“Innovative Digital Technologies for the Exploration of the Universe and Matter”
Status Report

Volker Lindenstruth, Kilian Schwarz, Raffaele Grosso, Paul-Niklas Kramp, **Serhat Atay**
Goethe University Frankfurt
GSI Darmstadt
11 March 2020

Work Packages

- Topic A
  - WP1) Tools for integration
    - Scheduling cloud jobs
    - Container technologies
    - Checkpointing
    - Access to experiment databases
  - WP2) Efficient use
    - Increase the efficiency of data-intensive applications on heterogenous resources using on-the-fly caches
  - WP3) Workflow control
    - Identification and control
    - In-Pilot Job Monitoring
    - Accounting
    - Optimization through data mining
- Topic B
  - WP1) Test of technology components

    Implementation and testing on different platforms of
    - Storage and caching solutions and
    - Virtualized services (databases, monitoring, accounting)
  - WP2) Job and resource management
    - Job distribution and monitoring in the environment of heterogeneous computing resources, including container virtualization
  - WP3) Virtualization user jobs
    - Recording the requirements
    - Determination and generation of the runtime environment
    - Creation of container metadata and
    - Checkpointing of container virtualization
  - WP4) Combined tests

    Testing of entire systems on different platforms
    - installation and maintenance effort
    - performance
    - scalability and
    - robustness

Milestones

- Topic A WP1 (A1):
  - M1 - June 30, 2019: successful integration of the container system in the SLURM scheduler used at GSI, functioning management system for images
  - M2 - June 30, 2020: Demonstration of singularity container system in the SLURM scheduler at Goethe-HLR
  - M3 - June 30, 2021: Generalization of the developed environment so that it can also be used by other experiments (e.g. CMS).
- Topic A WP2 (A2):
  - M1 - December 31, 2018: complete availability of the tested tools XRootD-Redirector-Plug-In, XRootD-Client-Plug-In and XRootD-Forward-Proxy.
  - M2 - December 31, 2019: first prototype of "Disk Caching on the fly" at the Goethe-HLR. (Postponed till June 30, 2020 in agreement with BMBF)
  - M3 - December 31, 2020: First prototype of "Disk Caching on the fly" at the University of Freiburg, integrated into the existing environment of the CMS experiment.
  - M4 - June 30, 2021: Expansion of "Disk Caching on the fly" at the University of Freiburg for operation with the ATLAS experiment.

Legend: Done, in progress

Milestones

- Topic B WP1 (B1):
  - M1 - December 31, 2019: Simplified installation (e.g. packaging) including documentation for the "Disk Caching on the fly" developed in work package A2 for heterogeneous resources.
  - M2 - June 30, 2020: Documentation of detailed performance measurements and the scalability of the disk caching system developed in A2. In particular, comparison of disk caching at Goethe-HLR with direct WAN access to the Luster cluster at GSI. (Postponed till Dec 31, 2020 in agreement with BMBF)
  - M3 - December 31, 2020: Simplified installation (e.g. packaging) including documentation for the container environment for grid jobs developed in work package A1.
  - M4 - June 30, 2021: Support for the commissioning of the disk caching system developed in A2 or the container environment developed in A1 in at least one other center for at least one other community.

Legend: Done

- Topic B WP4 (B4):
  - M1 - June 30, 2020: The singularity-based container system developed in work package A1 is to be used at the Goethe-HLR data center. The jobs sent should read the data to be analyzed from the "Disk Caching on the fly" set up in work package A2 at the Goethe-HLR. Performance measurements are to be carried out. (Postponed till Dec 31, 2020 in agreement with BMBF)
  - M2 - June 30, 2021: The container system developed in work package A1 is to be expanded to include an interface to HTCondor. The goal is a dynamic cloud created with HTCondor, whose jobs running in singularity containers extend over the batch systems of the HPC farms of GSI and the Goethe-HLR and their data either from the "Disk Caching on the fly" set up in work package A2 "at Goethe-HLR or directly from the Luster Cluster of GSI, if necessary via WAN access. Performance measurements are to be carried out.

Milestones

- A1:
  - M1 (June 30,2019) (singularity container at GSI, disk image management system)
  - M2 (June 30, 2020) (singularity container at Goethe-HLR)
- A2:
  - M1 (June 30, 2018) (availability of all tools for “disk caching on the fly”)
  - M2 (Dec 31 ,2019: postponed June 30, 2020) (disk caching on the fly at Goethe-HLR)
- B1:
  - M1 (Dec 31, 2019) (packaging and documentation for disk caching on the fly)

Legend: Done, in progress

Status of Milestones of A1

- M1: Singularity container at GSI, disk image management system
  - Singularity at GSI, running already via ALICE Grid Jobs
  - Disk image management system
- M2: Singularity container at Goethe-HLR
  - Reservation of four cluster nodes to test singularity at Goethe-HLR -done by the cluster admins
  - Singularity package installation at Goethe-HLR - installed on four reserved nodes and not tested yet due to cluster downtime in this week
  - Installing singularity to all nodes at Goethe-HLR - waiting for tests on four reserved nodes

Legend: Done, in progress

Disk Caching on the fly (vagrant setup)

- XrdProxyPrefix plugin employed by the client prefixes the address of the datamanager/ redirector
- The Datamanager/ Redirector checks the availability of the requested data on the cache
- The ForwardProxy retrieves the requested data from the external data server

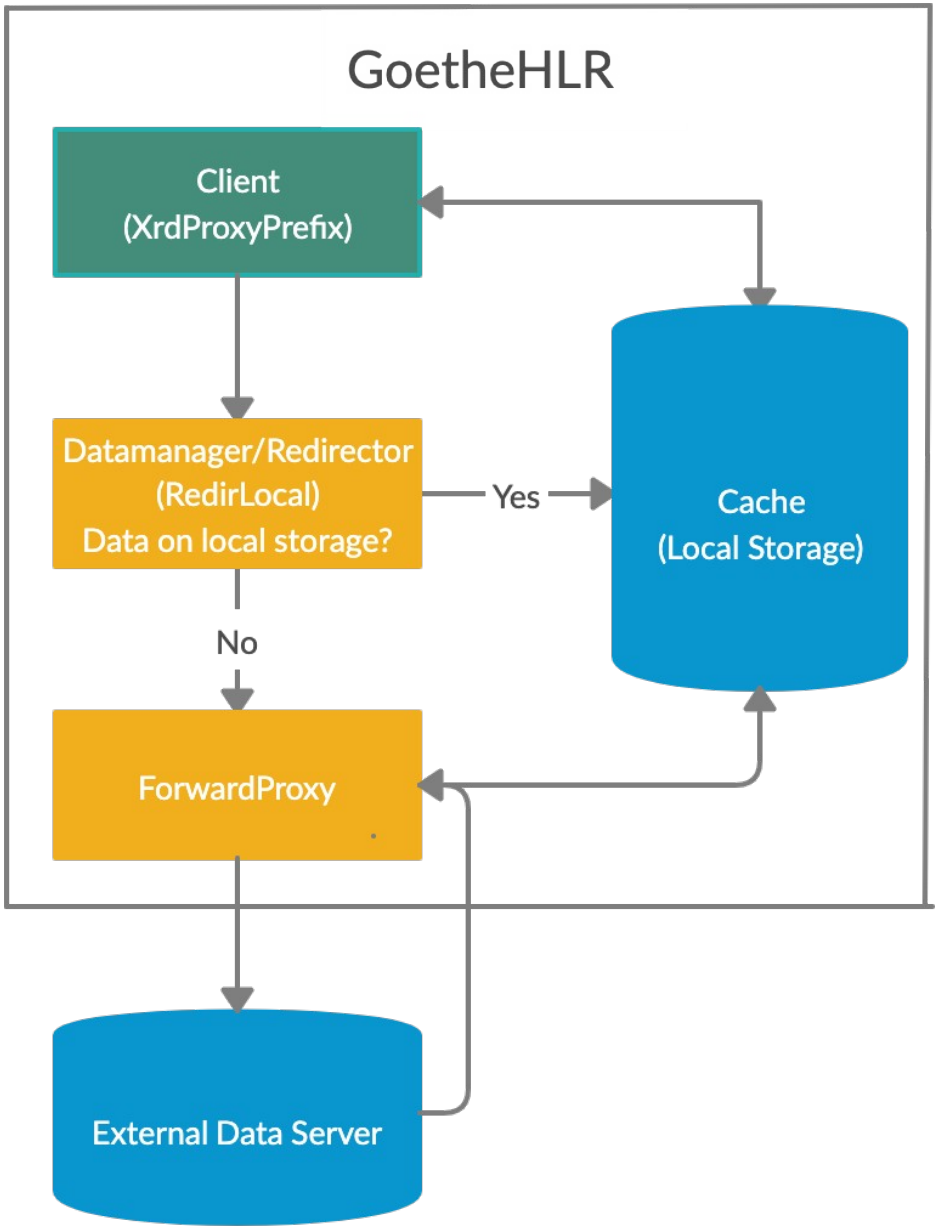

Status of Milestones of A2

- M1: availability of all tools for “disk caching on the fly”
  - Plugins
    - XrdForwardProxy -already exists at GSI GitLab repository
    - RedirLocal -already exists at GSI GitLab repository
  - Vagrant setup with five virtual machines on a private network to test “disk caching on the fly.” (client, data manager, forward proxy, external data server and data server)
- M2: disk caching on the fly at Goethe-HLR
  - Reservation of four cluster nodes to test “disk caching on the fly” at Goethe-HLR -done by the cluster admins
  - Test of xrootd servers
    - Installation of a VM software, ansible and vagrant on the cluster -installed virtualbox 6.0.6-6.0.18, vagrant and ansible on one of the reserved node
    - Installation the vagrant setup on one of the reserved node to check data caching -working fine
  - Setting up two virtual Xrootd servers at Goethe-HLR
    - Datamanager/Redirector -requested an IP address and a mount point to the shared filesystem
    - ForwardProxy -requested an IP address and a mount point to the shared filesystem

Legend: Done, in progress

Status of Milestones of B1

- M1: Packaging and documentation for disk caching on the fly
  - ForwardProxyPrefix and RedirLocal plugins
  - A vagrant setup which utilizes five virtual machine on a private network to test “disk caching on the fly”
    - Available for debian, centos and scientific linux machines via repository with documentation
    - https://git.gsi.de/dc/XRootD-Reps/xrootd-disk-cache-local-access/tree/ansibleCentos-xrootd-disk-cache-local-access

Legend: Done

Summary and Outlook

- A1-M1: Singularity container at GSI, disk image management system
- A2-M1: Availability of all tools for “disk caching on the fly”
- B1-M1: Packaging and documentation for disk caching on the fly
- Primary and urgent developments up to June 30, 2020
  - Singularity running at Goethe-HLR (A1-M2)
  - Setting up two Xrootd servers for “disk caching on the fly” (A2-M2)
- A1-M3: container jobs with CMS at KIT or bwFORCluster NEMO in Freiburg
- A2-M3: disk cache on the fly with CMS in Freiburg

Thanks

Computing Source Needs in Future

- Computing source needs >60x in next decade
- Flat budget (15-20%/year) ~6x-10x in next decade

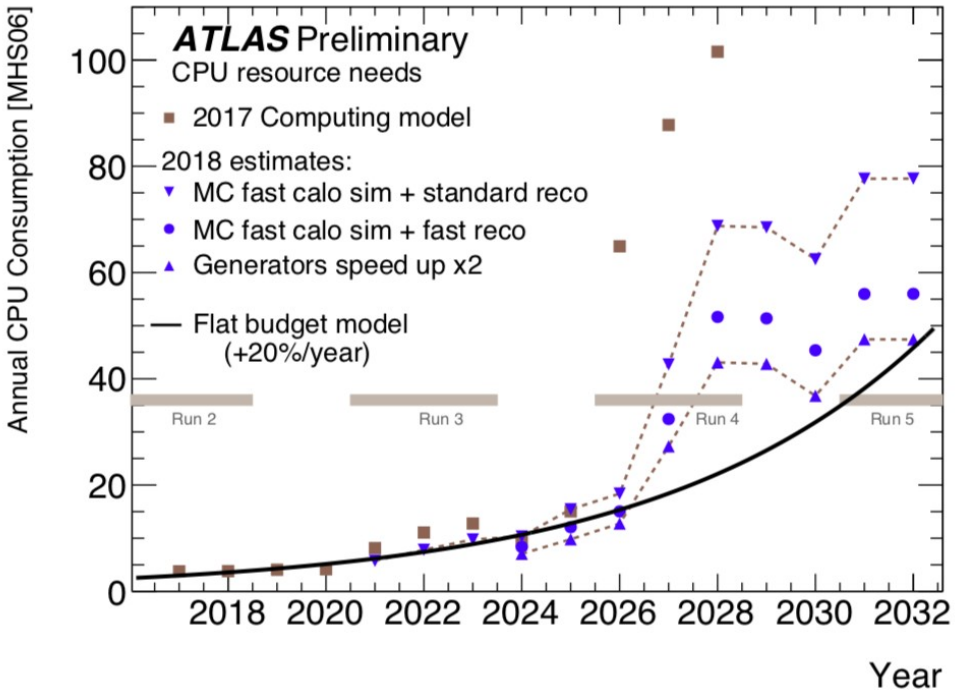

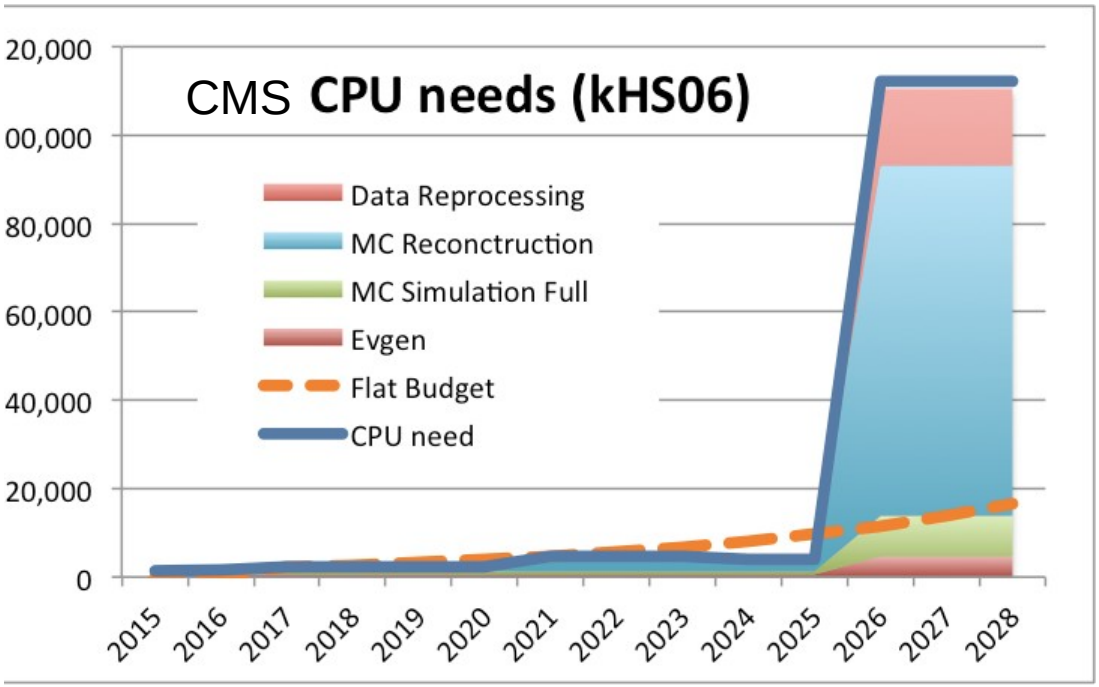

https://twiki.cern.ch/twiki/pub/AtlasPublic/ComputingandSoftwarePublicResults/cpuHLLHC_18.png

Motivation of the Project

- Cross-experimental solutions for the significant increase in the computing source needs
- Four topics of the project
  - Topic A
    - Virtualization and container technologies as well as new methods of data provisions
  - Topic B
    - The interoperability of the components, robustness and scalability
  - Topic C
    - Usage of deep neural networks for data analysis
  - Topic D
    - Exploiting the capabilities of modern processor architectures