SWAN: a Service for web-based analysis at CERN

Diogo Castro
On behalf of the SWAN team

https://cern.ch/swan

May 15th, 2024
CS3 JupyterHub Community Technical Workshop 2024

CERN

Introduction

SWAN's building blocks

Main user communities

- Physics analysis
  - Usually last stages of analysis
  - Interactive, exploratory
  - Collision event data, ntuple-like, columnar
  - More and more with Machine Learning
- Non-physics analysis (e.g. ATS)
  - LHC studies: extract machine measurements, query machine settings
  - Beam dynamics simulation
  - Query and process LHC logs distributedly via Spark
  - Query and plot monitoring data in experiment DAQ systems
- Education
  - Many schools/workshops use SWAN for teaching

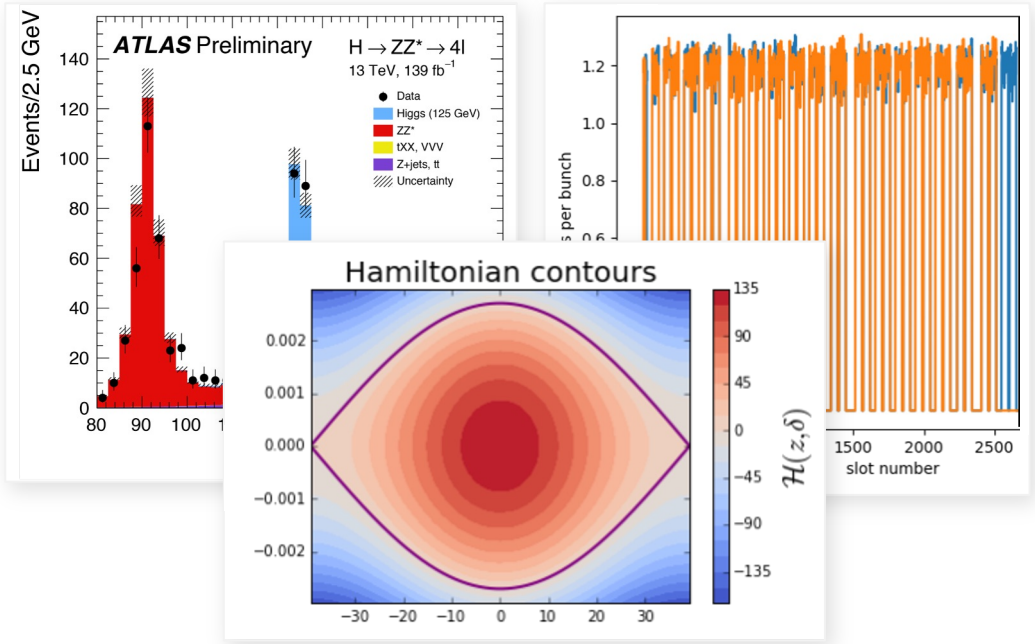

SWAN personalisations

- SWAN re-uses as much as possible from upstream projects
- We maintain modules created for Jupyter and JupyterHub
  - Thin layers to integrate with CERN resources and services
- These modules are released as open-source
  - e.g. Spark Monitor and Authenticator

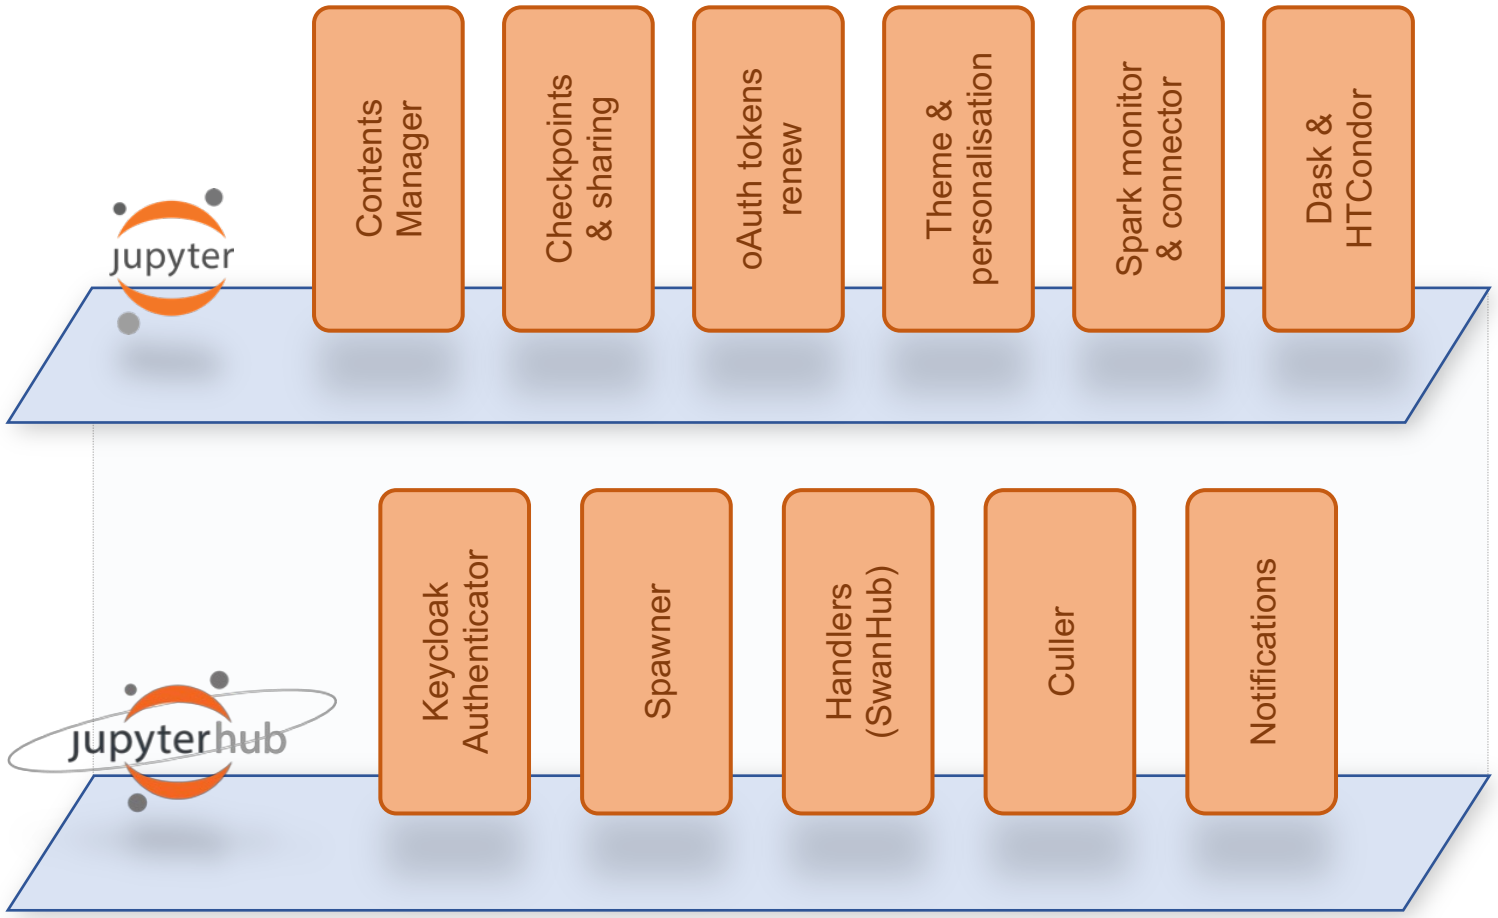

https://github.com/swan-cern/jupyter-extensions
https://github.com/swan-cern/jupyterhub-extensions

Classic UI (Notebook 6)

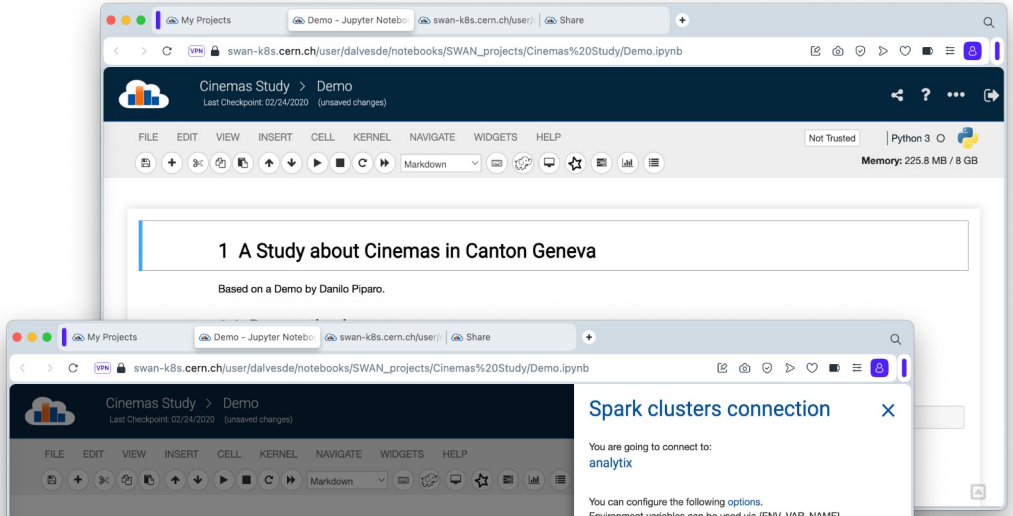

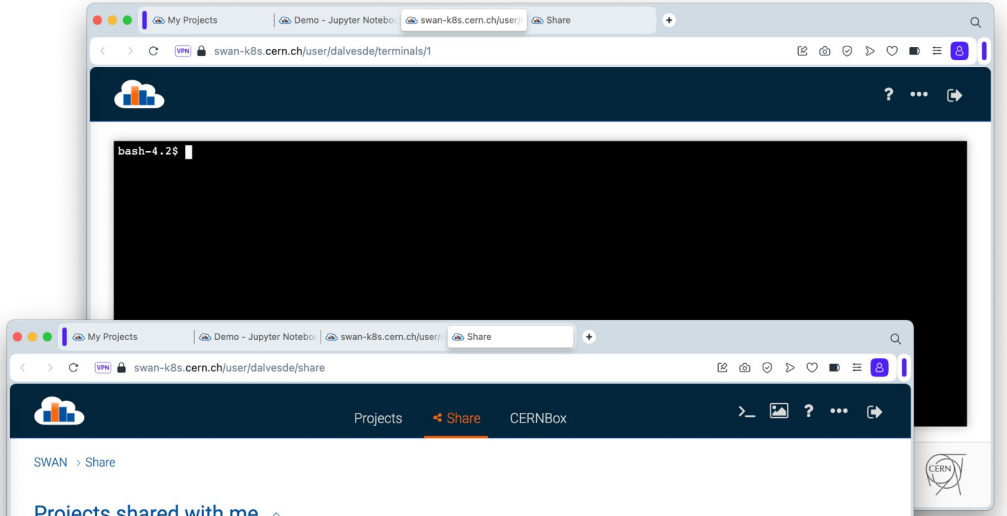

Migration to JupyterLab

- Deployed Jupyterlab v4
  - Extensions migrated to the new version
- Available as beta UI
  - Collection of user feedback underway
  - Users can use the classic UI in parallel
- Missing further integration with CERNBox
  - See next slide

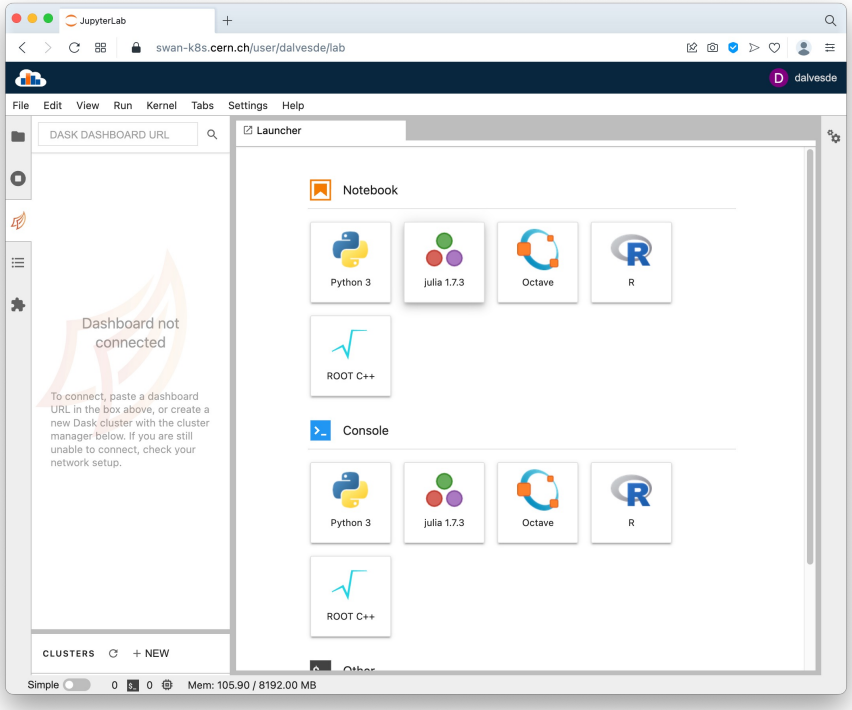

Architecture

Storage

- All the data our users need for their analysis
  - CERNBox as home directory
  - EOS (storage backend of CERNBox) Fuse mounted
  - Also experiments data available
- Sync&Share
  - Files synced across devices and the Cloud
  - Simple collaborative analysis
  - Users can share directly from SWAN's UI
- Lab Extension with full CERNBox capabilities under development

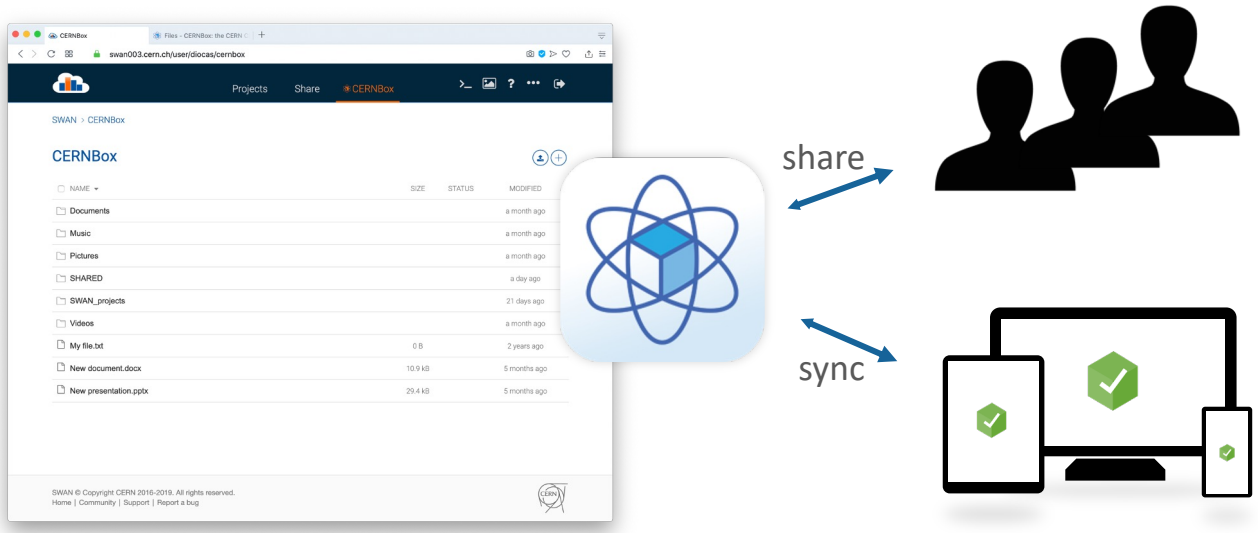

Consistent view across protocol boundaries

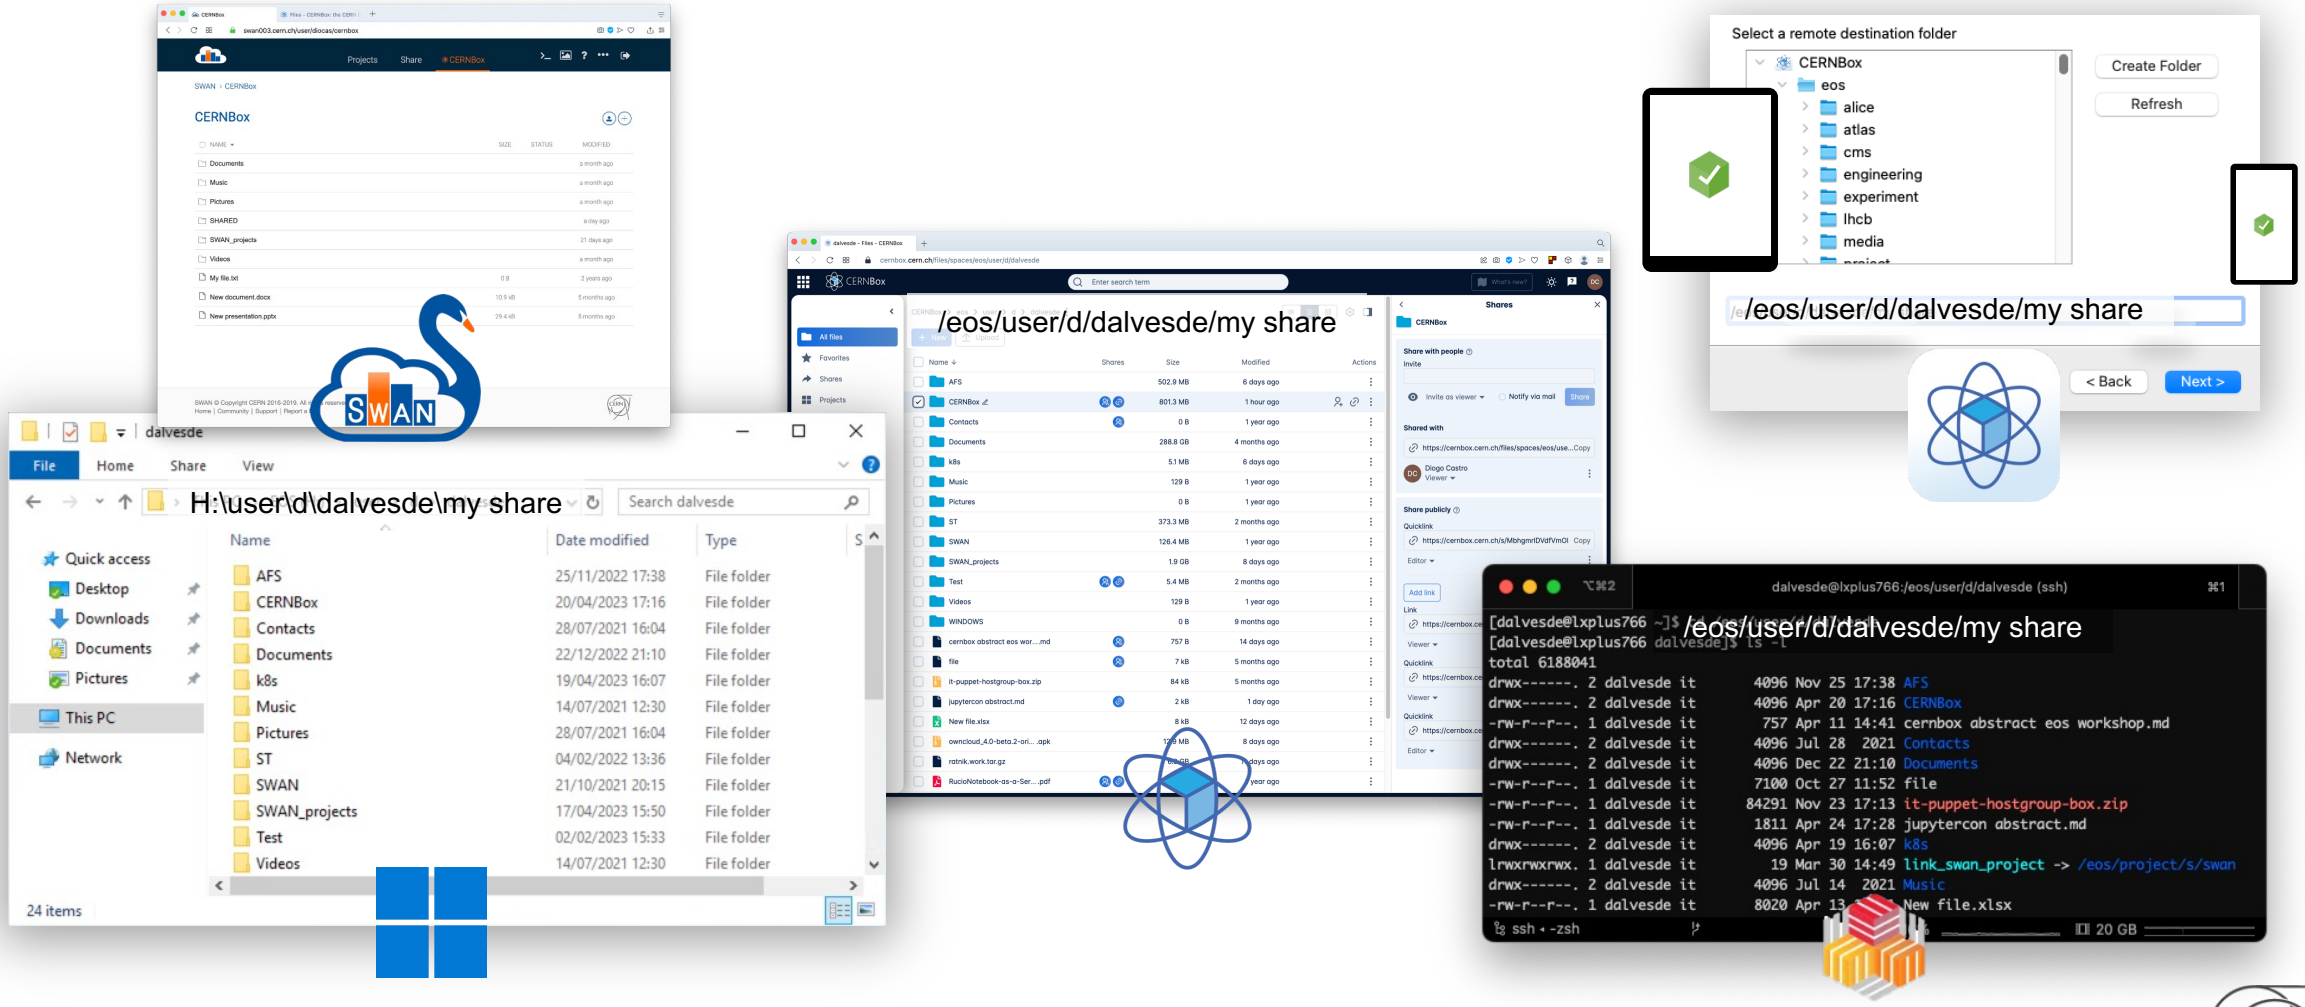

Software

- Software distributed through CVMFS
  - Distributed RO filesystem
  - Immutable software “stacks” maintained by librarians (called LCG Releases)
  - Lazy fetching of software
- Possibility to install additional packages on top of an LCG release
  - Stored on EOS/CERNBox
- WIP: custom software environments
  - Python environment independent of any LCG release
  - Picked during session startup
  - No plans for Binder integration due to security concerns

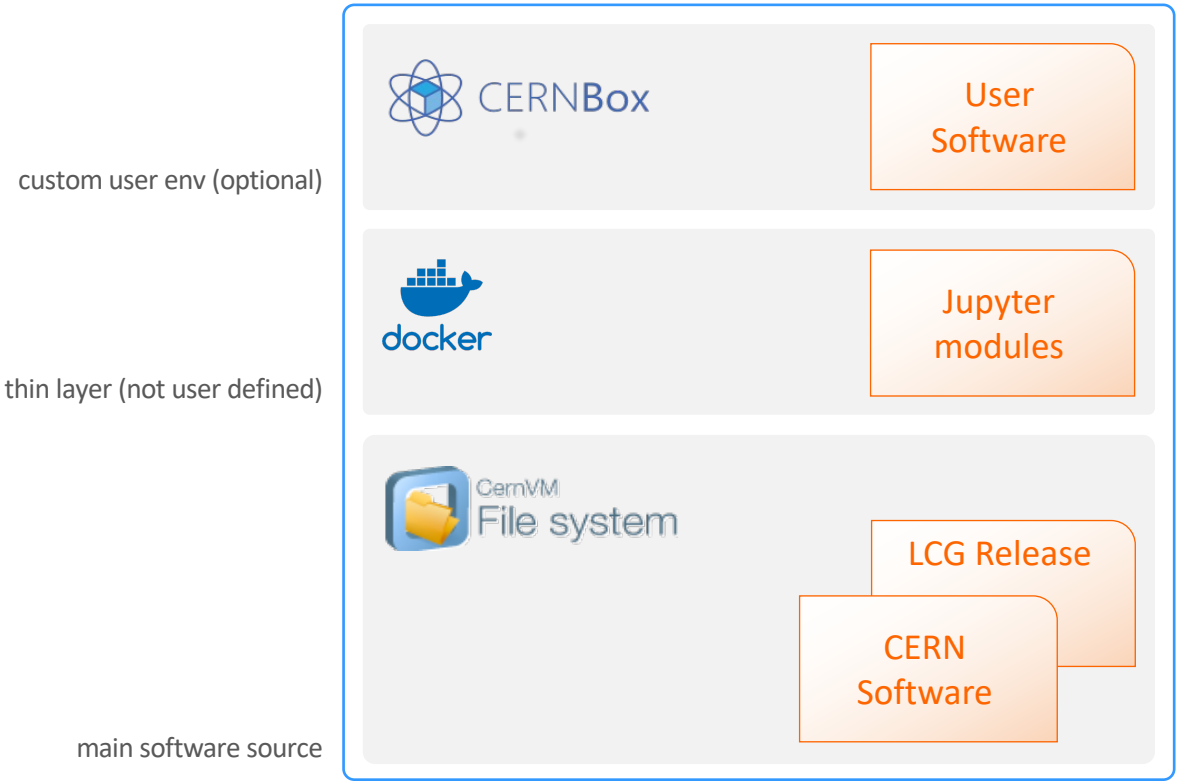

Infrastructure

- Fully running in k8s since this year
  - Initially, it ran in bare metal machines and later in a mix of bare metal + k8s cluster
- Helm charts based on *z2jh*
  - Before z2jh it ran on CERN-developed k8s yaml's
- 2 flavour of charts: *swan* and *swan@cern*
  - Allows deploying SWAN outside of CERN (ScienceBox)
  - Integrate the Python personalisations, EOS and CVMF
- JH 4.0.2

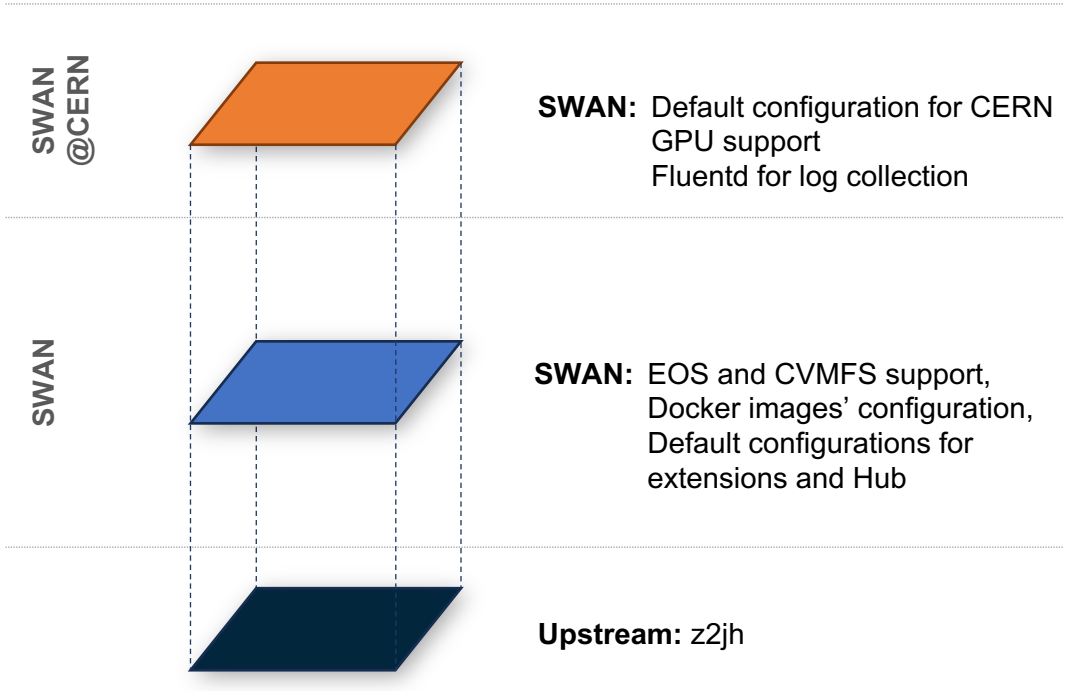

https://github.com/swan-cern/swan-charts

Container images: migration to Alma 9

- Key SWAN container images migrated
  - i.e. user session and JupyterHub images
- User images rewritten from scratch
  - Like upstream images, but Alma, not Ubuntu
  - Same entry points and configuration options
- Modular components' configuration
  - Independently configured on separate scripts
  - Easy to disable or add new components
- More runtime freedom
  - Opens the possibility to use the images in other contexts (e.g. CI)

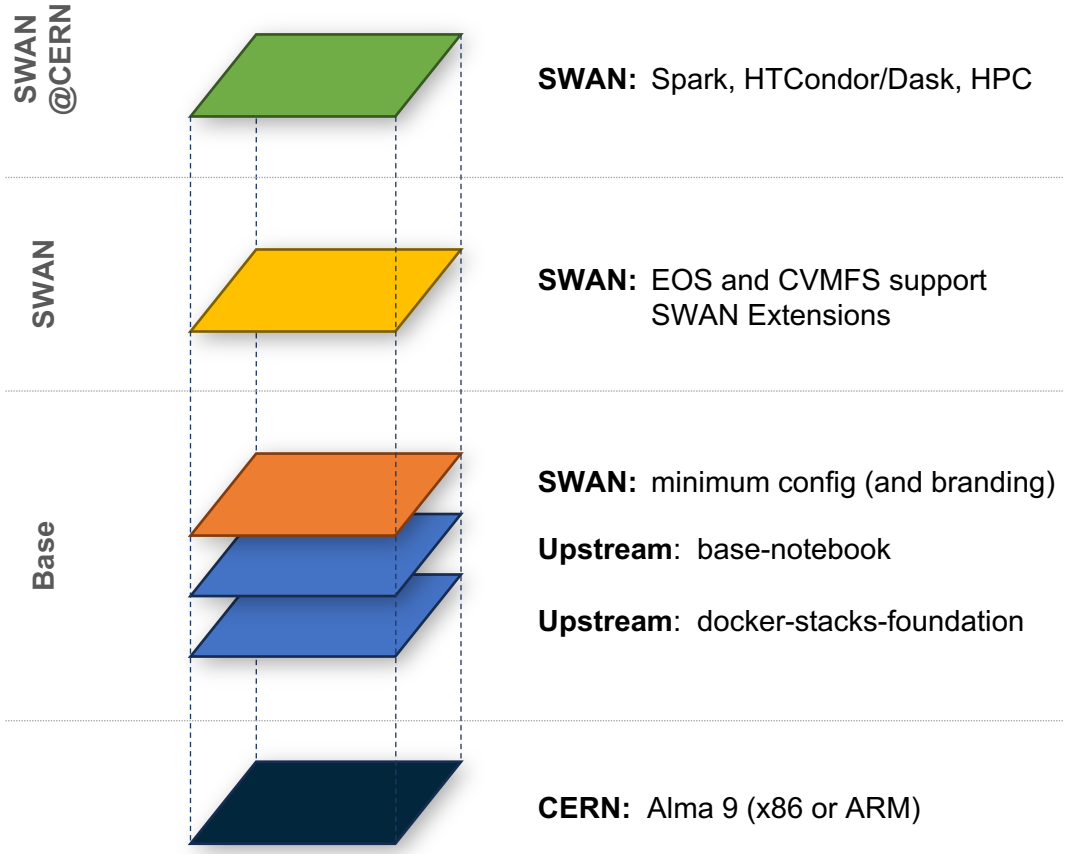

https://github.com/swan-cern/jupyter-images
https://github.com/swan-cern/jupyterhub-image

External computing resources

SWAN as entry point to computing resources

- A user SWAN session gets some resources (cores, memory) for running its notebook / terminal processes
  - 2 cores and 8 GB are the current defaults
- Additionally, SWAN can be used as an interface to access larger computing resources
  - Users launch computations elsewhere and inspect their results in the notebook
  - UI Extensions are provided to make it easy to connect to the external resources
- Current integrations
  - Batch/HTCondor pools
  - Spark/Hadoop clusters
  - HPC/Slurm clusters (work in progress, integration with CEPH FS)
  - GPUs (18 T4s, partitionable A100s for events, in the future access to a shared pool in CERN-IT)

Spark Connector and Monitor

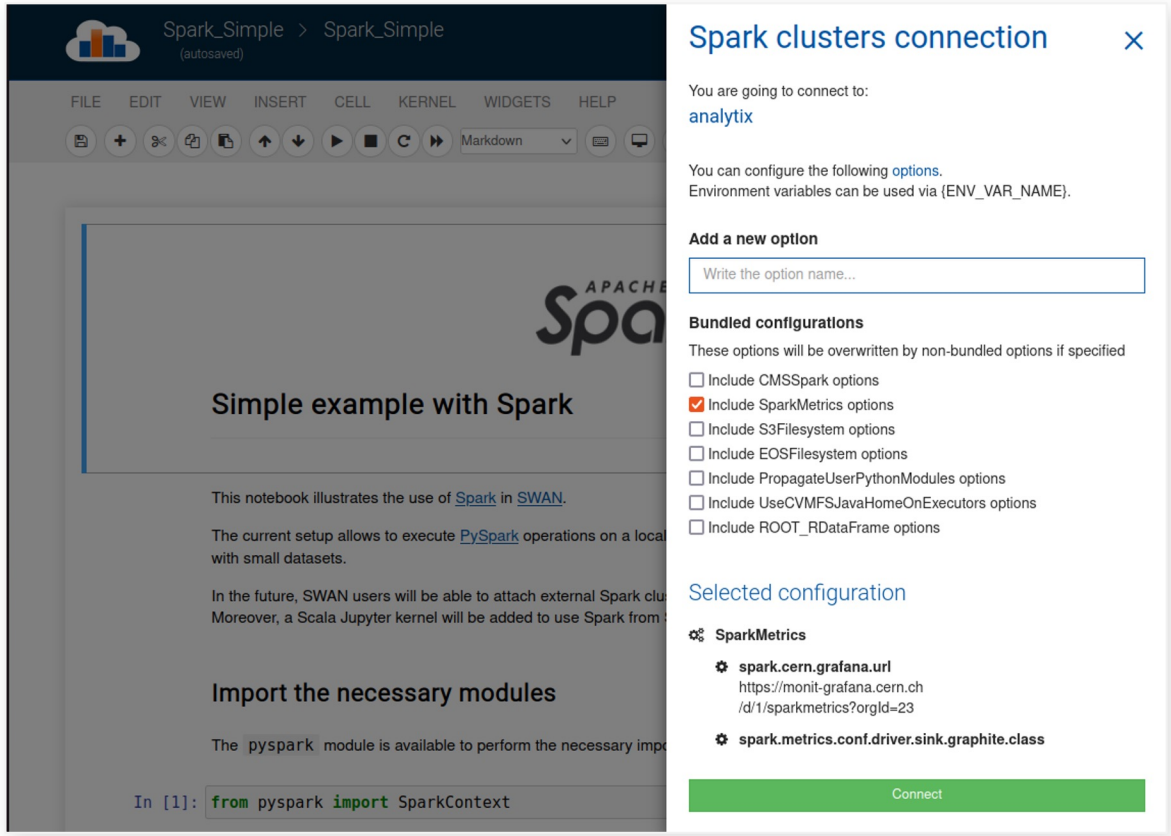

Google Summer of Code

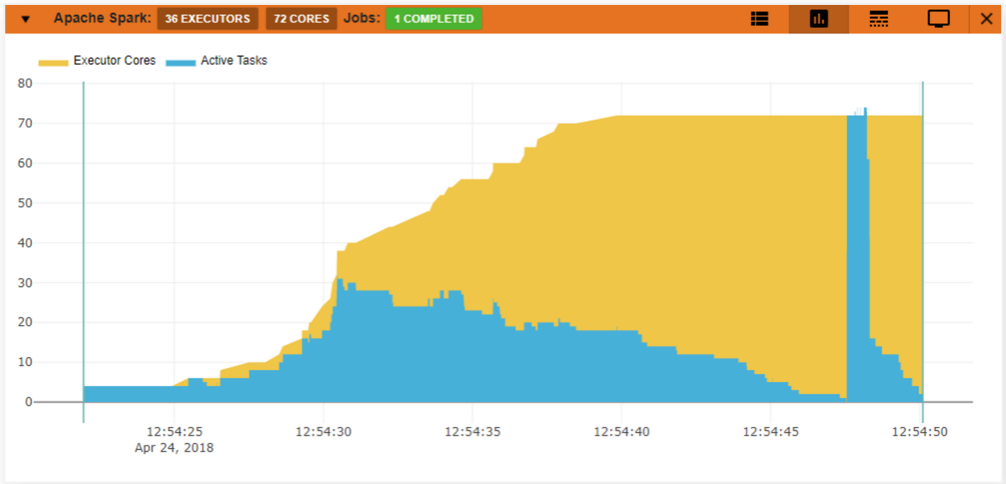

Analysis facility pilot

- Support interactive distributed analysis for High Energy Physics
  - Address the future analysis needs due to foreseen increase in data volumes.
- Dask as the connector to batch resources
  - The two main HEP analysis frameworks, ROOT and coffea, rely on Dask for running analysis distributedly
- For now, it uses overcommitted “static” slots on HTCondor
  - Optimizes usage of batch resources
  - A well-stacked batch farm with a good job mix can get to 80% CPU utilization
  - Known analysis jobs potential to stack nicely with other workloads to drive up utilization

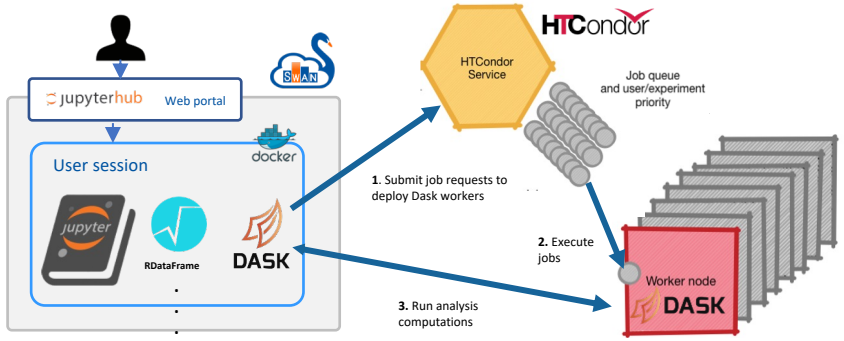

What next?

What next?

- SWAN in the TN (CERN's restricted Technical Network)
  - Includes work in custom software environments
- Finish Jupyterlab migration
- Deploy HPC integration and Rucio Lab extension
- Simplify operations
  - B/G deployment, GitOps, etc
- Investigate the world beyond notebooks
  - IDE

Contacts

Where to find us

- Contacts
  - swan-contact@cern.ch
  - http://cern.ch/swan
  - https://swan-community.web.cern.ch/
- Repository
  - https://github.com/swan-cern/
- Documentation
  - https://swan.docs.cern.ch
- ScienceBox
  - Install SWAN on-premises
  - https://sciencebox.web.cern.ch

Have you heard about REVA?

> https://reva.link

> CERNBox backend
  - Exposes APIs for the new ownCloud OCIS Web UI and sync/mobile clients

> Can have CEPH as storage backend (WIP)
  - Could add sync&share to your Jupyter CEPH FS mounted storage?

SWAN: a Service for web-based analysis at CERN

Thank you

Diogo Castro
diogo.castro@cern.ch

Backup slides

Architecture based on Kubernetes

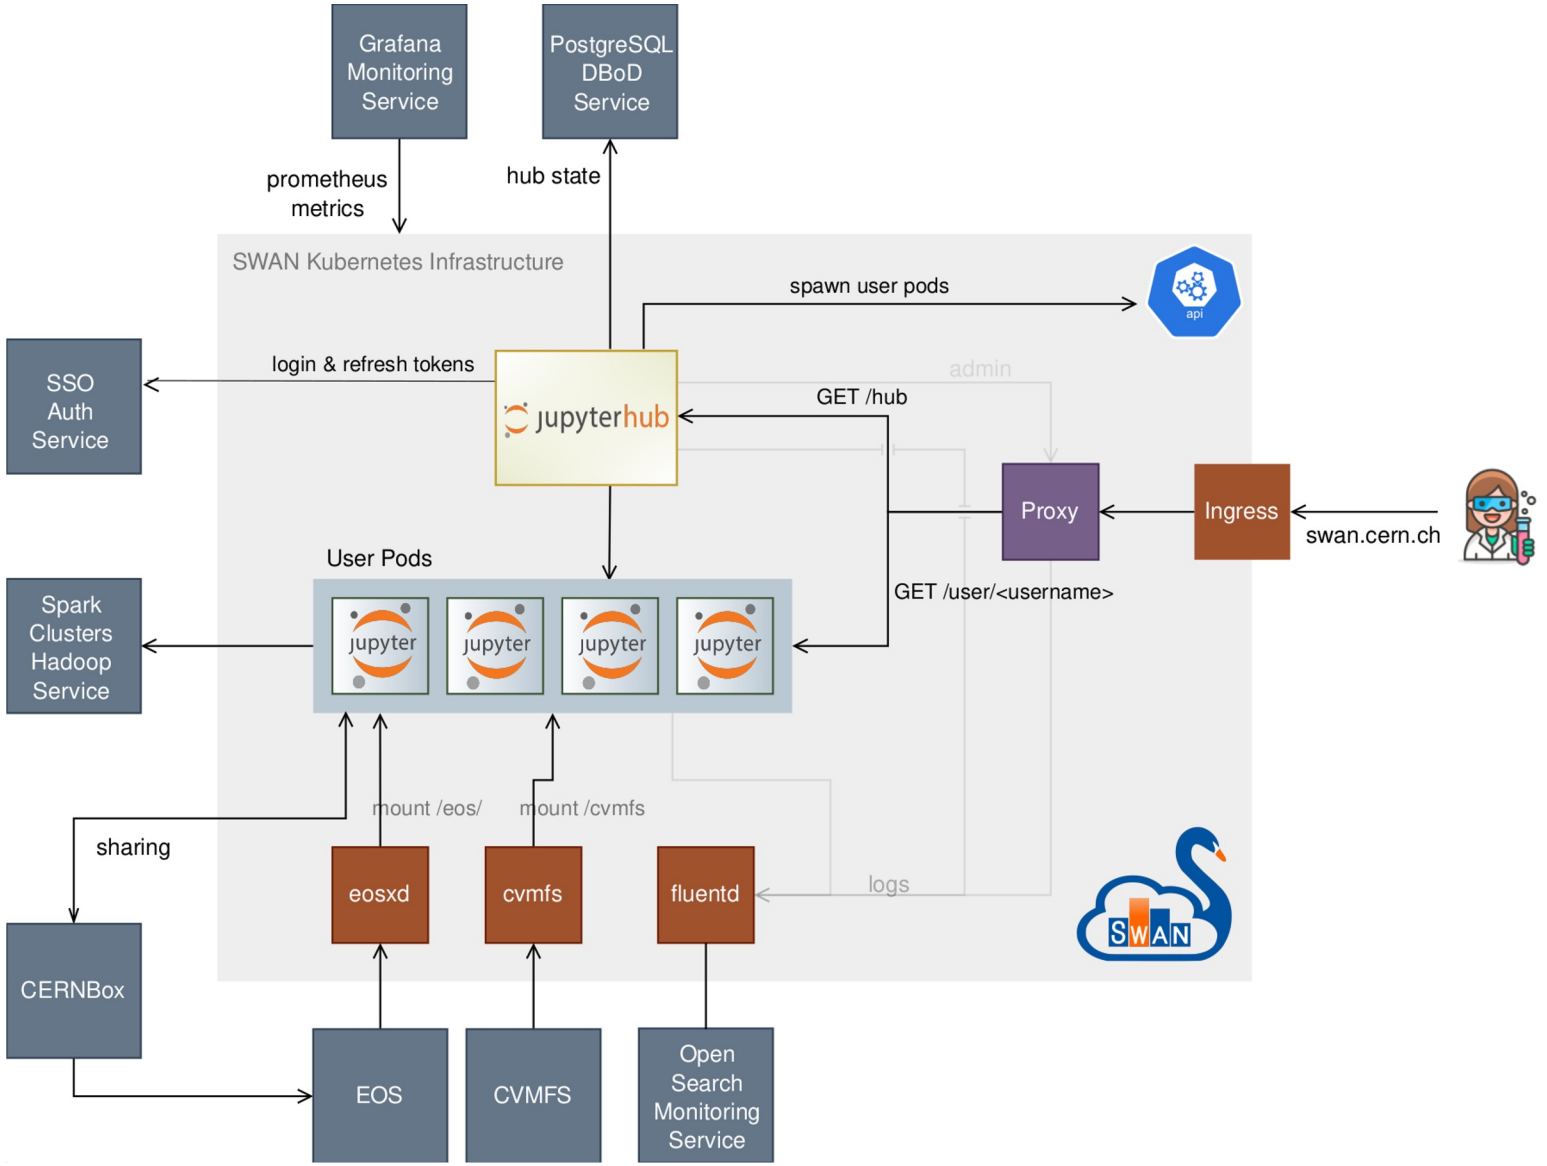

SWAN's Kubernetes setup
Based on Zero-to-JupyterHub

A note on collaboration

Current collaboration model for Jupyterlab

- In the beginning, notebooks could not be open in parallel
  - Conflicts would happen, especially on shared filesystems
- Now they can, and their data structures are synchronized
  - This looks awesome!
  - But optimal usage requires sharing the same Jupyter server and kernel (?)
- Jupyterhub proposes “collaboration accounts” instead
  - “Real-time collaboration without impersonation"

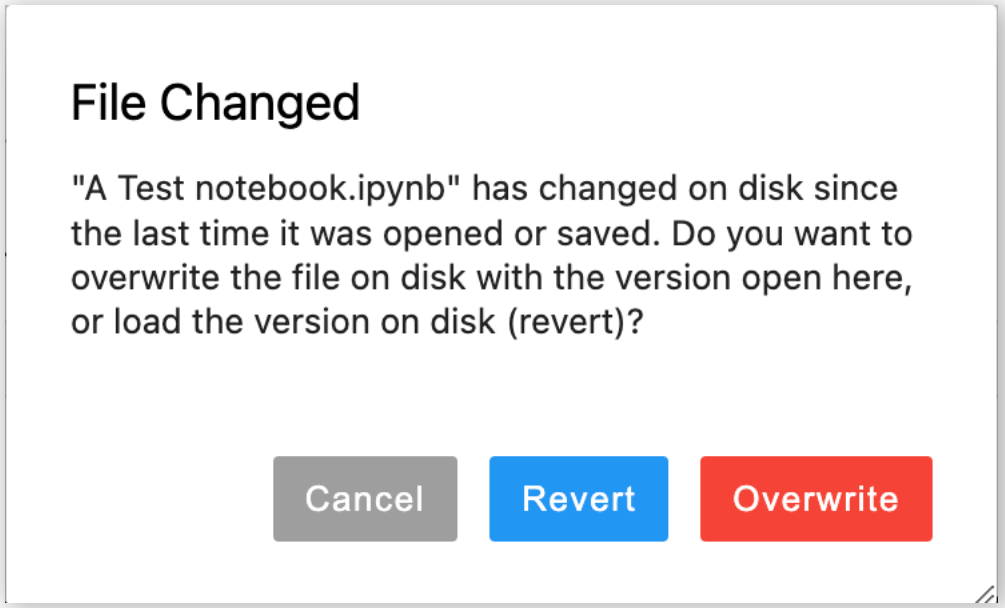

The problems of the current collaboration model

- A shared filesystem might mean access from different Jupyter servers
  - Or even other applications altogether
  - The concurrent editing does not work fully
- Collaboration requires coordination
  - This might not always be easy, especially if we don't know who is editing on the other side…
- Sharing the same server + kernel is risky
  - Full access to another user's account, storage, and permissions on many resources
  - Collaboration accounts help, but might be harder to coordinate or integrate with deployment
- We're not aware of use cases that would benefit from true concurrent editing

We proposed a complementary model better suited for large scale distributed environments

CERN

Collaboration model of the CS3Mesh project

- Same view as EFSS inside Jupyter
  - Access files, different mounts, shares, versions, etc.
- Sharing functionality
  - Share with users or public links
  - Same permissions everywhere
- Parallel access to notebooks
  - As alternative to concurrent editing
  - Opening the same notebook without creating conflicts (both locally or remote)
  - Execution environment independence

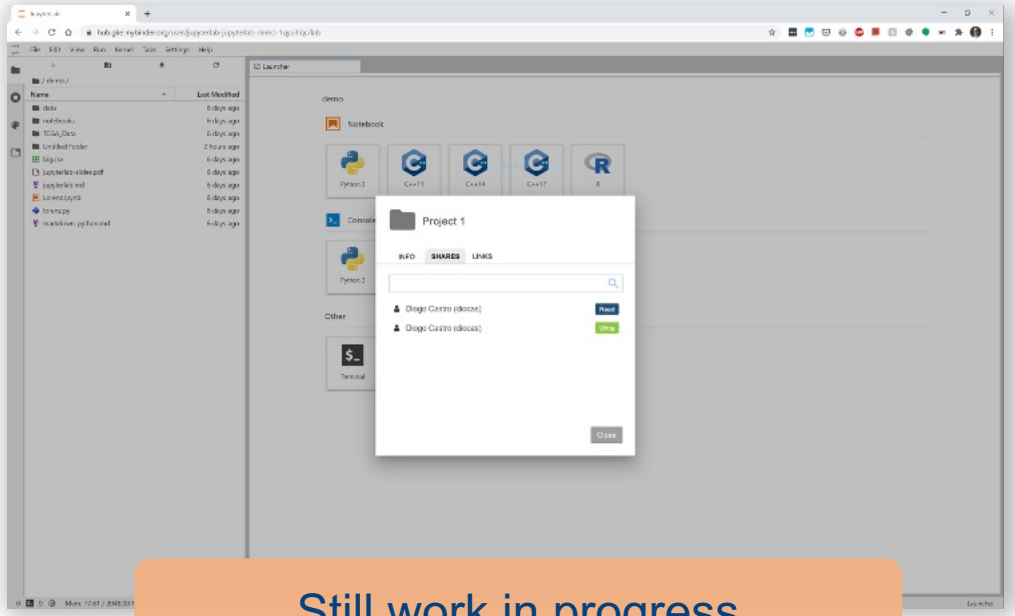

Still work in progress

https://github.com/sciencemesh/cs3api4lab