

---

# SSL Architecture

hybrid systems, deployments, developer support

Lincoln Bryant  
Enrico Fermi Institute  
University of Chicago



---

IRIS-HEP AS-SSL Blueprint @ NYU  
June 21, 2019



# Desired SSL capabilities



- Support a diverse catalog of deployment patterns & models
- Experiment patterns (scalability tests)
- Usability
  - Modality, Reservation
  - Metrics, logging, analytics
- Operation & Support
- Openness
  - to providers to contribute
  - to developers to conduct experiments
- Recording value
  - Analysis platform "blueprints":
    - Single site/region deployments
    - Multi-region deployment
    - Multi-cloud hybrids – e.g. SSL+GCP+CERN, etc..
  - Demonstrations & archival of demo artifacts





# An architecture that implements...

---

- a community platform
- supports groups and projects
- bespoke resources & configurations
- declarative & reproducible deployments
- services to build & manage artifacts
- scalable up and back down
- reduce cognitive load for developers and deployers





---

**Some ingredients..**



# Community platform

---

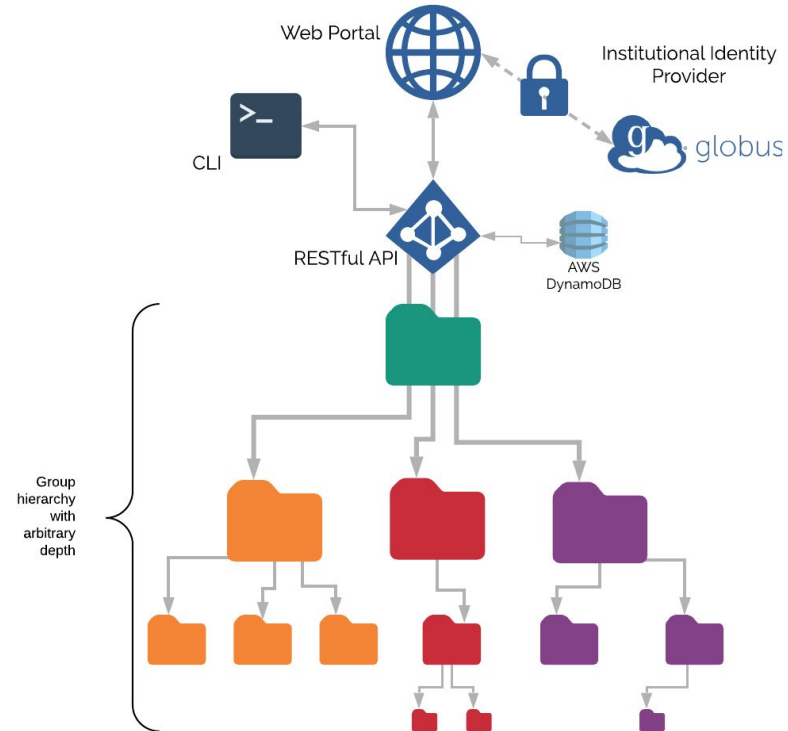


- Open to all working on software infrastructure in HEP
- CILogon, Globus to provide single-sign on and federated identity
- Lightweight user and group (project) management system
- Infrastructure itself composable, reusable



# Groups and projects

- Web and CLI interfaces for user management
- Groups organized as a tree structure with arbitrary depth
- Users can invite others, create sub groups, etc.
- Smart clients inspect the tree and implement appropriate provisioning of resources





# Resources

---

- Mix of bespoke dedicated resources and capability for users to bring allocations on others
- Container-based service orchestration on dedicated resources
- VC3-like technology to connect to HPC/HTC resources for batch
- Facilitate integration of commercial cloud resources when needed



# Orchestrating services in the SSL



- Need flexible infrastructure for supporting the workloads we expect from SSL
- Dynamically reconfigure existing hardware to be a HTCondor cluster today, Spark tomorrow, whatever is needed.
- Containerized services are getting a lot of attention in Industry right now– can we take advantage of the momentum?
- Want to “glue” clusters together, but abstract away infrastructure to whatever extent possible – clear a smooth road for the developers
- Potentially mimic cloud native groupings: e.g. create “zones” of resources





# CS 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
- Rebuilding as Kubernetes
- Explore federation to aggregate w/ others



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

# Kubernetes (k8s)

---



- Open source container orchestration platform
- Automate deployment, management, scaling
- Has origins in Google/Borg
- Supported/managed by Cloud Native Foundation
- Declarative model for deployments



# Declarative infrastructure

---

- Want infrastructure built under the SSL to be easily reusable and deployable to other sites
  - No more wikis with install guides!
- Declarative nature of Kubernetes is a good fit and gets us a long way down that road.
- SSL as an incubator for projects which then “graduate” to become full-fledged infrastructures that run on production resources.

# Federating platforms

---



- Expect users to outgrow the dedicated pool of resources we have now.
- Need an interface and mechanism to allow users of the SSL reach into resources at a heterogeneous collection of sites
- Many approaches in the Kubernetes community, waiting to see what survives & what will be most appropriate for us



# Extending into HPC/HTC

---

- Foresee workloads that require some service infrastructure in the SSL, but want to do batch computing elsewhere
- Want to facilitate by using technologies derived from VC3, HEPCloud and others
- Provision compute schedulers, data managers on SSL, schedule workers to HPC resources via overlays



# Artifact build and management

---

- Provide resources for building and registering containers, compiling software, etc.
- Off-the-shelf tools plugged into SSL resources with a little bit of glue.
  - Why wait 30 minutes for DockerHub to build your container?
- Is this obviated by CERN services? Perhaps more valuable for non-LHC experiments

# SSL "Glass"

---

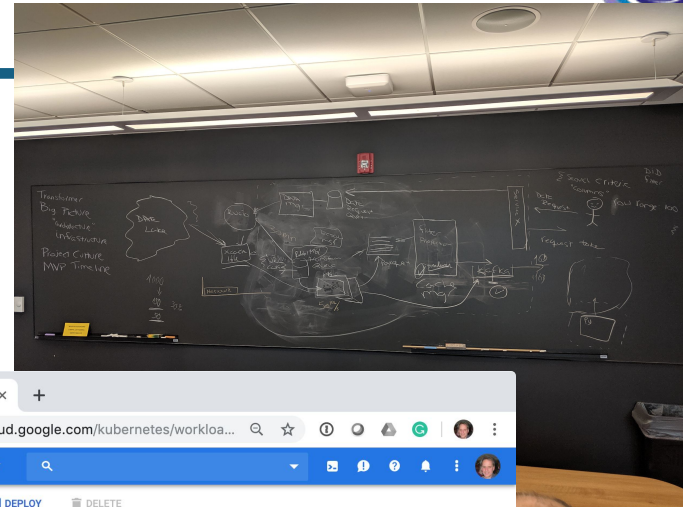


- Portal for visibility and organization
  - groups (projects)
  - resources
  - artifacts
- Metrics, logging, analytics
- Regional, national and international scopes

# 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







# Current status

---

- Group/identity bits are being developed for multiple projects, being repurposed for SSL.
- Kubernetes conversion of River cluster underway.
  - 4 nodes online, backfilling w/ OSG via SLATE
  - Brave early adopters come talk to me afterwards!
- Looking for partners to contribute infrastructure and a bit of effort – experiment with how to federate resources.



# Wrap up

---

- To briefly recap:
  - Institutional Identity and group management
  - Container-based, declarative software deployment and service orchestration
  - Mix of dedicated and non-dedicated resources
  - Exploring options for Federation
  - Building tightly integrated "pane of glass" for it all
- Integrate with industry best practices where practical!



# Discussion

---

- All of this is very nice, but we need to meet the needs of the community.
- We need input from Analysis Systems and others!

