HELMHOLTZ

RESEARCH FOR
GRAND CHALLENGES

Helmholtz Computing Activities for ErUM-Data and HL-LHC

DESY: Volker Gülzow
GSI: Thorsten Kollerger
KIT: Andreas Petzold, Achim Streit

DESY. JÜLICH Forschungszentrum GSI HI JENA HELMHOLTZ Helmholtz Institute Jena HIM HELMHOLTZ Helmholtz-Institut Mainz HZB Helmholtz Zentrum Berlin HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF Helmholtz-Zentrum Geesthacht Zentrum für Material- und Küstenforschung IPP KIT Karlsruhe Institute of Technology

The DESY Tier 2 Centre:

an integral Part of the existing and next WLCG Computing Model

- **Tier 2 status:**
  - Among the largest Compute Centres within WLCG
  - Dedicated Cluster for user analyses
- **LHC Run III**
  - Increased Pile-Up
  - Existing model sufficient
- **LHC Run IV**
  - High Luminosity LHC

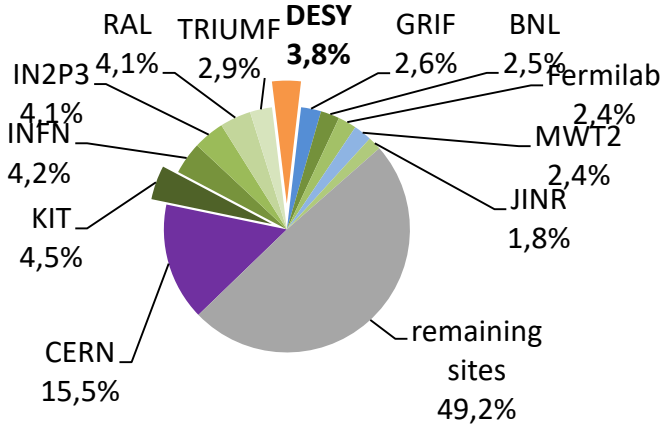

Share Worldwide LHC Computing Grid Sites (WLCG) (2019)

HL-LHC data rates forces revision of existing Computing Models

- Commitment: Tier 2 will be an integral part of the new WLCG model
- Provide resources for development of new model
- **continuing strong support for MU**

Strong commitment to High Energy Physics

Increasing support of computing for HEP

- **Strong Contributor in Storage Development**
  - Size and access to data are a major challenge for HL-LHC
  - Leading Laboratory in dCache development
  - Many Tier-1 sites use dCache
  - Important Part in the WLCG Data Lake

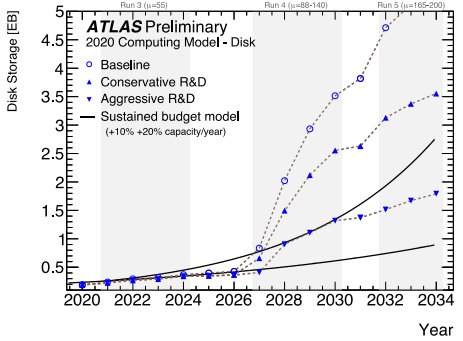

- **Major Commitment to Belle II**
  - Provide Services and Tools for the Belle II Collaboration
  - Provide Service for Belle II Membership Management
  - From 2021 providing Belle II as a Raw Data Centre one half of the German tape pledge

- **Use Synergies with other Communities on Site**
  - Share knowledge gained in HEP computing with other communities
  - Profit from shared knowledge with other communities (e.g. image based ML)
  - Share resources for development of new analysis frameworks

From Tier 2 to the Interdisciplinary Analysis Facility (IDAF)

Merging HPC and HTC Resources

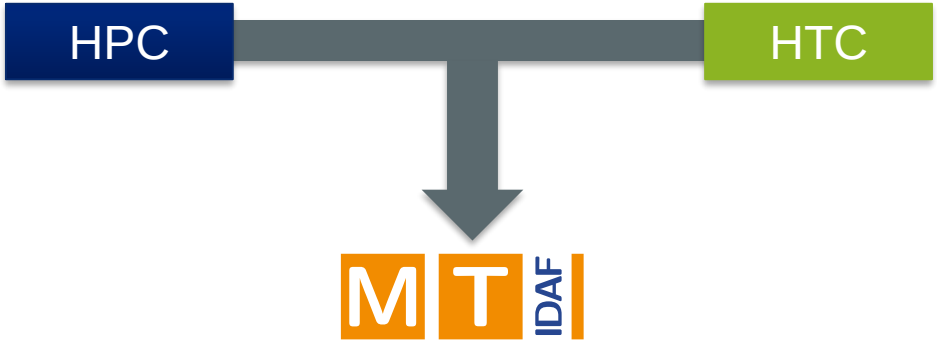

- **Support all three computing use cases:**
  - Grid large scale HEP production (HTC)
  - Individual user analysis (HTC)
  - Support simulations
  - HPC computing for Accelerator and Detector R&D, Photon Science data analysis, Machine Learning

- **CPU Resources (2019)**
  - HPC cluster: 31000 Cores
  - Tier-2 Centre (Grid): 22500 Cores
  - Tier-2 Centre (User): 11500 Cores
- **Storage (2019)**
  - Mass Storage: 25 PB
  - Cluster Storage: 2.6 PB
  - Tape Archive: 10 PB

Combining all use cases into a common IDAF makes it “win-win”

IDAF and Communities Beyond High Energy Physics

Beyond HEP – Emergence of additional Communities

- **Photon sources produce 10's of PB's per year**
  - EXFEL: AGIPD detector developed within MT-DTS
    ➡ **~1PB** of data taken within **a few days**
  - PETRA III beam-times:
    - Millions of images to be analysed in parallel
    - New detectors: **orders of magnitude increased data rates**

- **Next generation of Astro-Particle Observatories** ➡ cta cherenkov telescope array | IceCube South Pole Neutrino Observatory
  - Transfer of large amounts of data often from remote locations
  - Real Time data analyses to discover transient objects

- **Next generation accelerators**
  - Simulation ➡ increasing demand for HPC
  - Artificial Intelligence for predictive maintenance

GSI

Analysis and Data for Nuclear and Hadron Physics

Green IT-Cube

Compute Resources at GSI/FAIR

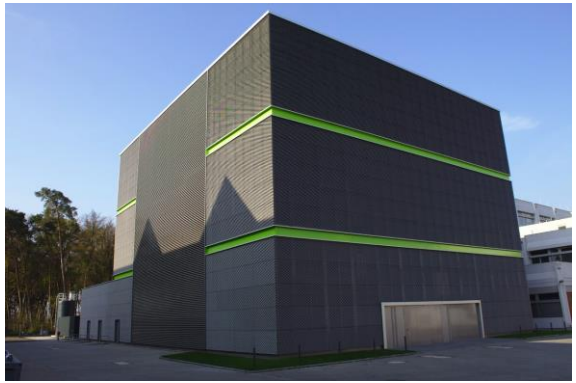

Environmentally friendly (PUE<1.07) and economic data center @ GSI

- 768 racks, 6 floors, 12 MW
- Technology commercialized

Software & Technology R&D

Enabling excellent science by easy access to cutting-edge technology

FairRoot: Framework for Analysis, Simulation and Reconstruction

- Large international user base (FAIR, ALICE/CERN, NICA, …)

Vc: performant and portable data-parallel programming

- Standardized: libstdc++

Data Management: high-performance, worldwide access to physics data, long-term availability

GSI Analysis System

Enabling excellent science by innovative systems

Online Farms (HLTs)

CBM, PANDA@FAIR

Novel concept: HLTs in common data center, fully integrated

- Cost-effective

ALICE Upgrade
O2-EPN@CERN
~1300 GPUs
several thousand cores

Data Processing, Simulation and Analysis

Common System for all use cases

- ~30.000 Cores
- ~600 GPUs (Use cases: LQCD, Accelerator, ML, Theory, HLTs)

Tier-2 for ALICE

Largest ALICE Tier-2

- Fully integrated into GSI-systems
- Optimized for high-performance data analysis

Recommendation by CERN-RRB: Use GSI Tier-2 as Analysis Facility for Run3+

GSI Data Storage Systems

Making world-class data available for science

Storing the Experiment Data at GSI/FAIR for analysis

Lustre-System
- >200 GB/s throughput
- 40 PB useable capacity
- HSM integration

Tape-Storage System
- 2 libraries at 2 locations
- 24 drives
- 10 PB used

Making the data efficiently available on a world-wide scale

Providing and developing systems to make the data available to regional, national and world-wide partners

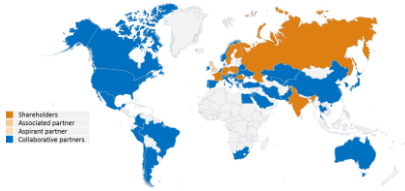

- Data access and distribution methods
- Distributed storage
- Network connections
- Efficient integration with compute resources

GridKa @ KIT

Data and analysis center for particle and astroparticle physics

- One of 13 WLCG Tier-1 centers
- 1/7 of global Tier-1 resources (= cornerstone of WLCG)
- Supporting all 4 LHC experiments, Belle-II, Auger and more (integral part of LHC data processing chain)

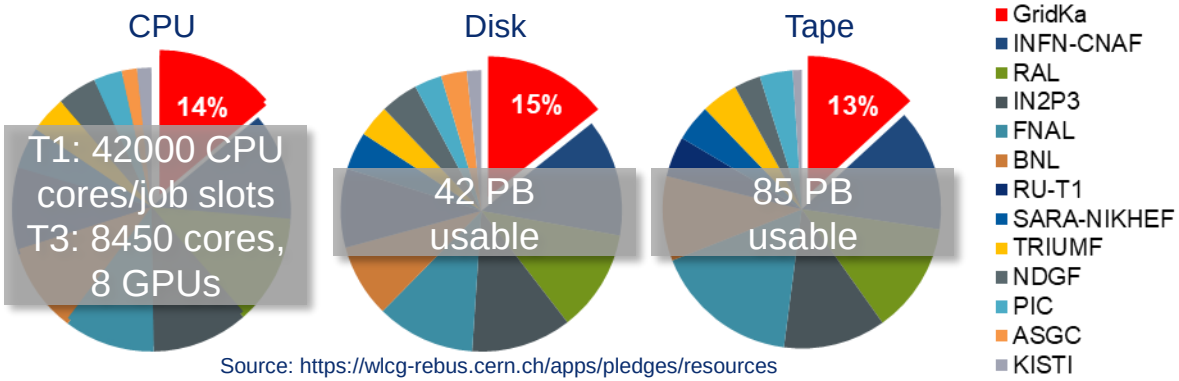

Source: https://wlcg-rebus.cern.ch/apps/pledges/resources

Reliability measured by the LHC experiments

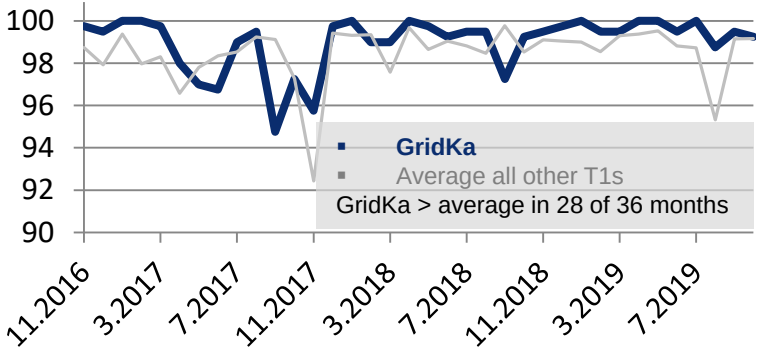

232 M core-hr
20 M jobs
57 PB in
110 PB out
0 downtime

Addressing Changing Computing Models

Leading-edge technology + in-house research and joint R&D

- **GridKa → Island in the Data Lake**
  - Requires massive scalability of storage and networks
- **Software Defined Online Storage**
  - Data access becomes less predictable
  - Increasing data access from remote compute sites
  - Dedicated sites (WLCG)
  - Opportunistically used sites (HPC, cloud)
- **Powerful Networks**
  - Redundant links to CERN (100 + 2x10 Gbit/s) and to DFN (2x100 Gbit/s)

Scalable online storage technology: throughput, IOPs, capacity

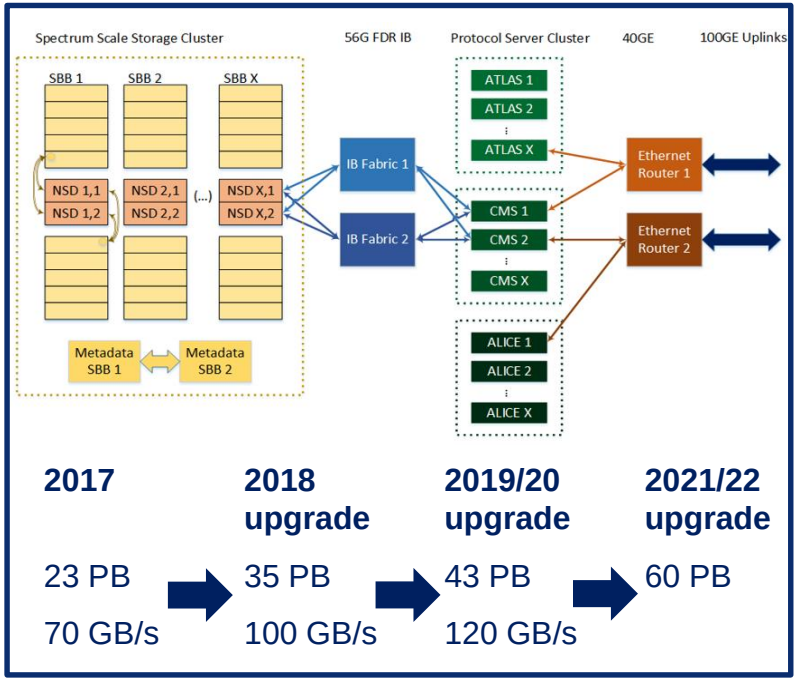

Addressing Changing Computing Models

Leading-edge technology + in-house research and joint R&D

- **Offline Storage on Tape – Today**
  - GridKa provides long-term archive
  - Focus on RAW data and simulation
- **Offline Storage – Future**
  - Cost of HL-LHC storage requirements and operations drive changes in computing models
  - Online storage as fast buffers and offline storage as permanent storage?
  - Experience gained during Run 1+2 → Performance increase for HL-LHC

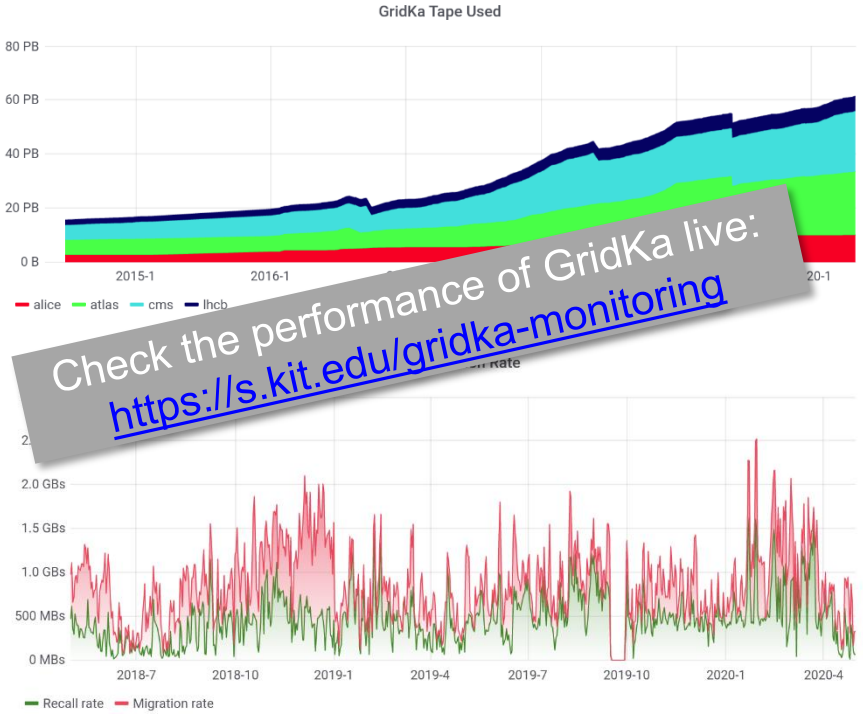

Addressing Changing Computing Models

Leading-edge technology + in-house research and joint R&D

- **Additional Opportunistic Resources**
  - Access to opportunistic (HPC, cloud, …) and heterogeneous resources (GPUs, …)
  - Hidden behind a single entry point
- **Optimized Resources and Increased Computing Efficiency**
  - Innovative ideas and improvements to speed-up analysis tasks
  - Optimized configurations of hard- and software
  - Sophisticated data and workload management

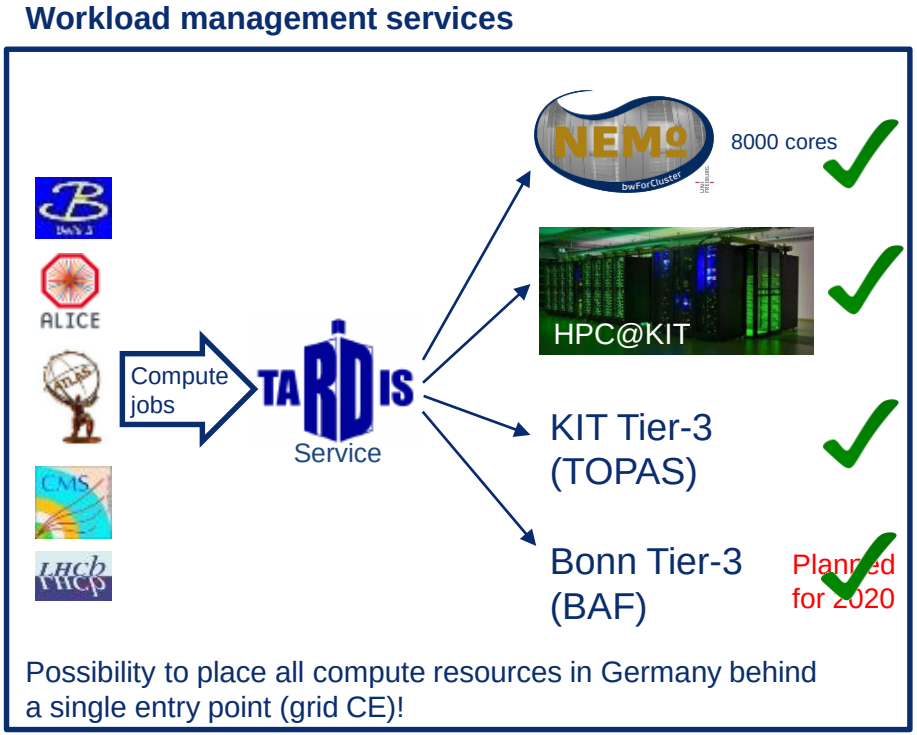

SUMMARY

- Excellent support for particle physics with leading-edge IT technology
- Scaling-up of computing services, resources and data center infrastructure according to requirements
- Interdisciplinary R&D in data management & analysis between experiments, Tier centers, MT DMA and RF Information
- Integration with federated infrastructures on national and international level (NFDI, EOSC, …)
- User transparent integration of HPC, cloud and other resources
- Indispensable research infrastructure for more than 10,000 physicists worldwide

- **Less Helmholtz Computing in Germany → Less science with HL-LHC and Belle-II**