Beyond High Energy Physics

Photon and accelerator science computing infrastructure at DESY

Speaker: Thomas Hartmann on behalf of Christian Voss

C. Beyer, S. Bujack, S. Dietrich, T. Finnern, M. Flemming, P. Fuhrmann, M. Gasthuber, A. Gellrich, V. Gülzow, T. Hartmann, Y. Kemp, B. Lewendel, J. Reppin, F. Schlünzen, M. Schuh, S. Sternberger, C. Voss, M. Wengert, S. Yakubov

Adelaide, 2019.Nov.05

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

DESY.

Interdisciplinary Laboratory

Research Communities supported by the Hamburg Site

- Storage and Analysis of Accelerator Data

FLASH. Free-Electron Laser FLASH European XFEL FF

- Data taken at Accelerator Development facility

SINBAD

- HPC simulations for future accelerators
- Storage and Analysis of test-beam data

Detector and Accelerator R&D

- Online and Offline Computing including user analysis

European XFEL PETRA III FLASH. Free-Electron Laser FLASH

- HPC Resources for dedicated interdisciplinary research centres on the DESY campus

CFEL SCIENCE CSSB Centre for Structural Systems Biology

EMBL Helmholtz-Zentrum Geesthacht Centre for Materials and Coastal Research

Research with Photons

- WLCG Tier II Centre for

CMS ATLAS EXPERIMENT LHCb

- Regional Data Centre and Hosting of Collaborative Services and Tools

Belle II

- Tier 0 Centre and VO Management

ilc international linear collider

High Energy Physics

Interdisciplinary Data Analysis Facility

Current Service Infrastructure

High Performance Computing

- Per node scheduling
- Powerful dedicated machines
- Fast interwoven network
- GPU resources

High Throughput Computing

- Batch submission
- Commodity machines
- Standard peer-to-peer network

Storage

Cluster/HPC

General Purpose

Tape Archival

Interdisciplinary Data Analysis Facility

Current Service Infrastructure

High Performance Computing

- Per node scheduling
- Powerful dedicated machines
- Fast interwoven network
- GPU resources

High Throughput Computing

- Batch submission
- Commodity machines
- Standard peer-to-peer network

Storage

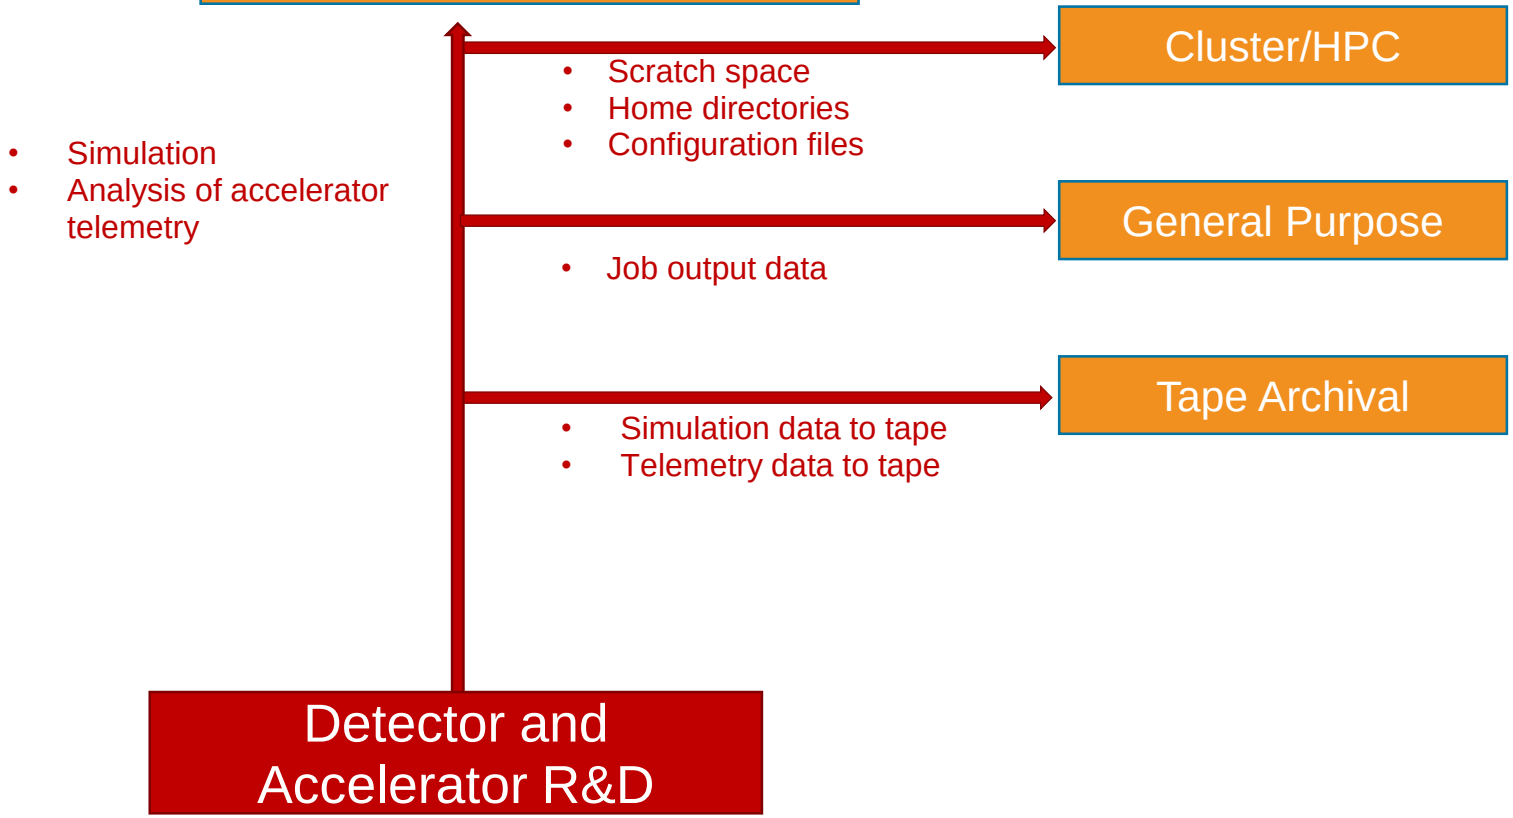

Interdisciplinary Data Analysis Facility

Current Service Infrastructure

High Performance Computing

- Per node scheduling
- Powerful dedicated machines
- Fast interwoven network
- GPU resources

High Throughput Computing

- Batch submission
- Commodity machines
- Standard peer-to-peer network

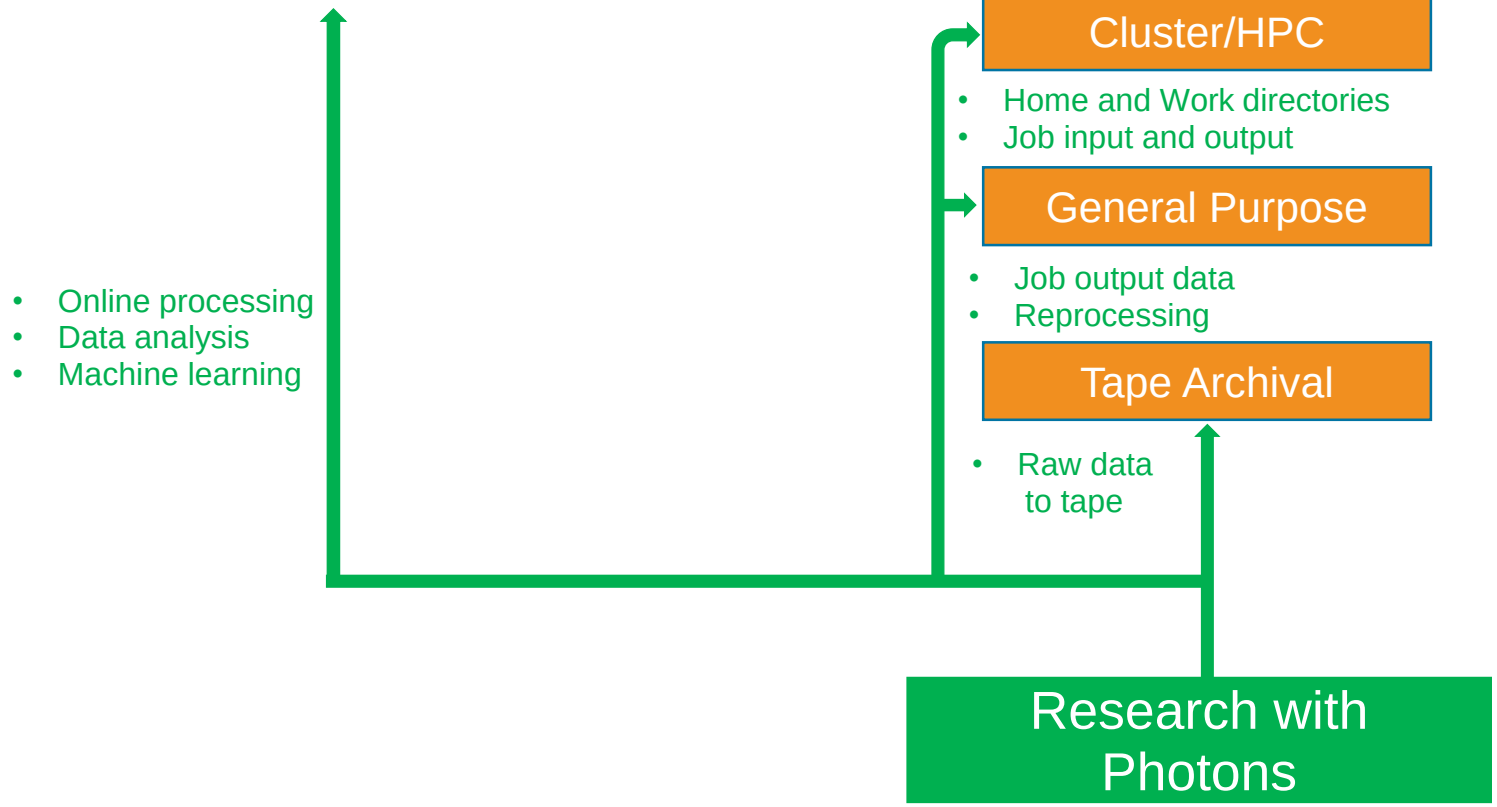

Interdisciplinary Data Analysis Facility

Current Service Infrastructure

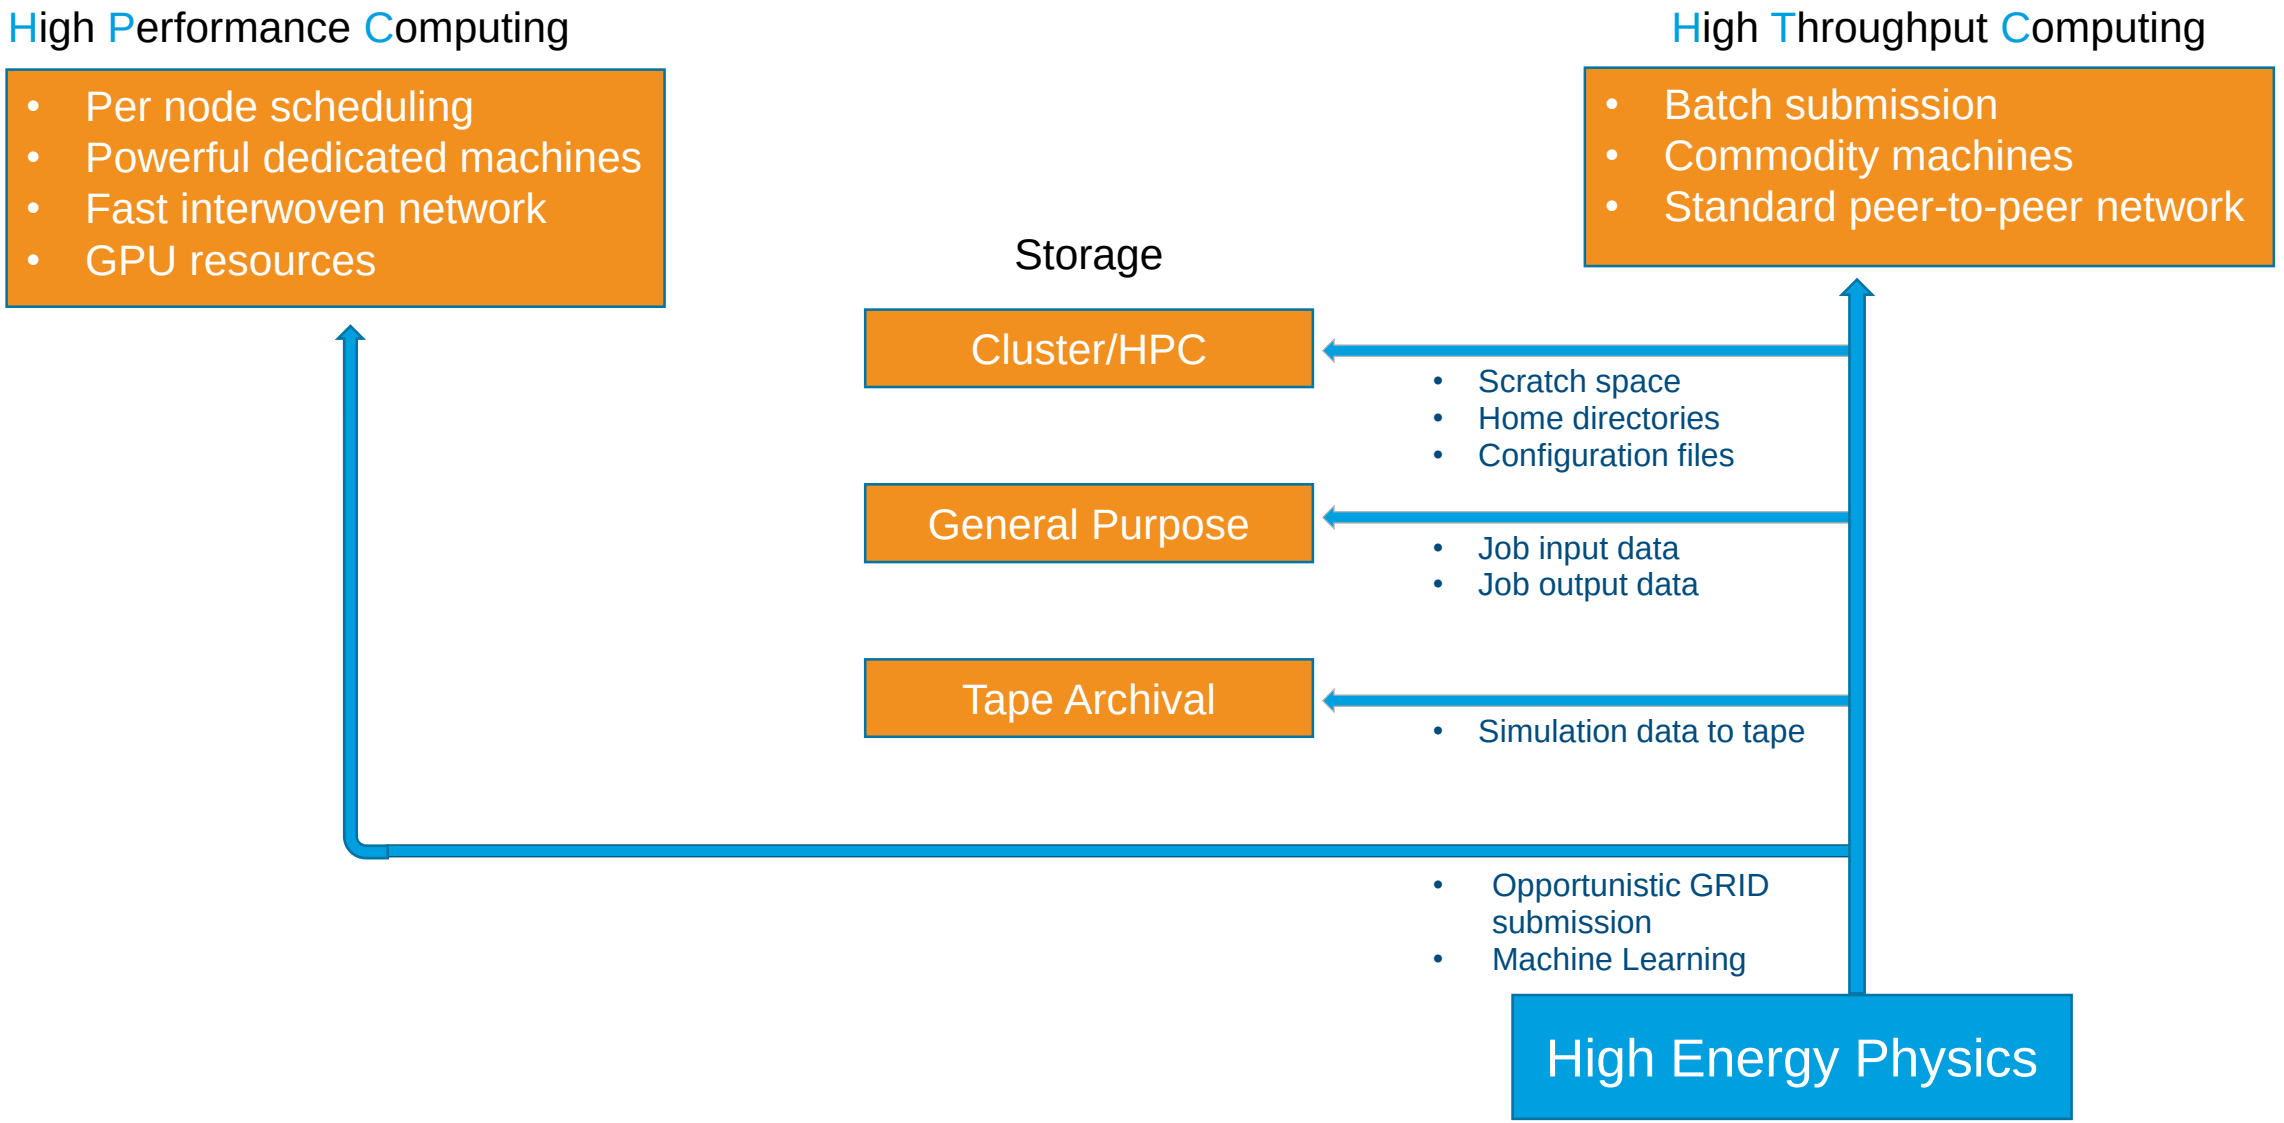

Interdisciplinary Data Analysis Facility

Overview

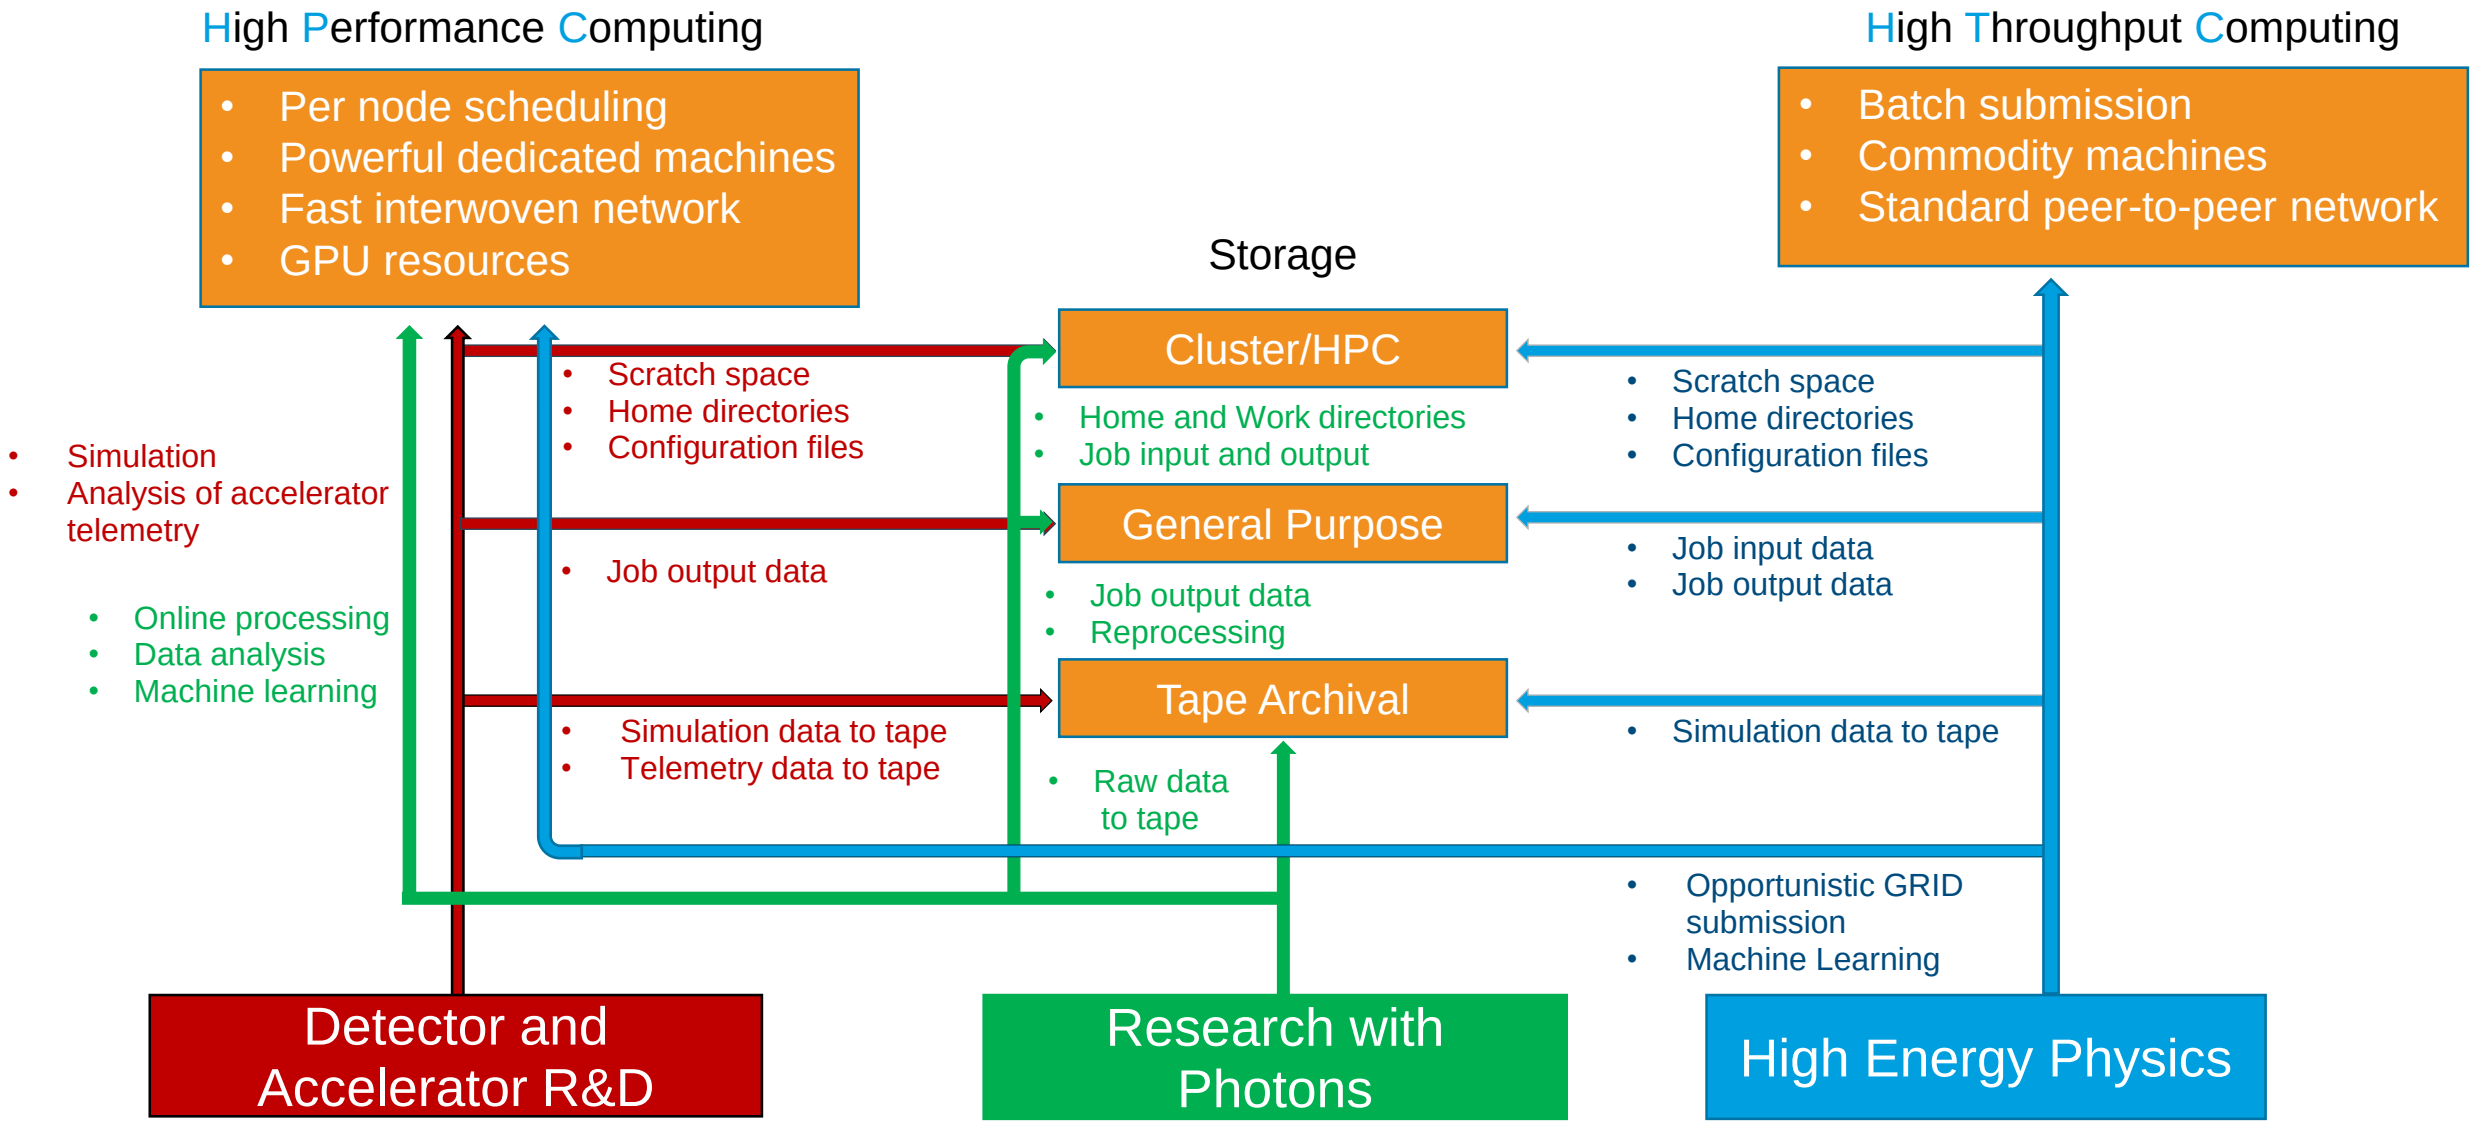

Different Needs of the DESY Research Departments

Different Computing patterns

- Apparent separation of user communities with respect to Computing
  - High Energy Physics community runs classic batch workflows on the GRID farms
  - Photon Science and Accelerator/Detector R&D uses predominately HPC resources

Why is that?

High Energy Physics

- Workflows opposite to classic HPC applications
- Embarrassingly parallelisable workflows

Accelerator/Detector R&D

- Simulations especially Plasma Wakefield calculations are classic HPC applications

Photon Science

- Applications needed inter-process communications due to exp. environments
- Data structure requires larger amounts of memory than HEP applications
- Higher demand for interactive analysis especially within online computing

Changes due to Evolving Experimental Environments

Changes to Computing Requirements

- Rise of new methods lead to adaptions of new concepts

What are these?

High Energy Physics

- Rise of Machine Learning and industry tools for data analysis
- Demand for GPU resources
- Decline of ROOT usage among HEP community

Accelerator/Detector R&D

- Demand on HPC will increase and local resources will follow this trend
- Expect no large changes to computing model

Photon Science

- Image processing in parallel suited for batch processing
- Workflows often turn out to be more HTC friendly

Adapting to Changing Requirements

Improving Direct Computing Support

High Energy Physics

- Long time integration into local HEP groups
- Consulting on storage/CPU and future investments in infrastructure/technology

- Can we achieve something similar for Photon Science?

Photon Science

- Every Customer at each light source needs to submit a proposal
- Reviewed and supervised by Beamline Scientist
- Add computing requirements and plans to proposal
- Consulting from Beamline Computer Scientists

↳ Recommendation:

Resources in High Performance Computing

Resources in High Throughput Computing

Adapting to Changing Requirements

User Experience

- Simply opening HTC and HPC resources to all users will not solve the puzzle
- Use different Batch systems on each cluster

- HPC Cluster
- Users: Photon/Accelerator Scientists

- HTC Cluster
- Users: HEP Scientists

- Difficult to understand and use one system well
- Do you want to enforce learning a second system?

Question of Access:

- Access to HTC System connected to membership of supported VOs (WLCG, Belle II or ILC/Calice)
- Access to HPC System on buy in basis:
  - Guaranteed access to your resources
  - Opportunistic access to remaining resources

Introduction of Cloud/Container Based Solutions

Improving User Experience

- Premise for our local Facilities for the Photon Science Community:

User Groups change, detectors persist

- Keep infrastructure for On/Offline analysis identical
- Only change the payload
- Users bring their payload

Prepared Cluster on Cloud/HPC/HTC

User Pay Load

Using of Cloud/Container Based Solutions

Advantaged for Free

User Pay Load

- Need infrastructure to build and ideally test these

- Industry and other labs have well tested solution: GitLab
- Naturally deployed on compute clouds
- Offers support for Continuous Integration and Continuous Deployment from a Browser-GUI

Proposed Workflow

- Scientists start with existing containers provided by us
- Scientists develop analysis using the most commonly used version control system
- Can automatically build and test there analysis
- Push their Analysis into a container repository
- Run their Analysis

Partitioning the Data Centre Using Cloud Technologies

Taking the Logical Next Step

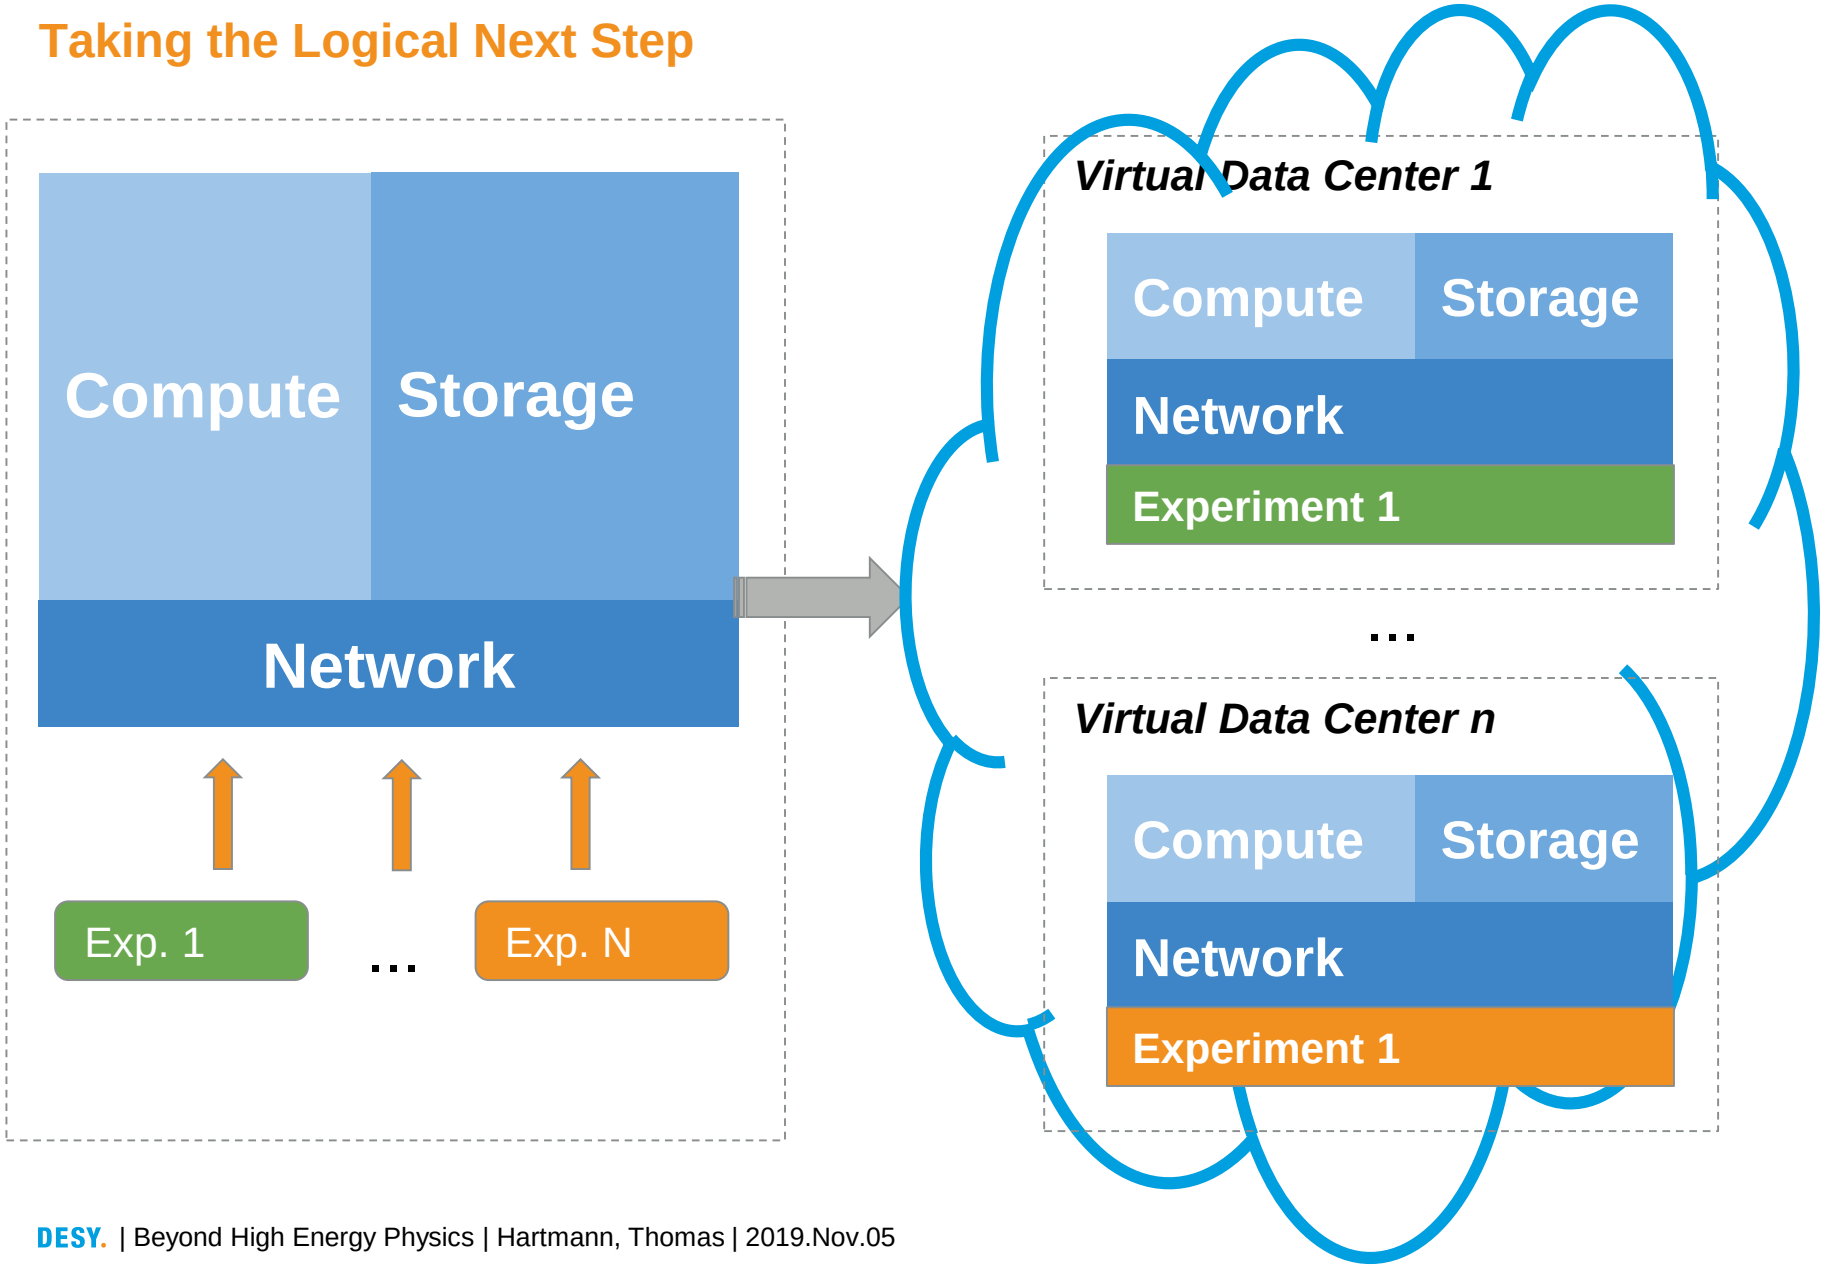

- Dynamic Provisioning

kubernetes

- Software Defined Networking
- Storage

How to Handle Data Access within this Paradigm

The Big Elephant in the Room

- Access and analysis of scientific data is most important service to provide
- Static namespace structure depending on the compute cluster
- Different storage systems depending on the compute cluster

General Purpose

- Access via NFS
- NFS layer ACLs
- Consistent Namespace

Cluster/HPC

- Access via NFS (HTC)
- Native Mount (HPC)
- NFS layer ACLs
- Different Systems for different clusters

- How to translate the ACLs into the Container?
- How to make all storage systems universally accessible?
- Is a global namespace what we need or even want?

Thank you

Contact

DESY. Deutsches
Elektronen-Synchrotron

www.desy.de

Name Surname
Department
E-mail
Phone