
SSL

activities and towards year 2

Rob Gardner
Enrico Fermi Institute
University of Chicago



IRIS-HEP EB meeting
6/10/2019





SSL program of work

- WBS 6.1 – creating and operating scalable cyberinfrastructure
- WBS 6.2 – establishign devOps patterns through blueprint meetings & workshops
- WBS 6.3 – participate in functional testing
- WBS 6.4 – provide database services for metrics aggregation and dashboards
- WBS 6.5 – as needed, provide backend cyberinfrastructure for training



WBS 6.1: acquiring & operating SSL cyberinfrastructure



- Recall the SSL has no dedicated equipment
- We'll need to improvise:
 - leveraged – use existing LHC computing facility and other resources for dedicated tests
 - repurpose – adoption from finished projects
 - purchased – smallish cloud resources for development
- Andrew Chien providing CS research cluster from his lab we can repurpose to base platform



River cluster – SSL base platform services



- Repurposed UChicago CS research cluster
- Vintage but nice: (~50)
 - CPU: 2 x Intel Xeon E2650 v3 12-core processor, 2.3GHz, 30MB cache
 - DRAM: 16 x 16GB TruDDR4 Memory 2133MHz, 256GB
 - Disks: 2 x 800GB SATA MLC SSD, 1.6TB
 - 10G NICs
- 2x40 Gbps to SciDMZ
- Rebuild as Kubernetes
- Explore federation to aggregate w/ others



Federated ID access (institutional, CERN account), edge services hosting, Unix account provisioning, LHC software env.



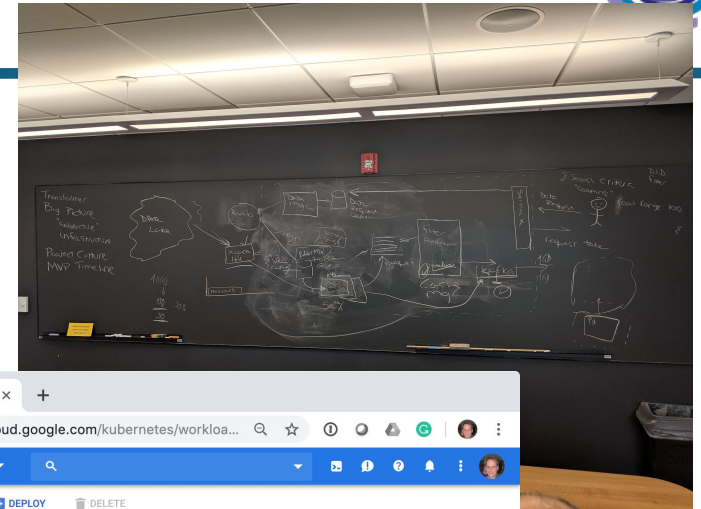
WBS 6.2: DevOps patterns

- Goal here is align with industry devops best practices where its practical.
- We discussed reproducible "deployments" of orchestrated at the Kickoff.
- Will have opportunity to collect input at Analysis Systems-SSL blueprint at NYU

WBS 6.3 Functional Testing



Until SSL base platform operational we can use GKE for testing. Early deployments for iDDS/ServiceX



Name	Location	Cluster size	Total cores	Total memory	Notifications	Labels
analysis	us-central1-a	3	3 vCPUs	11.25 GB		Conn
servicex	us-central1-a	3	3 vCPUs	11.25 GB		Conn

Name	Status	Type	Pods	Namespace	Cluster
atlas-base	OK	Deployment	1/1	servicex	servicex
did-finder	OK	Deployment	1/1	servicex	servicex
invariant-mass-analysis	Running	Pod	1/1	default	servicex
kafkacat	Running	Pod	1/1	kafka	servicex
servicex	OK	Deployment	1/1	servicex	servicex
servicex-kafka	OK	Stateful Set	3/3	kafka	servicex
servicex-kafka-zookeeper	OK	Stateful Set	3/3	kafka	servicex
testclient	Running	Pod	1/1	kafka	servicex
transform-cli	Running	Pod	1/1	servicex	servicex
transformer	Does not have minimum availability	Deployment	1/1	servicex	servicex



WBS 6.5 cyberinfrastructure for training



- <http://codas.slateci.net/>
- JupyterLab machine learning platform for CODAS-HEP last year – available again
- Leverage NSF projects: PRP, CHASE-CI, and SLATE
- Coordinate with SSC for detailed requirements

CODAS PLATFORM
Supporting Computational and Data Science for High Energy Physics

Purpose
A computational platform optimized for machine learning applications, supporting the second school on tools, techniques and methods for Computational and Data Science for High Energy Physics (CoDaS-HEP), 23-27 July, 2018, at Princeton University.

Elements
The platform provides hosted JupyterLab instances for students as well as access to GPU resources of the Pacific Research Platform (PRP) and CHASE-CI. Elements of the SLATE Kubernetes platform and effort from the VC3 project were used in the development.

Links

- [CODAS-HEP.org](https://codas-hep.org)
- [2018 School Program](#)
- [HEP Software Foundation](#)

[GitHub \(this portal\)](#)

[SLATE: Services Layer at the Edge](#)

[Pacific Research Platform](#)



Towards Year 2

from the annual report



WBS 6.1: Creating and operating a scalable cyberinfrastructure platform for IRIS-HEP R&D: the first SSL research cluster will be operational and supporting both DOMA and Analysis Systems prototypes. Year 2 goals include supporting G2.16 Move prototype of Analysis Systems Components to SSL.

WBS 6.2: Through blueprint workshops and prototyping efforts, best practice DevOps patterns will be established and documented. We will continue process begun in Year 1 associated with G2.5, the first IRIS-HEP Blueprint workshop on Analysis Systems R&D on Scalable Platforms at NYU on June 21 & 22, 2019.

WBS 6.3: Participate in functional and scalability testing of data lake and analysis system prototyping efforts. In Year 2, the SSL will work to support DOMA goals:

- **G3.14:** IDDS prototype that can automatically transform / optimize data formats.
- **G3.15:** With SSL, scale test a data lake prototype across multiple sites.

WBS 6.4: Provide database services for metrics aggregation, visualization and analytics for IRIS-HEP testing campaigns on the SSL. In Year 2 the SSL will provide a dashboard for iDDS/ServiceX.

WBS6.5 Provide the Institute with user-facing and back-end cyberinfrastructure components as needed for education, training and outreach events, supporting two training events during the year. In Year 2, this will include CoDaS school and FIRST-HEP events as needed.

