Contribution ID: **430** Type: **Oral**

Opportunistic data locality for end user data analysis

Wednesday 12 October 2016 11:45 (15 minutes)

With the LHC Run2, end user analyses are increasingly challenging for both users and resource providers.
On the one hand, boosted data rates and more complex analyses favor and require larger data volumes to be processed.
On the other hand, efficient analyses and resource provisioning require fast turnaround cycles.
This puts the scalability of analysis infrastructures to new limits.
Existing approaches to this problem, such as data locality based processing, are difficult to adapt to HEP workflows.

For the first data taking period of Run2, the KIT CMS group has deployed a prototype enabling data locality via coordinated caching.
The underlying middleware successfully solves key issues of data locality for HEP:

- Caching joins local high performance devices with large background storage.
- Data selection based on user workflows only allocates critical data to optimize throughput.
- Finally, transparent integration into the batch system and operating system reduces compatibility issues for user software.

While the prototype has sped up user analyses by several factors, the scope has been limited so far.
Our prototype is deployed only on static, local processing resources accessing file servers under our own administration.
Thus, recent developments focus on opportunistic infrastructure to prove the viability of our approach.

On the one hand, we focus on volatile resources, i.e. cloud computing.
The nature of caching lends itself nicely to this setup.
Yet, the lack of static infrastructure complicates distributed services, while delocalization makes locality optimizations more complicated.

On the other hand, we explore providing caching as a service. Instead of creating an entire analysis environment, we provide a thin platform integrated into caching and resource provisioning services. Using docker, we merge this high performance data analysis platform with user analysis environments on demand. This allows using modern operating systems, drivers, and other performance critical components, while satisfying arbitrary user dependencies at the same time.

Tertiary Keyword (Optional)

Virtualization

Secondary Keyword (Optional)

Computing middleware

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Author: FISCHER, Max (KIT - Karlsruhe Institute of Technology (DE))

Co-authors: HEIDECKER, Christoph (KIT - Karlsruhe Institute of Technology (DE)); KUHN, Eileen (KIT - Karlsruhe Institute of Technology (DE)); QUAST, Gunter (KIT - Karlsruhe Institute of Technology (DE)); GIFFELS, Manuel (KIT - Karlsruhe Institute of Technology (DE)); SCHMITT, Marcus (KIT - Karlsruhe Institute of Technology (DE)); SCHNEPF, Matthias Jochen (KIT - Karlsruhe Institute of Technology (DE))

Presenter: FISCHER, Max (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing