



Platform for Facility R&D

In order to meet the demands of the HL-LHC era and beyond, computing facilities will need to evolve. Here we present the IRIS-HEP Scalable Systems Laboratory, a new platform for developing next-generation facilities.

Industry-Standard Tooling

Flexible infrastructure is needed to support a plethora of workloads beyond traditional batch computing



SSL leverages "Cloud Native" technologies such as Docker, Kubernetes and Helm for managing services and batch workloads.

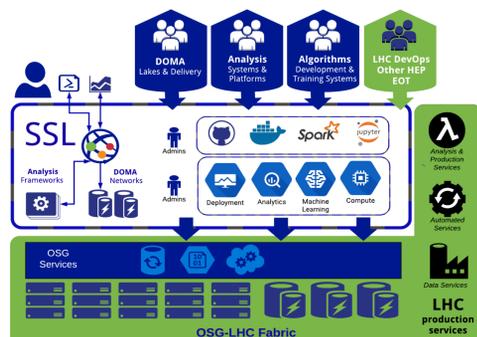


kubernetes

These tools allow users to simply and easily deploy workloads onto SSL-like clusters from the comfort of their own laptop in a declarative way.

The Path to Production

The Scalable Systems Laboratory allows users to deploy novel services, analysis frameworks, data transformers, batch workloads, etc.



IRIS-HEP users can use the SSL to experiment and incubate software ideas, scale their workloads to thousands of cores, and ultimately package their code to run on production resources elsewhere, such as CERN.

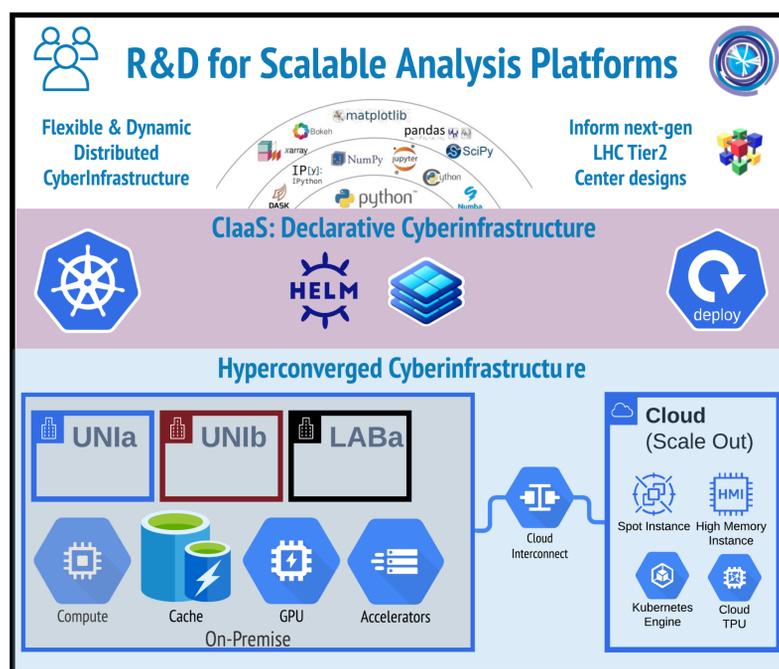
Building a Federated Platform

The SSL leverages expertise from the Pacific Research Platform, SLATE, and WLCG communities



SSL Team Experts are available to..

- Guide users to appropriate sites in SSL, such as the UChicago RIVER cluster or the Pacific Research Platform
- Help other sites get started with standing up their own Kubernetes resources

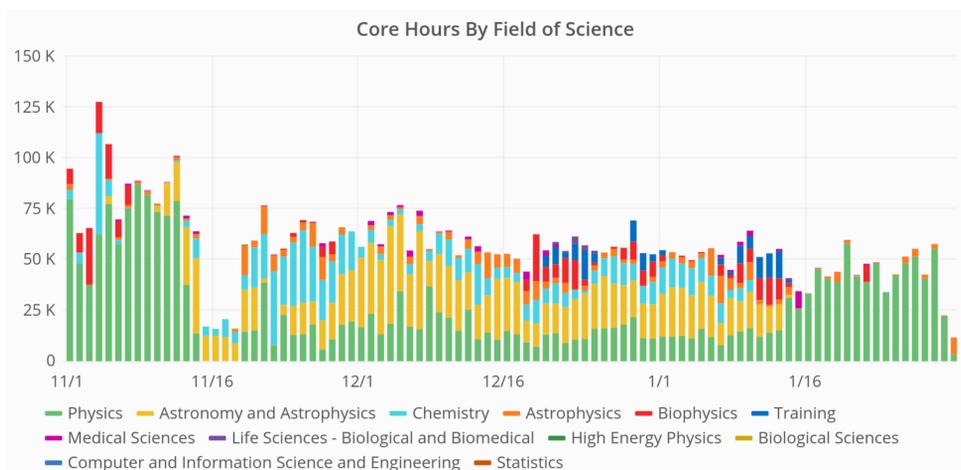


First Piece of the Substrate: RIVER Cluster @ UChicago

- 70 nodes
- 48 cores, 256GB RAM, 2x800GB SSD, 10Gbps
- 2x40Gb to ScienceDMZ
- Kubernetes v1.16

Users and Success Stories

DOMA::ServiceX	Data transformation and delivery service for LHC analyses
Frontier Analytics	Analyze and improve data access patterns for ATLAS Conditions Data
perfSONAR Analytics	Network route visualization based on perfSONAR traces
Parsl / FuncX	Parallel programming in Python, serverless computing with supercomputers
Large-Scale Systems Group @ UChicago	Serverless computing with Kubernetes
DOMA::Skyhook	Programmable storage for databases, scaling Postgres with Ceph object store
REANA	Reusable Analysis Service
CODAS Platform	JupyterLab notebooks, access to GPU resources on the Pacific Research Platform for annual summer CoDaS-HEP training event
SLATE & OSG	Backfilling otherwise unused cycles on SSL with work from the Open Science Grid using the SLATE tools



Nothing wasted:
Nearly 1.5M CPU hours backfilled by OSG in one month!

