Scalable Systems Laboratory (SSL)

Area Lead Chats

Rob Gardner

April 23, 2020

Chat Organization

- Responses to questions
- Referral to SSL slides at the NSF review (avoid duplication)

Questions

1. Summarize the set of projects/activities and associated effort for your area
2. Are there internal or external collaborations associated with each project or activity? For external collaborations, is IRIS-HEP leading, contributing or simply “connecting/liaising”?
3. Which project/activities/goals are making progress and which are not? (Area lead’s opinion) For those that are not, what is impeding progress?
4. How are each of these projects/activities connected to, being informed by or planning on delivering (eventually) to the experiments? Are there relevant blueprint meetings or workshops that should happen to make progress?
5. What would be potential Year 3 milestones for each of the projects? (First ideas, to be iterated with PIs and the whole team as this process moves forward.)
6. What “grand challenges” would be useful to organize involving your area during Year 3 of IRIS-HEP? How would these challenges depend on efforts from other areas of IRIS-HEP, the US LHC Ops programs or the experiments?
7. Are there new opportunities where effort from IRIS-HEP can make an impact? Is the alignment of the focus areas in IRIS-HEP appropriate?
8. How are projects currently managed in your area? What tools are being used? How is progress measured? How are risks recorded, identified and mitigated?
9. Are the metrics being used to measure success clearly defined? How well do metrics in your area measure progress, success or impact? Where can the metrics be improved or refined to better measure progress, success or impact?

Summarize the set of projects/activities and associated effort for your area

This was the WBS for the SSL area in years 1, 2:

- WBS 6.1 - creating and operating scalable cyberinfrastructure
- WBS 6.2 - establish 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

IRIS-HEP SSL effort at UChicago* (~ 1 FTE):

- 7.5 % - Gardner (area lead)
- 33.3% - Bryant (DevOps, K8s deployment)
- 30% - Weinberg (testing ServiceX)
- 2% - Chien (Faculty CS research, student supervision)
- 25% - Chen Zou grad, Neha Langareddy undergrad (research)

Refer to 18 month NSF review slides

Are there internal or external collaborations associated with each project or activity? For external collaborations, is IRIS-HEP leading, contributing or simply “connecting/liaising”?

ACTIVITY

- Re-purposed existing cluster to form a 'kernel' of a distributed Kubernetes service
- Provided testbed for DOMA **ServiceX**, **Skyhook**
- Deployed of **CoDaS-HEP** training platform; **FuncX**
- Supporting ATLAS **Harvester** edge containerization
- **REANA** Helm dev & deploy
- Additional activities:
    - *ATLAS and WLCG analytics services (**perfSONAR viz, Frontier-Squid** analytics)*
    - *Configured the SSL to 'backfill' with **OSG Connect** or ATLAS **PanDA***
    - *SSL big contributor to **COVID-19***
- Partnership established with Pacific Research Platform to share expertise and access to CPU and GPU resources

→

COLLABORATIONS

- WLCG Kubernetes Working Group (L. Bryant co-chair)
- OSG-LHC, SLATE (containerization, deployments)
- DOMA (deployments, functional testing, scale testing - ServiceX, Kafka)
- CoDaS-HEP instructors, US ATLAS Ops, Pacific Research Platform
- REANA development team (sharing unpriv deployments)
- ATLAS ADC (Analytics deployments on SSL), ATLAS PanDA (Harvester containerization)

Which project/activities/goals are making progress and which are not? (Area lead's opinion) For those that are not, what is impeding progress?

- All areas are making progress. The facility acceleration R&D had a slow start, but we have recent progress (c.f. here) and have recruited a new CS grad student (incoming this Fall) to work in this area.

- The SSL has aggressively identified a course for flexible, declarative and dynamic facility infrastructure based on new, fast-moving open source technology, not precisely aligned with scientific computing infrastructure

- So this has brought a number of challenges:
  - *Expanding the Kubernetes substrate (involving more resource providers)*
    - Challenge of lack of Kubernetes expertise in our community
    - Effort available to train, on-board, and collaborate
  - *Support and management of "applications" and "users" of the SSL*
    - Lack of expertise of IRIS-HEP developers with service deployment
    - Integration with external, dependent services (e.g. storage, caches)

How are each of these projects/activities connected to, being informed by or planning on delivering (eventually) to the experiments? Are there relevant blueprint meetings or workshops that should happen to make progress?

- The delivery to the experiments happens in a number of ways
  - *Helping IRIS-HEP research areas with deployments and testing*
  - *Delivering to LHC computing resource providers*
  - *Helping experiments with containerized service deployment*
  - *Providing education and training infrastructure*
- Blueprint meeting with Analysis Systems for scalable platforms
- SSL-TEAM monthly telecons
- Chicagoland k8s-HEP meetup
- Leadership in WLCG working groups
  - *WLCG Kubernetes WG (co-chair)*
  - *WLCG Federated Operations Security WG (co-chair)*

What would be potential Year 3 milestones for each of the projects? (First ideas, to be iterated with PIs and the whole team as this process moves forward.)

1. Realize a **lightweight**, three-site "k8s substrate" which has the following ingredients and capabilities:
   a. *SSL single sign-on and minimal user/group management*
   b. *Distributed storage solution based on Rook (Ceph)*
   c. *HTCondor, Spark and Ray task spawners from interactive notebooks*
   d. *Deliver a* ***k8s deployment pattern*** *that can reproduced at a single site (e.g. a Tier3, or a local network of 'gaming' -Nvidia cards- laptops in an LHC analysis group)*
2. Create a scalable federated notebook service to support interactive machine learning education and analysis
   a. *Based on recent discussions w/ AS, CERN IT*
3. Analysis Challenge, and 10 PB day long Grand Challenge (next slide)

What “grand challenges” would be useful to organize involving your area during Year 3 of IRIS-HEP? How would these challenges depend on efforts from other areas of IRIS-HEP, the US LHC Ops programs or the experiments?

As discussed at the NSF Review ("**analysis challenge**", c.f. response to review questions), and a grand challenge currently being discussed by CMS (**10 PB One Day Challenge**, cf here)

Are there new opportunities where effort from IRIS-HEP can make an impact? Is the alignment of the focus areas in IRIS-HEP appropriate?

Yes. Three opportunities we're looking at:

1. R&D to evolve the LHC computing facilities to be more flexible, declarative, dynamic, and responsive to LHC physicits for scalable analysis platforms.
2. Supporting the Snowmass effort with an analysis platform (JupyterLab access to GPUs, parallel frameworks: HTCondor/Ray/Spark).
3. For external, broader impacts, working with computing coordinators from other domains (e.g. from MMA) with similar analysis workflows. Recall interest from Patrick Brady at review (yet to be followed up)

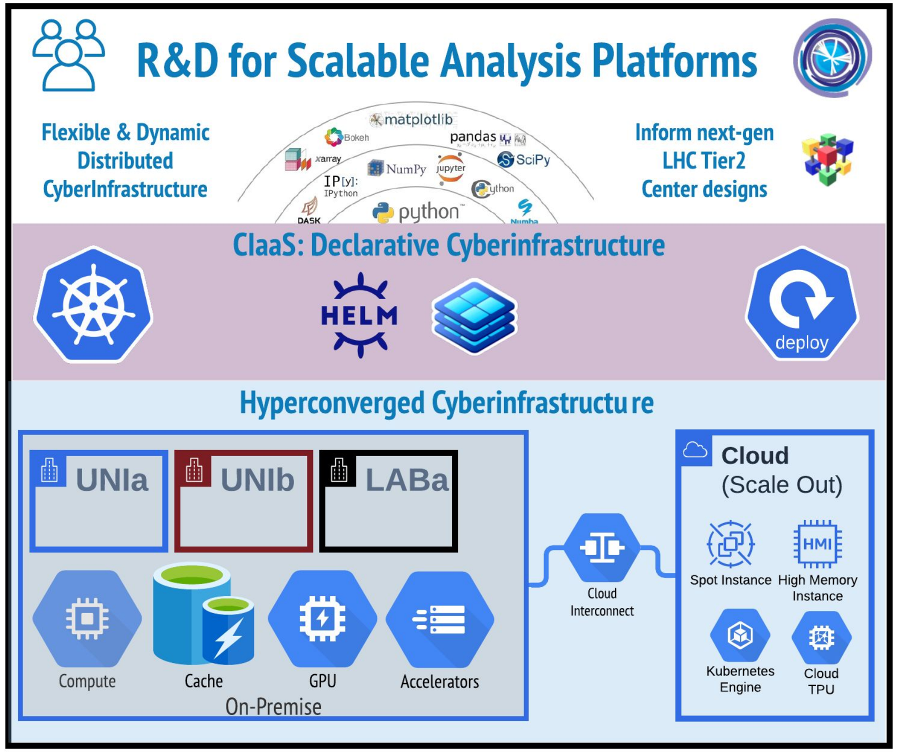

Refer to 18 month NSF review slides

How are projects currently managed in your area? What tools are being used? How is progress measured? How are risks recorded, identified and mitigated?

We have so little effort informal communication is being used, currently.

Progress has been "measured" primarily through myriad project reporting required by IRIS-HEP.

Are the metrics being used to measure success clearly defined? How well do metrics in your area measure progress, success or impact? Where can the metrics be improved or refined to better measure progress, success or impact?

They are not particularly well defined. Suggestions welcome.

Refer to 18 month NSF review slides

EXTRA SLIDES

Copy of NSF Review slides follow

Focus Area: Scalable Systems Laboratory (SSL)

- Vision: Community innovation platform & facility R&D, path to production
- Work with IRIS-HEP research areas, Blueprint, LHC software and computing teams, OSG-LHC
- Utilize industry stand tooling (cloud native Linux container packaging and orchestration: Docker, Kubernetes, Helm)
- Leverage related NSF CI projects SLATE, Pacific Research Platform and others

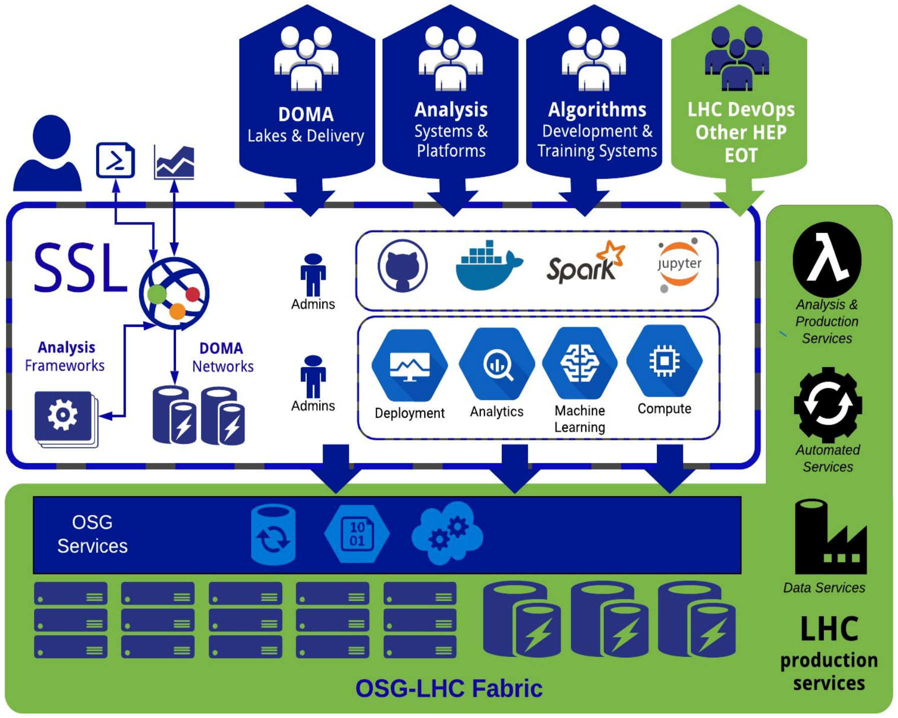

Major Activities

- Re-purposed existing cluster to form a 'kernel' of a distributed Kubernetes service
- Provided testbed for DOMA ServiceX development and scalability testing
- Deployed of CoDaS-HEP training platform
- Supporting ATLAS Harvester edge containerization
- Deployed REANA service
- Additional activities:
  - *Provided various ATLAS and WLCG analytics services*
  - *Configured the SSL to 'backfill' with OSG Connect or ATLAS PanDA*
- Partnership established with Pacific Research Platform to share expertise and access to CPU and GPU resources

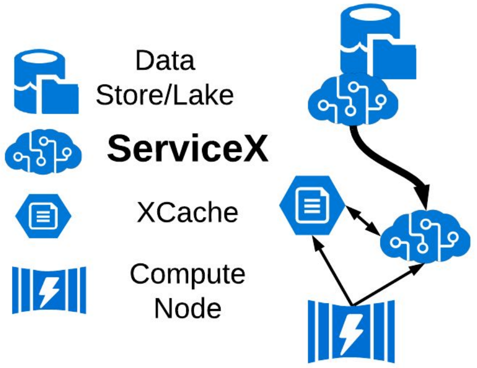

reana

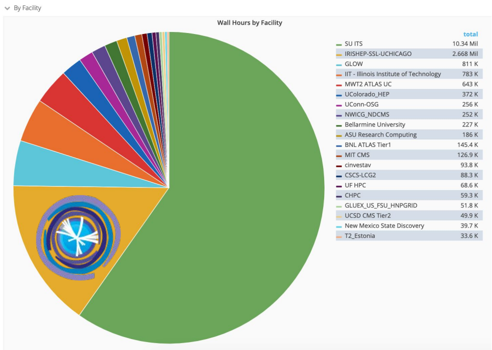

OSG CONNECT

Summary of SSL Deployments

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

SSL Highlight

- Provisioned Jupyter machine learning environment for 55 students attending the CoDaS-HEP school
- Scheduling backend to GPU resources from the NSF Pacific Research Platform and CHASE-CI
  - *More in the following presentation on Training*

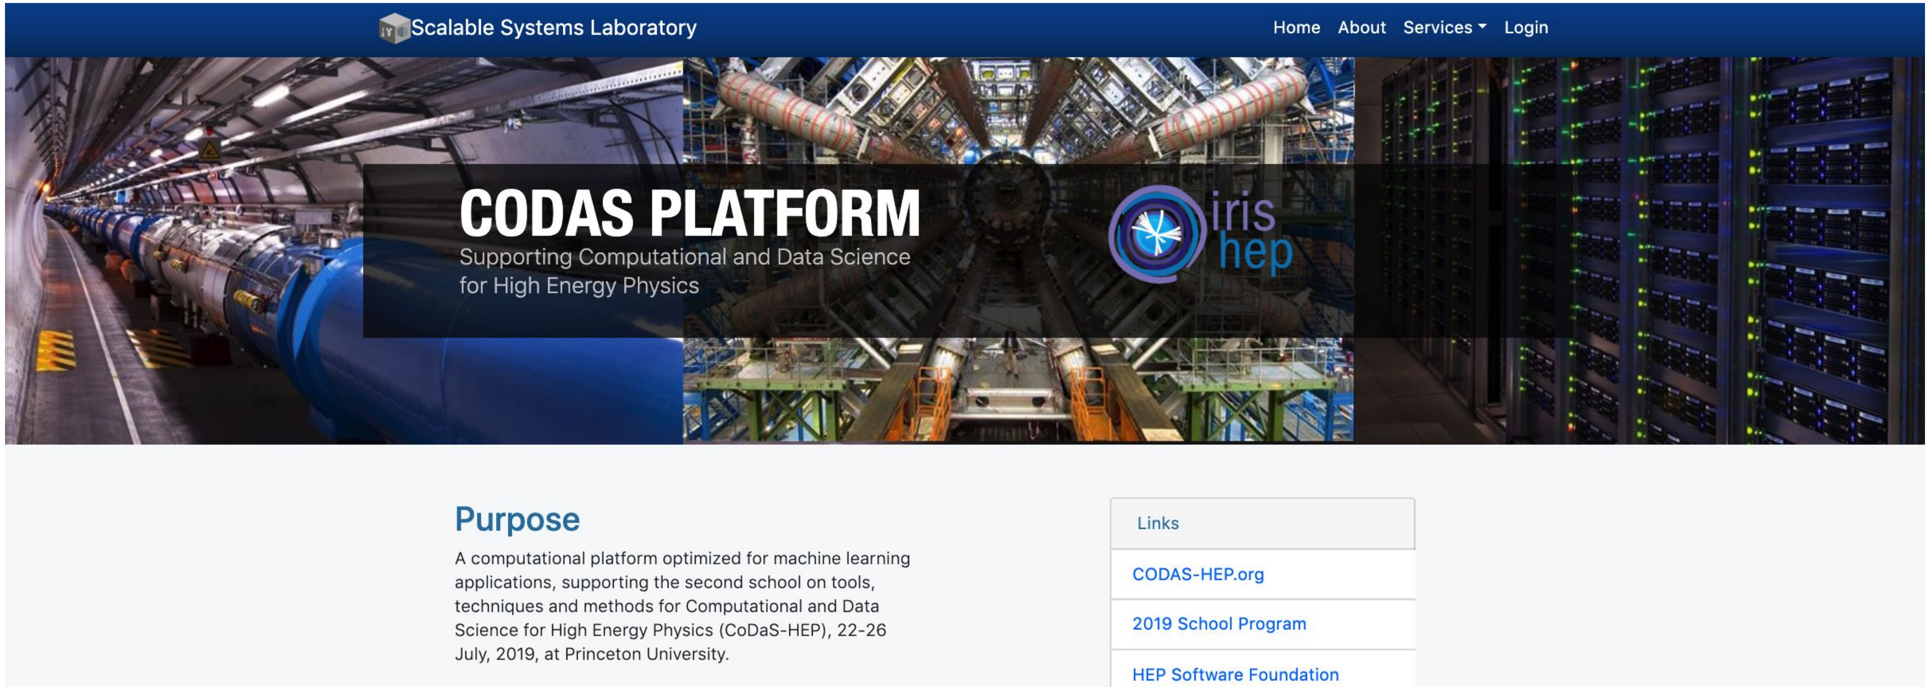

Community Building

- Presentations at ATLAS Software and Computing workshops
- Analysis Systems - SSL blueprint meeting
- Organized Kubernetes BoF at CHEP2019
- Talks at Kubernetes pre-GDB (Grid Deployment Board) at CERN, December 10, 2019
- Monthly SSL call on Kubernetes topics, open to community

Evolving LHC computing facilities

- Facility R&D towards more flexble systems to support HTC batch & low-latency interactive
- Emphasis on "declarative, distributed CI" for reproducible patterns

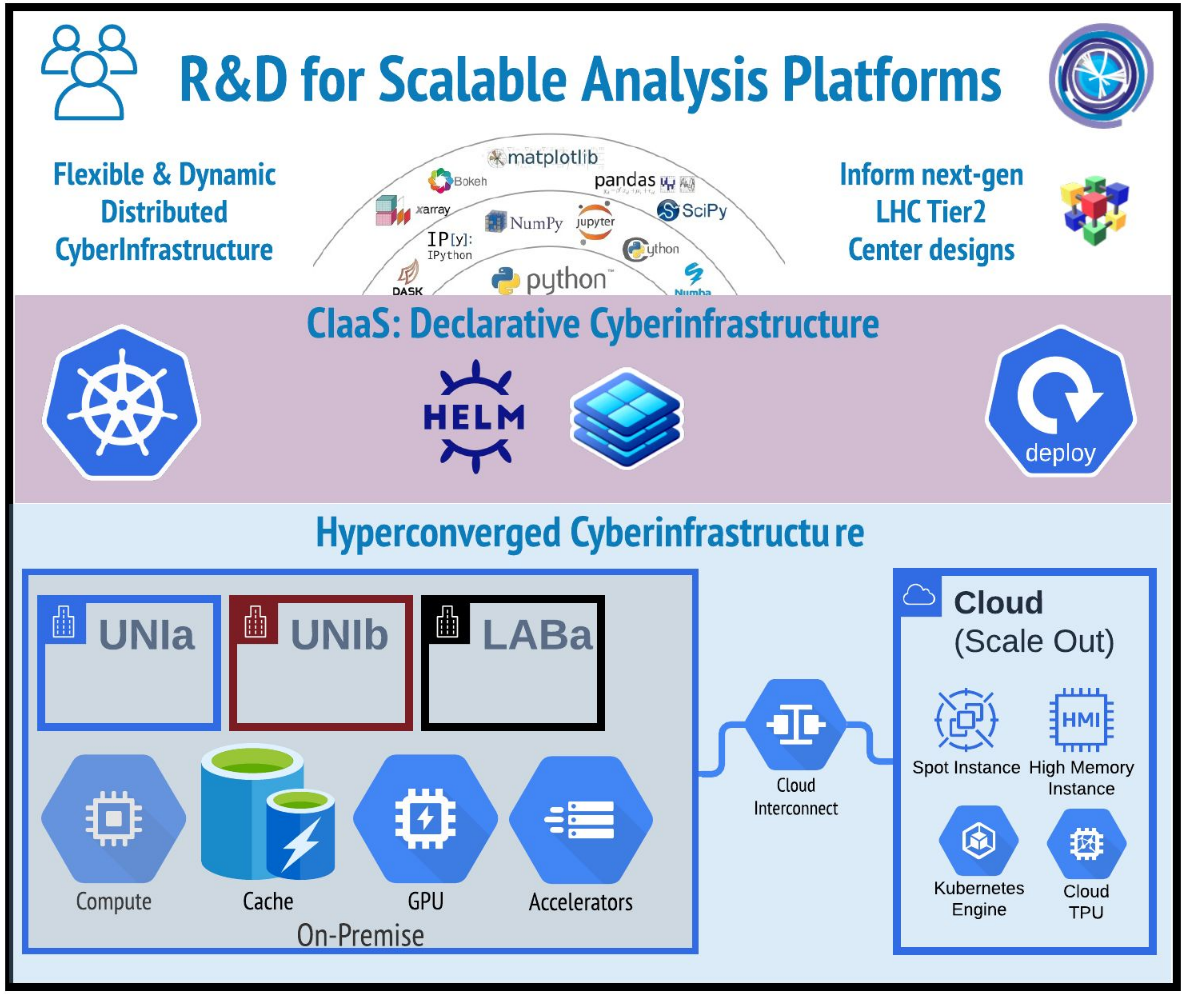

Metrics

Metric	Target	Current Value
M.6.1: Number of deployed and operated platform services	Design phase targets taken from initial DOMA and AS estimates. A suitable metric for the execution phase will be desgned.	12
M.6.2: Number of SSL user groups	Determined by number of developers & other devops requiring resources	>9
M.6.3: Number of engaged institute areas per quarter	5 Institute areas were targeted for engagement	4 (DOMA, AS, Training, OSG-LHC)
M.6.4: Number of integrated SSL cluster resources	Determined by scale demands requirements of development teams, both within the Institute and outside	2 (1 dedicated, one shared)

Area Summary

- The SSL has been established as a flexible testbed and integration platform
- The SSL team has deployed dozens of services supporting DOMA and AS R&D, ATLAS analytics and WLCG networking services and the OSG-LHC
- The SSL has supported a training event and has directly leveraged other NSF cyberinfrastructure projects
- Looking forward:
  - *Involve additional sites*
  - *Continue to collect SSL infrastructure requirements for all areas*
  - *With the larger Kubernetes community (inside HEP, CERN and beyond), R&D on multi-tenancy, scheduling and federation to provide scalable, multi-site platforms*
  - *Use research results to inform design of next-generation of LHC computing sites, including more flexible Tier-2 and analysis centers*