OSG Technology Update

Brian Bockelman

State of the Union

- OSG Technology has drastically evolved over the past five years:
    - **CE philosophy transition** from “job submission” to “resource acquisition” (pilots).
    - Underlying **CE technology transition** changed from Globus GRAM to HTCondor-CE.
    - **Information services** are transitioning from LDAP to HTCondor-based.
    - **Accounting system** is being modernized.
    - VO-installed software (when needed) migrated to **OASIS**.
    - **Storage stack** is being simplified.
- The Technology Area additionally collaborates closely with operations for OSG’s expanded portfolio of services.

Storage Simplification

- **SRM retirement**:
  - Goal to retire SRM endpoints from one USATLAS and one USCMS site this year. **One USCMS site done**; actively partnering with USATLAS.
  - Remove bestman2 from release series (April 2017?). Would imply dropped support by fall 2017.
- **Stash / StashCache**: Responding to the fact that storage management / storage element paradigm is **too complex** for the non-LHC VOs,
  - **Stash**: Effort by User Support team to provide a single high-performance storage element for OSG VO users.
  - **StashCache**: Effort by OSG Technology to provide a caching layer for *any* VO, tuned for working set sizes O(10TB).
  - The single-SE / caching paradigm requires far less investment from VOs to utilize.

OASIS / OSG-Storage

- OASIS is the name for our CVMFS infrastructure:
  - HTTP-based content distribution network for repository contents.
  - Single, shared repository at GOC. Good for VOs with little / no support and OSG-internal activities.
  - Key-signing infrastructure for VO-hosted repositories. Good for VOs with active support teams (such as FIFE).
  - OASIS provides an install-once, read-almost-everywhere semantics for VO software, **when needed**. To maximize portability, users & VOs are encouraged to use simpler techniques (HTCondor file transfer) where applicable.
- OSG Storage is a new extension for OASIS. Allows VOs utilizing StashCache to provide a POSIX interface to StashCache.

HTCondor

- HTCondor now provides the base for our information system, CE, OSG VO-hosted service, and glideinWMS service.
    - Having a common base software stack allows us to concentrate our expertise and minimize our dependencies on external teams.
- To some extent, there's a never-ending treadmill of needed scalability improvements and, for the CE, improvements of the batch system integration.
    - This is mostly delivered by the HTCondor Flightworthy team.
- OSG's leadership in this area is reflected in increased partnerships with the European HTCondor community.
    - CERN is steadily migrating from LSF / CREAM to HTCondor / HTCondor-CE. This means LHC VOs must maintain a high level of compatible with these shared components!

GRÅCC

- Our accounting system, Gratia, has been on minimal-maintenance-only for several years. The central Gratia collector has reached a breaking point.

- We are integrating a new service, GRÅCC (pronounced “grok”), that reuses many standard components that are currently Gratia code:

    - **RabbitMQ** for message distribution.

    - **ElasticSearch** for the backend database.

    - **LogStash** for uploading records to the database.

    - **Grafana** / **Kibana** for analytics and visualization.

- On top of these three components, we have various integration scripts to transform & replay this data. It’s being run as a new service, not a software product.

- Importantly, this allows us to re-route records to alternate backends (such as XDMod). Goal is to have each piece of functionality be pluggable: never again replace all at once!

- **Status**: Basic functionality has been demonstrated. Between here and September, plan is to flesh out more functionality and integrate with various accounting scripts (e.g., uploading to WLCG). *Goal is to have the option of turning off Gratia by December 31, 2016.*

Future-looking Projects

- Improved isolation of payloads. Investigating use of *singularity*, a software project from LBL, to provide the same level of isolation as glexec *without x509 certificates*.
    - On future platforms (RHEL8), this can be done in completely unprivileged mode. Isolation could be done across the OSG with **no site support necessary**!
- Modernize the authentication / authorization infrastructure. The software (GUMS, edg-mkgridmap, VOMS-Admin) and processes (authz/authn template generation) are nearly obsolete.
    - No current activities beyond planning. **I don't expect any software we use today for auth{z,n} to be used in 5 years**.

The Next Five Years

- The next five years hold many challenges:
  - Finish off the many ongoing transitions!
  - Improve integration with non-OSG resources: HPC facilities, non-WLCG sites, commercial clouds.
  - Slowly expand our storage capabilities from the current OSG-Storage offerings. Particularly, **we need an external software partner** if we want revolutionary work here.
  - Increasingly decouple our user authentication & authorization scheme from the “traditional grid model”.