

# GRACC REVIEW 2016

---

12 DECEMBER 2016 / 13:00 / FCC117

## PROJECT STATUS

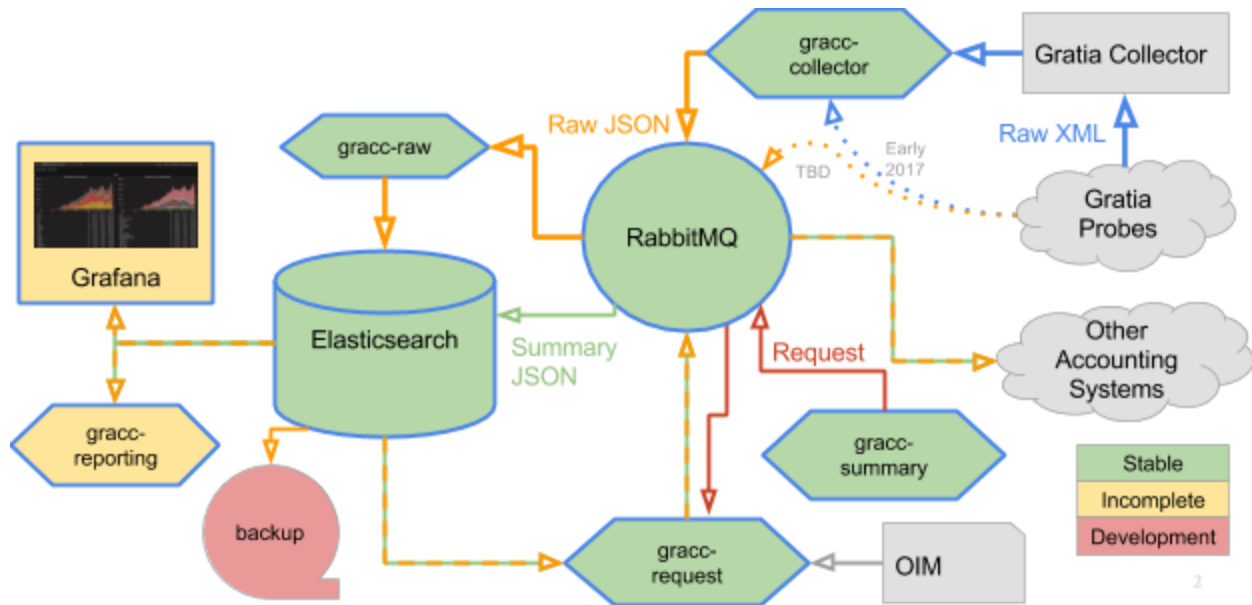
Review of what we set out to do, and where we are.

### Project Goals

[https://docs.google.com/document/d/1ojQdFk2YyjF3J\\_Q-PFHDpxkcaY6q6YBEsMJ0jdFR4tM](https://docs.google.com/document/d/1ojQdFk2YyjF3J_Q-PFHDpxkcaY6q6YBEsMJ0jdFR4tM)

- Retire GRATIA database and collector by 12/31/2016!
- Keep all the existing probes without modification
- All historical data must be transferred to the new infrastructure.
  - the granularity of information kept online may be reduced to monthly aggregates for anything that is older than 1 year
  - Any plot for more than two weeks does not have to be with an hourly granularity
- The display.grid.iu.edu does not change.
- The look and feel of Gratia displays may change at both gratiaweb and myosg.
- Careful attention must be paid to guarantee that the safety of the accounting data is as assured in the future as in the past.
- A transition like we do now is easier in the future.
- New system should be designed with long term maintenance in mind:
  - no more java or alike
  - flexible to schema evolution
- “workflow data” on workflows that run on OSG should be available to the wider CS audience.
- Allow study of joint of perfSonar, HTCondor, XRootd, gridftp, network transfer statistics

## Status



- Five-node Elasticsearch cluster and auxiliary services/agents running on HCC Anvil cloud at UNL. RabbitMQ (ITB) at GOC.
- Main OSG Jobs stream live since July
  - Raw record pipeline:
    - Gratia collector (replication)
    - gracc-collector
      - transforms to JSON
      - Sends to RabbitMQ
    - gracc-stash-raw
      - Checksums & timestamps receipt
      - Quarantines old records
      - Indexes into Elasticsearch
  - Summary pipeline:
    - gracc-summary
      - makes request through RabbitMQ
      - Systemd cron job for regular summarization (past week every 15 minutes)
      - CLI for on-demand
    - gracc-request
      - Aggregates raw records from Elasticsearch
      - Gets field corrections from Elasticsearch, updates summary records

- Collects data from OIM, adds to summary records
    - Sends to RabbitMQ
  - gracc-stash-summary
    - Checksums & timestamps receipt
    - Indexes into Elasticsearch
- Gratia MasterSummaryData loaded into Elasticsearch, included in gracc.osg.summary alias.
- Grafana
  - Basic dashboards for viewing Pilot & Payload usage by VO, Site, Project, etc.
    - <https://gracc.opensciencegrid.org/dashboard/db/pilot-jobs-summary>
    - <https://gracc.opensciencegrid.org/dashboard/db/payload-jobs-summary>
    - <https://gracc.opensciencegrid.org/dashboard/db/osg-project-accounting>
  - Fifemon also utilizing GRACC for some displays, either directly or with rollups in Graphite.
    - <https://fifemon.fnal.gov/monitor/dashboard/script/experiment-computing-summary.js>
    - <https://fifemon.fnal.gov/monitor/dashboard/db/fife-history>
- Reports
  - FIFE production reports using GRACC
- ITB and Transfer streams being collected
- Active monitoring and alerting of:
  - gracc-collector
  - RabbitMQ exchanges and queues
  - <https://gracc.opensciencegrid.org/dashboard/db/gracc-osg-monitor>
  - Elasticsearch
    - <https://gracc.opensciencegrid.org/dashboard/db/grace-elasticsearch-monitor>
  - VMs
    - <https://gracc.opensciencegrid.org/dashboard/db/grace-system-monitoring>
- Total active data: 4.5 TB, 1.4 billion documents
  - <https://gracc.opensciencegrid.org/dashboard/file/grace-indices.json>

# ROADMAP TO PRODUCTION

What is **required** before GRACC can be considered production-ready?

## Main OSG Jobs Stream

- gracc-collector needs to handle probe record bundles (GRACC-42)
  - Heartbeats from probe?  
<https://github.com/opensciencegrid/gracc-collector/issues/6>
- Full history must match Gratia (GRACC-40)
  - Still some anomalies
  - Gratia summaries need OIM information added, need Njobs aggregation in summary (GRACC-41)
- VO, Project, and Site Name correction (GRACC-44)
  - Handle site name corrections using existing field correction capability in request agent?
- OIM Integration (GRACC-7)
  - Need apel normalization factor for APEL reporting
- Special snowflakes:
  - Comet (GRACC-42)
  - IceCube (GRACC-27)
- Incomplete/incorrect records, e.g. Ga. Tech recently (GRACC-43)
- Update Raw records mappings & reindex all records (GRACC-33)
  - Expect weeks to do reindexing! May be faster with Elasticsearch 5.
- Agent bug fixes
  - Request agent crashing  
<https://github.com/opensciencegrid/gracc-request/issues/38>
  - Summary agent crashing  
<https://github.com/opensciencegrid/gracc-summary/issues/8>

## OSG Transfer & Other Streams

- Collector handle StorageElement and StorageElementRecord records  
<https://github.com/opensciencegrid/gracc-collector/issues/5>
- Transfer summaries (GRACC-32)
  - Need Njobs remapped (GRACC-33)
- ITB summaries
- Fermilab streams: Fermicloud, AWS, AAF

- Need collectors
- Custom summarizers
- Host on GRACE or ...?

## Grafana Dashboards & OSG Display

- “Making pretty” the graphs.
- Usability tests with enthusiastic volunteers (certainly not unwilling undergrads)
- Grafana Templating lacking ad-hoc filter capability of GratiaWeb
  - Send power users to Kibana?
  - Create Kibana dashboards to facilitate? (wishlist?)
- What is missing that is on GratiaWeb & do we need it?
- OSG Display (GRACC-18)
  - Where will it be deployed? Requires Python 2.7 (EL7)
- Allow study of joint of perfSonar, HTCondor, XRootd, gridftp, network transfer statistics

## Reporting

- Email summary reports (GRACC-4)
- OSG-Connect Project (GRACC-37)
- Flocking (GRACC-36)
- Per-Site Reports (GRACC-26)
- WLCG (GRACC-25)
- XSEDE (GRACC-20)
  - Mats expects to have test version done by end of Dec.
- APEL (GRACC-19)
- Opportunistic users (GRACC-39)
- Review of what reports we have (what is missing?)
- What reports are using raw records & why?

## Services

- Elastic stack upgrade to v5+
  - Full cluster shutdown
  - Kibana and Grafana compatibility?
- Elasticsearch co-hosting
  - OSG and LIGO history (50% of indices, 15% of storage)
  - FIFE (20% of indices, 5% of storage)

- Production RabbitMQ
- Access restrictions
  - Service certificates to restrict writing to Elasticsearch (currently using certs issued by private CA)
  - Basic auth for kibana access is not maintainable - options?
  - Don't request certs from frontend users (GRACC-29)

## Backups

- Tape backup & recovery of raw records
  - Maintain tape backups in Fermilab Enstore? Use AAF?
- Elasticsearch snapshots

## Monitoring and Alerting

- Need to decide between InfluxDB/Telegraf and Prometheus (Snap, other?) to monitor systems and services. Both support alerting, plus Grafana (see #gracc-notifications in Slack).
- Agents need to publish telemetry; Prometheus client libraries make this easy [https://github.com/prometheus/client\\_python](https://github.com/prometheus/client_python)
- What should we alert on?

# PRODUCTION

We've turned off Gratia; what next? (besides the five stages of grief)

## Deployment & Operations

- Gratia transition. How long do we keep collector up?
- MOU with HCC and GOC; support? Who operates production GRACC/GRACE? (also see "Funding" below)
- Raw record retention policy (GRACC-16)
- Name corrections maintainer
  - Summarization operations
- Service packaging. Containerized deployment? (GRACC-12)

## User Access

- Kibana access for power users; public Kibana? (GRACC-14)
- Another analytics platforms - Zeppelin, Jupyter?

## Wishlist

- What reports do we want?
- Raw record "nearline" archival, e.g. HDFS
- Probes reporting directly to RabbitMQ
- New fields in OIM that improve reporting?
- Tag summary records with processing errors (GRACC-31)

## Funding

With OSG Fermilab funding ending in April, 2017 who will be in charge of running the services? This is a question to Brian and ET. I don't think we could solve it at the workshop but they should be addressed.