Leibniz Supercomputing Centre

of the Bavarian Academy of Sciences and Humanities

Tools and systems for HPC and AI at LRZ

22.7.2019 | Luigi Iapichino

About the presenter

Dr. Luigi Iapichino

- Astrophysics and Quantum Computing Application Specialist
- High Performance Systems Division, LRZ

- Team leader of the Application Lab for Astrophysics (LRZ AstroLab)
- Lead of Quantum Computing @ LRZ
- Expert in computational astrophysics and simulations
- Member of the PRACE High-Level Support Team

Email: Luigi.Iapichino@lrz.de

Thanks to Nicolay Hammer and David Brayford (LRZ) for having provided many of the next slides

Leibniz Supercomputing Centre

Garching, Germany

Leibniz Supercomputing Centre

Bavarian Academy of Sciences and Humanities

We are the computing backbone for advanced research science in Bavaria

250
employees
approx.

56
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

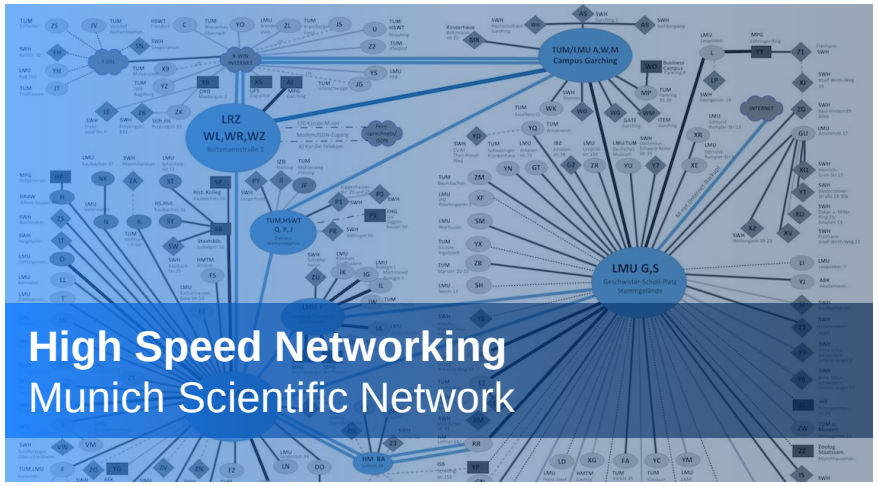

High Speed Networking
Munich Scientific Network

High Performance Computing
SuperMUC-NG, Linux cluster

Big Data
Bavarian State Library Digital Archive

Virtual Reality and Visualisation
V2C (CAVE, Powerwall)

For German Research

Gauss Centre for Supercomputing (GCS)

Alliance for Germany's Tier-0
high performance computing centers

- LRZ | Munich | SuperMUC-NG
- HLRS | Stuttgart | Hazel Hen
- JSC | Jülich | JUWELS

Founded 13. April 2007

For European Research

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

LRZ Systems and Access

National and International

PRACE

GCS

Local and Regional

Munich TUM and LMU
~30% of SuperMUC usage

***Students**
Training future experts

Bavarian projects
<1 million CPU hours

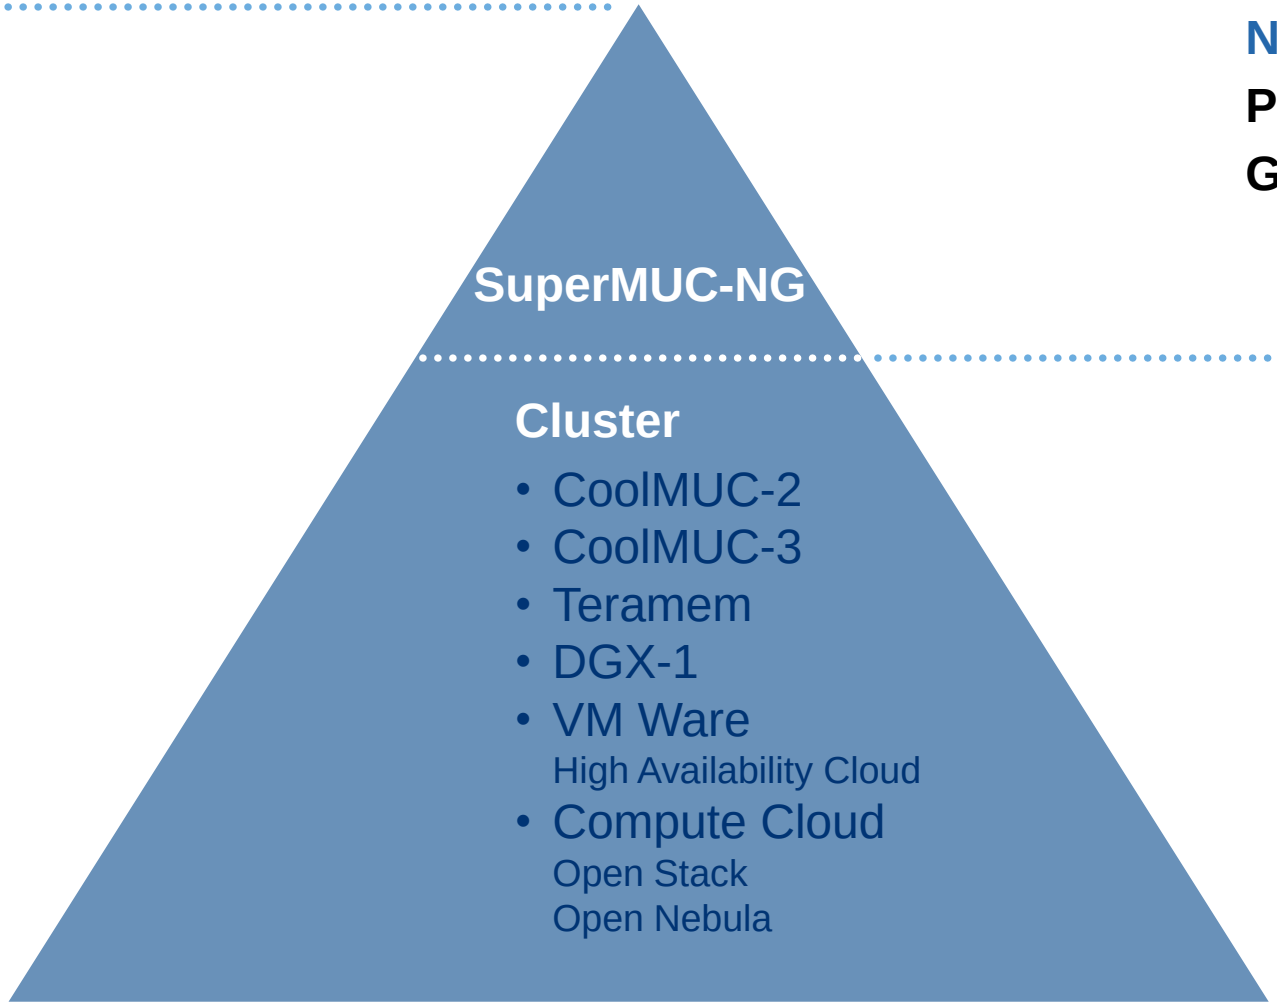

High Performance Data Computing at LRZ

Emerging Communities
Increasing **computing** demands

HPC User Communities
Increasing **analytics** demands

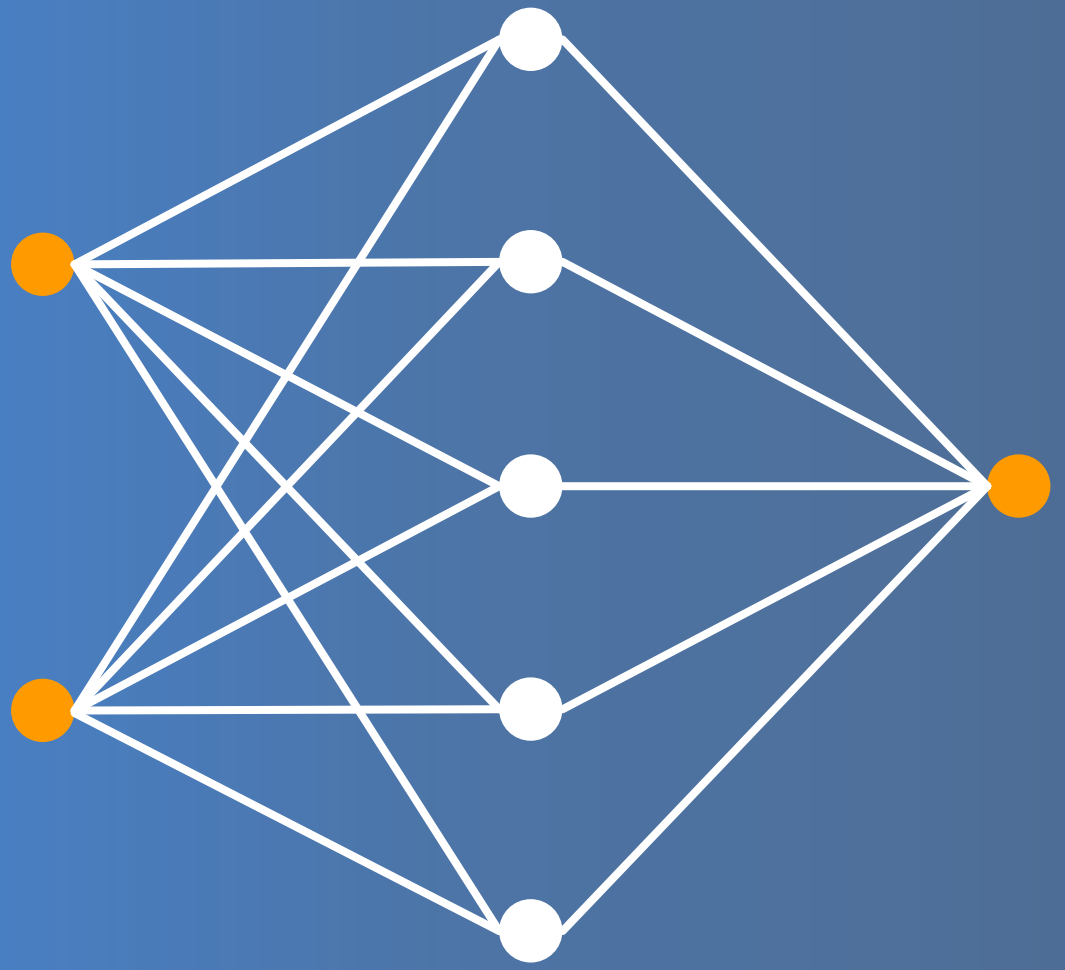

Data Intensive Computing, Data Analytics & AI

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

New User Communities with New Workflows

Ability and Expertise to Target Large Scale Problems

THE COMPUTING INNOVATION CYCLE

Advanced computing and huge volumes of data creates new opportunities for information insight.

HPC

Modeling & Simulation (M&S)

Natural World ▹ Hypothesis ▹ Equations ▹ Algorithms ▹ Computing ▹ Data ▹ Analysis

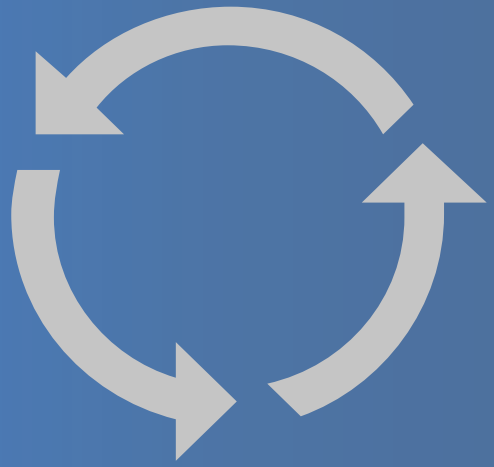

AI & Machine Learning

Data ▹ Algorithms ▹ Computing ▹ Pattern Recognition

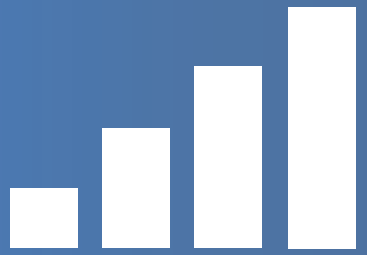

Big Data

THE COMPUTING INNOVATION CYCLE

Advanced computing and huge volumes of data creates new opportunities for information insight.

HPC

Modeling & Simulation (M&S)

Natural World ▹ Hypothesis ▹ Equations ▹ Algorithms ▹ Computing ▹ Data ▹ Analysis

AI & Machine Learning

Data ▹ Algorithms ▹ Computing ▹ Pattern Recognition

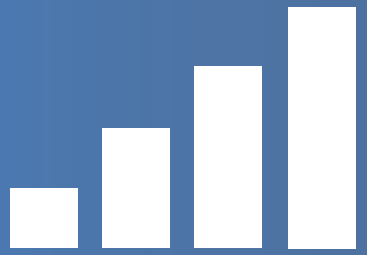

Big Data

High Performance AI (HPAI) in a **Container**

Transition AI algorithms from the
laptop to supercomputer
with minimal effort

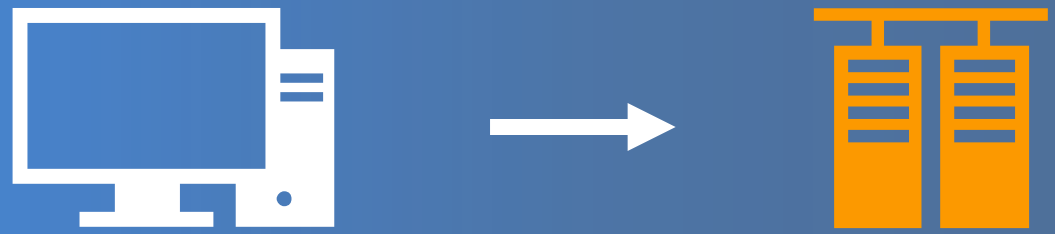

“It just works”

HPAI =

M&S

- Equation based on model
- Computing driven
- Numerically intensive
- Creates simulations
- Monte Carlo
- Larger problems
- Iterative methods
- PDE

+

- Linear algebra
- Matrix operations
- Iterative methods
- Compute intensive
- Data transfer
- Predictive
- Probabilities
- Stencil codes
- Calculus
- Pattern recognition
- Graphs

Analytics

- Finds patterns
- Correlations in data
- Logic driven
- Creates inferences
- Knowledge discovery
- Graphs
- Data-driven science
- Predictions
- CNN
- RNN

Differences Between HPC & AI

AI	HPC
Large number of small files	Small number of large files
Large memory nodes (+1TB)	Memory per node (32/64GB)
Single node	Multiple nodes
Single GPU/accelerator node	Distribute compute over many nodes
Local node storage	Typically diskless systems (no local storage)
Data transfer within a single node. (PCI bus)	Data transfer between multiple nodes
Matrices are typically small	Medium to large matrices
Root privileges	User privileges

Requirements for AI on HPC

Compute intensive hardware

Optimized AI frameworks (TensorFlow, Caffe)

Optimized software (numerical libraries, Python)

HPC specific software (distributed computing, workload manager)

Method of deploying the AI software in a simple, straightforward and flexible way

Need to get to: “It just works”

Compute System Diversity
with fully integrated Central Data Silos

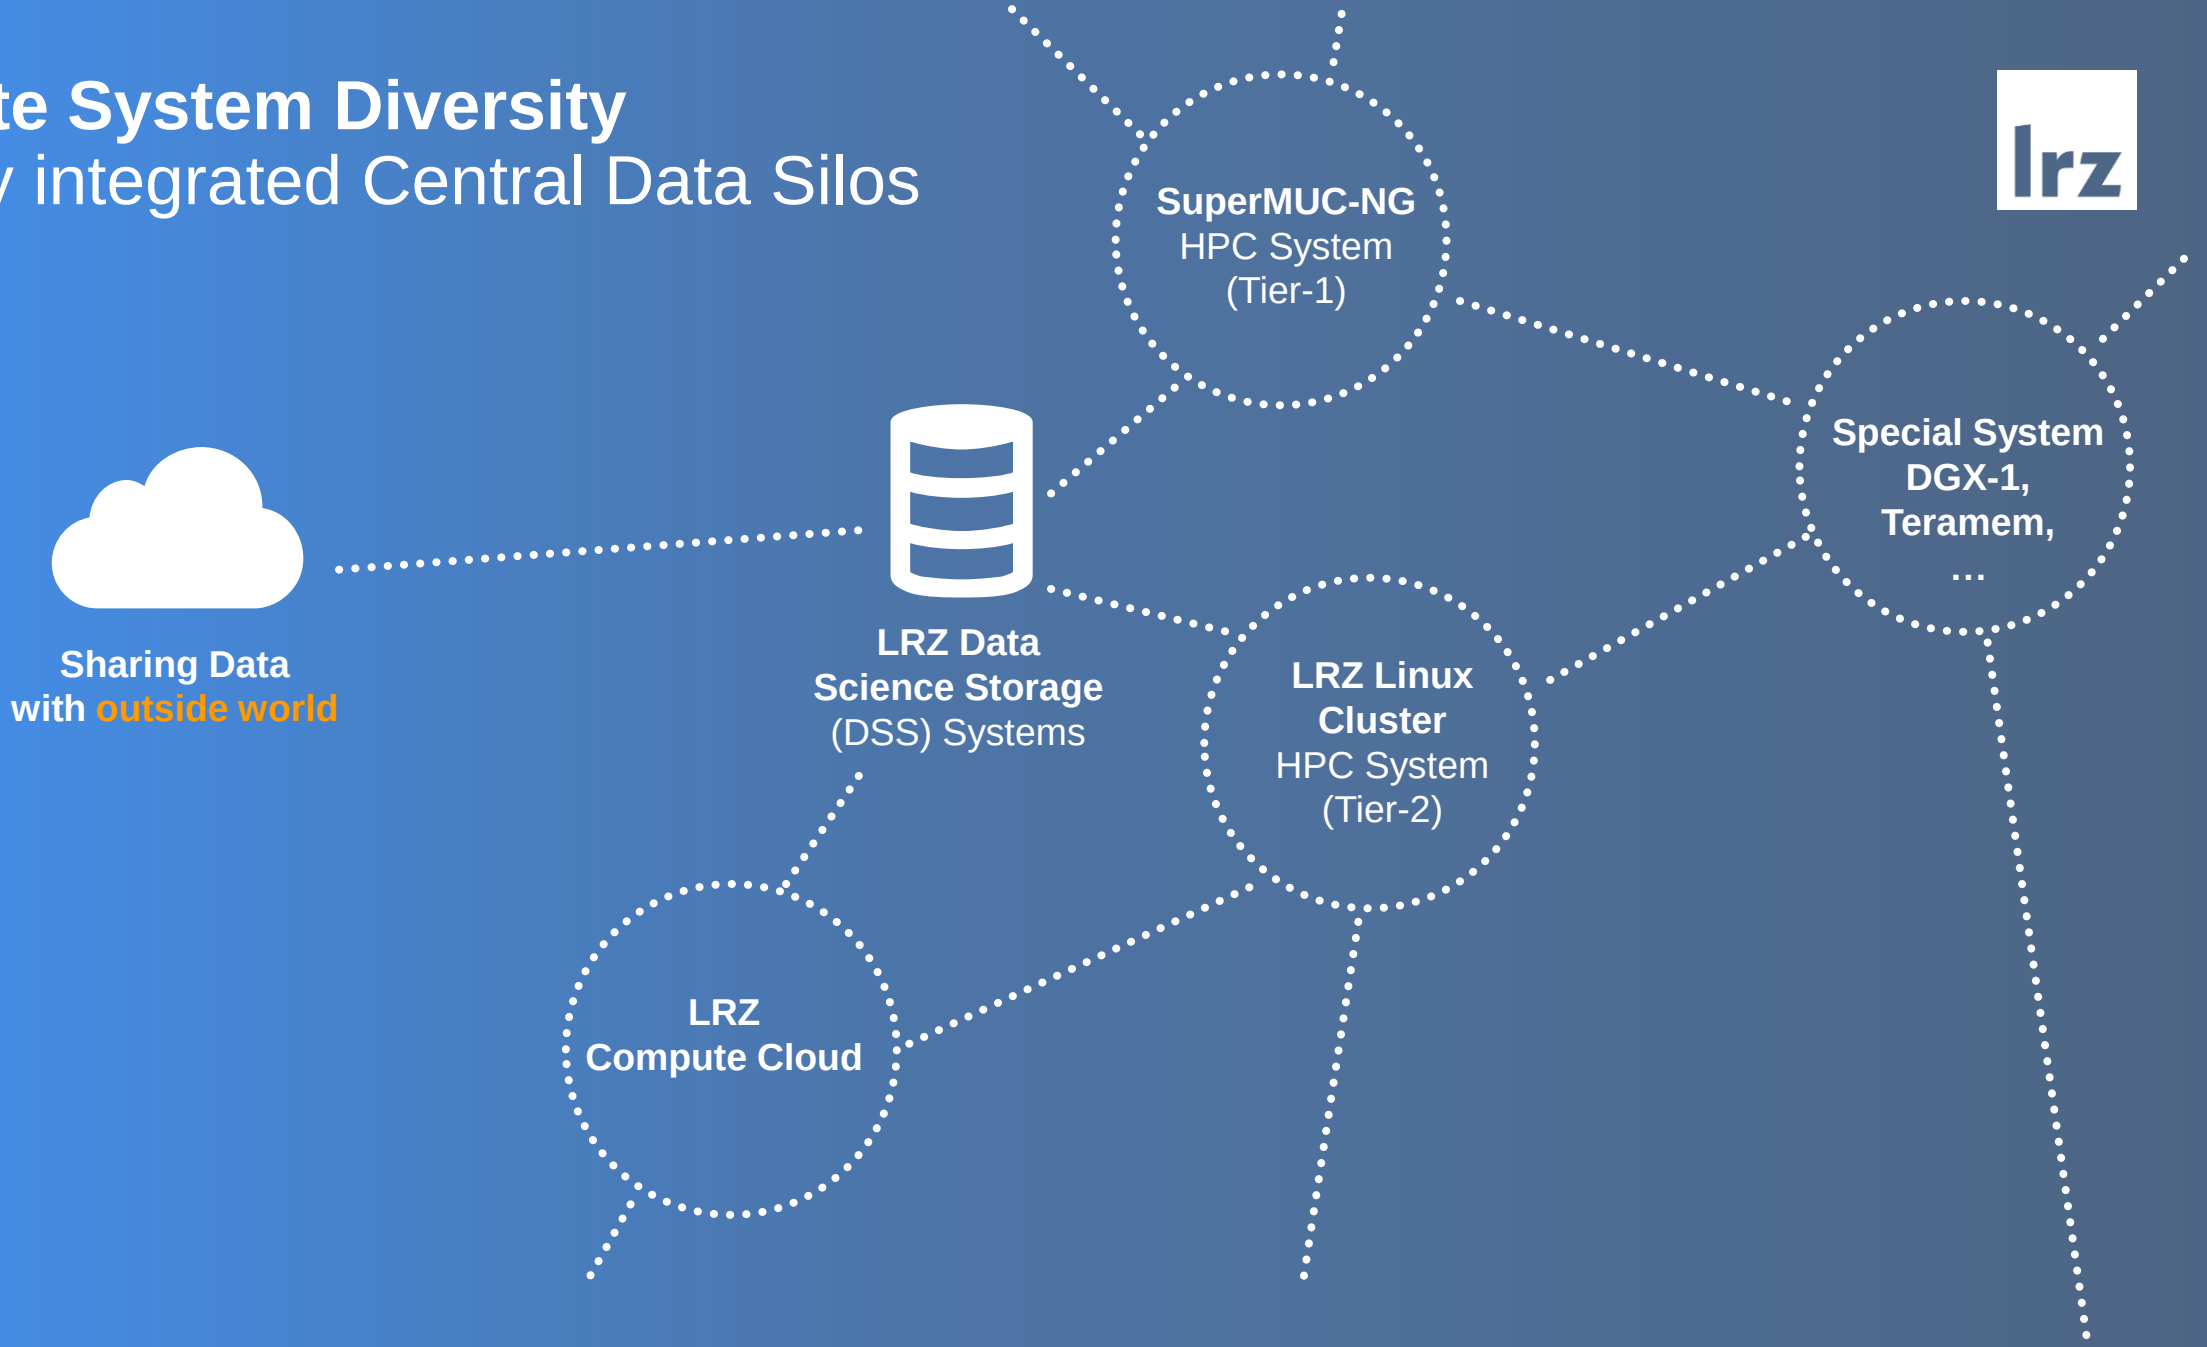

Our training system: the LRZ Compute Cloud

A new service for LRZ users

It allows to upload and use your own virtual machines

Hardware overview:

- 82 nodes: Intel® Xeon ® Gold 6148 (40 cores) @ 2.40GHz, 192 GB memory
- 32 nodes: Intel ® Xeon ® Gold 6148 (40 cores) @ 2.40GHz, 768 GB memory, each with 2x Nvidia Tesla V100 16 GB RAM
- 1 huge memory node: Intel® Xeon® Platinum 8160 (192 cores) @ 2.10GHz, 6000 GB memory