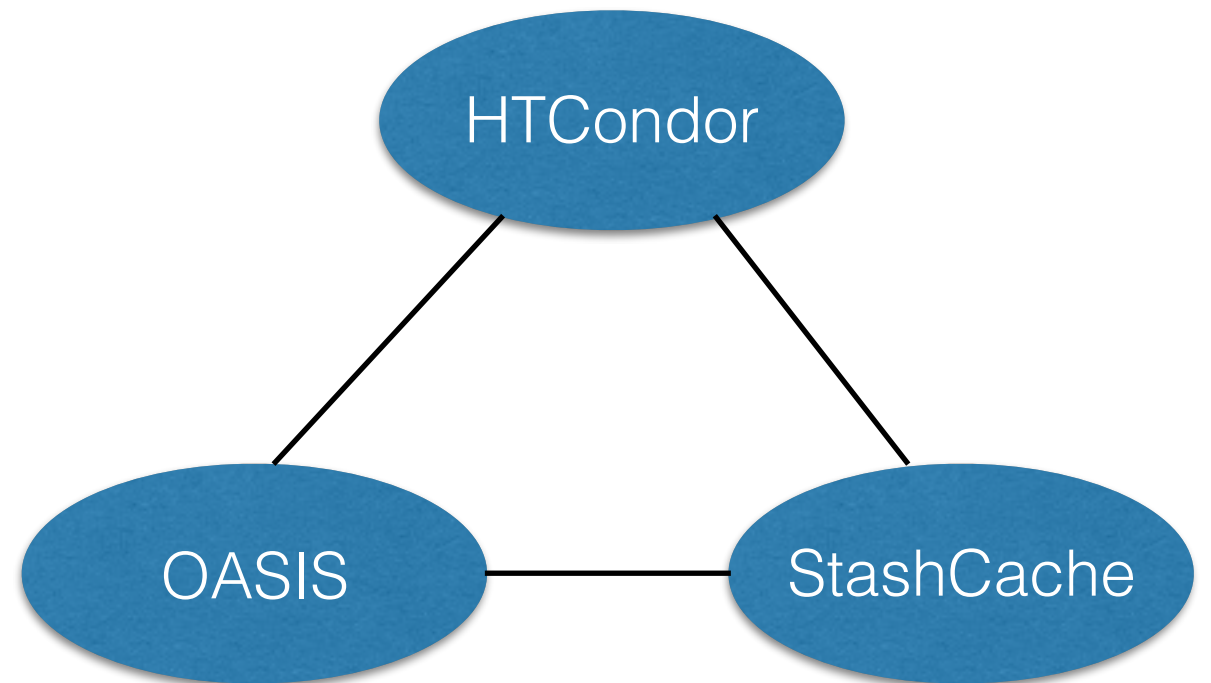


Technology Investigations Update

Brian Bockelman
OSG Staff Retreat, May 2016

Slide from Year 4: 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!

How did we do?

- From the deliverables listed in last year's presentation; items finished:
 - HTCondor-CE: Improve non-HTCondor backend support.
 - First version of htcondor-ce-bosco (with Gratia accounting).
 - Increase monitoring coverage of Stratum-1s.
 - Auto-updating client configuration for OASIS.
 - Have StashCache export HTTP.
 - Use StashCache to serve a CVMFS repository (stash.osgstorage.org, nova.osgstorage.org).
 - Variety of minor HTCondor improvements: <https://htcondor-wiki.cs.wisc.edu/index.cgi/tktview?tn=4882>
- Things partially finished:
 - GlideinWMS: allocation-based provisioning.
 - OASIS available at more HPC sites.
 - Migrate code from OASIS to upstream CVMFS.
 - Improve payload/pilot isolation.
- Things that did not really start:
 - Work with HTCondor team to do their own pilot creation.
 - Site authorization cleanup.

Documentation **Technology**

- Will have a longer discussion in the operations section, not just oriented at technology.
- We are looking at using a common document generation package (mkdocs) and workflow (github-based).
 - I think this workflow better matches the software release process.
- (Short demo to follow if we get projector working...)

Gratia V2 - GRÅCC

- Gratia deploy has hit a critical state.
- Working to break monolithic collector into a smaller standalone pieces.
 - Will keep compatible with existing probes.
 - Communication between standalone services happens via GOC's messaging service.
 - Looking for July milestone to show basic functionality.
- Backend database is ElasticSearch.
 - Will allow all raw records to be kept online and queried.
 - Goal is to extend the record to the full HTCondor ClassAd.
- Resulting service can separate the *raw* records from the *summary* records. This means we can send the either record type to multiple databases - such as XDMOD.

Monitoring

- Per yesterday's discussion, will be looking at revamp of how we do monitoring:
 - Move RSV-like functionality into the base CE (currently in CE-View).
 - Investigate extending our "stashcache" approach - use condor to manage and monitor a transfer daemon - to GridFTP.
 - Current status of services gets sent to the OSG Collector; a monitoring service pushes this to a database.
 - Work with Operations to determine how to display this data.

LIGO

- LIGO was certainly the surprise of 2015. They forced movement on the following technology topics:
 - Utilizing XSEDE project IDs (access to Stampede).
 - Access to VIRGO and LDG sites.
 - Actual RHEL7 use.
 - Data access.
 - (Hopefully soon) identity management.
- Other VOs (IceCube, SBGrid) are starting to push us on GPUs.

CVMFS & StashCache

- Working on providing script to bootstrap trust roots (CAs, CRLs, VOMS).
 - Bootstrapping is straightforward; determining semantics on failure appears difficult.
 - This is our mechanism to get out of the business of distributing CRLs.
- Push upstream to incorporate better support for OASIS in CVMFS. Otherwise, would like to keep it a quiet year for OASIS.
- Deploy and refine the “secure CVMFS” features to meet LIGO’s needs.
 - Push the data federation features to production.

Anti-Browser Campaign

- I'd like to start a campaign:
 - ***No user should need to load a certificate into a browser for a OSG website.***
- It's a sad indictment of our project that users ***don't need certificates*** for:
 - Running jobs.
 - Transferring files.
- But they ***need certificates for:***
 - Joining a VO (FNAL VOs)
 - Reading twiki pages
 - (site / VO admins) Making changes in OIM.
 - Visiting <https://www.opensciencegrid.org>.

CILogon 2.0

- Proposal:
 - With each OIM user, associate a *federated identity* (brian.bockelman@example.com). No passwords!
 - Use existing CILogon integrations to manage authentication with services.
 - Remove client certificates from wiki and OIM by December 2017?
- Long-term: CILogon 2.0 will integrate COmanage:
 - Self-service group management for VOs.
 - SSH key management.
 - Identity management workflows: registration approvals, expiring users, etc.
 - All offered as a software service.

Security Infrastructure

- I remain unhappy with most aspects of our authentication / authorization infrastructure.
- I think everything from software used, to documentation, to generation of default templates needs to be completely overhauled.
- This area does not have the effort to do this during this year.

SRM

- We support SRM endpoints by maintaining the BestMan implementation.
- Maintenance is a significant cost.
- Stakeholders are effectively just USATLAS / USCMS.
- What's a viable mechanism to deprecate this in 2017? I feel I understand the USCMS aspect but not USATLAS side.

Containers

- A few small items on containers:
 - At Nebraska, we're working to allow users to submit pilots requesting arbitrary containers (using HTCondor-CE and Docker Universe).
 - We currently distribute "the OSG" via RPMs, tarballs, and CVMFS. I'd like to add Docker containers as another mechanism.
 - Will keep track of evolution of "unprivileged containers" in RHEL7. Would like understand their capabilities and be aware once they can provide the isolation features we are interested in.

Staffing

Notes:

- Sum is 2.5 FTE.
- Zvada is moving some effort to gWMS Ops
- Edgar's LIGO and IceCube time difficult to classify

Name	Institution	Level	Notes
Bockelman	UNL	0.5	
Zvada	UNL	0.25	Xrootd expert
Hover	BNL	0.10	Blueprints
Caballero	BNL	0.40	CVMFS Stratum-1
Dykstra	FNAL	0.25	CVMFS core, StashCache
Weitzel	UNL	0.5	BOSCO, GRACC
Fajardo	UCSD	0.5	gWMS LIGO