


GridKa Tier-1 Status and Future

Andreas Petzold




STEINBUCH CENTRE FOR COMPUTING - SCC



Today

- we have come a long way since 2001



Helmholtz Data Federation

July 1, 2001

Requirements for a Regional Data and Computing Centre in Germany (RDCCG)

Peter Maltzacher, GSI, ALICE
Andres Sandoval, GSI, ALICE
Lutz Köpke, Universität Mainz, ATLAS
Alois Putzer, Universität Heidelberg, ATLAS
Dankfried Lamske, RWTH Aachen, CMS
Günter Quast, Universität Karlsruhe, CMS
Volker Lindenstruth, Universität Heidelberg, LHCb
Michael Schmelling, MPI Heidelberg, LHCb

Marcel Kunze, Universität Bochum, BABAR
Klaus R. Schubert, TU Dresden, BABAR
Thomas Müller, Universität Karlsruhe, CDF
Lars Schmitt, TU München, COMPASS
Peter Mättig, Universität Wuppertal, D0

Peter Braun-Munzinger, GSI, for KHK
Rolf Heuer, Universität Hamburg, for KET
Rainer Mankel, DESY, for DESY

Abstract


A Regional Data and Computing Centre in Germany (RDCCG) is proposed by the German Particle and Nuclear Physics Community. The steps towards the production phase are described, including hardware requirements, expected manpower support from the experiments, deliverables, and milestones. A proposal is made for the management structure of this centre.



First 70TB for
GridKa

Users

- GridKa: particle and astro-particle physics



- LSDF: systems biology, climatology, synchrotron radiation ...

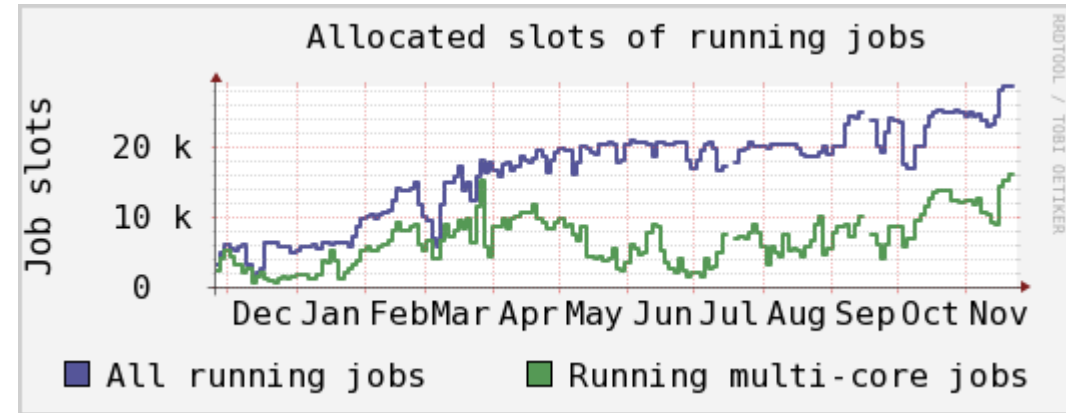


GridKa Team

- ~20 people for hardware, services, batch system, online storage, offline storage, network, experiment contacts
- many are active in GridKa and other projects
- evolution towards integrated service teams
 - online storage team responsible for online storage in GridKa/LSDF/SDIL/...
 - offline storage team responsible for tape storage for GridKa/LSDF/bwDataArchive/...
- activities in HEPiX/WLCG/EU projects
 - CPU benchmarking
 - AAI activities in AARC, Indigo
 - HNSciCloud

GridKa Resources – CPU

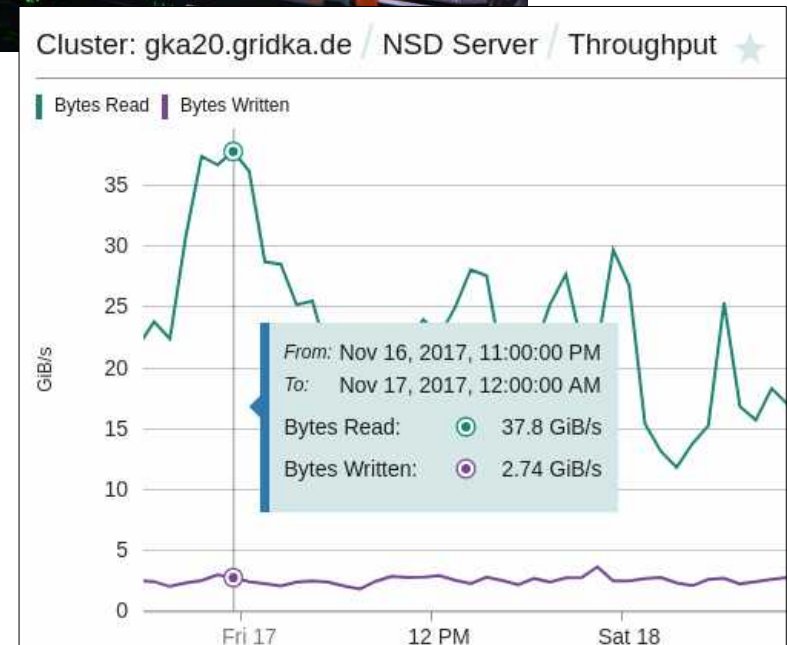
- HTCondor
- 850 WNs, ~24k job slots, 320kHS06
- today all WNs still on premise
- KIT CMS group runs jobs at Freiburg
 - >10k slots during testing last year
- HNSciCloud
 - KIT is testing dynamic cloud extension on 3 clouds
 - next phase starting in 2018 will provide production quality resources on 2 clouds



GridKa Resources – Disk




- new storage installation started in 2016
 - 13.5PB old storage systems are being retired
- GPFS – software defined storage layer
 - used for almost everything @GridKa&LSDF
- 21PB used in file systems @GridKa
- 44PB installed since last week for GridKa+LSDF
- dCache/xrootd/NFS/S3 on GPFS
 - future dCache with shared file system support for pools?
 - CMS T3 will serve as testing ground for new possibilities: SRM-less, GPFS on WNs...
- Exploring cloud storage integration in HNSciCloud



GridKa Resources – Tape

- Switch to enterprise technology great success
 - much improved performance
 - continuing investments in enterprise drives and media
- 35PB LHC data on tape
- TSM current capacity 85PB (T10K-D technology)
 - drives: 17 T10K-C, 12 T10K-D, more to be purchased
- HPSS migration planned (GridKa&LSDF&others)
 - future proof system for capacity and performance
 - big effort to migrate data and dCache/xrootd interface
 - 3 libraries at KIT Campus North (255PB T10K-D)
- BelleII will start using GridKa tape from 2020



Network

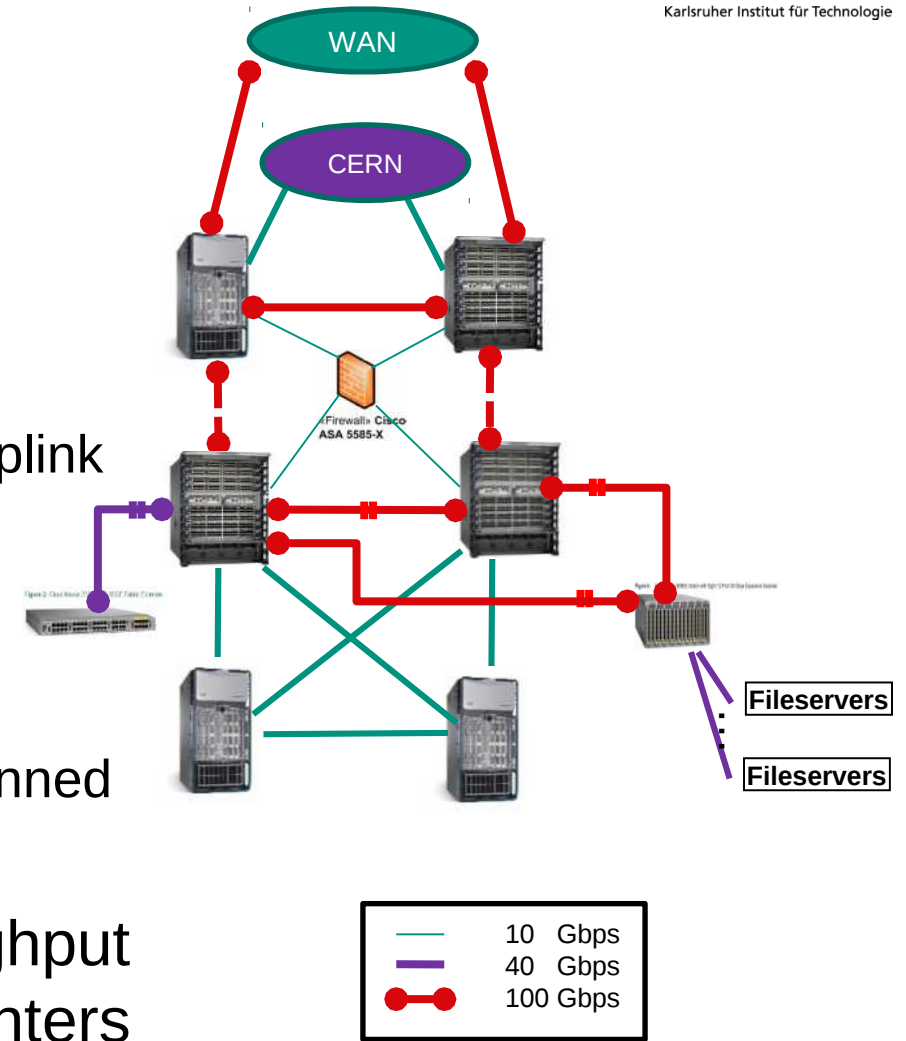
LAN

- 4 fabric routers, 80-200Gb/s interconnect
- file servers: 40Gb/s to TOR switches
- TOR switches: 400Gb/s to fabric routers
- WNs: 10Gb/s to fabric extenders, 20/40/80Gb/s rack uplink

WAN

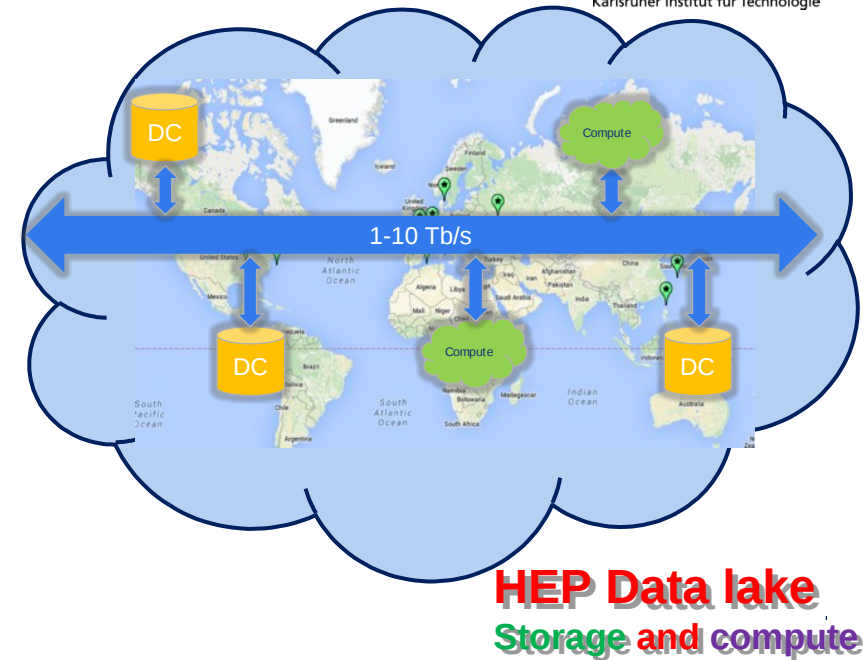
- 2 border routers
- dedicated 20Gb/s to CERN, next year 100+20Gbit/s
- 100Gb/s to X-WIN/LHCOPN/LHCONE, 2x100Gb/s planned
- connections to cloud providers via DFN

- GridKa network is prepared for increasing throughput required by remote data access from existing centers and clouds



Towards Run-3 and HL-LHC

- expect increasing WAN bandwidth
 - 100Gb/s, planned 200Gb/s for LHCONE/OPN/X-WIN
- expect increasing peak loads on storage and greatly increasing capacity requirements
 - large investments in tape as cheap high latency storage
 - GPFS ideal for exploring tiered storage solutions
 - transparent cloud integration
- expect increasing peak requirements for CPU resources
 - cloud bursting capability with HTCondor&ROCED
- sustainable operations
 - KIT power&cooling plant provides inexpensive cooling for GridKa
- KIT future strategy: focus on energy, mobility, information
 - GridKa is treated as information infrastructure and has full support of KIT president and VPs





We Are Hiring!

Contact

andreas.heiss@kit.edu

andreas.petzold@kit.edu