

Technology Investigations Update

Brian Bockelman
OSG Staff Retreat, May 2015

Staffing

- Would like to continue previous year's trend toward keeping effort in Software area and pulling together effort for time-limited projects.
- Exceptions:
 - **Zvada**: Take advantage of site-locality.
 - **Dykstra**: Already assigned to special projects.


Name	Institution	Level	Notes
Bockelman	UNL	0.5	
Zvada	UNL	0.5	Xrootd expert
Hover	BNL	0.25	Propose move to Software
Caballero	BNL	0.25	Propose move to Software
Dykstra	FNAL	0.25	CVMFS core, StashCache
Illingworth	FNAL	0.5	StashCache

Review from Last Year

- Major items mentioned last year:
 - **Finish HTCondor-CE transition:** Mostly this was about the hand-off to software (plus ce-collector). *Done.*
 - **OASIS:** Enable external repo support; “*Install locally; run globally*”. *Done.*
 - **Common Input Data:** Became StashCache in January. *In progress.*

Vision for OSG TI

- **Jobs and Resources:**
HTCondor.
- **Software Distribution:**
OASIS.
- **Data Distribution:**
StashCache.



During Year 4, I want to strengthen and refine these three pillars!

Jobs and Resources: HTCondor

- Includes job execution (HTCondor), provisioning element (HTCondor-CE), and resource provisioning (GlideinWMS).
- Goals for the year:
 - HTCondor: Continue to manage steady stream of small improvements.
 - See: <https://htcondor-wiki.cs.wisc.edu/index.cgi/tktview?tn=4882>
 - Having HTCondor team own pilot creation.
 - HTCondor-CE: Focus on improving non-HTCondor-backend support
 - <https://htcondor-wiki.cs.wisc.edu/index.cgi/tktview?tn=5045>
 - GlideinWMS: Monitoring items (holdover from last year).
 - Investigate allocation-based provisioning.

osg-ce-bosco*

- **Goal:** Provide a BOSCO-based backend to the htcondor-ce.
- A new sub-package of the osg-ce, this will focus on allowing User Support to quickly deploy new endpoints. Minimize number of “knobs” required.
- Development needed - particularly for Gratia accounting through BOSCO.

* still working on the name.

Software Distribution: OASIS

- Got an OASIS release out the door! All the hallmarks of “first release” - probably 4-5 issues that were overlooked during testing.
- Goals:
 - Increase monitoring coverage of Stratum-1 and external repo status. Particularly for items missed by last week’s upgrade.
 - Migrate as much as possible from OASIS configuration script to OASIS software. Migrate as much as possible from OASIS software into CVMFS upstream.
 - Engineer an auto-updating client.
- Will be working to get this enabled at HPC sites (particularly, NERSC and SDSC).

Data Distribution: StashCache

- **StashCache**: A set of strategically-placed caches throughout the OSG available for VOs to distribute data.
 - Currently, data is exported via xrootd protocol using the xrootd server.
- Goals for this year:
 - Open StashCache to three customers: NoVA, OSG-Connect, and OSG-XD.
 - Investigate alternate protocols such as HTTP; look into using condor_cached to manage cache contents.
 - Investigate using StashCache to distribute a CVMFS repository.

Other Items

- Improved payload/pilot isolation (“certificate-less glexec”): Will likely help / advise security on this.
 - Can cover specifics during Mine’s talk.
- **Site authorization cleanup:** Switching default deployment from GUMS to simpler mkgridmap (or variant).
- Network datastore: Believe our piece of this done. Not currently seeing any follow-up necessary from the Technology team.

Top Concerns

- As usual, most of our work depends on external / non-OSG teams for completion.
 - To compensate, will continue quarterly F2F meetings with TimC & TimT to help re-plan and re-allocate effort.
- Seems difficult to estimate “velocity” when keeping most of the effort outside the area.
- Riskiest part of the workplan is StashCache items - only large “new” piece of technology.