U.S. CMS Operations Program

CMS and OSG and beyond

Ken Bloom

OSG Council

4 October 2017

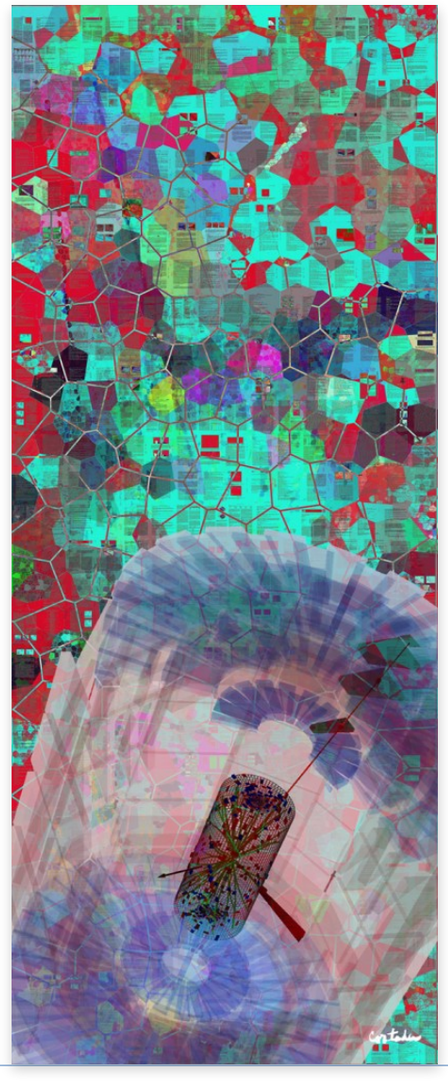


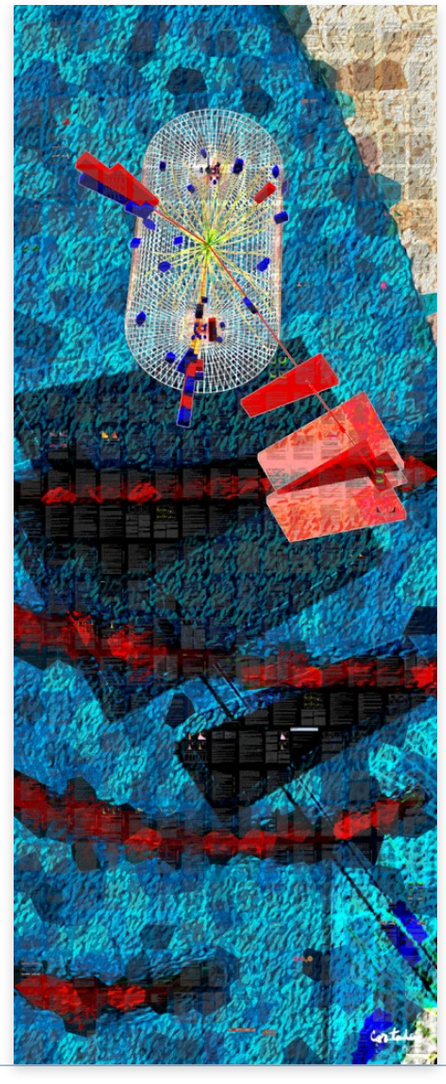


U.S. CMS and OSG

- OSG has worked very well for U.S. CMS!
  - And thus it is important that its services are maintained
- CMS has made excellent use of resources on the OSG
  - Now ~40M hours/month, dwarfing current annual allocations at HPC centers
- OSG provides an integrated software infrastructure, support and technical evolution for many U.S. CMS activities, e.g. T2 and T3 centers
  - Software support, grid operations, network support, access to resources "across the campus," support for T3s, CMS Connect service, T3 in a box, accessing non-owned resources and HPC, etc.
- OSG, HTCondor, U.S. CMS and U.S. ATLAS form an ecosystem of organizations that advances distributed high-throughput computing (dHTC) to meet the needs of the LHC scientific program through continual development efforts
  - Fermilab, Nebraska, UCSD, Wisconsin are institutes with strong involvement in both U.S. CMS and OSG

Examples of past collaboration

- Many benefits from R&D and joint projects involving OSG!
- HTCondor and OSG worked closely with global CMS, CERN, BNL and FNAL to improve HTCondor scalability and robustness
  - Led to adoption of HTCondor as the default batch system for most T1s worldwide and CERN
- OSG and HTCondor developed the HTCondor-CE (Compute Element = job submission interface to site) as the replacement to GRAM and CREAM
  - The formerly-dominant globus-based technologies in both US and EU lost funded support years ago
  - CERN has adopted HTCondor and HTCondor-CE
- OSG now working with U.S. ATLAS and U.S. CMS on drastic simplification of the infrastructure software
  - BDII, BestMan, SRM, GRAM, Gratia replacements begun/completed
  - gLexec, VOMS-Admin, GUMS, RSV transitions still pending
  - Has reduced effort required from T2 and T3 sites to run services
  - OSG provides intellectual leadership WLCG-wide on future-facing areas such as adoption of containers
  - Leads WLCG in retiring niche “grid technologies”, making room for modern techniques with the same principled approach

Cataloging critical services for OSG-LHC

- Working with OSG and U.S. ATLAS, prepared a thorough catalog of OSG services that are critical to CMS, and FTE required to support them, in conjunction with OSG ET
  - Note: this discussion only covers operational services, and not the R&D pieces of OSG that have driven e.g. projects on previous slide
- Extended discussion with U.S. ATLAS S&C leaders to agree on definitions of services and understand where services are shared between the experiments, in full or part
- Executive summary: U.S. CMS requires 9.4 FTE to support the critical services, of whom 5.2 can be shared with U.S. ATLAS
  - This represents current services; strictly restricting this exercise to the currently existing scope

Detailed assessment of work breakdown

- Have estimates of effort required on individual tasks at the level of 0.1 FTE
- Agreed to by U.S. CMS and U.S. ATLAS
- Also have some understanding of current key personnel working on each task

			Sharing assessment		
			Shared FTE 2019++	ATLAS FTE 2019++	CMS FTE 2019++
Infrastructure Software maintenance and integration		HTCondor	0.00	0.00	0.10
Infrastructure Software maintenance and integration	Compute Element (CE)	HTCondor-CE	0.25	0.00	0.00
Infrastructure Software maintenance and integration	CE	SSH-based Compute Elements (HTCondor-CE-BOSCO)	0.25	0.00	0.00
Infrastructure Software maintenance and integration	Storage Element (SE)	HDFS	0.00	0.00	0.50
Infrastructure Software maintenance and integration	SE	dcache	0.00	0.25	0.00
Infrastructure Software maintenance and integration	SE	Xrootd including caches	0.25	0.25	0.25
Infrastructure Software maintenance and integration	SE	Object Store	0.00	0.25	0.00
Infrastructure Software maintenance and integration	SE	Load-balanced GridFTP / Storage Element	0.10	0.00	0.00
Infrastructure Software maintenance and integration	Software Distribution	Squid	0.10	0.00	0.00
Infrastructure Software maintenance and integration	Software Distribution	CVMFS server	0.10	0.00	0.15
Infrastructure Software maintenance and integration	Worker Node	Singularity	0.10	0.00	0.00
Infrastructure Software maintenance and integration	Worker Node	Singularity image(s) including worker node client	0.00	0.10	0.10
Infrastructure Software maintenance and integration	Connect services to NSF HPC and university campus clusters	ATLAS Connect, CMS Connect	0.50	0.00	0.25
Infrastructure Software maintenance and integration		T3-in-a-box	0.00	0.00	0.20
Infrastructure Software maintenance and integration		Parrot	0.00	0.00	0.10
Infrastructure Software maintenance and integration		Network services	0.20	0.00	0.00
Infrastructure Software maintenance and integration		End-to-End testing and release of data transfer components, worker node components, etc, to maintain coherence between middleware/software components	1.00	0.00	0.00
Total			**2.85**	**0.85**	**1.65**
CVMFS service operation		Stratum-1	0.00	0.10	0.20
CVMFS service operation		central Squid fallback proxies	0.00	0.10	0.20
CVMFS service operation		Maintain cvmfs key signing, CVMFS Stratum-1 monitoring (OSG operates 2 Stratum-1s; BNL and FNAL Stratum-1s get configs from OSG)	0.10	0.00	0.00
Total			**0.10**	**0.20**	**0.40**
Accounting, Registration, Monitoring		GRACC (database and dashboards)	0.10	0.00	0.00
Accounting, Registration, Monitoring		myOSG / OIM (REBUS, network meshes, and authoritative naming for WLCG)	0.10	0.00	0.00
Accounting, Registration, Monitoring		transition to AGIS/CRIC	0.00	0.20	0.00
Accounting, Registration, Monitoring		Batch system reporting to WLCG (APEL reports)	0.10	0.15	0.15
Total			**0.30**	**0.35**	**0.15**
Job submission infrastructure operation		glideinWMS Factories	0.00	0.00	0.50
Job submission infrastructure operation		PanDA pilot factories	0.00	0.50	0.00
Job submission infrastructure operation		CMS Connect login host and associated services.	0.00	0.00	0.50
Job submission infrastructure operation		AutoPilotFactories	0.00	0.25	0.00
Job submission infrastructure operation		Harvester	0.00	0.75	0.00
Total			**0.00**	**1.50**	**1.00**
Cybersecurity infrastructure	Security	OSG CA (Host certificate CA/RA, distribution and updates, PKI tools). Currently OSG provides "host certs" (user certs can come from CERN). We need some provider for host certs that will work with our global grid use-cases. Incommon CA certs, because they are part of IGTF, will work but require InCommon membership. Other options like Lets Encrypt could also provide such certs but we are not sure of how globally trusted they are.	0.00	0.00	0.00
Cybersecurity infrastructure	Security	Software: Globus GSI infrastructure, VOMS trust roots, LCMAPS authorization framework, fetch-crl updater, IGFT CA bundle	0.25	0.00	0.00
Total			**0.25**	**0.00**	**0.00**
Ticketing and front-line support	Security	Security advisories and trusted alerts	0.20	0.00	0.00
Ticketing and front-line support	Security	24/7 CSIRT-esque responses for the most basic level of our computing infrastructure	0.30	0.00	0.00
Ticketing and front-line support		24/7 ticketing support (ticket exchange mechanism with EGI GGUS and XSEDE)	0.00	0.00	0.00
Ticketing and front-line support		Front-line support for grid-related site issues (US ATLAS sites and opportunistic sites, mainly through factory support)	0.00	1.00	1.00
Ticketing and front-line support		Aggregation, archival, analysis and access to network performance data (PerfSonar)	0.20	0.00	0.00
Ticketing and front-line support		Provide ticketing capabilities for software and support	0.50	0.00	0.00
Total			**1.20**	**1.00**	**1.00**
Coordination		Ensure compatibility of software and approaches worldwide	0.10	0.00	0.00

List of critical services

- Infrastructure Software maintenance and integration
  - integration and release of the OSG software stack, including HTCondor-CE, ssh-based CEs, HDFS, Xrootd, storage element, squids, CVMFS, Singularity, CMS Connect, T3 in a box, packaging
- CVMFS service ops for software distribution
  - Stratum 1 operation, squid proxies, key-signing, monitoring
- Accounting, registration, monitoring
  - dashboards, site registration, reports to WLCG
- Job submission infrastructure
  - glideinWMS pilot factories, CMS Connect login host
- Cybersecurity infrastructure
- Ticketing and front-line support
  - ticketing capabilities, support for sites, security advisories/responses, network issues
- Coordination
  - representation to external entities

Summary of FTE by function

Category	Shared with U.S. ATLAS	CMS specific	Total
Infrastructure software maintenance and integration	2.9	1.7	4.5
CVMFS service operation	0.1	0.4	0.5
Accounting, registration, monitoring	0.3	0.2	0.5
Job submission infrastructure operation	0.0	1.0	1.0
Cybersecurity infrastructure	0.3	0.0	0.3
Ticketing and front-line	1.2	1.0	2.2
Coordination	0.5	0.0	0.5
TOTAL	5.2	4.2	9.4

Mapping onto U.S. CMS needs

- Categorizing how these services and this effort support U.S. CMS needs
  - ~2 FTE for OSG “production services” that enable resource sharing for data processing, data management, etc.
    - ★ Strong overlap with other users, both U.S. ATLAS and other elements of the U.S. HEP program, i.e. intensity frontier which is relying on OSG sharing infrastructure for ~50% of their computing resources
  - ~2 FTE for U.S. CMS production job submission infrastructure
    - ★ glide-in factories and site support needed for site interoperation
  - ~5 FTE to enable U.S. CMS T2 and T3 infrastructure
    - ★ Both development of the technology stack and operational support
    - ★ Some easily shared with U.S. ATLAS, some not
- But note again that these 9 FTE do not include OSG technology evaluation and integration efforts that ultimately benefit the U.S. LHC collaborations

Recent events

- December-January: In the course of the U.S. LHC ops programs external review, becomes clear that agencies will not necessarily support OSG in its current form
  - Many questions about what the essential services from OSG to U.S. LHC are, and their costs (hence all the previous slides)
  - Understood by agencies that these services are critical to the LHC science mission, but they might choose to have them provided in a different way
- We have tried to make clear that we would like to maintain OSG in something like its current form
  - OSG services are outside the core expertise of our program
    - ★ We would probably not do as good a job as OSG in providing them
  - Having many communities receive these services via OSG provides efficiencies of scale
    - ★ It would probably be more expensive for us to provide them
  - Any transition would be disruptive and incur short-term costs
  - Could lose shared infrastructure that enables campus participation and our ability to access opportunistic resources
  - Would need a different mechanism to drive technology evolution
    - ★ Model of finding partners to work on projects of common interest has worked well for us and the entire community

"OSG" in the longer term

- OSG services for the LHC experiments are being discussed in the context of a "Scientific Software Innovation Institute" (S2I2)
  - Meant to bring together computer scientists and domain scientists
  - Solicitation would certainly be suitable for the technological work carried out by OSG, not as clear on operational efforts
- There is an active S2I2 Conceptualization grant (PIs Elmer, Neubauer, Sokoloff) that is funding a series of workshops this year targeted at producing a well-defined strategy for R&D on software and computing models for HEP, with HL-LHC as a particular target
  - To prepare for an S2I2 Implementation proposal that would establish the institute
  - U.S. based effort, but working closely with the world-wide Community White Paper (CWP) effort initiated by WLCG, goal is to deliver the CWP this fall
  - Institute could be proposed in 2018 and funded for 2019-23

Can an S2I2 include OSG?

- Exact plan for the institute still in development
  - Original motivation for institute was not OSG continuation; more focused on software sustainability, using new technologies, and in general responding to the strong software and computing challenges provided by the HL-LHC
    - ★ Estimating a ~5-10 gap between naive expectation of HL-LHC computing needs and available hardware budget!
  - However, it is not clear what other options there are with NSF
  - U.S. CMS has tried to make it clear that OSG functions need to be considered as part of the conceptualization process
- The output of the conceptualization effort will be a strategy document that should be the basis of a future NSF solicitation for an actual institute proposal
  - Or maybe multiple solicitations, one of which will be explicitly OSG-like?
  - But given that the institute would be targeted at HL-LHC challenges, NSF support might be for an LHC-OSG, and not “the rest of OSG”
- No solicitation yet, so hard to imagine that funds could be available before 2019

Managing transitions

- This is a known risk within the U.S. CMS Operations Program and we are trying to be prepared for it
- Can probably help manage a gap between OSG funding periods
  - Understood as a potential call on management reserve in 2018
  - Do not want to lose experts whose efforts will be critical to meet the challenges of the HL-LHC in the coming decade
- U.S. CMS will make every effort to make this work, but would need to work with funding agencies and OSG partners

Outlook

- U.S. CMS has derived great benefits from our partnership with OSG, and we want to sustain it
  - We are getting good value out of it, in both operations and technology evolution
  - A transition would be extremely disruptive
  - The “broader impacts” of OSG are compelling
- Given that OSG services are needed for our science mission, we expect that the U.S. LHC collaborations will find a way forward with help from funding agencies
- It is not clear that other organizations will be supported!
- We need a strategy for interaction with the agencies that will provide a sustainable path forward, and we will need to work together to implement it