


# LSDF – Large Scale Data Facility at KIT

Marcus Hardt, J.v.Wezel, R.Stotzka, T.Jejkal, V.Hartmann, M.Sutter, A.Garcia, A.Hammad | December 2, 2010

STEINBUCH CENTRE FOR COMPUTING (SCC)



- **Build Germany's largest Data Storage** to support data intensive science
- Attract new demanding users to exploit synergies
  - Understand, develop and deploy of community specific services
  - Collect and analyse use-cases
  - Honour community specific techniques and systems
  - Integrate and develop generic methods and tools
- World wide secure access to data
  - International collaborations
  - Ensure transparent access
  - Integration with established access / authentication methods
- Added value and tools for data processing
  - Integrate well connected clusters
  - Meta-data management and tools
- Long term data storage
  - Store the data itself (legal reasons)
  - Provide Methods to access/understand the data
- Funded within the national HGF programme "SuperComputing"



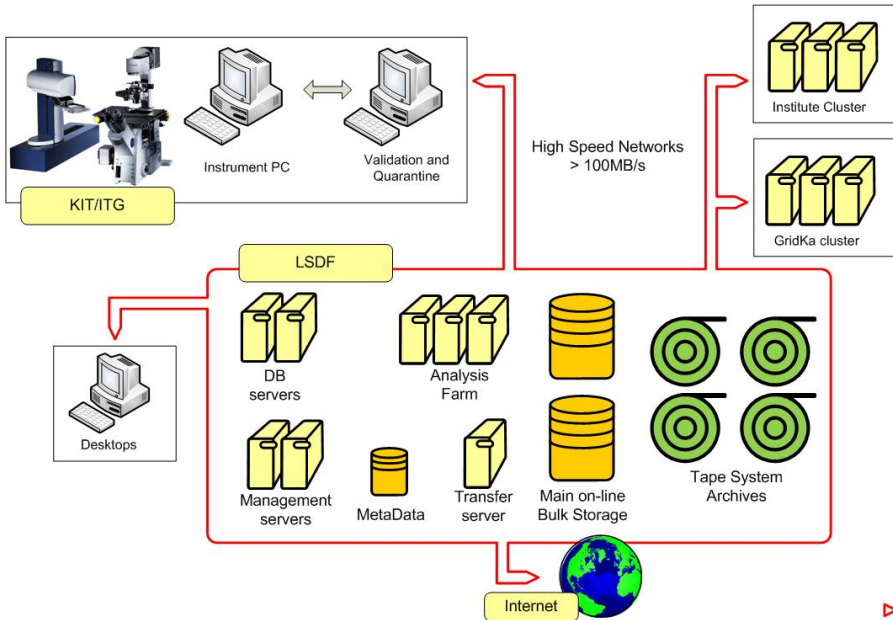
- 2 high grade disk systems
  - 550 TB [Data Direct Networks]
  - 1,2 PB - 3 PB (2011) [IBM]
  - Tape backend for archive and backup
  - Fibre Channel attached
- Dedicated storage servers
  - GPFS on top of each storage system
  - exported as GPFS, NFS, CIFS, webdav
- Directly attached to data sources / processing clusters
  - 10 Gb/s dedicated redundant backbone
  - 10 Gb/s dedicated links to prime partners, 1Gb/s others
- Well connected inside Germany:
  - 10GE to all data-centres in Germany
  - 40GE to HLRS-Stuttgart
  - 100GE connectivity to FZJ planned and tested

- Experiments should level-0 process data locally
  - i.e. where the data is created
  - 15 days to transfer 1 PB over ideal 10Gb/s link
- Small processing cluster:
  - 58 nodes with 8 cores, 36GB memory
  - Directly attached to storage (GPFS)
  - Hadoop environment
    - 110 TB HDFS, Hadoop native filesystem
    - Good scalability on commodity hardware
    - However, use-case must match map-reduce paradigm
  - Cloud environment OpenNebula
    - Users can deploy own dedicated data-processing VMs
    - Reliable, highly flexible, and very fast to deploy



OpenNebula.org



# LSDF birds-eye view



- Implementing / integrating experiments
  - e.g.: High throughput microscopy workflow
    - Images are stored to hadoop cluster
    - Preprocessing right there
    - Storing output on LSDF
- Metadata:
  - Data that cannot be found (in a few seconds) is nonexistent data
    - accessibility increases the data value
    - simple access (input and retrieval) increases acceptance by communities
  - Higher level access layer vs. indexed searching
- Dissemination to new – potentially interested – communities
- Add support for IPv6

# Some experiments at KIT/LSDF

- USCT
  - for breast-cancer diagnosis
  - 0.3 PB per year
- Immunogenetics Institute Charité Berlin
  - Computer tomography of sea urchins
  - "several hundred" TB per year
- ANKA
  - Synchrotron Beam Source
  - 240 PB per year (2013: 1 PB per year)
- Robotic high throughput Microscopy
  - Biology / Genetics
  - 1 PB per year
- Large Hadron Collider (LHC)
  - Currently uses 5 PB at dedicated facility
  - 1 PB per year
- Yearly growth, **currently** summing up to
  - 150 TB: 2010
  - 1 PB: 2011
  - 2.5 PB/a: 2012 - 2014
  - 5 PB/a: 2015 +



# To conclude

- First hardware up and running
- First software services available
- First data stored
- First experimental data processed
- Focus on user requirements
  - Added value services on top of large storage
- Many scientific communities interested and getting involved
- Future activities
  - Grow beyond KIT, HGF and build international collaborations
  - Involve new experiments
  - Approach new communities
  - Explore new techniques, integrate/develop new services

Scaling to terabit networks and exabyte storage must start today



Thank you.

- Provisioning of storage and archives in exabyte scale
- Development of software technologies for distributed data management and archiving
- Development of efficient transport protocols from the experimental facilities, e.g. robotic microscopes, to the LSDF
- Development of technologies to handle the special requirements of experiment data (e.g. 3D image stacks) of various research communities (e.g. systems biology)
- Development of open standards and implementation across computing centre borders
- Provisioning of compute resources for data analysis
- Development and integration of data analysis services
- Specific support for users with data intensive applications
- Development of data and meta data models for specific user groups
- Optimized data organization for specific user groups

- Hardware
  - Storage capacity (planned deltas)
    - 10Q4 2 PB
    - 11Q2 4 PP (SONAS on 2 PB IBM storage)
    - 12Q2 6 PB
  - 10 Q4 Dedicated tape storage
  - 11 Q2 Improved network connectivity and services
    - Dedicated 10 Gb/s backbone for 3 more institutes
    - Initial support for IPv6 in Q2 2011
  - 11Q3 40 Gb/s Link to BioQuant/Heidelberg
- Services
  - 10Q4 Enable direct access storage for first experiments, Q3 2010
    - stop gap measure until tools are available
  - 11Q1 iRods software operational
  - 11Q2 Additional communities integrated
    - ANKA (synchrotron radiation ring)
    - IMK (meteorology and climate research)
    - Geophysical data

# LSDF Network

