

## T 25: Grid Computing

Time: Tuesday 17:00–18:30

Location: H-HS V

T 25.1 Tue 17:00 H-HS V

**Simulating distributed caching for end user analyses** — RENÉ CASPART, R. FLORIAN VON CUBE, ●TABEA FESSENBECKER, MAX FISCHER, MANUEL GIFFELS, CHRISTOPH HEIDECKER, MAXIMILIAN HORZELA, EILEEN KÜHN, GÜNTER QUAST, and MATTHIAS SCHNEPF — Karlsruhe Institute of Technology

With the start of the High-Luminosity LHC era in the near future, the massive amount of data processed in HEP user analyses results in challenges for network and storage infrastructures. In order to face these challenges developing concepts to use existing computing infrastructures more efficiently becomes essential.

Bringing data close to the computing resources is a promising approach to overcome throughput limitations and improve the overall performance. One way to achieve this is the placement of caches close to the computing resource. In order to efficiently use caches in a heterogeneous, distributed computing infrastructure an adapted operation concept needs to be developed.

Our goal at KIT is to develop a concept for facilities focused on end-user analysis. Therefore, we monitor our system and create a simulation to study the effects of different caching concepts. This allows to investigate the possibilities that distributed caching offers to computing facilities designed for end-user analysis.

In this contribution, we give an overview of the simulation tool and discuss different cache operation concepts of distributed caching.

T 25.2 Tue 17:15 H-HS V

**XCACHE for ATLAS production** — ●NIKOLAI HARTMANN, GÜNTER DUCKECK, RODNEY WALKER, and CHRISTOPH ANTON MITTERER — Ludwig-Maximilians-Universität München

High-level caching provides a promising and lightweight storage solution that can replace or complement permanent storage clusters. The XRootD project includes a proxy caching server (XCACHE). Integration with the ATLAS data management system Rucio is added through a plugin.

We have setup such an XCache server at the LRZ-LMU Tier-2 and tested several ATLAS workflows. In particular, we ran an ATLAS production queue that processes all input via the XCache server. In our initial setup several problems showed up when operated longer-term under heavy load in production conditions. The latest XCache version proved to be more stable. We also evaluated different hardware configurations of the caching server, comparing the performance of individual disks and RAID 6.

T 25.3 Tue 17:30 H-HS V

**Performancesteigerung in HEP Arbeitsabläufen durch koordinierte Caches** — RENE CASPART, R. FLORIAN VON CUBE, TABEA FESSENBECKER, MAX FISCHER, MANUEL GIFFELS, CHRISTOPH HEIDECKER, ●MAXIMILIAN HORZELA, EILEEN KÜHN, GÜNTER QUAST und MATTHIAS SCHNEPF — Karlsruhe Institute of Technology

Der stetig steigende Bedarf an IT-Ressourcen durch eine explosiv wachsende Menge an Daten und die zur Analyse benötigten Simulationen insbesondere in Hinblick auf den High Luminosity LHC ist nur mit neuartigen Computing-Konzepten zu decken.

Ein vielversprechender Ansatz um dem Ressourcenbedarf nachzukommen, bietet neben der Bereitstellung zusätzlicher Ressourcen, eine Effizienzsteigerung in den Arbeitsabläufen. Die Effizienz der Arbeitsabläufe in der Hochenergiephysik ist oft durch den Datentransfer limitiert. Um dennoch die volle Rechenleistung auszunutzen sind für verteilte Rechensysteme optimierte Caching-Lösungen, bei der das Workflow Management mit dem Batchsystem koordiniert wird, ein wichtiger Beitrag zur Effizienzoptimierung.

Dieser Beitrag umreißt grundlegende Konzepte für koordinierte Caching-Lösungen.

T 25.4 Tue 17:45 H-HS V

**Benchmarking of GRID resources and infrastructure using**

**the HammerCloud service** — ●DAVID HOHN, BENOÎT ROLAND, BENJAMIN ROTTLE, and MARKUS SCHUMACHER — Albert-Ludwigs-Universität Freiburg

The high luminosity LHC is a major upgrade to the current accelerator and will provide an order of magnitude more high energy physics data in the coming decades than exists now. To enable the analysis of this wealth of data the computing resources need to follow suit and be upgraded as well, both in storage and processing capabilities.

Reliable performance information is the foundation for planning of upgrade efforts as well as efficient operation of the complex computing systems. To this end, benchmarks provide essential insight. This talk will present fast benchmarks of CPUs which can enable prompt matching between available and required resources, as is mandatory to incorporate opportunistic and heterogeneous resources. Furthermore, benchmarks of remote storage access via several different protocols, e.g. WebDAV via Dynafed, direct WebDAV and Xrootd, will be presented. Dynafed could allow for aggregation and more efficient usage of storage resources.

Both sets of benchmarks are performed by the HammerCloud service within the ATLAS GRID computing network. This constitutes an extension of the HammerCloud use cases from functional to additional performance testing.

T 25.5 Tue 18:00 H-HS V

**Scheduling heterogeneous resources for federated usage** — RENÉ CASPART, ●R. FLORIAN VON CUBE, TABEA FESSENBECKER, MAX FISCHER, MANUEL GIFFELS, CHRISTOPH HEIDECKER, MAXIMILIAN HORZELA, EILEEN KÜHN, GÜNTER QUAST, and MATTHIAS SCHNEPF — Karlsruhe Institute of Technology

To overcome shortcomings in computing resources, the opportunistic usage of additional resources, not specifically dedicated to HEP, can help. Such resources might be HPC centers, or commercial and institutional cloud providers. An overlay batch system is used to provide a single point of entry to this plethora of resources and enables transparent usage for the user.

A novel on-demand resource management system COBaID/TARDIS has been developed and extensively tested at KIT to allocate and integrate such resources. Small computing centers are enabled to contribute computing resources to the grid in a very light-weight manner using those services. Such resources can be connected, allowing for jobs to move freely between different batch system instances.

In this contribution we present the latest software developments and the progress made in applications in terms of federating resource pools.

T 25.6 Tue 18:15 H-HS V

**Performance monitoring of opportunistic resources at ATLAS-BFG** — ANTON J. GAMEL, ●STEFAN KROBOTH, BENJAMIN ROTTLE, and MARKUS SCHUMACHER — Albert-Ludwigs-Universität Freiburg

The workload of computing clusters is typically unpredictable and tends to alternate between over- and under-utilization of the available resources. Tools such as COBaID/TARDIS provide an easy way to opportunistically make under-utilized resources of one site available to other sites. Fine-tuning of the involved software in order to optimize efficiency and user experience needs to be performed in a production environment and is therefore difficult without continuous monitoring of logs and meaningful metrics. In this work we present the current situation at Freiburg University where resources of the NEMO cluster are used to extend the ATLAS-BFG cluster in an opportunistic fashion using COBaID/TARDIS. The talk covers the tools involved in the collection and analysis of logs and metrics acquired from different sources within the local and opportunistic resources. Examples of how the aggregation of logs and the monitoring of metrics aids decision-making are shown. Besides fine-tuning of the involved tools, this setup can also be used to detect problems and anomalies early on.