

Gratia Review

<https://indico.fnal.gov/conferenceDisplay.py?confId=6158>

1

First of all we want to **thank You, the reviewers,**
Stu Fuess, Ian Fisk, Maria Girone, David Swanson
for making the time (in what we know are very busy
schedules) and travel (in winter with full and delayed
flights) to help us, the Gratia sponsors, team and
management, think about and better plan for the future.

Agenda :

Most of the meat is in Tanya and the Stakeholders talks. I
am setting the scene and background.

We are especially looking for your thinking and wisdom –
we hope we have allowed plenty of time for questions and
discussion.

Gratia Background

<https://indico.fnal.gov/conferenceDisplay.py?confId=6158>

2

The Fermilab Computing Sector is committed to the ongoing maintenance and support of the Gratia software.

- The software is public domain with the code hosted on sourceforge
- Stakeholders are (and new ones can be) supported under agreements (SOW, MOU, Sector annual budgeting).
- The Computing Sector commits to provide the technical leadership and core development effort, with operations, first line of support and contributions coming from and in collaboration with stakeholders.


We all appreciate that Accounting is a management service – not core to the processing and analysis that gets the scientific results and discoveries out, “boring” to many. Be that as it may:

- Accounting is deemed the one Service OSG Sites are required to install and report back to central operations
- Reporting is required under the WLCG MOU
- Accounting/usage provides an easy to present summary of the progression and growth of a Facility
- Accounting provides “eyes” into what is actually happening in large distributed, non-deterministic system. This is very valuable and difficult to come by; and this has resulted in a regular string of new requests.

Gratia transport and repositories used for Resource Testing and Availability History



Approachable, easy to understand plots of growth and use.



Why a Review Now?

- Understanding and assessment of the return on and the business case for investments by the organization.
- Timely to revisit expectations, planning and commitments
 - ▣ The Fermilab Computing Sector recently reorganized, with the end of Run II and the stable operations of CMS new initiatives with competing priorities are coming rapidly to the table.
 - ▣ Steve Wolbers – head of the Scientific Computing Facilities quadrant in the Scientific Computing Division has asked for the review given the leadership, core development team and main operations effort are in his organization.
- Hearing from stakeholders and reviewers of the value of and continued need for the project – both the need to sustain the maintenance and at what kind of level to continue support development – is essential to justify continued investment.
- Ongoing challenges faced by the project in coordinating contributions from the contributors and prioritizing requests from stakeholders.
- The Gratia team wants to take stock and understand whether the current approach, and technologies, are still correct and still warranted. For example, Is there a “silver bullet” new implementation out there we could adopt with less effort.
- Technical leadership has changed hands several times. Tanya, as the latest lead, is doing an excellent job. Your time and efforts will help her to continue to and increase her success.

Goals of the Gratia Project



Goal of the project defined early on:

“The Gratia Project designs and deploys robust, scalable, trustable and dependable grid accounting, publishes an interface to the services and provides a reference implementation.”

“grid” has been treated as generic and current/
updated goal would replace it with
“cyberinfrastructure”

Charge to the Reviewers

As reviewers you are asked to consider the 1-5 year time frame for the Gratia project to:

- Understand and comment on the scope, status and plans.
- Comment on the priorities and appropriateness of the pending requests, requirements and planned evolutions.
- Offer advice and guidance to the management and sponsors of the project for their continued support.
- Recommend whether the current approach, and technologies, are still warranted.
- Advise whether and/or when the project should look into whether there exists any new implementations out that we could adopt and maintain effectively (with fewer resources in the longer term).
- Suggest areas of investigation and work that would enhance the value of the project.

Added recently by presentation reviewers:

- Guidance on how to prioritize and what the prioritization of pending requests.
- Guidance on how to best coordinate contributions.

Potted History

- The Gratia project was started in 2005 as a PPDG extension activity.
- Development, deployment and software support effort cost has been ~10FTE years (~2FTEs for ~3 years; , ~0.5 FTE for ~1 year, ~1 FTE for ~3 years.)
- It was recognized that the initial requirements led to a “data acquisition” type architecture, with many sensor Probes, a high performance, resilient transport layer – with buffering to overcome transient transport problems and downtimes, and a high transactional repository.
- The transaction rate, the number of distributed sources, the size of the data store(s), the architecture of distributed/communicating data stores were all requirements that could not be met by other existing implementations at the time.
- The core focus was on ensuring the performance of the data transfer and scalability of the acquisition, storage and reporting could be sustained.
- Core technologies have had a couple of redesigns and many extensions to meet changing needs.
- The transport layer and repository service have proven reusable for extensions in the information collected and stored – job accounting, availability (RSV), and then storage and data transfer.
- Less emphasis, and less success has been had, with the native reporting and GUI tools.
- The Gratia accounting service has provided “easy to present and understand” summary of progress of the Facilities that have adopted it.
- The list of requests for the project, developments as a result of changing environments and services.
- The resources needed for maintenance and support have not shrunk.

Current Stakeholders

- Fermilab
 - ▣ Contributes: Provides core Gratia development/support team; develops probes, new capabilities for Clouds; Provides test systems to developers and integrators.
 - ▣ Users: FermiGrid, FermiCloud, HPC Scientific Facilities
- Open Science Grid
 - ▣ Contributes: Software Team efforts, probes, web reporting interface, expert debugging,
 - ▣ Users: Executive Team, Production (including Operations, Campus Grids), User Support, VOs
- US ATLAS, US CMS, ALICE-USA
 - ▣ Contributes: Reporting scripts to the WLCG
 - ▣ Users: US CMS and US ATLAS Tier-1s.
- University of Nebraska HCC
 - ▣ Contributes: extensions to the software, new probes, expert testing and debugging
 - ▣ Users: All HCC distributed infrastructure
- WLCG
 - ▣ Contributes: Interfaces, integration and coordination between OSG and EGI accounting services.
 - ▣ Users: Receives information on behalf of US LHC and creates necessary reports
- XSEDE users with Allocations on OSG
 - ▣ Transparent to the XSEDE users - Gratia interfaced into XSEDE AMIE reporting so that reports through XDMOD and XSEDEDB.

Options and Decisions for the Future – Continue Development?

- Should we “freeze” development saying original goal has been met (or not – but if not yet, why should we expect this to be achieved in the foreseeable future)?
- Do we consider updating the technologies – looking for new developments from elsewhere?
- Given we continue development, how should we prioritize requests across the stakeholders?

Options and Decisions for the Future – Management, Planning issues

- The management and coordination of the many threads is complex.
- The probes are developed outside of the project and then maintained by the developers. The OSG Software Team maintains and supports some of the probes and is responsible for the releases and distribution through the VDT.
- All needs to be integrated and work well as a system.
- Issues related to distributed responsibility/authority for probes include:
 - ▣ How it is decided to develop a particular probe.
 - ▣ Meeting the commitments for maintenance, support and documentation.
- The OSG Software Team is a key contributor and important interface for the Project.
 - ▣ They are responsible for the distribution of all the Gratia services.
 - ▣ Perhaps need to clarify and pay higher level attention to interface to OSG Software Distribution Team which distributes all the s/w through VDT.
- Would perhaps benefit from a periodic board across the stakeholders and effort contributors that reviews and prioritizes the various request

Options and Decisions for the Future – OSG Thrust into the Campus ?



- Is it appropriate (for OSG) to account for usage on or from the Campus that uses (OSG) CI services?
- Do we have a good set of requirements for this and if not how do we acquire them?

Options and Decisions for the Future – One accounting system for all?

- OSG PI/Technical Director statement: “Why is there more than one accounting system for National CIs?”
 - ▣ Regular attempts to collaborate with other US CI (TeraGrid/XSEDE).
 - ▣ Several accounting projects developed (finished/adopted/dropped) in the EU.
- OSG PI/Technical Director response to announcement to the ET of the review in October:
 - ▣ Does the project (Gratia) have a design/vision/mission document?
 - On web page and requirements document
 - ▣ Do we have an understanding of the "marketplace" for such accounting systems?
 - Current stakeholders and possibly XSEDE
 - ▣ Should we consider a separation between that information collection process and the storage and reporting part?
 - Possibility if collaborate with XSEDE? Is there system useful?
 - ▣ Can we look for more clients/users of Gratia?
- Expect Clouds are resources in production within ~2 years.

Options and Decisions for the Future – Another attempt to collaborate with XSEDE

- Currently trying again as part of Lothar/Miron/John initiative.

Usage Accounting and Analysis: There has been significant discussion on this to date and a document has been developed (see attached). While this has made good progress on defining the technical problem and a general approach to addressing it, this now needs to have an associated plan put together.

The group should be clear on a set of short-term deliverables to address the most relevant accounting and usage monitoring needs between XSEDE and OSG. Both projects individually have mature usage accounting systems in place. The group should assess relevant analysis queries and information displays that would serve to track and monitor usage across XSEDE and OSG, in particular the opportunistic use of CPU cycles.

Once relevant queries are established, a gap analysis will likely be required to identify what concepts and data are missing in the existing systems to connect the information between the different information providers on XSEDE and OSG and allow cross-projects analyses and monitoring. In case of gaps the group should develop the required concepts and approaches and define a possible work plan.

- OSG representatives: Ruth Pordes, LATBauerdick, Chander Sehgal
- XSEDE representatives: Steve Quinn, Dave Hart
- TAS representatives: Tom Furlani (or designee)

We look forward to your input and guidance..

