OSG Technology Area

Brian Bockelman

OSG Internal Review

Mission

- Improve the understanding of OSG processes and technologies and improve the state-of-the-art of Distributed High Throughput Computing.

- Understand the technology landscape surrounding OSG.

- Provide the best possible technology stack to the OSG and the OSG Production grid.

OSG Technology

- OSG Technology is split into three sub-areas:
  - **Architecture**: provides broad guidance about how OSG technology direction.
  - **Investigations**: Projects to effect change in the software stack.
  - **Software**: Tasked with the producing the OSG Software distribution and managing software lifecycles. *Covered in the next presentation*.

OSG Architecture

- **Primary activities**:
  - **Organize the Blueprint process**: Quarterly meetings which outline how we believe the OSG technology stack should be constructed.
  - **Maintain documents**. Do updates to the Blueprint document as necessary. Publish Blueprint meeting notes.
  - **Communication and Outreach**: Participate in and contribute to the WLCG Grid Deployment Board meetings (this is a technical coordination body for the WLCG), other
    .

OSG Investigations

- **Primary Activities**:
    - *Investigating new software components*: Identify technological components that potentially disrupt the OSG in the medium-term. Understand if and how these would apply to the OSG.
    - *"Expert level" consulting*: Provide intellectual contributions or hands-on debugging to other areas or external software.
    - *Targeted software contributions*: Contribute bugfixes and features necessary for OSG to software maintained outside this area.
- This program of work is mostly organized around a series of time-limited projects.
- "Expert level consulting" includes debugging software issues that cannot be solved by either the software team or upstream software.

Effort Allocations - Year2

Sub-Group	Personnel	Institution	FY14 FTE effort	Comments
OSG PI	Livny	UW	0.3	
Management	Bockelman	UNL	0.2	
Architecture	Hover	BNL	0.5	
Investigations	Bockelman	UNL	0.3	
	Caballero	BNL	0.5	
	Zvada	UNL	0.5	Recent hire
	Sfiligoi	UCSD	0.2	
Software	(Detailed in next talk)	UCSD, UW, UNL. BNL	6.6	Excludes 0.5 TBN at UNL

Total effort (minus software): 2.5 FTE

Architecture and Investigations Year 1 Project Summary

- Investigations:
  - HTCondor-CE: Investigated whether HTCondor can be used as a gatekeeper as a part of the CE. Ultimately, decided this would be the basis for the next-gen CE.
  - Information services: Looked into ActiveMQ as a replacement for BDII for information services. Decided not to pursue further.
- Expert-level consulting:
  - HTCondor: Helped design condor_tail, condor_ping, schedd audit log, python bindings, and transfer throttling.
  - GRAM: Served as intermediary and triage for GRAM5 issues between OSG Software and Globus Toolkit.
  - BOSCO: Assisted debugging scalability issues.

Architecture and Investigations Year 1 Project Summary

- Targeted software contributions:
  - JGlobus: Temporarily managed JGlobus project in order to finish SHA-2 upgrade.
  - HTCondor: Contributed python bindings, Globus and Condor-C fixes to support HTCondor-CE, improve packaging of the procd.

Architecture and Investigations Year 2 Summary

- Architecture:
  - Draft OSG Provisioning Services document.
- Investigations:
  - Cross-CE: Have an OSG-run CE which is a virtual pool backed by OSG VO glideins; this was ultimately not selected to continue.
  - APF: See if APF can reasonably launch HTCondor-based glideins.

Architecture and Investigations Year 2 Summary

- Targeted software contributions:
    - HTCondor-CE: Continuing from Year 1, implemented HTCondor bug fixes and performed validation for PBS backends.
    - HTCondor: shared port scalability, condor_q scalability, TCP keepalive improvements, remote history, scalable classad updates from worker nodes.
- Expert-level consulting:
    - Improve scalability of HTCondor-G/GRAM5 through finding previously unknown system tuning issues.

Stakeholder Impact

- OSG Technology Area (excluding software) feeds mostly into other OSG areas (primarily Software and Operations), which are then released to stakeholders.
    - Other than documents we write, there is no direct “stakeholder deliverable”.
- In the next two slides, I highlight items that have almost direct impact and indirect impact.

Direct Impact to Stakeholders

Item	DOE or NSF Impact	Affected Group	Description
JGlobus	Both	bestman2-based sites	Improved JGlobus to be SHA-2 compliant
OASIS	Both	Small sites, OSG VO, and Intensity Frontier	OASIS provides a mechanism
Debugging	DOE	USLHC Tier-2 sites	Solve GRAM5 issues on large sites
Debugging	DOE	USLHC Tier-2 sites	Expert level debugging.

Indirect Impact

- **HTCondor-CE**: Improves scalability of the CE and ease-of-operations. More importantly, decreases our reliance on external projects and hence reduces risk.

- **Information Services**: The decision to not support GLUE2 puts us at a disadvantage in terms of interoperability with WLCG, especially pieces taken from EGI.

  - However, this saves us significant effort in implementation.

Key Findings

- We still struggle with finding proper avenues for disseminating OSG Blueprint (both document and knowledge put together during the meetings).

- We feel OSG provides a reasonably holistic approach for workflow management in combination of HTCondor and glideinWMS. However, despite several attempts, there is no holistic approach to data management.

  - Data management tends to be done with a series of best practices.

- In the first two years, significant technology improvements mostly impacted sites on the production grid.

Upcoming Challenges and Plans

- In Year 3, we need to pivot from being “site focused” to focus on the issues facing users.
    - At last Blueprint meeting, identified work items related to improving the User Support pool.
    - Would like to help User Support design a OSG Submit Host package.
- Lacking a ‘magic bullet’ for data management, we hope to continue evolutionary improvements on current best practices through targeted contributions to external software.