

OSG Operations

OSG Planning Retreat 2018
Jeff Dost

Effort

- FKW effort numbers:

IRIS-HEP	Nebraska	Operations			Derek Weitzel		0.50
CESER	UCSD	Operations			Diego Davila		0.30
OSG N5Y	UCSD	Operations			Edgar Fajardo		0.30
OSG N5Y	UCSD	Operations			Jeff Dost		0.30
OSG N5Y	Nebraska	Operations			John Thiltges		0.50
IRIS-HEP	Nebraska	Operations			Marian Zvada		0.50
OSG N5Y	Nebraska	Operations			Marian Zvada		0.25
OSG-NP	NERSC	Operations			TBD	X	0.50
TNRP	UCSD	Operations			TBD "Edgar"	X	0.20
OSG N5Y	UChicago	Operations			TBD "Suchandra"	X	0.70

– 4.05 FTE Operations total

Service Breakdown

- Before we can accurately assign efforts we needed to survey all of the OSG services to understand what exists
- Operations team has started cataloging all of the services in a spreadsheet:
 - <https://docs.google.com/spreadsheets/d/1cIu0YcNAT8Kiy5u2xBz8ZzG2zImVK3pRX9iFHeglJuk/edit?usp=sharing>

Understanding Ownership

- 4 types of services in OSG:
 1. OSG owned and operated
 2. OSG owned, operated by non-OSG staff
 3. externally owned, OSG staff assists in operations
 4. externally owned, no OSG staff involved (e.g. site services: CEs, XRD Caches, etc)
- I claim OSG Operations is responsible for 1-3
- I also claim before transition, OSG Ops had a very limited picture of the services external to IU (2 and 3) and the effort to run them
 - See slide 1 of OPS report from 2017 retreat
 - <https://indico.fnal.gov/event/14739/contribution/11/material/slides/0.pdf>

Breaking down Effort

- Ideally I would know how much FTE it takes for every service
 - But our inventory so far has ~40 services!
 - Getting the numbers is still **work in progress**
- Factory Ops is an exception because it is strictly shift based
- Next slide shows Factory Ops, then the rest of effort breakdowns are by person

Glidein Factory Effort

- Breakdown:
 - Marian: 0.25 FTE OSG-CORE
 - Jeff: 0.6 FTE external (CMS)
 - Marco: 0.5 FTE external (CMS)
 - Edita*: 0.5 FTE external (CMS)
- Total: 1.85 FTE; 0.25 OSG-CORE; 1.6 external
- Shift based:

	M	T	W	Th	F
EUR	MARCO	EDITA	EDITA	MARCO	MARCO / EDITA ALT
US	jeff	JEFF	JEFF	marian	MARIAN

– UPPERCASE: active shift

– lowercase: cold shift

* Edita is CERN CAT-A, can only do CMS related work

Effort by person

- First, we go through FKW's list
- Then I move into people who operate OSG services that I think are missing

Jeff Dost

- 0.30 FTE OSG-core:
 - Operations AC Coordinator
- 0.60 FTE external (CMS):
 - Factory ops discussed previously

Derek Weitzel

- 0.50 FTE IRIS-HEP
- XRootD Monitor Collector
- GRACC:
 - Message Broker
- CE Collector

John Thiltges

- 0.50 FTE OSG-core
- Stash Cache
 - Origin: LIGO (proprietary data)
 - * externally owned but operated by OSG staff

* I believe service list for John is incomplete, more on this next slide

Marian Zvada

- 0.50 FTE IRIS-HEP:

- OASIS:

- HCC CVMFS repo
- OASIS

- Stash Cache:

- Top level redirector

- GridCF Repo

- GRACC:

- frontend and datanodes
- APEL accounting service

- PerfSONAR:

- perfsonar collectors

- Web Pages: Display, Map

- 0.25 FTE OSG-core:

- Glidein Factory

* not all of Marian's services under IRIS-HEP look like they should be there. Could Glidein Factory be called IRIS-HEP instead to open up another 0.25 of OSG-core? Or should some of these services be offloaded to John Thiltges?

Edgar Fajardo

- 0.30 FTE OSG-core
- Glidein Frontends:
 - Gluex
 - Ligo
 - UCLHC
 - Glow (provides support as needed)
- Stash Cache
 - Origin: Flatiron Institute
 - * externally owned but operated by OSG staff
 - Caches: Chicago POP, Kansas City POP, Manhattan POP

People not in FKW Ops Table

Mats Rynge

- ?? FTE ?? funding
- Glidein Frontends:
 - OSG Flock
 - XD Login
- What other effort? any OSG Connect / User Infrastructure beyond those listed above?

UChicago

- ? FTE external (UChicago Tier2 / SLATE / VC3)
- User Infrastructure
 - Freshdesk * in flux
 - Globus endpoint
 - NFS for /home
 - OSG Connect
 - Stash Cache:
 - Origin: OSG Connect
 - Hosted CEs (Infrastructure only)

* I realize some of this will be picked up by the TBD “Suchandra” hire at UChicago. how will the Chicago Tier2 role / effort change once that happens?

Note on Hosted CEs

- Was managed by UChicago
- Jeff and Marco Mascheroni agreed to take temporary ownership until new effort in UChicago can pick them back up
- Small % effort “borrowed” from Jeff and Marco's Factory ops combined 1.1 FTE

Brian Lin

- ?? FTE ?? funding
- Topology (?)
- Software infrastructure:
 - Jira
 - Software integration testbed
- Web pages:
 - DNS (cloudflare / godaddy)
 - GitHub Doc Pages

Other Software team ops

- Carl Edquist: ?? FE ?? funding
 - OSG Repo
- Matayas Selmecci: ?? FE ?? funding
 - Koji

Fermilab ops

- ?? FTE (external)
- Mailing lists
- DocDB
- Indico

Shawn McKee?

- I'm not clear on the potential overlap between network pipeline / personar team and operations team

Accomplishments

(since I started as OPS AC in June 2018)

- Re established weekly ops meetings:
 - <http://opensciencegrid.org/operations/#meeting-minutes>
When: Mondays 1:00 pm Central
URL: <https://unl.zoom.us/j/183382852>
Phone: +1 669 900 6833 or +1 408 638 0968 or +1 646 876 9923
Meeting ID: 183 382 852
- New email list for ops team: operations@opensciencegrid.org
- Started service catalog already discussed:
 - <https://docs.google.com/spreadsheets/d/1clu0YcNAT8Kiy5u2xBz8ZzG2zImVK3pRX9iFHegljUk/edit?usp=sharing>
- UNL GRACC disaster recovery exercise

Concerns

- Getting the big picture of all the OSG services right, and understanding the effort required for each
 - (already discussed in detail)
- Effort required to keep Topology information up to date
 - Had one joint meeting with Software team, need to have more to continue understanding the effort involved
- Our ticketing system that disappeared with IU, and whether Freshdesk is the right replacement

Freshdesk, ctd

- Works great for one way support (think OSG Connect):
 - customer → service
- Ops needs multi-directional support, to support things like:
 - VO ↔ factory ops ↔ site
 - Adding external help is the norm (OSG software, gWMS, HTCondor and other dev teams)
- Ticket threads other than Security incidents should be readable by everyone, not just ticket agents

Ticket triage

- Not just Freshdesk issue, used to have dedicated IU staff effort to triage ops tickets
- For now, operators are forced to triage their own tickets (and not doing a great job)
- We should be focusing on the problems contained in the ticket, not on tickets expiring / violating arbitrary SLA rules
 - Having dedicated staff for this reduces overhead for the people working on the issues

Plans for Coming Year

Internet2 and StashCache

- Complete integration of Internet2 caches with existing stash infrastructure
- Begin adapting kubernetes orchestration as official operations service to manage caches
- Dedicated effort once TBD's are hired:

CESER	UCSD	Operations			Diego Davila		0.30
OSG-NP	NERSC	Operations			TBD	X	0.50
TNRP	UCSD	Operations			TBD "Edgar"	X	0.20

New Singapore site

- Involves Operations for:
 - Hosted CE deployment and maintenance
 - Cache deployment using Kubernetes
- Some combination of UChicago and TBD effort
- Would be a good proof of concept of Kubernetes infrastructure proposed in prev slide

Validate GRACC reporting

- Some combination of effort from Marian and Derek to help understand discrepancies
- Come up with some kind of automated checks for data integrity / misconfigured probes
- Factory ops to help cross reference data mined from pilot logs

IRIS-HEP deliverables

1	Timeframe (months)	Description	WBS x-ref	Risk Register	OSG Owner
9	3–6	Update OSG Service Level Agreements (SLAs) with the LHC experiments	7.1	R.7.1	Operations
10	3–6	Review the operations transition from OSG to OSG-LHC for completeness	7.1, 7.4	R.7.1	Operations
18	6–12	Report to OSG Council on the completed review of operations transition at the OSG AHM Council meeting on 21 March 2019	7.4, 7.5	—	Operations
19	Deliverable	6–12 Summarize past year of the monthly accounting reporting to WLCG, including all issues with reporting and their resolutions	7.5	-	Operations
29	Deliverable	12–18 Evaluate operations of DOMA services by US ATLAS and US CMS for common use	7.2, 7.5		Operations

- Review transition – in particular, any balls dropped?
 - The only major one in my opinion is the ticketing system situation already discussed
 - The other action item is to better understand effort involved to maintain Topology as replacement for OIM