GridKa LK-II

Status Report

Andreas Petzold, Achim Streit

GridKa

Data and analysis center for particle and astroparticle physics

ALICE · ATLAS EXPERIMENT · CMS · LHCb · Belle II · IceCube Neutrino Observatory

WLCG Worldwide LHC Computing Grid

~10,000 physicists worldwide
~ 1,300 from Germany

Global Effort → Global Success

Results today only possible due to extraordinary performance of accelerators – experiments – Grid computing

Observation of a new particle consistent with a Higgs Boson (but which one…?)

Historic Milestone but only the beginning

Global Implications for the future

R-D Heuer

- Tier-1 center: cornerstone of the Worldwide LHC computing Grid
- Integral part of the LHC data processing chain

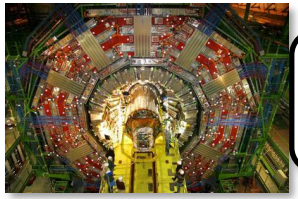

Tier-0

CERN

RAW Data

GridKa Tier-1

Data archive

- Event reconstruction
- Event selection
- Simulation
- ...

Physics datasets

Compute – Disk – Tape – Network

Tier-2s

M T IDAF

Analysis jobs

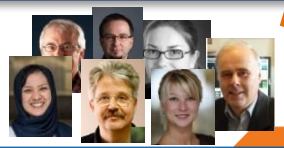

Compute – Disk – Tape – Network

Batch System

- Smooth operations of ~50 k logical cores
  - 31.3 M jobs & 325 M core-hours in last 12 months
- Several rolling reboots for security patches
- 56 GPUs (NVIDIA V100) available
- Additional (opportunistic) resources managed through COBalD/TARDIS @ GridKa

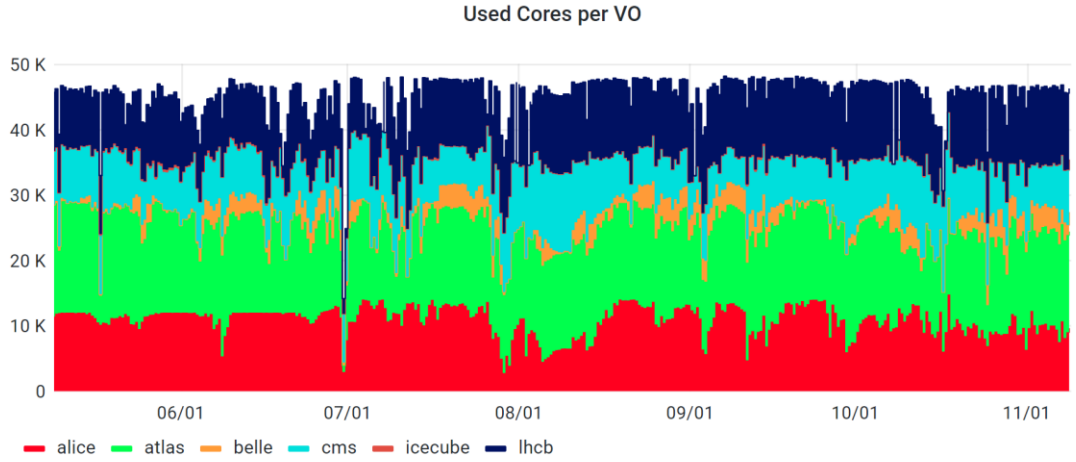

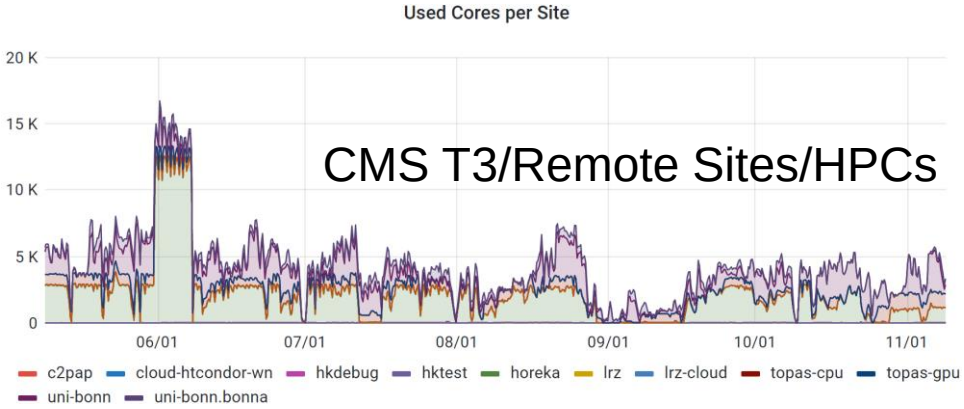

Compute – Disk – Tape – Network

New Worker Nodes

- 196 new nodes installed in 1st week of November
  - Dual AMD EPYC 7742
  - 570 GB RAM, 7 TB NVMe
- 50k logical cores
- 550+ kHS06
- Burn in running
- 900 WNs to be retired in 2022

10:00 → 12:30

Online Storage

- Smooth operations of 46.4 PB scalable online disk storage
  - 90+ GB/s bandwidth to CPUs
  - 80 PB data written, > 500 PB read in last 12 months
- Smooth operations of dCache and xrootd
  - Upgrade to dCache 6.2
  - Intensive preparation for tape challenge

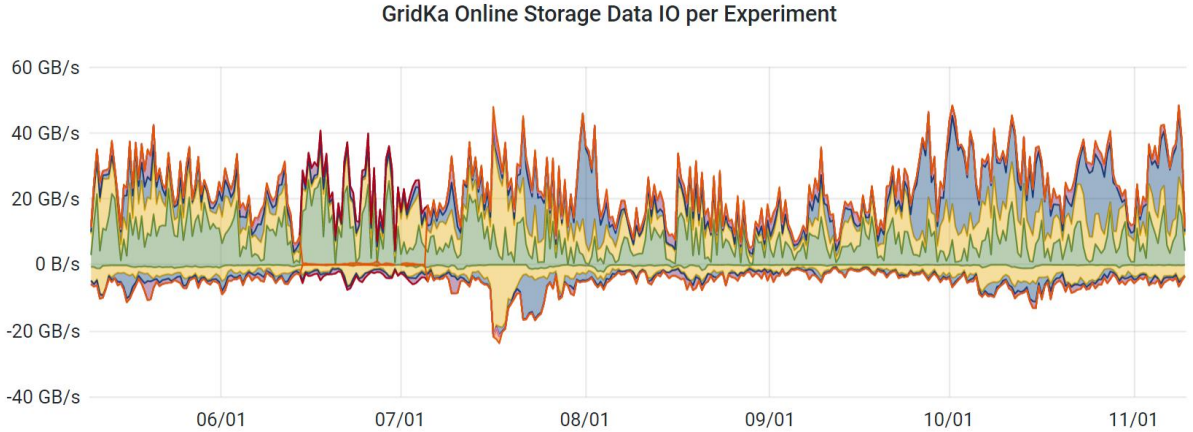

Online Storage

- Call for tender successfully finished
  - 71 PB usable storage
  - 44 Seagate Enclosures
    (dual controllers & 106 18 TB disks)
- 22 storage & 48 protocol servers
- Final setup and deliveries still under discussion, but looking ok
- Replacement Spectrum Scale Metadata servers to be delivered in 2021
- 10 PB and 20 protocol servers retired, data migrated

Reliable Offline Storage

- Smooth operations (67.7 PB occupied)
- Migration to HPSS
  - 23 PB finished for CMS and Belle-II
  - Preparation for LHCb (5.7 PB), ALICE (10 PB) and ATLAS (27 PB)
- Preparation for Tape Challenge
  - dCache pools increased, switched to latest storage agent HW
  - Extra tape classes according to experiment requirements

Spectra TFinity Tape Library KIT Campus South

Tape Challenge for Run 3

- Targets reached and surpassed for all VOs
  - LHCb delayed by broken 40 G transceiver
  - ALICE should not have been part of the challenge @KIT

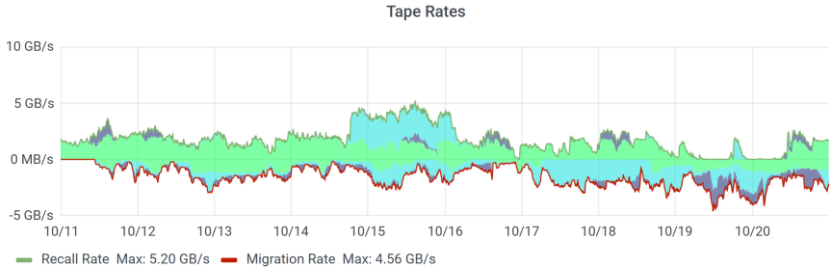

- Able to utilize all drives
  - But: dCache config prevented optimal drive utilization (migration to tape)
  - → understood and fixed
- Tape buffers will need to grow!

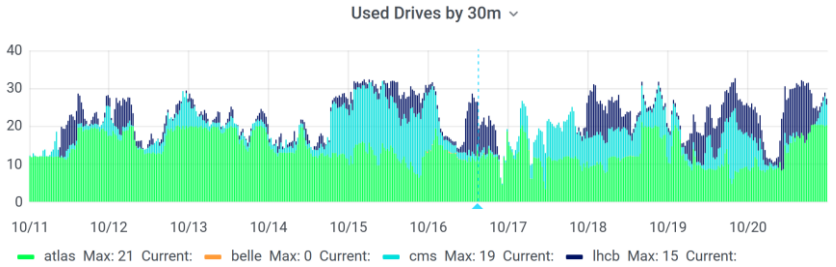

Compute – Disk – Tape – **Network**

Status & Challenge

- Smooth operations of 400 Gb/s WAN connection
  - 2*100 Gb/s to CERN/LHCOPN
  - 2*100 Gb/s to Internet/X.WIN/LHCONE
  - no changes in 2022/2023
- LAN upgrade to 400 Gb/s in 2022/2023
  - New backbone routers with X*400Gb/s interconnect
  - upgrade storage switches to 2*400 Gb/s uplinks

- Network challenge
  - T0→T1 highest peak rate for KIT

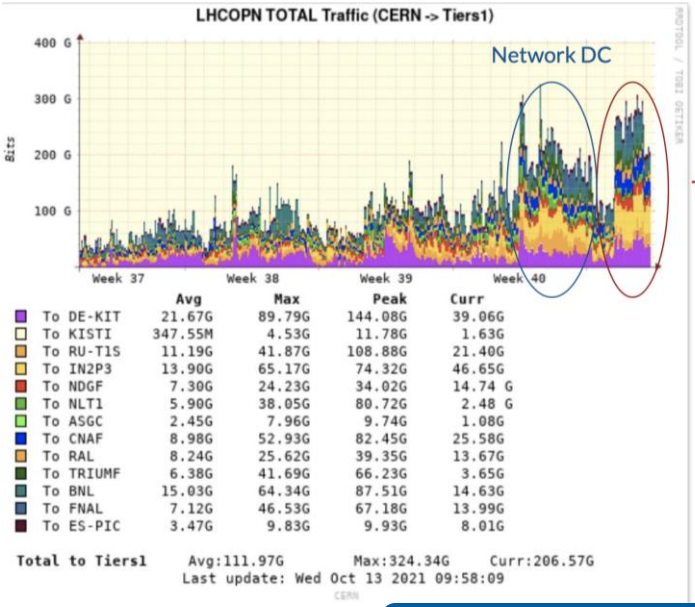

Tape Challenge

Monitoring

http://s.kit.edu/gridka-monitoring

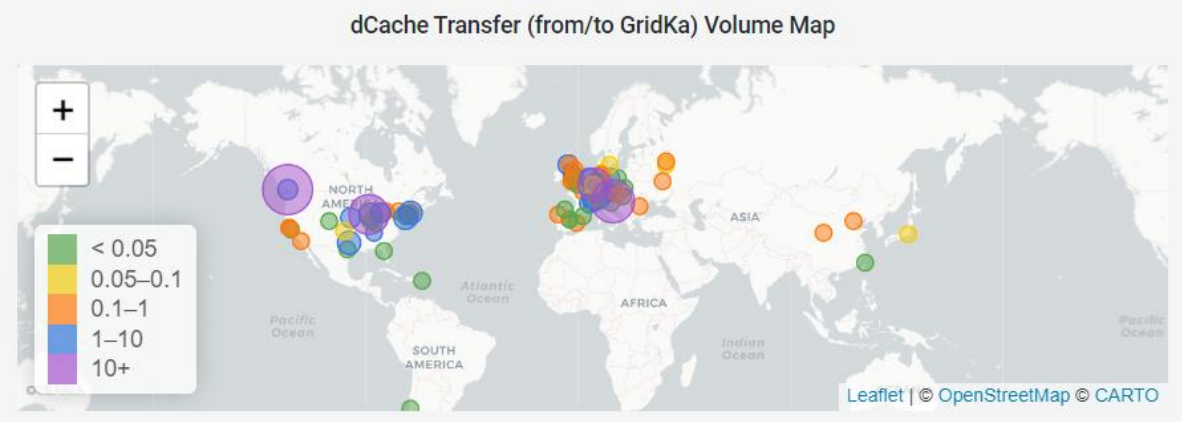

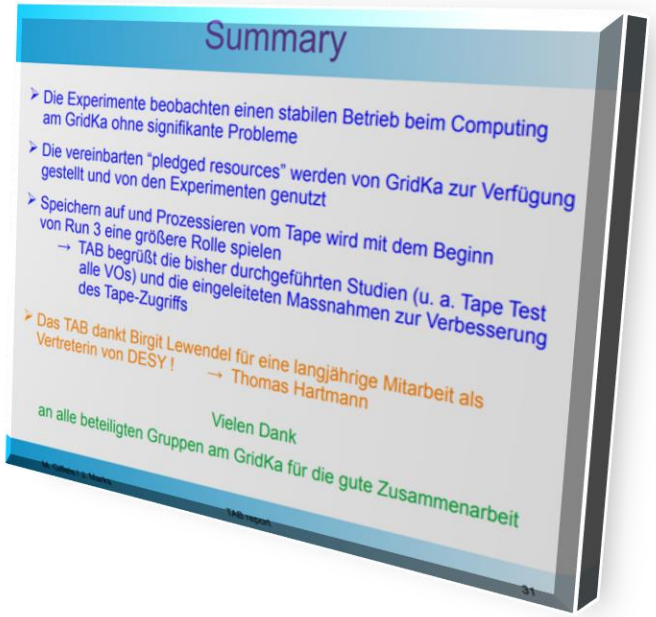

NO operational restrictions due to Corona, some challenges during delivery and commissioning of new HW

Chip/Material Crisis:
We were lucky until now!
But expect situation to get worse before it gets better!