Contribution ID: **566** Type: **presentation**

GRACC: GRid ACcounting Collector

Tuesday 10 July 2018 11:15 (15 minutes)

The OSG has long maintained a central accounting system called Gratia. It uses small probes on each computing and storage resource in order to usage. The probes report to a central collector which stores the usage in a database. The database is then queried to generate reports. As the OSG aged, the size of the database grew very large. It became too large for the database technology to efficiently query to generate detailed reports.

The design of a replacement required data storage that could be queried efficiently to generate multi-year reports. Additionally, it requires flexibility to add new attributes to the collected data.

In this paper we will describe the GRACC architecture. GRACC uses modern web technologies that were designed for large data storage, query, and visualization. That includes the open source database Elasticsearch, message broker software RabbitMQ, and Grafana and Kibana as data visualization platforms. It uses multiple agents that perform operations on the data to transform it for easier querying and summarization.

Authors: WEITZEL, Derek John (University of Nebraska-Lincoln (US)); ZVADA, Marian (University of Nebraska Lincoln (US)); BOCKELMAN, Brian Paul (University of Nebraska Lincoln (US))

Co-authors: RETZKE, Kevin Michael (Fermi National Accelerator Lab. (US)); BHAT, Shreyas (Fermi National Accelerator Laboratory)

Presenter: BOCKELMAN, Brian Paul (University of Nebraska Lincoln (US))

Session Classification: T3 - Distributed computing

Track Classification: Track 3 –Distributed computing