

Rob Gardner

Scalable Systems Lab

Facility R&D and Integration with LHC production environments

w/ Lincoln Bryant & Fengping Hu

Facility R&D

- ❖ Facility R&D broadly refers to activities related to the exploration and innovation of systems, services and physical infrastructure that provide platforms suitable for HL-LHC service environments and runtime ecosystems.
 - These can be purely **local facilities** (platforms deployed within a local area network) or **distributed**, in the sense of interoperating services over wide area networks.
 - c.f. IRIS-HEP 2.0 strategic roadmap
- ❖ IRIS-HEP and the LHC experiments are adopting cloud-native application management methods for infrastructure and services
 - Helm charts and Fluxcd as a technology to standardize deployment
- ❖ Flexible strategies are being tried out in many contexts
 - Scalable Systems Laboratory (SSL) at UC and at UNL
 - UChicago Analysis Facility
 - These patterns can be adopted at other sites

Facility Integration

- ❖ Facilities Integration refers to activities relating to the installation, configuration, and operation of development and pre-production software components.
- ❖ Where possible development services are deployed in realistic contexts, **alongside existing systems and production infrastructure**.
 - This is the most immediate path to understanding which components in the overall cyberinfrastructure present the most significant scalability challenges.

Aligning with IRIS-HEP 2.0

- ❖ We encourage US LHC Operations programs to adopt approaches with the needed flexibility
 - Concepts in IRIS-HEP 1.0 (K8s substrates, GitOps services) now in production (analysis facilities)
 - Inform the next generation Tier2 and Analysis Facilities -- where relevant
 - Must figure a way to evolve what's below to **provide needed flexibility** above
- ❖ Scalable Systems Laboratory Resources
 - For the past four years, a dedicated K8s cluster at UC to support R&D
 - "RIVER" has been an important space for prototypes and even called into production use for analytics and hosted CE services
 - UNL has K8s deployed at scale for Coffea-Casa development
- ❖ Feedback needed
 - for on-going development
 - for production facility "perturbations"

Facility R&D Goals

- ❖ We plan to keep River operating for the next year to provide an **SSL "reference platform"**
- ❖ **Constructing analysis facilities:** participate with the broader WLCG community in identifying best practices, infrastructure management patterns, and recipes forged on the SSL leading to reliable and scalable service deployments capable of supporting hundreds of users at Run3 scales in the near term, but with proven capabilities to scale up to Run4 during LS3.
- ❖ Evolving **Tier-2 infrastructure management:** retrofitting Tier-2 facilities with substrates and higher level infrastructure management tools and services, utilizing industry standard solutions where possible.
- ❖ Capturing **facility patterns and blueprints:** Creation and curation of GitOps charts and repository actions to provide reliability and reproducibility such that both infrastructure and applications can be easily replicated across facilities.

Facility R&D Goals II

- ❖ **Identifying infrastructure (and service) bottlenecks:**
 - A critical activity during data and analysis grand challenge exercises, and experiment demonstrators, are load tests of representative software workloads that accurately mimic planned infrastructure and HL-LHC use cases.
 - This will require **instrumenting services and infrastructure** with the needed monitoring hooks, streaming the resulting metrics to analytics dashboards, identifying and characterizing any bottlenecks that might exist, and measuring overall performance.
- ❖ **Improve resource sharing across tasks on facilities**
 - Kubernetes itself cannot match the full spectrum of fair share scheduling capabilities for CPUs, ServiceX transformers, Dask workers for Coffea-Casa, HTCondor job slots for traditional batch processing.
 - This must be dealt with during Run3 for analysis facilities which are providing new and old analysis environments.

Wide Area Infrastructure

- ❖ Remote (central) management of infrastructure to reduce cost, improve efficiency
- ❖ Examples are distributed caches
 - In US ATLAS we centrally manage networks of XCache, Squids, Varnishes & are eager to work with others doing the same (e.g. OSDF)
- ❖ And embedding services in the network
 - "Neighborhood Caching" (ESnet caching project)
 - Server Side Data Delivery
- ❖ Neighborhood caching could extend to other services

Early Deliverables

- ❖ Establish infrastructure and orchestration services necessary for a sustained reference SSL. December 2023
- ❖ ServiceX deployed inside FABRIC node at CERN. December 2023
- ❖ Releases of Analysis System pipelines supporting distributed analysis deployed on SSL. Continuous
- ❖ Provide a deployment package for upgraded software distribution and conditions data caching servers. May 2024
- ❖ Evaluate distributed workers. May 2024
- ❖ Curate and publish production AF deployment patterns including integration of traditional systems in use during Run3 and forward-looking analysis systems. December 2024
- ❖ Support demonstration of running a full analysis suitable that uses machine learning. December 2024
- ❖ Evaluate technologies, services and infrastructure management patterns for next generation LHC Tier-2 facilities and publish for the community. December 2025

Related Breakouts this Week

- ❖ K8s / Future Facilities monitoring
- ❖ ServiceX next steps this afternoon
- ❖ Future Analysis Facilities R&D tomorrow 9:30
- ❖ OSG-LHC session tomorrow 11:30

refine plans based on discussions here..

Questions?