ServiceX Detroit Retreat

July 19 & 20, 2025
Previous Retreats:

- Nov 2024 in Austin, TX (document)
- March 2025 in Seattle, WA (document)

Logistics

- Meet at 9:30 am on Saturday and Sunday morning
- Location: Courtyard Detroit Downtown
    - Piquette Room 4th floor
-

Reference Materials

- ServiceX Issue Planning Board

People Attending

Please fill out when you are arriving/departing, etc.

- Gordon (arrive 9 am Saturday Morning, need lift/uber to Lansing on Sunday night)
- Ben - Fly to Seattle on Sunday 4:45
- Kyungeon
- Peter
- Roger (Arrive 10pm Friday, Leave for Ann Arbor Sunday Night)
- Artur
- Dmitry: arriving Saturday after lunch, leaving Sunday noon
- Matt

Desired Outcomes

-

To-do Items (todo at workshop)

Topic Brainstorming

Add a topic here, along with your name!

- IRIS-HEP Review (done)
    - IRIS-HEP still does not know when it will happen. But it sounds like NSF employees got asked for travel needs, so there are signs that travel may becoming possible again.
    - Prefer a Zoom review
        - Bring in outside experts for advice
        - NSF still can't travel
        - Maybe in DC? Some problems with this
        - No way before October
    - IRIS-HEP Retreat in October
        - Make sure we have a good ServiceX presence
        - Usage in experiments
        - Need milestones
- Marketing (done)
    - Conferences
        - ACAT - Sept
            - Ben Poster on transactions
            - Gordon poster on using this for ML data extraction
            - Talk in Track 1- Artur (general)
            - Good conference to getting at end users
        - CHEP - 2026
            - Good conference for getting at end users and facilities people
            - Who is submitting!?!?!? Deadline is currently Dec 19th
            - One abstract aimed at facilities people, and one at end-users
        - ROOT Workshop (November - Gordon is going as an IRIS-HEP rep)
            - Emphasize RDataFrame support (someone from UT)
            - Deadline: 15 Sept 2025
        - ATLAS S&C in late Sept (Gordon)
            - Gordon has a plenary talk there…
        - Other physics or ATLAS/CMS workshops and conferences?
            - The AF committee was the right venue for outreach
            - Goal: Higgs -> \mu\mu framework - can you get SX work there without massively changing (or at all) the framework
            -
    - Analysis Facilities (done)
        - Do they overlap
        - How could someone outside of the US access ServiceX
        - IRIS-HEP would not buy hardware ex-US
        - Could we take advantage of the AF's mailing list
            - Offer demos or tutorials
        - Need fresh perspective on AF/ServiceX relationship
        - What are the buy-in costs for an AF to deploy ServiceX?
            - Must have Kubernetes

- Could we reduce the barrier to deployment
    - Would it make it easier to deploy at an AF if we move the transformer jobs to HTCondor?
    - Could we substitute xrootd/EOS for S3?
        - ACTION: xrootd target to write outputs
    - Should have a decent pre-reqs doc that we can hand to an AF
    - ACTION: Talk to AF owners/stakeholders about barriers to installation and use

- Tier2s
    - Northeast
    - Midwest
    - Great Lakes
    - Southwest
- Examples and outreach (done)
    - Landing page has links to ServiceX deployments
        - Github Issue
    - Simple examples to share and get people interested in ServiceX
        - Gordon's examples are useful but still complicated
            - Get mc particle list
            - Get ML training data (this is not simple)
        - We do have some simple examples
        - Peter will put in a better example
        - Replicate some current ATLAS examples using ServiceX to show difference
        - Utilities (like MC tooling, etc.)
- What is ServiceX good for? (done)
    - Get ServiceX on TopCP FastFrame page!
        - Start from PHYS or PHYSLITE
        - Get systematics
        - For now this is just experimental
        - Useful for R&D
    - EasyJet
    - Making a large dataset manageable
    - Exploring a dataset
        - Much easier to do these things
    - ATLAS Tutorials reimplemented
    - Testimonials
    - Fast turnaround for R&D
    - Near-interactive
    - Things that are just too complicated to make the grid worthwhile
    - Simple env setup on local machine
        - Only one pypi package + yaml token to work with AODs DAODs Ntuples Parquet (many availble releases)

            - Instead of: ssh to lxplus, setup ATLAS, Grid credentials, setup rucio, setup panda, setup TPCT or Fast frames or Athena, run on grid & pray (very release dependent!)
        - ACTION: Roger is going to take a stab at turning this into a few sentences which we can post/user on our servicex home page
            - **Accelerate your analysis cycles** — go from query to usable data in hours, not days.
            - **Cut down dataset sizes early** — filter and transform centrally, so you only download what you need.
            - **Simplify analysis onboarding** — minimal learning curve for new students or collaborators.
        - Artur is putting this into his tutorial
        - ACTION: We need to update the servicex landing page: ServiceX | Institute for Research and Innovation in Software for High Energy Physics
            - Matt is taking this on!
- ATLAS S&C Scrubbing - outcomes affecting us (done)
    - AF's, postdocs, etc.
    - There is a question of whether the non-uchicago ones will be supported, Claims that Chicago doesn't have enough resources and people wait for their jobs, but they are most used.
    - So, funding looks like it will remain stable for this.
- External Marketing
- List of "uses" in the "wild" (by group and not by group)
- Demo Day's coming up we should use?
    - Previous demo day talks:
- CMS/Purdue update/status
    - Working with a big analysis - able to reduce number of columns by 50%
        - 5-6 Minutes CMSfor original dataset
        - Filtered: 4 minutes
        - Compression level didn't have big impact
    - Need some way to recover source file for an output file (this output file came from this input file) to track down "bugs" and "weird" behavior.
    - Where will the bottlenecks be?
        - Location of input files is very important
        - Will our Rucio nearest location lookup work correctly with CMS implementation?
        - Want to run over entire Run2 dataset?
            - Couple of PB
            - For 200GBs we were using 3,000 cores
        - Maybe use interlink to burst out into Slurm/Condor cluster
- DUNE study
    - Had a student over the summer standing up an instance
    - But they have gone to tokens (which doesn't work with us)

        - They are trying locally and will come back to us to chat
    - HSF/WLCG Workshop Talk
        - Talk, lots of good questions, discussion. Not sure any hard “leads” came out of it, however.
    - Tutorial next week (Artur)
        - On git (asap)
    - ATLAS induction day / software tutorials ? (artur)
    - ATLAS Software Docs
        - Should we have add something like that.
        - Roger and Artur will come up with ideas of pages to submit to atlassoftwaredocs web site. Will present at future ServiceX meeting.
- Staffing/Supportability (done)
    - List everyone and where they plan to concentrate in the next 6 months?
        - Ben
        - Roger
        - KyungEon
        - Artur
        - Matt
        - Peter
        - Gordon
        - Dmitry - CMS Contact
    - Python 3.9 decision approaching fast
        - Relevant because this is the default Alma 9 version (although e.g. lxplus will have more recent versions available for use)
        - In general, should we push specific solutions to Python, etc. version reproducibility, or should we rely on users to learn how to set things up for their specific environments?
- New Features/Physics (partial - continue next time?)
    - RNTuple (all)
        - Input
            - xAOD func_adl can’t read it, topcp can’t, etc.
            - Uproot based things have no trouble read
            - ACTION: check flat rntuple can be read by uproot func-adl
            - RDF based read no trouble
        - Output
            - Uproot-raw knows how to do it
            - RDF is the most recent version can do it
            - Func_adl uproot does not know how to do it yet
            - Func_adl xAOD does not yet do it
            - TopCP is the same thing.
            - What about conversion in the side-car?
                - Yes, please - we need this for transformers that can’t write RNuples

- Making ServiceX-local useful for everyone (Gordon)
  - There are two different approaches
  - Gordon Method
    - Drop in replacement to the `deliver` function
    - CodeGen step
    - Science step - either from container or in local python env
    - `Deliver` splits the execution path based on the locality of the file.
    - Useful for developing code gen and science container
    - Uses docker and in-process (or wsl) to run
    - Should it be part of the servicex library?
      - Small number of package dependencies
      - Few people actually use it
      - Keep separate package for now
    - Needs a new name (ServiceX - local implies running servicex locally rather than servicex light, servicex-direct)
  - Matt Method
    - Develop ServiceX locally
    - Mount local dirs into container
    - Uses minikube
    - Overmind
  - Ben Method
    - Use docker desktop/ORB to run kubernetes
    - Scripts to port forward RabbitMQ/Postgres/Minio to desktop
    - Use pycharm with special app config to connect to these port forwarded services
  - Going Forward
    1. Building a small easily installable ServiceX that anyone can run on laptop or server to test out SX
    2. Develop scripts for above to allow things like running the servicex_app natively, etc.
- Stand-alone commands based on uvx, etc. (Gordon)
  - Good idea
- Using LLM's to write SX code (Gordon)
  - See testing site
  - Hint files
  - Putting hint files into packages (Roger)
- FuncADL unified frontend (Roger) (done)
  - Want the user to be able to pass a query, dataset, etc., and then it directs automatically to the correct backed.
    - We need to be careful about being too clever
    - People will need access to the deeper layers
    - But - as an intro to get into things started, this is perfect.
    - Could MCP/LLM help here?

        - Roger should read up about Analysis Description Language (Harrison Prosper)
    - Config.yaml support for FuncADL xAOD (Roger) (done)
        - Routing data through uproot backend to process funcADL specific cuts
        - Diagram
        - What features would we need to test this outside servicex?
            - Read from S3 - a new DID finder (beware signed uRLS expiring, etc.)
            - Side car could sign url just before handing to science image
        - Would this help with joins?
        - Should be very similar to the TopCP codeGen
    - Dataset discovery options (Dmitry) (done)
        - CMS Rucio/DAS - probably needs new DID finder
            - ACTION: Dimitry and us (Peter) need to have a sitdown to understand the way CMS uses rucio
        - XRootD w/ wildcards (not working at Purdue EOS)
            - ACTION: bug report (authentication)
    - Options to store outputs (Dmitry) (done)
        - Handle just via client downloads to target area? (client cooperates with S3 by default to purge downloaded datasets?)
            - Plan to iterate on a plan
            - Directory structure
        - Can work now: user's run script can download files to local final position with complete flexibility.
    - Query optimizer (analysis_utils)
        - Awkward type tracer running as a highlevel array with same form
        - Must run once through all the code to catch needed branches
        - Issues:
    - Deliver function Error handling (done)
        - Guardlist parsed to deliver output vs runtime error?
        - Empty list issue 585
            - See proposal here
    - Deliver command line error handling (done)
        - The dump when there is an error - error and then 100's of successful files, you can't find what the error
        - **ACTION: Clean up the file list**
    - Could we use the Kibana API to fetch errors? (done)
        - Github Issue
    - Systematic Errors
    - Simplifying servicex yaml fetching/login
        - ACTION: `servicex login` workflow
        - Issue to make this as easy as Globus CLI
    - User-specified compression

        - ACTION: uproot-raw query support. Secondarily: support in other transformers. Unlikely to do : support in sidecar.
    - New file formats
        - ACTION: investigate support for feather v2 in our libraries. Discuss file format enum and extension support. Restructure sidecar translation to go through common awkward format.
    - CLI review
        - ACTION: review commands, whether options are matched to user experience
    - What is wrong with SARA if anything? Xrootd transfers slow. ACTION: studies.
        -
- Technical:(done)
    - Codex Cloud (Gordon)
        -
    - Out Of Memory
        - Should we be programming this for a 2 GB limit?
        - Ephemeral files are a separate limit that gets set, and is memory backed.
        - Func_adl only seen - but why only R25
        - We are pretty sure we see this on prod and af
        - Can we set a limit of 10Gb on ephemeral storage
        - Should we compile this once and then distribute it to all the pods?
- Reliability/Supportability:
    - Current status of OOM errors in prod
    - UX of serviceX backend errors (Roger)
    - FuncADL documentation update (Roger)
        - Making documentation more findable
        - Add links First Examples - ServiceX 3.2.1 documentation
        - ROger will come up with a proposal for modifying the servicex-frontend documentation with easy to discover links to other documentation materials for the queries
        - Also front page of services-frontend
        - Roger and KyungEon - write a new “ecosystem” markdown page
    - Instrument ServiceX deployments with monitoring tools
    - Cleanup Service, take 2
    - What are expected/observed performance bottlenecks of ServiceX? (Dmitry)
    - Make -int work, and reload it when develop containers change
        - Run integration tests against -int?
    - Further integration tests for non-uproot-raw?
        - Certainly catches problems on prod
        - Need to test `develop` backend
        - Servicex-int is broken
        - It would be good to get new `develop` images in -int whenever they change

        - **ACTION: Figure out how to respawn pods when a docker image updates** Can this be done with flux? Maybe make postgres volume persistent and delete the helm deployment
        - Can we extend the tests to check other stuff like xRootD globbing/xAOD FuncADL
        - Clear the datasets to insure we are testing DID Finder
    - Stress Testing
        - Install locust via helm chart
        - Works locally, but problem locust uses gEvent which is incompatible with async
        - Result is that we can't run multiple workflows in a single pod
        - Need some example ServiceX transforms to run in this framework
        - **TODO: Provide Matt with working examples**
        - **TODO: Create an official frontend docker image** useful for integration testing
    - ServiceX Deployment guide is obsolete
- Project Management
    - When should the next meetup be?
        - Will co-locate with IRIS-HEP retreat
        - Next one is UIUC -
- Lessons from the recent ("not latest") ServiceX use case for physics analysis?
    - Interesting effect for one dataset - it was very slow - but a rucio copy rule got the full 1.1 TB to MWT2 in 5 minutes… The protocol used was likely very different from the xrootd used by SX - should we be looking at protocol prioritization by site?
    - This is TopCP - and thus a very practical way to run it. Systematics were run, but since this is BLS, it was a very limited set (e.g. muons).
- ServiceX documentation URLs (done)
    - servicex.readthedocs.io -> servicex-backend.readthedocs.io?
    - servicex-frontend.readthedocs.io -> servicex.readthedocs.io?
    - Yes (keep servicex-frontend valid & redirect?)
- Fix the frontend github action to enable it to publish commit with new version (done)
    - ACTION: do this
- Func_adl xAOD backend produce a new xAOD? (done)
    - Output a skimmed xAOD file that could then be used for other things
    - Specifically TopCP on a skimmed dataset
- New IRIS-HEP integration challenge (done)
    - Goals
        - Part of the AGC challenge - Alex and Oksana leading
        - No specific throughput target like the 200gbps challenge
        - Large scale analysis
            - ~ AGC pipeline **upgrade with systematics**, ML inferencing…
            - Benchmark of our utilities
            - CMS & ATLAS

        - Moving away from a single notebook - to something a bit more realistic (frameworks, separate tools, etc.)
        - Framework/infrastructure can be called from notebooks to do quick single variable studies
            - For simplistic idea of this see the sx data fetcher notebooks.
        - Coffea with virtual arrays (NanoAOD / Physlite schemes)
        - Full analysis where everything has been figured out
        - TopCP on grid for systematics -> ServiceX for fine-grain cuts and skimming
    - What new features can SX showcase?
        - Output
        - R&D using SX to explore dataset
        - Create NTuples from output of TopCP
    - Schedule:
        - aiming for talk(s) at CHEP
        - First working version at end of September
    - How SX can be integrated?
        - ATLAS
            - Runing TopCP for ntuple maker
                - Systematics?
            - Skimming and delivering Ntuples (post TopCP on grid)
        - CMS (Dmitry)
            - NanoAOD skimming (recompression?)
                - Can we use gitOps to manage the generation of these skims?
                - A' la pangeo-forge
                - Use flux
                - Run this as a K8s job
                - Where does the output go and how do we tell user
                - Allow per-request location change for downloads
                - ACTION: Ben will find out what the filelist length limit is
                - **Actions: Fix the permissions problem with xRootD DID Finder**
                - **Action: Get xRootD writes to work in sidecar**
            - Using SX instead of CRAB with MiniAODs?
                - Main use case would be joins
                - Officially deprecate miniAOD code Gen
                - ACTION: cmssw transformer -> enabled: false

- Milestone designation sprint?
-