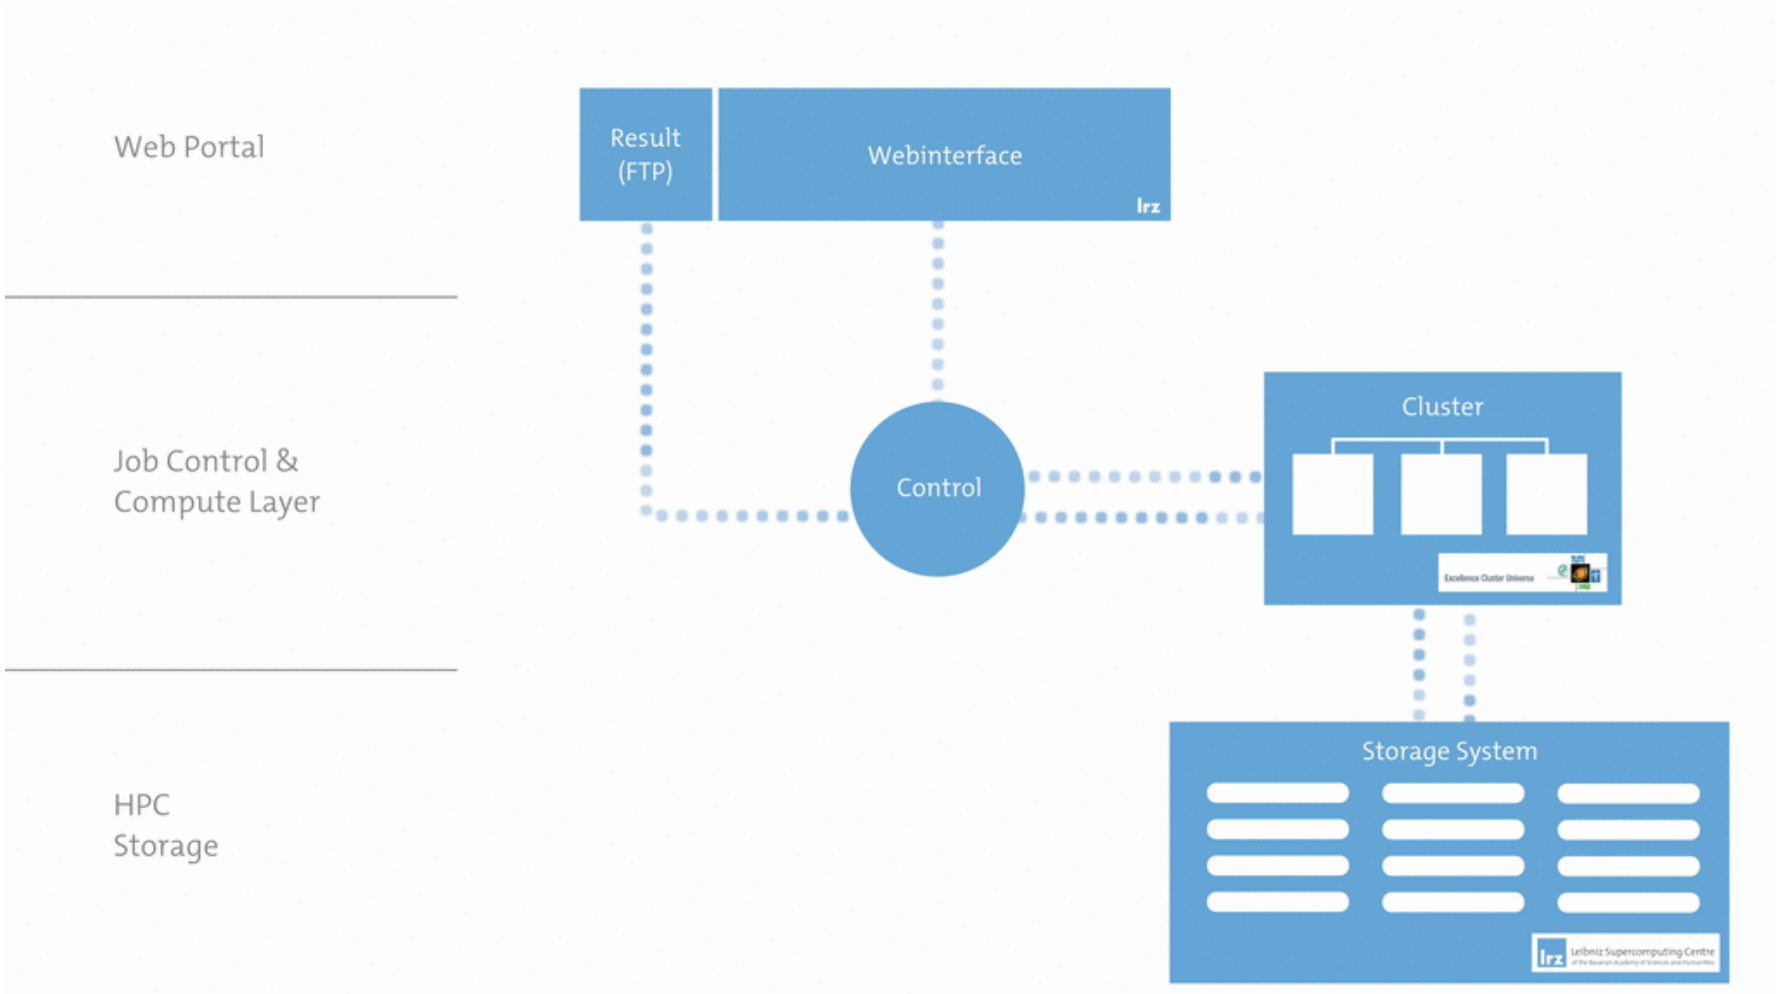

https://c2papcosmosim.uc.lrz.de/

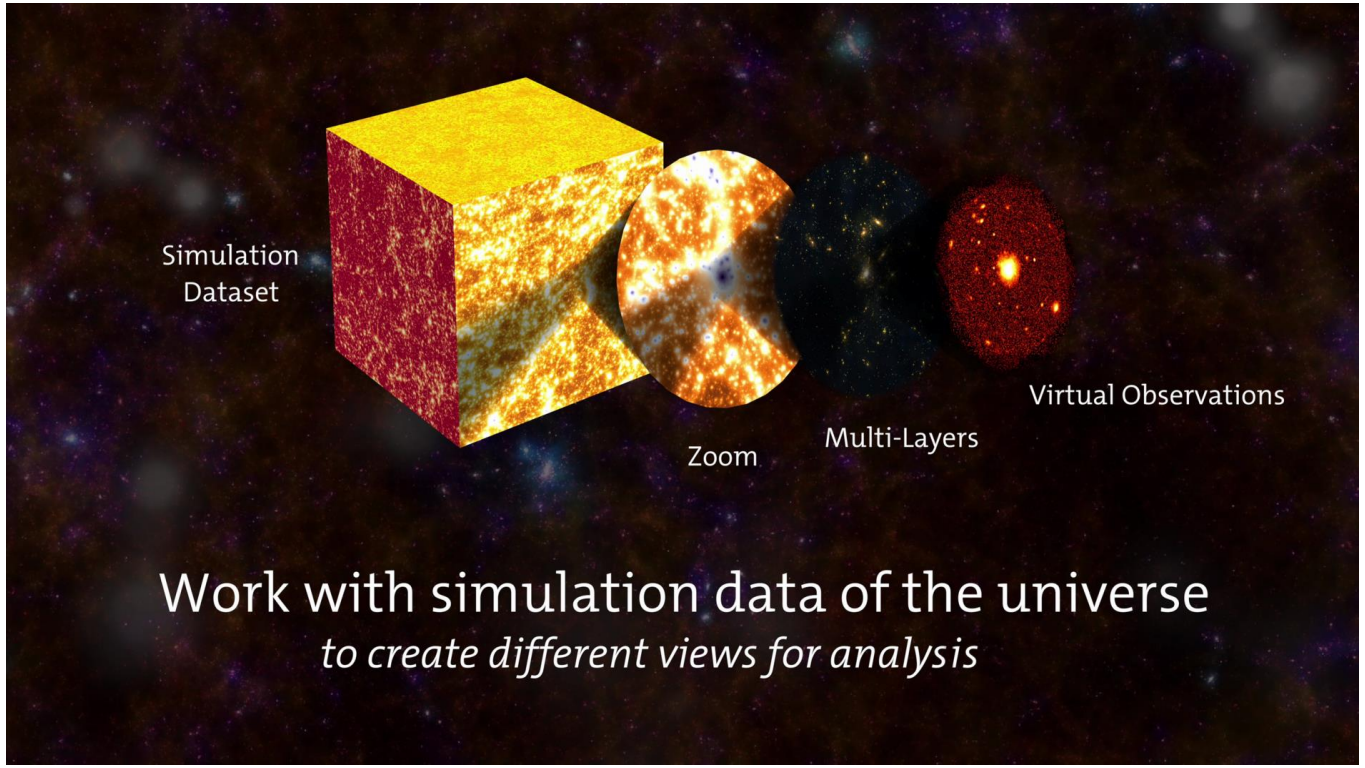

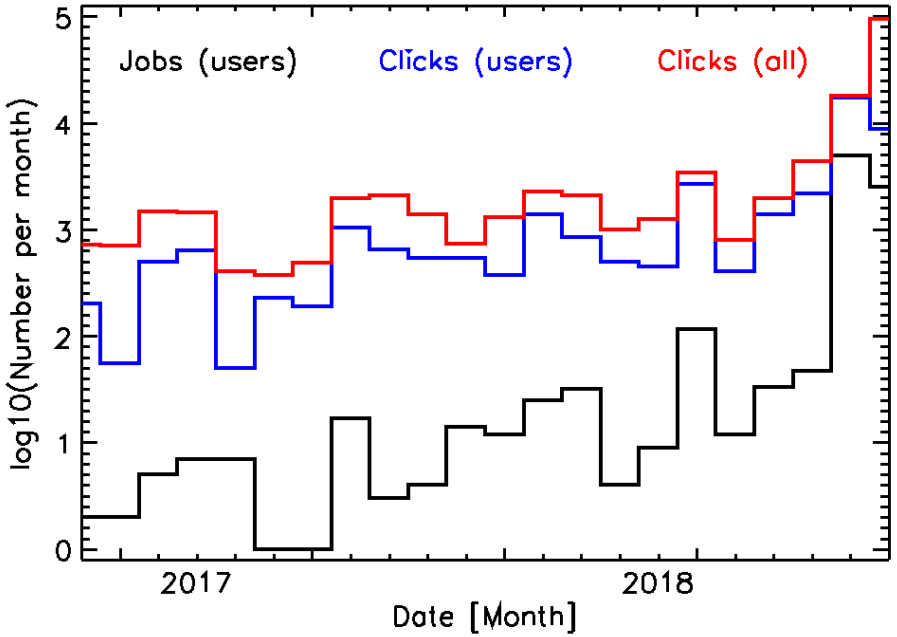

As of beginning 2019

- 130 registered users;
- first scientific publications with analysis performed on the Web Portal.

FAIR Supercomputing Data at LRZ

Towards a new Research Data Management Service

Research Data Management – Why?

External Scientific Users

Research Data Search Engines

Metadata Store + **lots** of „glue“

Medium-volume systems

- MWN Cloud Storage
- LRZ GitLab

Large-volume | HPC

- Data Science Storage (DSS)

Large-volume | Archive

- IBM Spectrum Protect (Tapes/Disks)

[1] e.g. Wilkinson M.D. et al., 2016, Scientific Data, 3, 160018

‘Let the Data Sing‘ - A lightweight, Microservice-based RDM Layer

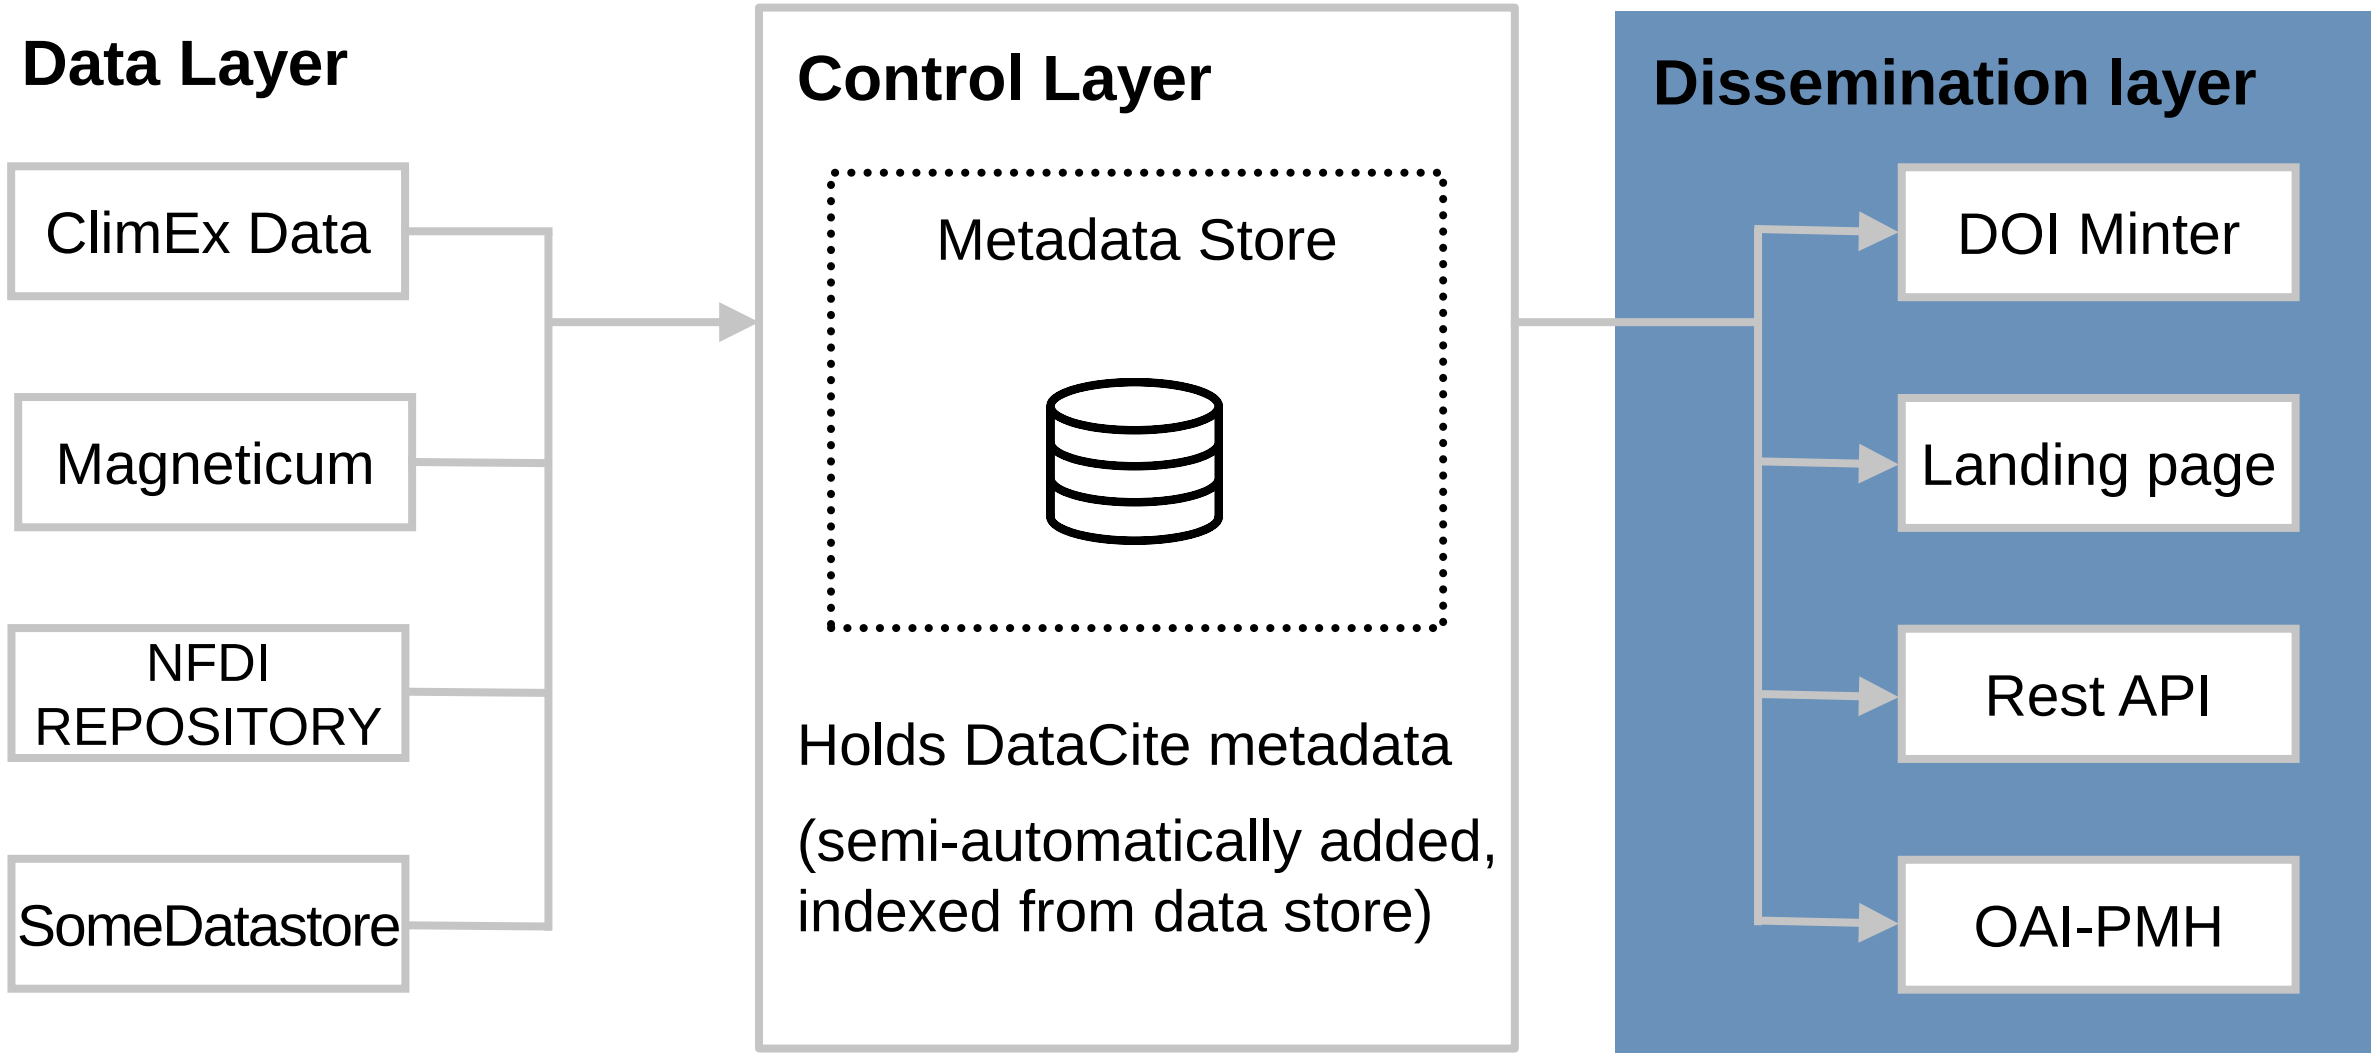

Vision

FAIR data (in NFDI etc.), and „best-of-breed“ analytics solutions

Users

Semi-Automatised Data Analytics

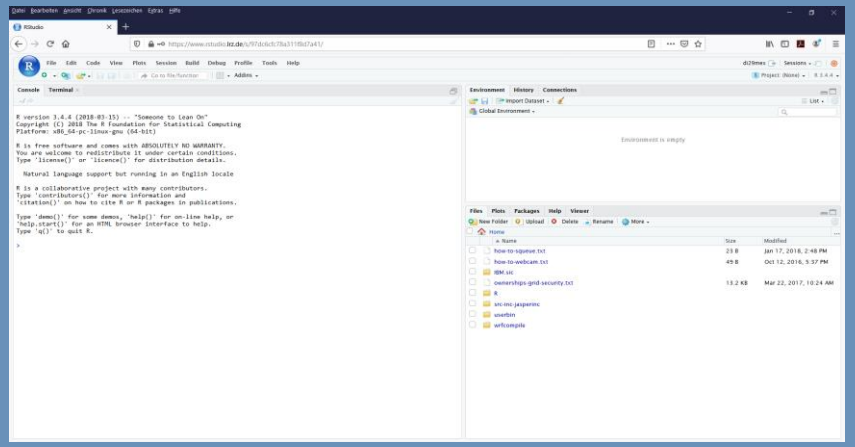

- Solutions as JupyterHub (JSC, GCS), www.rstudio.lrz.de
- Supported by persistent identifiers & metadata

Research Data Search Sites

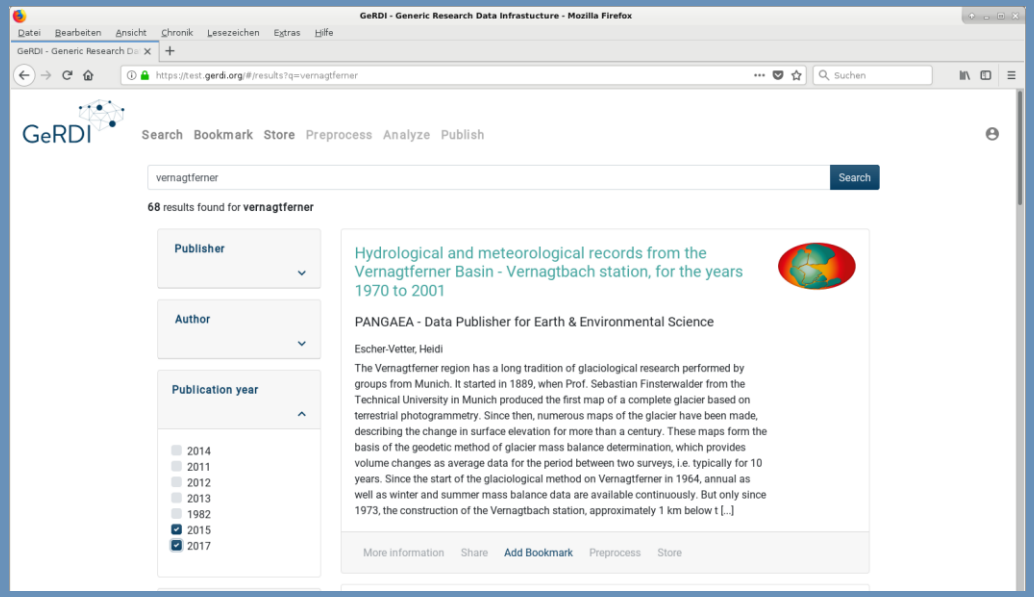

- Picture: GeRDI (Generic Research Data Infrastructure, DFG project including LRZ)
- EUDAT-B2FIND
- BASE
- Google Dataset Search, …

Thank you!

Luigi Iapichino – LRZ AstroLab & Quantum Computing lead – iapichino@lrz.de
LRZ AstroLab – astrolab@lists.lrz.de

Stephan Hachinger – LRZ RDM Team lead – hachinger@lrz.de
LRZ RDM Team – rdm@lists.lrz.de