SSL

News, Status, Plans

Rob Gardner
Enrico Fermi Institute
University of Chicago

IRIS-HEP Retreat
Fermilab
September 13, 2019

SSL: Path to Production

Provisioning of software **environments** and development tools.

Distributed platforms materialized with tools like containerized edge services.

Integration point with the OSG and LHC experiment services (data, analysis).

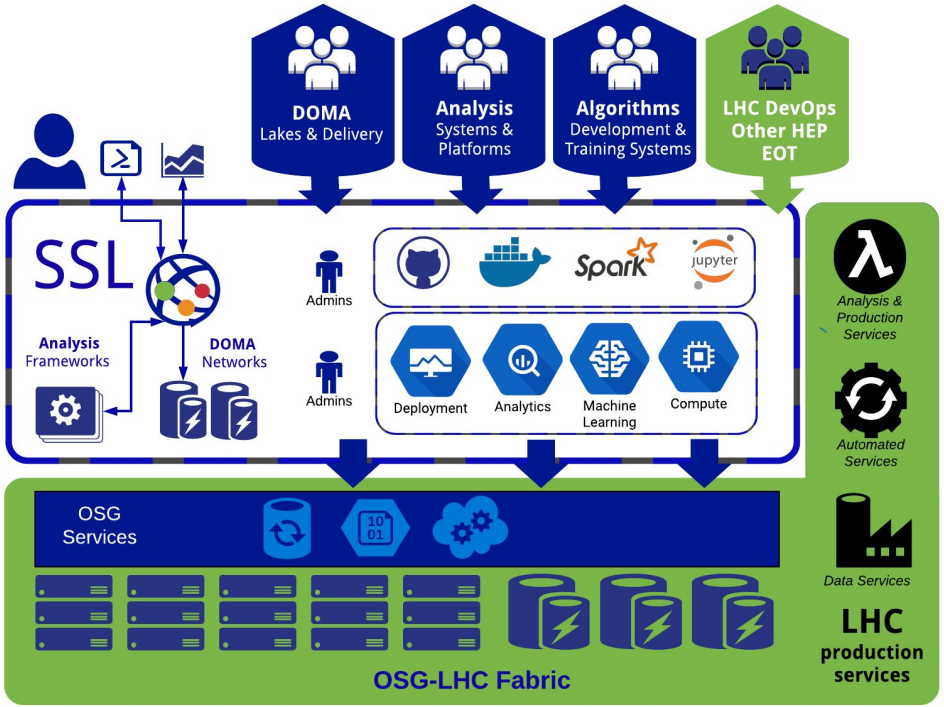

SSL People

- Core group to support base environment
- CS research on accelerated services
- Leverages effort from R&D areas, external projects & research computing staff
- Organizes contributed resources for scalability testing

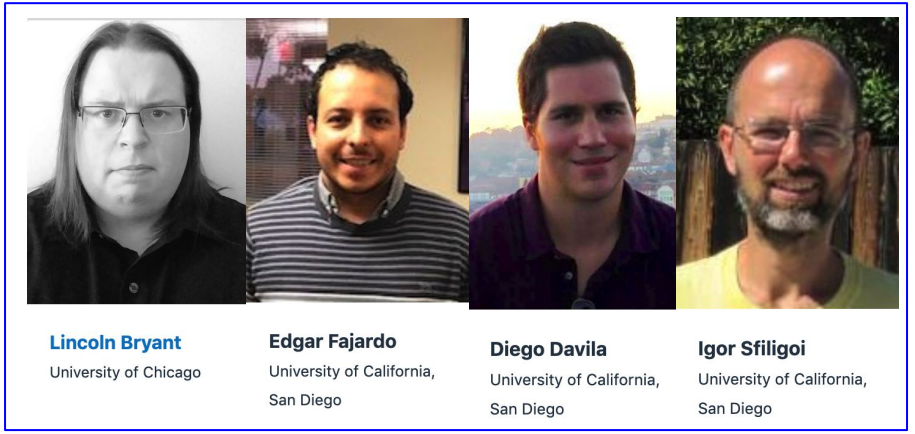

Lincoln Bryant
University of Chicago

Edgar Fajardo
University of California, San Diego

Diego Davila
University of California, San Diego

Igor Sfiligoi
University of California, San Diego

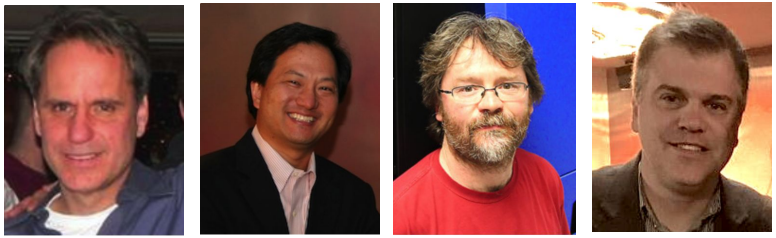

Rob Gardner
University of Chicago

Andrew Chien
University of Chicago

Frank Wuerthwein
University of California, San Diego

Mark Neubauer
University of Illinois at Urbana-Champaign

Analysis Systems & SSL Blueprint

- June 21–22, 2019 @ NYU
- Engagement with AS team, SSL team and representatives from NCSA, SDSC, NYU Research Computing, Google and RedHat (26 participants)
- Key outcomes (doc):
  - **Kubernetes as "common denominator"**
  - **Federated "substrate" project identified**
  - Industry partnerships explored

SSL Successes thus far

ServiceX on SSL

DOMA iDDS development on Google Cloud reproduced on IRIS-HEP SSL with Docker containers and Kubernetes

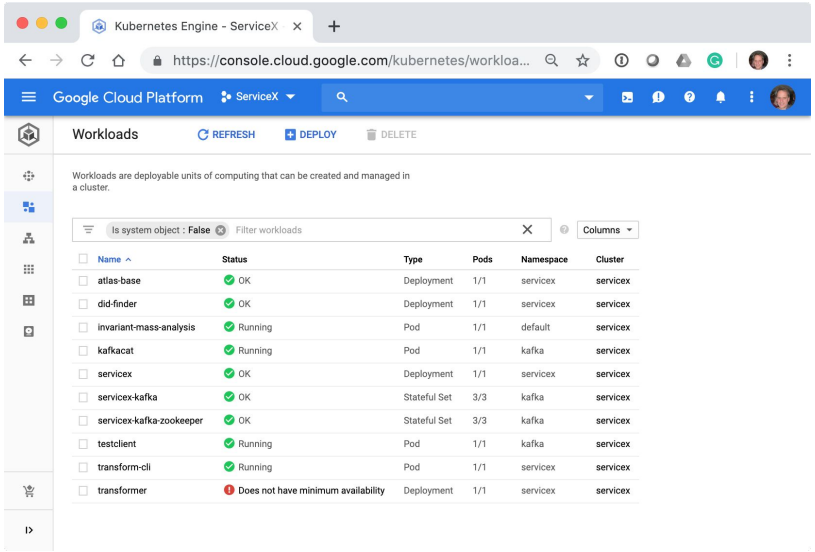

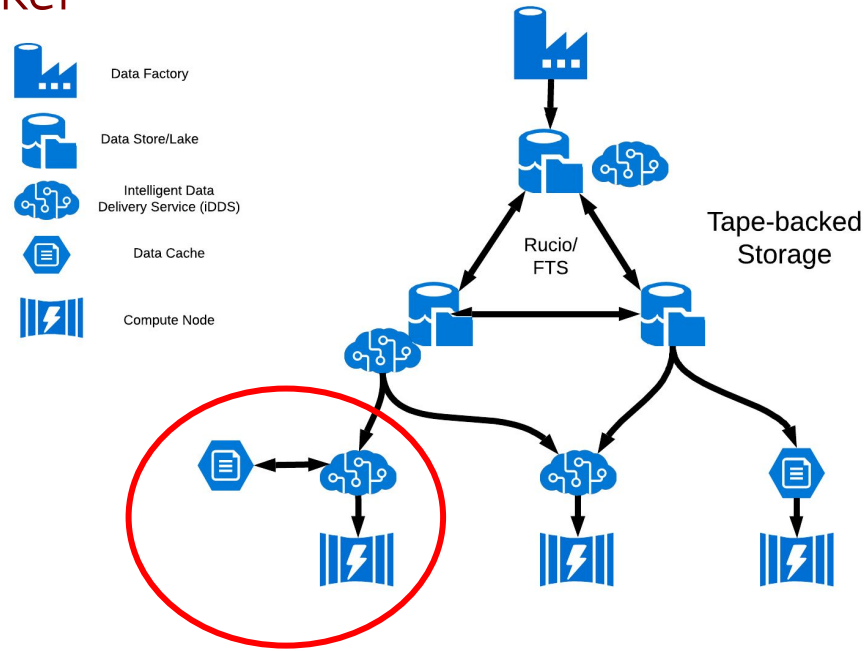

Service Metrics & Dashboards

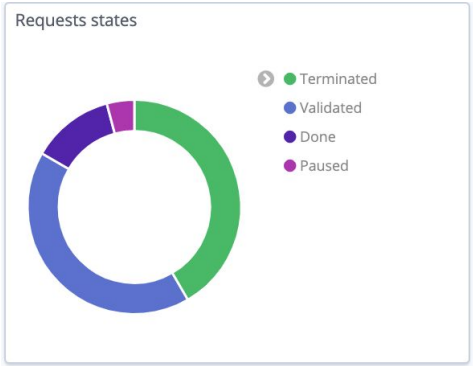

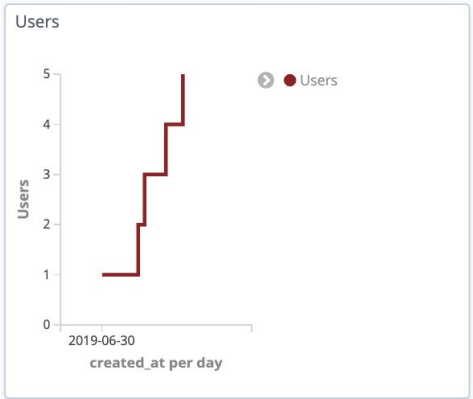

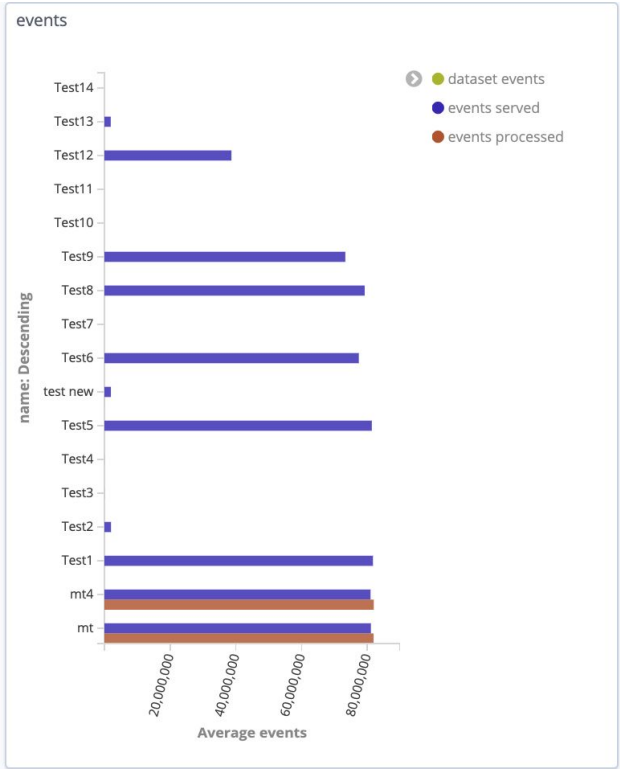

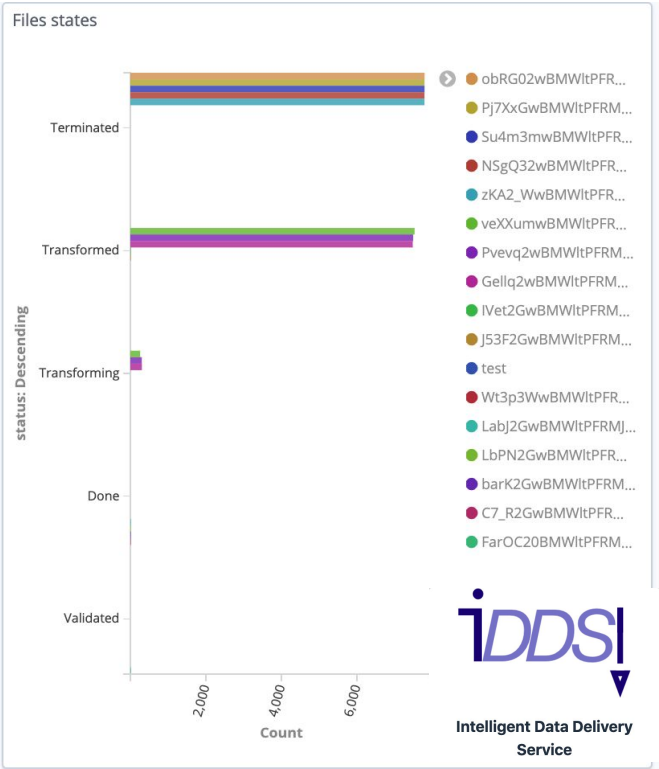

iDDS
Intelligent Data Delivery Service

REANA on SSL

- UChicago CS Student (Neha Lingareddy) was able to deploy the REANA framework on SSL over the summer
- Started with no knowledge of Kubernetes, REANA, etc – was able to produce ATLAS "Recast demo" plots in a couple of days

Neha Lingareddy is a rising Third Year in the College at UChicago majoring in computer science.

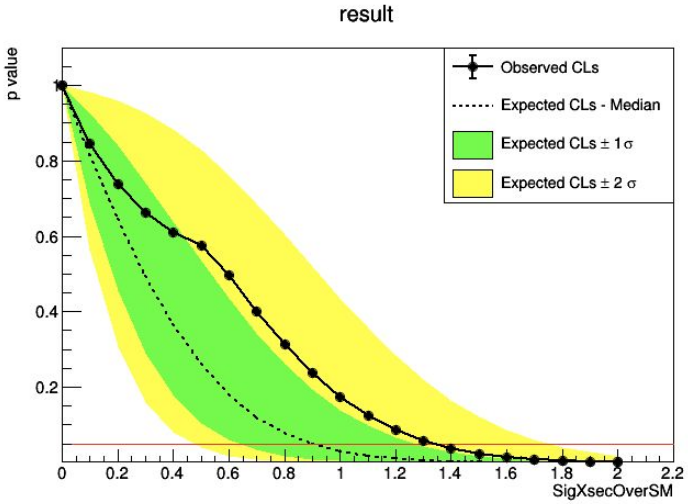

Backfill from OSG

- Deployed OSG-like worker containers via SLATE
  - Most of the usual trimmings: CVMFS, glidein startup scripts (no singularity yet)
  - Flocking from OSG Connect
- Large contributor to OSG!
  - 2.83M core-hours over last 30 days

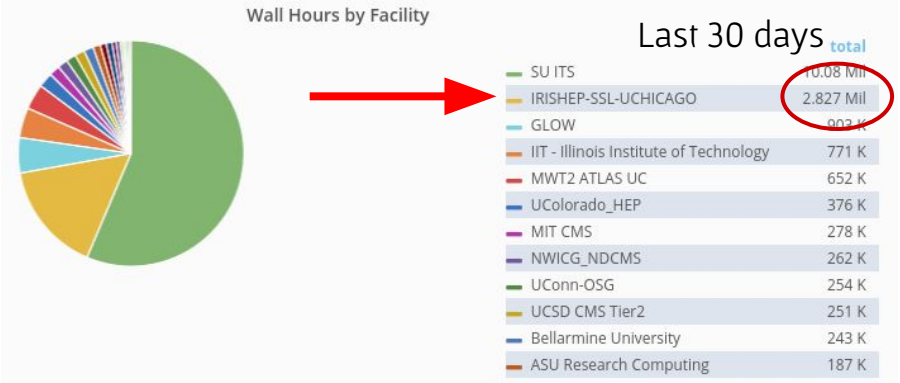

Cyberinfrastructure for training on "SSL"

- JupyterLab machine learning platform for 55 CoDaS-HEP students @ Princeton
- Use Kubernetes from Pacific Research Platform – short term GPU access

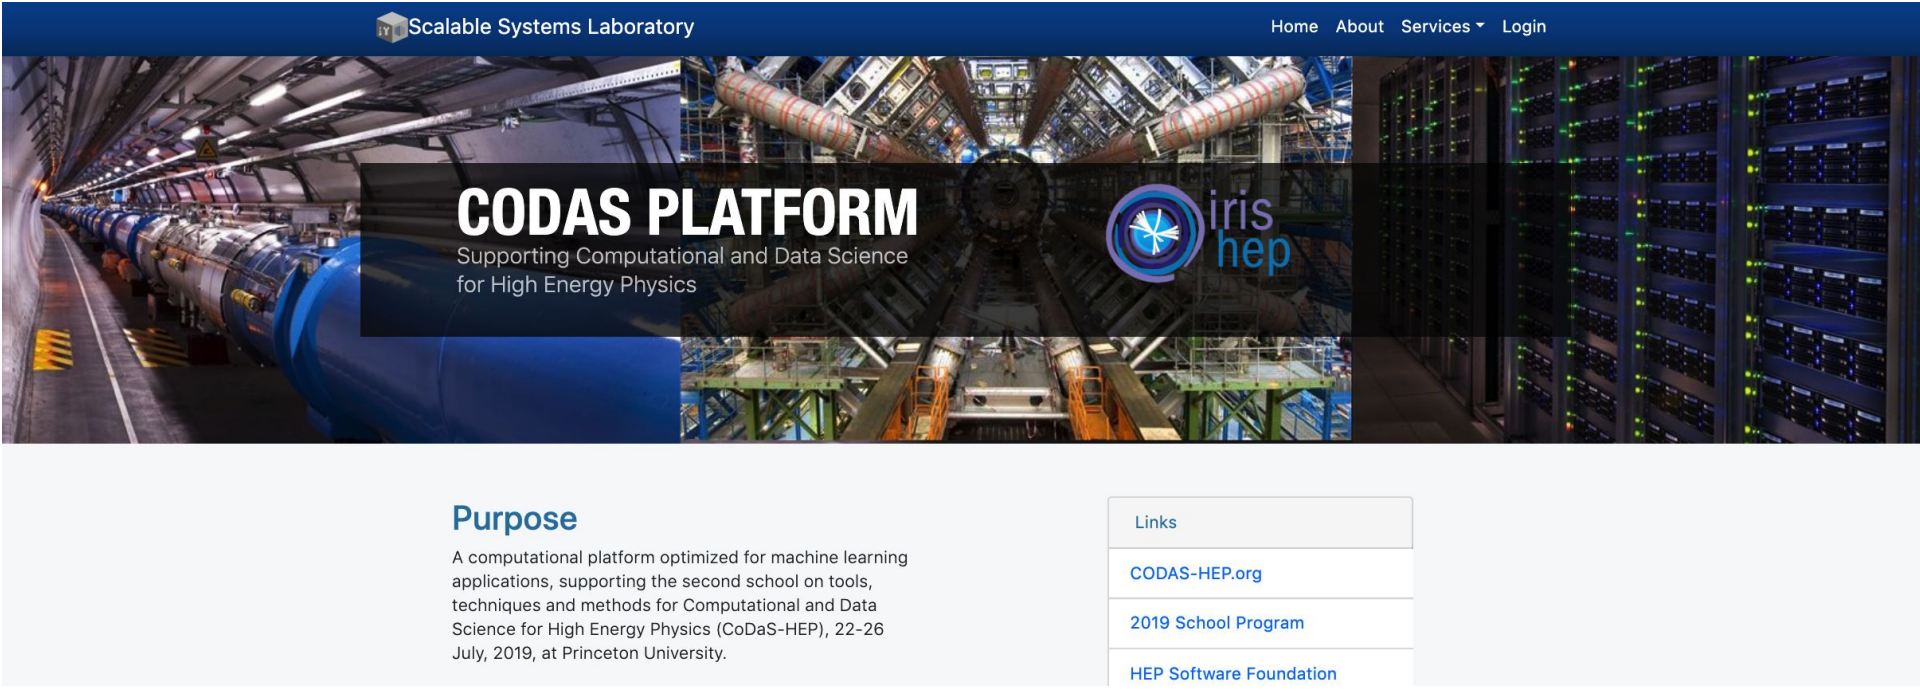

First piece of the substrate

- Repurposed R&D cluster
  - Vintage but nice (3k cores)
  - 2x40g to campus 100g SciDMZ
  - Kubernetes for flexibility for services and compute; SLATE enabled
- Deployed REANA application
- Deployed DOMA ServiceX components
- Deploy HTCondor workers and backfill with OSG Connect

Federated ID access (institutional, CERN account), edge services hosting, Unix account provisiong, OSG–LHC sofware environment

Moving forward

- Adding more resources to SSL
  - clearing house to k8s resources & building the community
  - some federation-like activity where it makes sense
- Providing additional services to the SSL
  - For user access and storage volumes for services
- Supporting DOMA milestones
  - Data Lake+ServiceX+Skyhook, iDDS
- Supporting AS milestones
  - Deployment & benchmarking of existing systems, focus on RECAST/REANA

Thank You