

SSL Session at the IRIS-HEP Retreat

September 4-6, 2024

University of Washington, Seattle

IRIS-HEP Institute Retreat: <https://indico.cern.ch/event/1374962/>

SSL Focus Area talk Wednesday, Sep 4 ( 2024.09.04 SSL Focus Talk at Retreat)

SSL Session is Thursday, Sep 5, 13:30 - 15:00 PDT

Overall planning document:  IRIS-HEP Retreat Session Proposals - 2024

Project Planning folder:  Year 7 SSL Planning ; Last year:  Year 6 SSL Planning

The goals of the retreat are to:

1. Checkpoint the status of the IRIS-HEP efforts to date (Year 6 of the project)
2. Clarify the gaps between where we are now and what will be needed for the HL-LHC startup
3. Update our plans for delivery of IRIS-HEP products to our partners (experiments, ops programs and the LHC and HEP community)
4. Elaborate a vision for what IRIS-HEP can achieve in Year 7 and beyond.

SSL Session Topics

Supporting Training Events

- << David and Jim will be at the first part of the session and will offer direct feedback>>
- From Slack convos:
 - From David Lange: would it be useful/interesting for you to present/discuss your work to support training events (eg, in India and Codos) in the training parallel session? I think I see that most of you are planning to come in person
 - From Jim Pavarski: I'd say: **authentication, startup time of large images, and GPU instances**. These have been the three pain points so far. In a conversation, I'm willing to learn new best practices—the solution is not necessarily technological. But we do need to collectively find a better way to do these things.
 - Lincoln: I think these are great topics. I could certainly give updates on **Authentication** - I'm happy to report that we should be able to have users go to a particular web page and request a group nowadays. No more importing

of a CSV of users whose emails might not match the identity asserted by the provider. I am working on making it even simpler / nicer for users. As for large images, that's a tricky one but I think we could at least quantify what is a large image, what is a long startup time, and what are our expectations. I am not sure what the GPU instance problems are, but we should definitely discuss

Jim Pavarski Tutorial Feedback

- See slides in Indico: [Tutorial compute platform](#)
- Having a prominent ATLAS IAM button is a problem - students click it anyway
- CILogon openness is good, but there are issues.
- Want the platform to be allowed long after the event.
- The allow list - doing it manually
- How about an open enrollment period - say 10 minutes.
- Discussing by-passing the Binderhub interface; strongly want to avoid having to "teach" the binder interface, and expose the users to more pitfalls.
- The large blue start server button - which they should not press.
- Large image handling:
 - also too many images
- GPU instance problems:
 - Wifi network instabilities
- Saving work
- December is the next event
- Can layering help? A base image with 80% of the software, and then students install remaining part on top
- Pre-pulling a base layer
- NRP has an image caching service
- Could we have someone on shift during the workshop, to cover emergencies?

Lincoln on Authorization

- Slides:
 - [2024.09.04 Simple AuthZ for Training Events](#)
- Jim: concerned about time to approve - pressing the button too time consuming
- Keycloak page - confusing; which one provider did I choose?
 - Can skip this page to present
- Can you delegate approval to trainers? Ans: yes
- Brian L - are there other tools available in Keycloak? Another point about managing multiple email addresses.
 - Also - Reporting: per workshop list of attendees is very helpful.

Fengping on Binderhub Improvements

- [Binderhub misc topics](#)
- Large images. Continuous image puller.
- Binderhub w/ Dask

- Model GPU selection to choose sites - shouldn't need to do that

Next major even is in December - schedule a mock training event in October to try out features and iterate on feedback.

Other Topics

- ServiceX DevOps best practices
- Multit-User IDAP demonstrator Planning
- Evolving towards Federated Analysis Facilities
- **Monitoring:** show AF usage Accounting (GRACC, APEL, CRIC...). UNL has been an alpha tester for the OSG's new kubernetes probe.
- **Federation:** Not only with UNL (which now has GPUs and ARM server, but I think this is done), but with NRP (done!). What other resources would it be useful to federate with? Could we federate with Expanse in the future? Federate with the k8s cluster at NET2?
- Major goal during the retreat is to map discussions into the Y7 plan below.

Y7 Project Planning

We need to identify Y7 deliverables and Milestones where the SSL is needed.

- W6.1 SSL Infrastructure
 - Providing equipment and service orchestration services necessary for a sustained facility R&D laboratory for IRIS-HEP
 - **HL-LHC Gaps:** the major gap identified during the IDAP demonstrators was network bandwidth limitations and lack of throughput through XCaches.
 - **Benefits and Impacts:**
 - **Milestones:**
- W6.2 Supporting Data Delivery
 - Helping ServiceX become an declarative, interactive data interface for HL-LHC analysis.
 - **HL-LHC Gaps:** There are several dependencies on the analysis models, data formats, and
 - **Benefits and Impacts:**
 - **Milestones:**
- W6.3 Analysis Facilities
 - Integrating IRIS-HEP Coffea Casa & other interactive technologies into production HL-LHC analysis environments.
 - **HL-LHC Gaps:** Adoption
 - **Benefits and Impacts:**

- **Milestones:**
- W6.4 Tutorial & Training Event Infrastructure
 - Deploying persistent resources in support of tutorials and training events (T&TE). This includes creation of declarative deployment patterns that can be reproduced elsewhere.
- W6.5 SENSE (Diego)
 - Slides: [2024.09.05 SSL SENSE milestones](#)
 - Good to get an ATLAS site participating
 - Q: Lincoln - what is needed at the site? dCache/K8s/switch config? Whats the high level checklist?
 - Ans: don't use your production infrastructure. There are recipes for k8s. IPv6. Need three subnets we can use to get to your infrastructure.
 - Justice: need switch access and control BGP. FRR is a clone of router, so you don't need access to the switch at the site.
 - Lincoln: we can check all the boxes except the BGP requirement.
 - Justice - will have FRR running on a server at Purdue
 - Q: Shawn - wasn't Eduardo (NET2 which is at MGHPC) working with the group already? (Rafael) Gen-4 machine sent. Frank W is in the loop.