

# GRACC

## New Generation of the OSG Accounting



Kevin Retzke — [kretzke@fnal.gov](mailto:kretzke@fnal.gov), Tanya Levshina — [tlevshin@fnal.gov](mailto:tlevshin@fnal.gov), Bo Jayatilaka — [boj@fnal.gov](mailto:boj@fnal.gov),  
Chander Sehgal — [sehgal@fnal.gov](mailto:sehgal@fnal.gov), Shreyas Bhat — [sbhat@fnal.gov](mailto:sbhat@fnal.gov) • Computing Sector, Fermilab, Batavia, IL  
Brian Bockelman — [bbockelman@cse.unl.edu](mailto:bbockelman@cse.unl.edu), Derek Weitzel — [dweitzel@cse.unl.edu](mailto:dweitzel@cse.unl.edu) • University of Nebraska — Lincoln  
Robert Quick — [rquick@iu.edu](mailto:rquick@iu.edu) • Indiana University Frank Wuerthwein — [fkw@ucsd.edu](mailto:fkw@ucsd.edu) • University of California — San Diego

### Introduction

- Throughout the last decade the Open Science Grid (OSG) has been fielding requests from user communities, resource owners, and funding agencies to provide information about utilization of OSG resources.
- With the rapid expansion and transformation of the OSG resources and access to them, we are faced with several challenges in adapting and maintaining the current accounting service, Gratia, developed in 2006.
- The newest changes include:
  - acceptance of users from numerous university campuses, whose jobs are flocking to OSG resources,
  - expansion into new types of resources (public and private clouds, allocation-based HPC resources, and GPU farms),
  - migration to pilot-based systems,
  - migration to multicore environments.

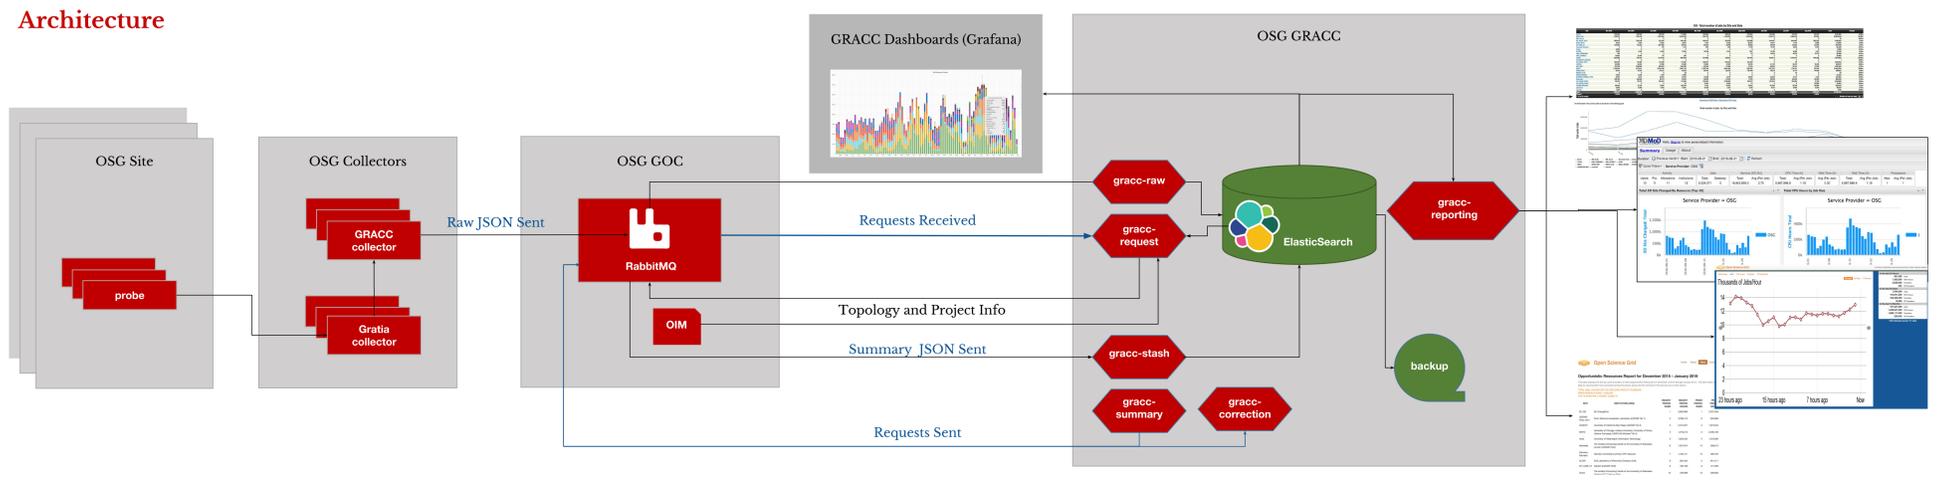
### GRACC Overview

- Next-generation OSG accounting service, based on open-source technology and compatible with Gratia.
- Consists of swappable, independent components:
  - Data Storage: Elasticsearch
  - Data Handling: Logstash: data handling
  - Data Exchange: RabbitMQ
  - User Interface: Grafana
  - Analytics: Kibana
  - Data Collection: custom probes, including Gratia probes
- Reports are generated by custom scripts accessing ElasticSearch.

### Requirements

- All historical summary data needs to be preserved.
- Generate daily summary:
  - of job wall duration, cpu usage per site, VO, Project, DN (with role), user, exit code for Batch and Payload resources.
  - transfer summary (size and wall duration) per storage site, vo, user, direction, exit code.
  - raw and logical usage of storage space per site and vo.
- Accepts input from existing Gratia probes.
- Ability to aggregate at different levels (site versus cluster versus CE) and rename site aggregation.
- Should be able to use WebUI to extract the accounting information.
- Continues to interface EGI and XSEDE accounting services and provides information in accordance with existing agreements.
- Allows new metrics to be reported and summarized.

### Architecture

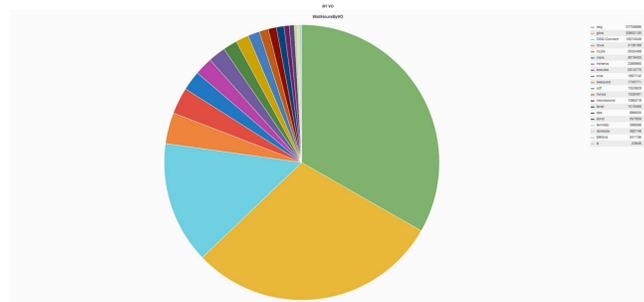
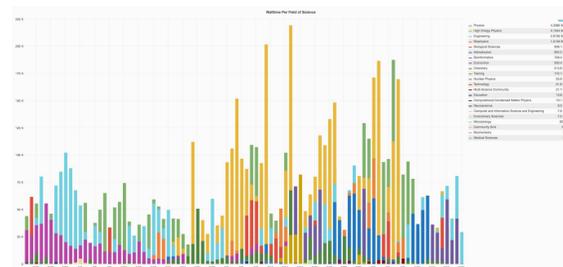
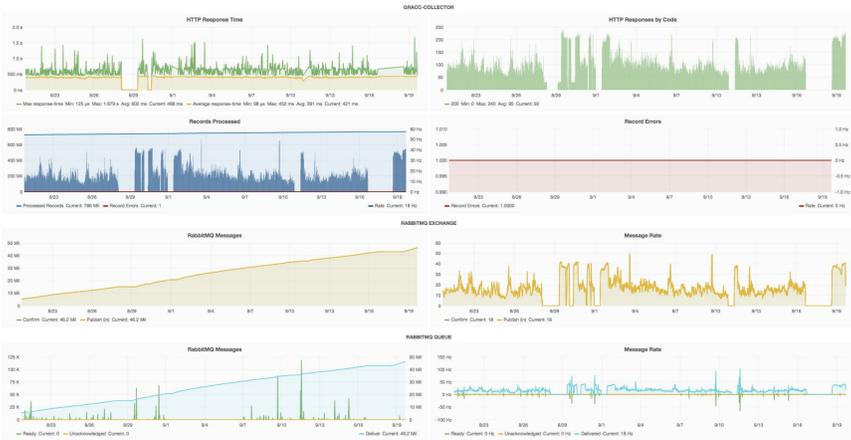


### Deployed Components

- ES Cluster at Nebraska (HCC) on “Anvil” cloud
  - head node and 4 data nodes.
  - each node has 8 CPUs and 64GB of RAM.
  - 5 TB Ceph-backed network storage per node
  - Four data/master nodes, one “client” node running collector and agents
- RabbitMQ at Indiana University (GOC)
- Collector, Forwarder (at Fermilab)

### WEB UI

- Uses Grafana’s dashboards to display batch, payload jobs and file transfer metrics.
- Provides means to create Kibana plots for savvy users.
- Integrated with the OSG Information Management Service (PIs, Science, Topology Information).



### Statistics

- 1.1 billion documents in ElasticSearch.
- 3.5 TB current disk usage for summary & raw documents.
- 1 million new job records per day.

### Planned Development

- Integration with WLCG and XSEDE Accounting.
- Summarization of Storage and Transfer records.
- Full integration with the OSG Information Management Service.
- Implementation of backup and restoration procedure.
- Generation of email reporting.
- Production: Spring of 2017.