CHEP 2013
Poster ID 221

GRID ACCOUNTING SERVICE:
STATE AND FUTURE DEVELOPMENT

Tanya Levshina — *tlevshin@fnal.gov*, **Chander Sehgal** — *cssehgal@fnal.gov* • Computing Sector, Fermilab, Batavia, IL
Brian Bockelman — *bbockelman@cse.unl.edu*, **Derek Weitzel** — *dweitzel@cse.unl.edu*, **Ashu Guru** — *aguru2.unl.edu* • University of Nebraska — Lincoln

Introduction

- **Accounting is a crucial component of Grid Services.**
- **An accounting service collects resource utilization information including detailed user identities.**
- **It is important for**
  - **verifying pledged resource allocation per particular groups and users.**
  - **providing reports for funding agencies and resource providers.**
  - **understanding hardware provisioning requirements.**

Overview

Gratia is a resource accounting service adopted by the Open Science Grid, used at Fermilab, University of Nebraska - Lincoln and several other institutions.

- It is a robust and scalable service that is designed to collect various metrics from individual sites using local probes.
- The available metrics include daily data collection of wall and CPU utilization, as well as storage capacity and data transfer by user, VO and facility.
- The Gratia framework provides means to generate various reports that allow graphical and textual access to the data. It allows access to the data via several web portals.
- It offers various ways of data filtering, replication and forwarding.
- The available metrics are used by users, VO and site administrators, OSG management, as well as funding agencies.

Statistics

- 200,000 records per hour is the maximum observed rate for processing
- 1 billion individual usage records collected since 2005
- 900 TB accounted data transferred per day
- 130 sites reporting usage
- 16 types of probes for usage and data collection

Architecture

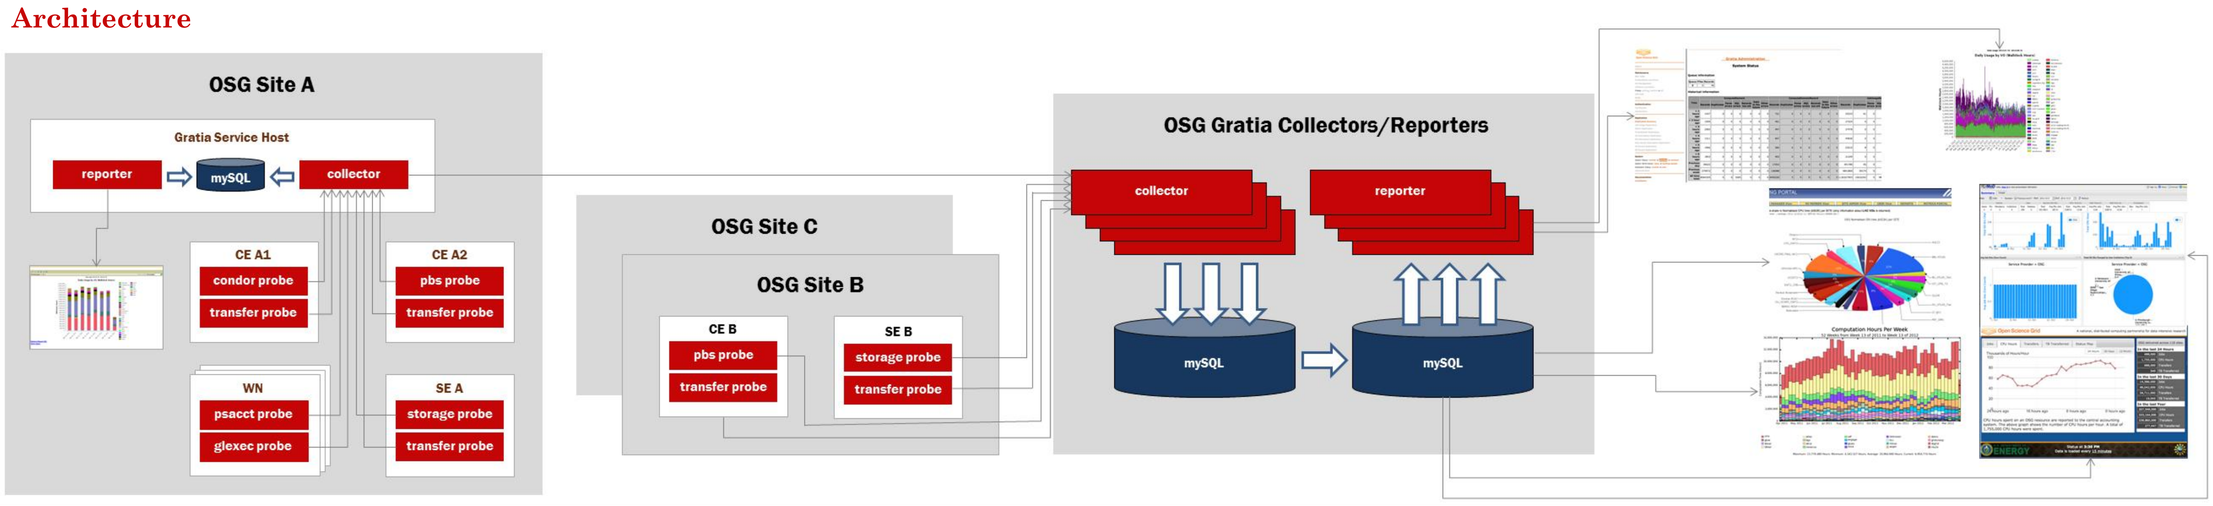

Campus Accounting

A Campus Infrastructure provides a means to access shared resources. It enables researchers to transparently use various computational resources that include resources from other campuses as well as resources from the national cyber-infrastructure, such OSG or XD.

The primary concern for Campus Grids (CG) accounting is the accurate recording of usage of the OSG resources by CG users. We need to account for jobs that have been flocked to OSG resources from CG submitters.

A local campus user usually does not use a certificate. Several modifications were made to Gratia so probes can account for a local user's usage on campus and on OSG.

We currently require that system administrators that maintain CG infrastructure register a new group of researchers in the OSG Information Management System (OIM) by providing minimal information about the nature of research (Field of Science), the name of a Principal Investigator and the preferred Project Name. The project name is set by users during the submission of the job to the CG.

The jobs records, submitted by CG users that contained specific Project Names are collected in Gratia. This way the resource usage could be traced back to the individual group.

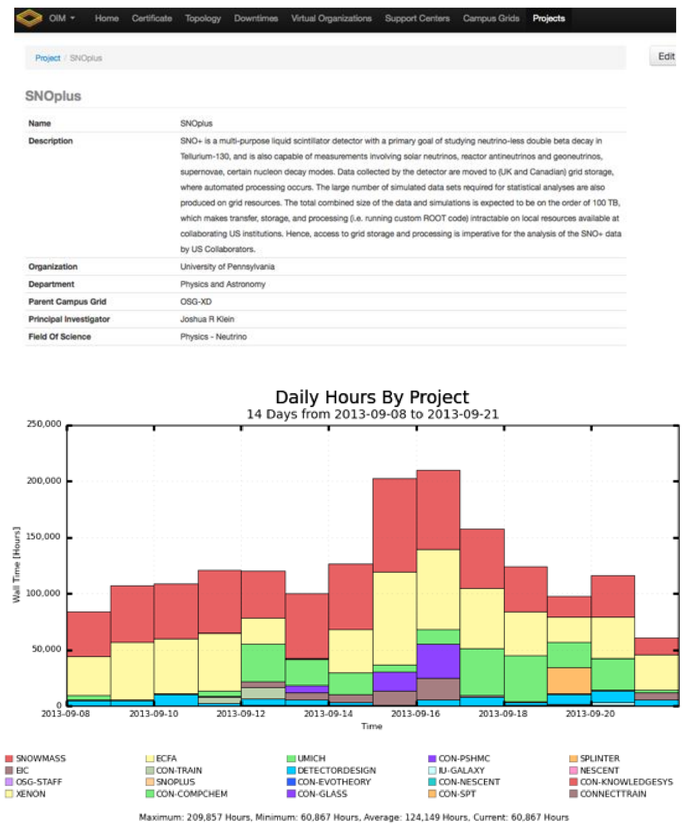

Cloud Accounting

FermiCloud Service is an infrastructure-as-a-service facility that provides the access to virtual machines on demand without sysadmin intervention. It is used by developers, integrators and testers. It also allows to start on-demand services such as gridftp or mysql servers.

The accounting information is collected by the Gratia cloud probe. Currently we have implemented a probe that interfaces with Open Nebula 3.2. All information about the VM's owner, physical machine, states, start and stop times, memory and cpu provisioning could be acquired from interfaces provided by Open Nebula.

The OGF Usage Record structure, used by Gratia, was sufficient to represent the VM accounting information. No changes in Gratia collector code were required.

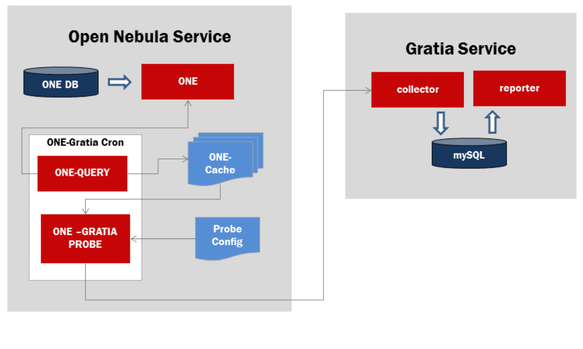

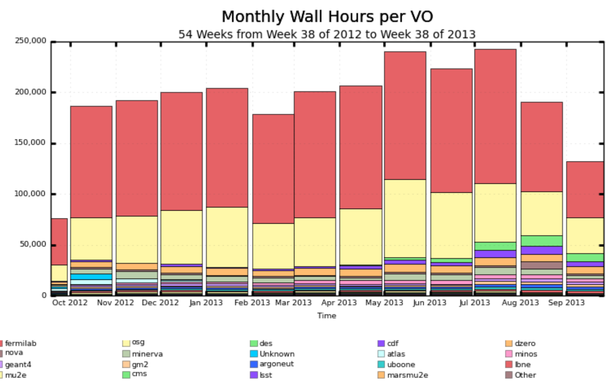

Integration with other CIs accounting services

XD(AMIE)
OSG has been an XD Service Provider since April, 2012. OSG provides a Virtual Cluster that shields XD users from the necessity to be aware of the OSG infrastructure. Jobs submitted into this Virtual Cluster are executed on the OSG physical clusters. The account records collected in Gratia, summarized by allocation (projects) are reported to the XD account management system (AMIE) by using AMIE API.

EGI(APEL)
Accounting records for US-LHC Tier1 and Tier2 facilities collected by Gratia are forwarded to the EGI accounting system (APEL), in accordance with signed MOU agreements. An external service parses and analyzes accounting records for a resource, scales wall time and cpu usage, and then forwards this data to the APEL Server using the EGI GOC.

Future Development

Expansion of Collected Metrics:

- Collected data is focused on CPU Usage. We are considering means to collect memory usage, as well as WAN and LAN network activities.
- Currently cloud probe only interfaces with Open Nebula. We will be working on development of a generic interface to various other Cloud Management Services such as Open Stack, Eucalyptus, and others. We should be able to incorporate into accounting records not only the static attributes and metrics of VM but also dynamic usage of the resources such as cpu, memory, disk usage and networking activities. We are planning to account for persistent storage used by a VM as well as elastic storage used only during the VM lifetime.

Accounting & MultiGrids Infrastructure:

- For the time being, means for information exchange between various accounting services are limited. We are planning to explore the possibility to define a common interface that if accepted by various accounting services would allow to pull information across different distributed computing infrastructures and present a unified view of resources utilization.

Acknowledgments: **Philippe Canal** — Gratia architect /developer, **John Weigand** — OSG/EGI reporting, **Mats Rynge** — OSG/XD reporting, **Soichi Hayashi** — Project/OIM registration, **Brad Hurst** — GratiaWeb development
Fermilab is Operated by the Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.