Facilitating Collaborative Analysis in SWAN

E. Tejedor, D. Castro, D. Piparo, P. Mato
E. Bocchi, J. Moscicki, M. Lamanna

https://swan.cern.ch

July 4th, 2018
CHEP 2018, Sofia (Bulgaria)

Introduction

SWAN in a Nutshell

- Analysis only with a web browser
  - No local installation needed
  - Based on Jupyter Notebooks – interactive computing
  - Calculations, input data and results “in the Cloud”
- Support for multiple analysis ecosystems
  - ROOT, Python, R, …
- Easy sharing of scientific results: plots, data, code
- Integration with CERN resources
  - Sofware, storage, mass processing power

The Pillars of SWAN

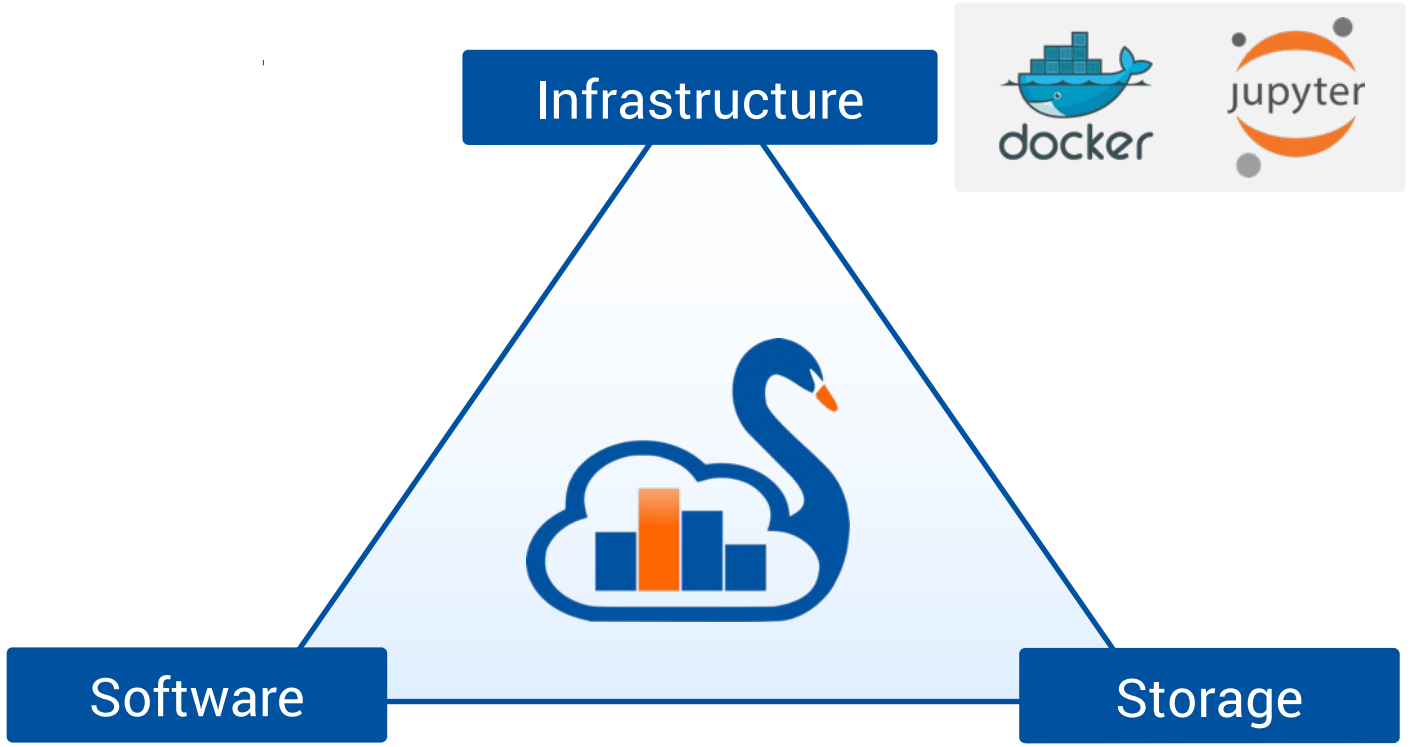

docker

jupyter

CernVM File system

CERNBox

EOS

CERN

Text

Code

Graphics

Simple_ROOTbook_py (unsaved changes)

Control Panel | Logout

File Edit View Insert Cell Kernel Help | Trusted | Python 2

Code

Displaying graphics

We can now draw the histogram. We will at first create a canvas, the entity which in ROOT holds graphics primitives. Note that thanks to JSROOT, this is not a static plot but an interactive visualisation. Try to play with it and save it as image when you are satisfied!

In [5]:
```
c = ROOT.TCanvas()
h.Draw()
c.Draw()
```

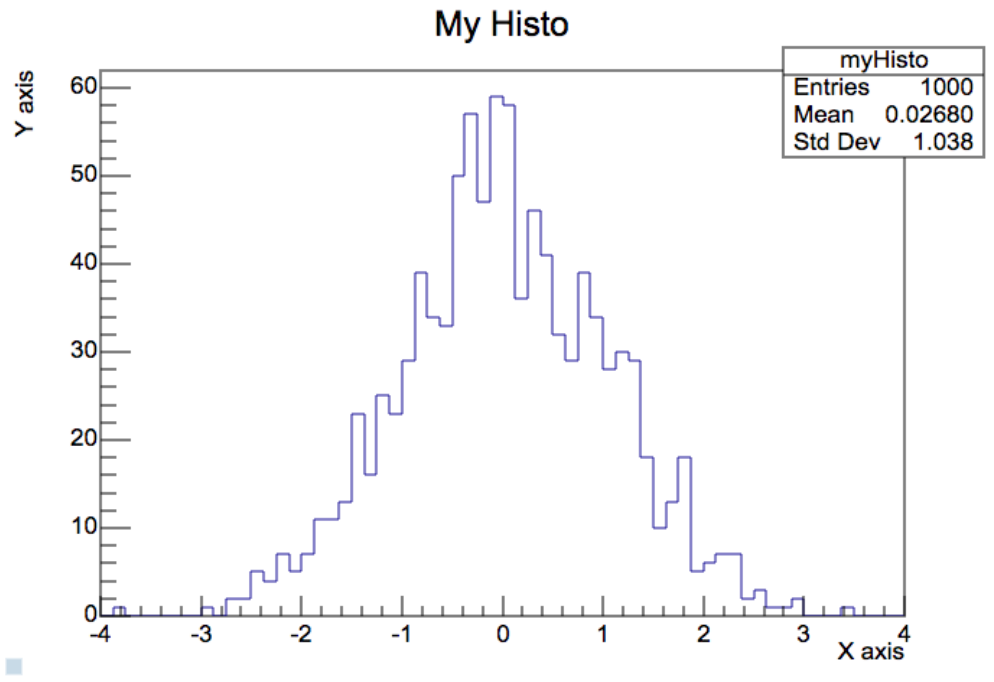

Collaborative Analysis

Cloud Storage as your Home

- CERNBox is SWAN's home directory
  - Storage for your notebooks and data
- Automatic synchronization
  - Files synced across devices and the Cloud
- Provides foundations for sharing
  - Collaborative analysis

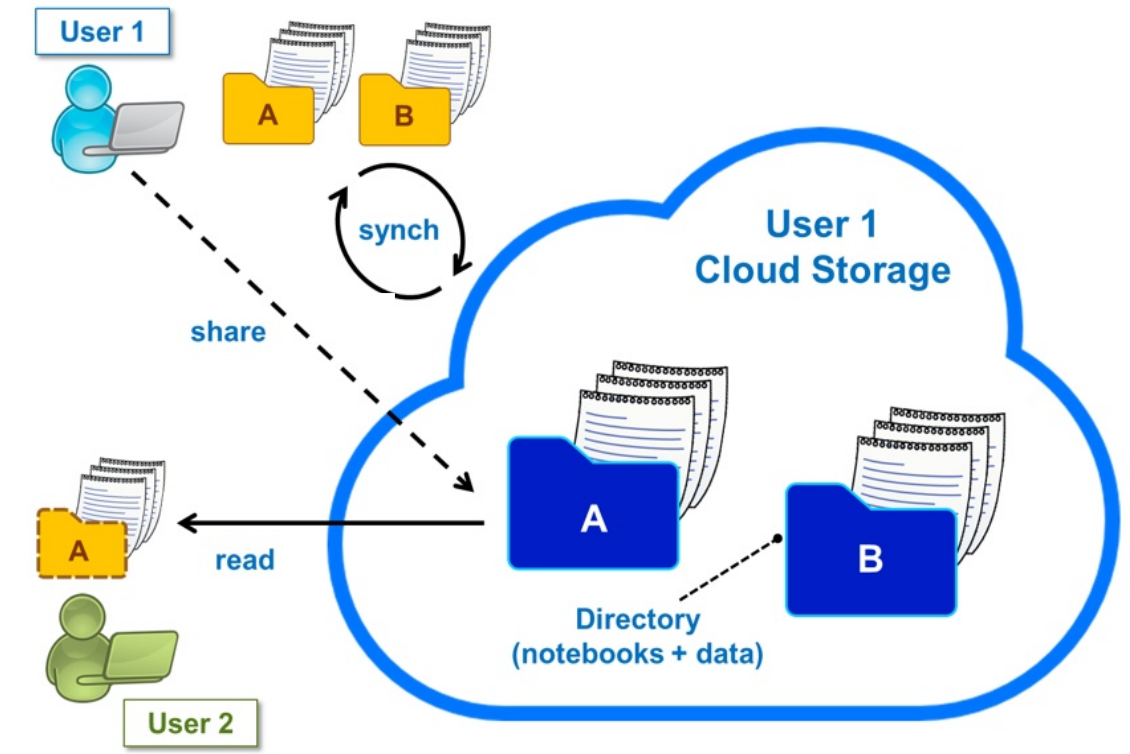

New User Interface

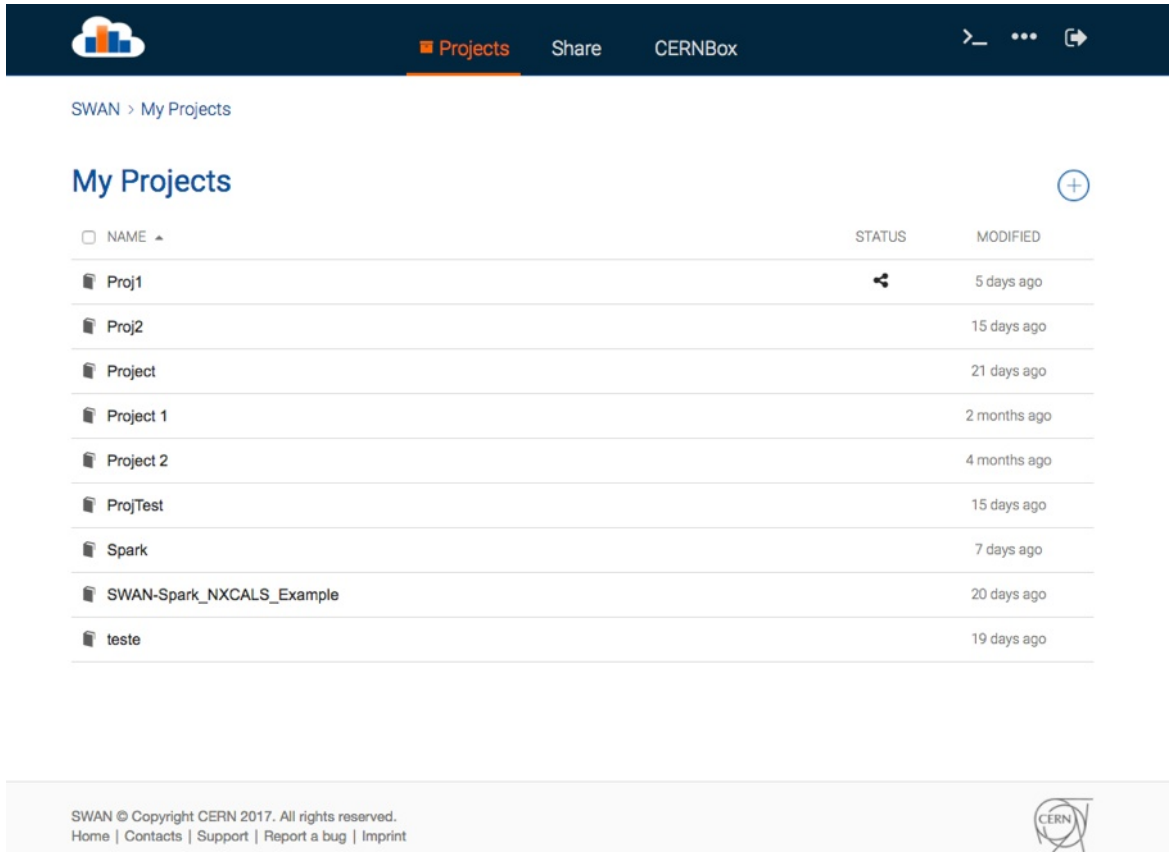

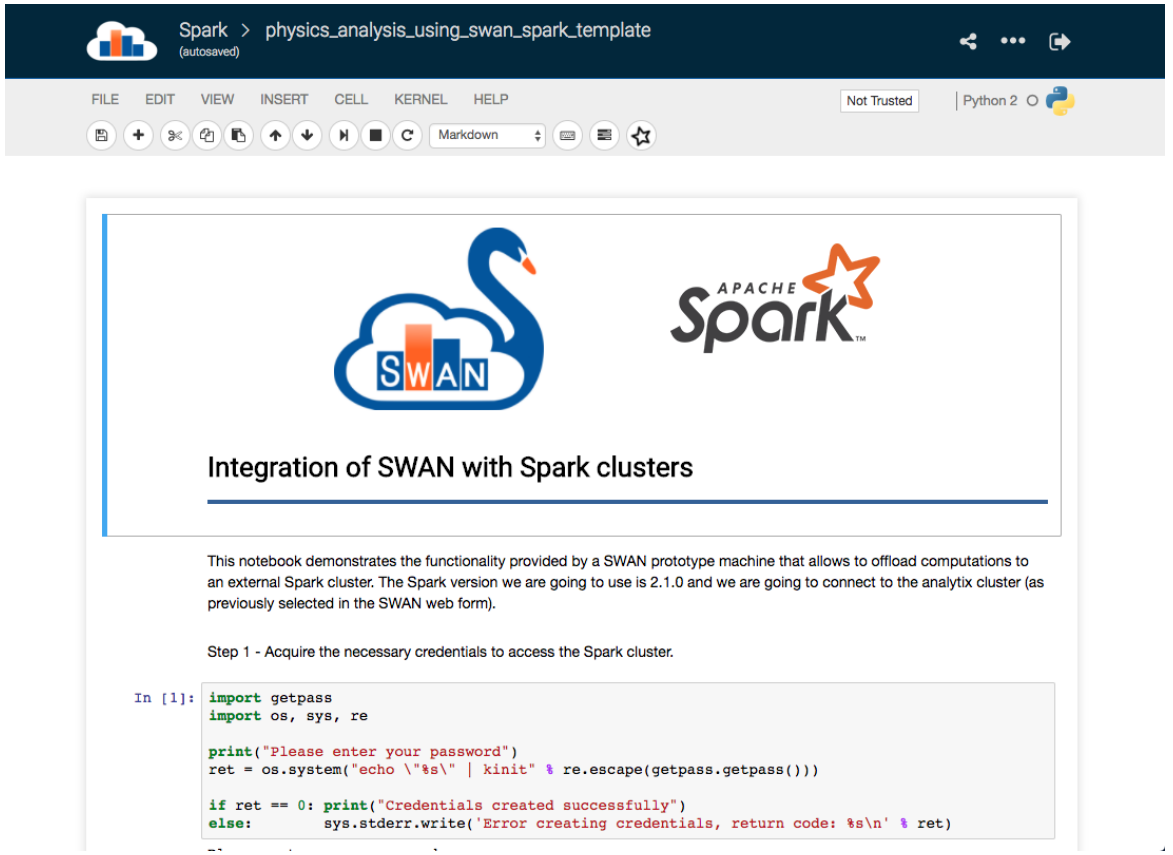

Sharing Made Easy

- Sharing from within the SWAN interface
  - Integration with CERNBox
- Users can share “Projects”
  - Special kind of folder that contains notebooks and other files, like input data
  - Self contained

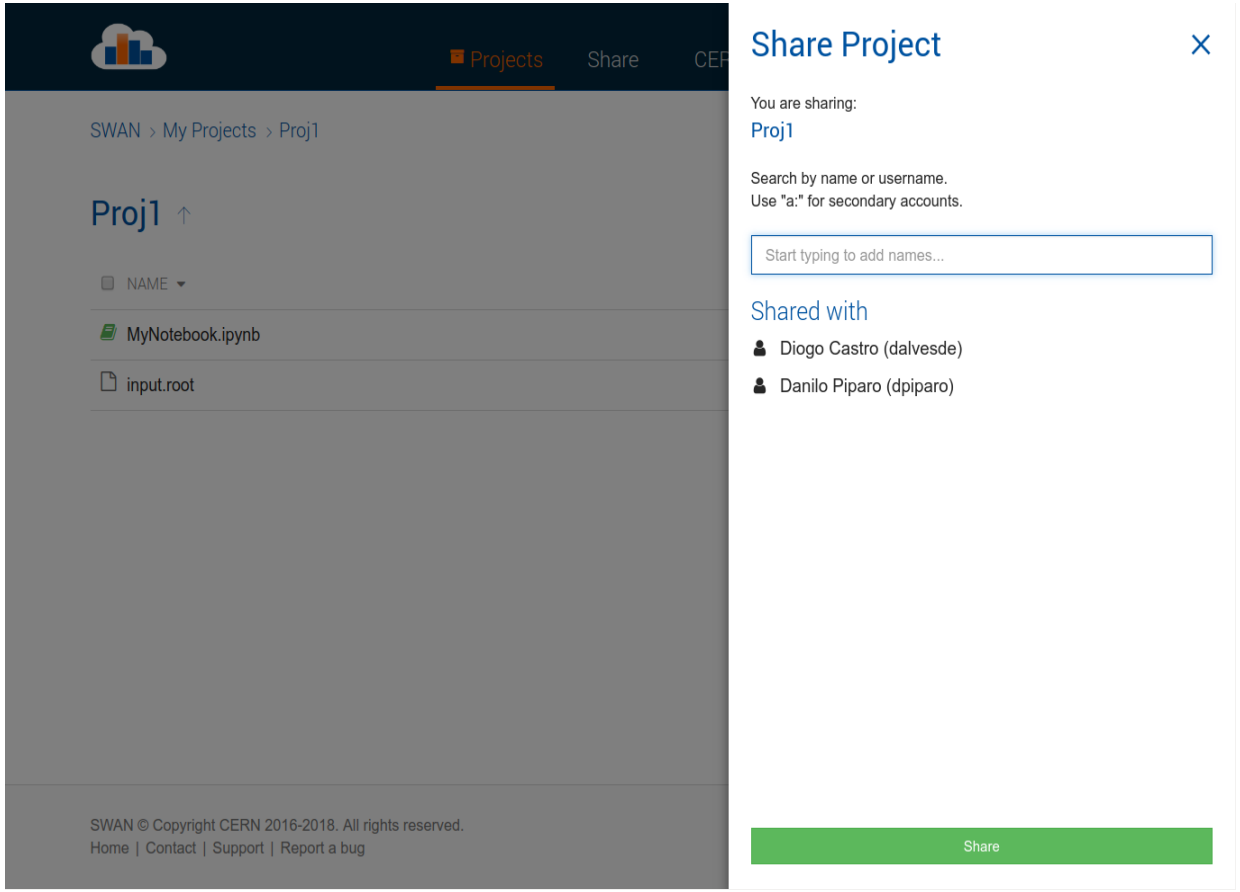

The Share Tab

- Users can list which projects...
  - they have shared
  - others have shared with them
- Project information
  - Sharer
  - Size
  - Date

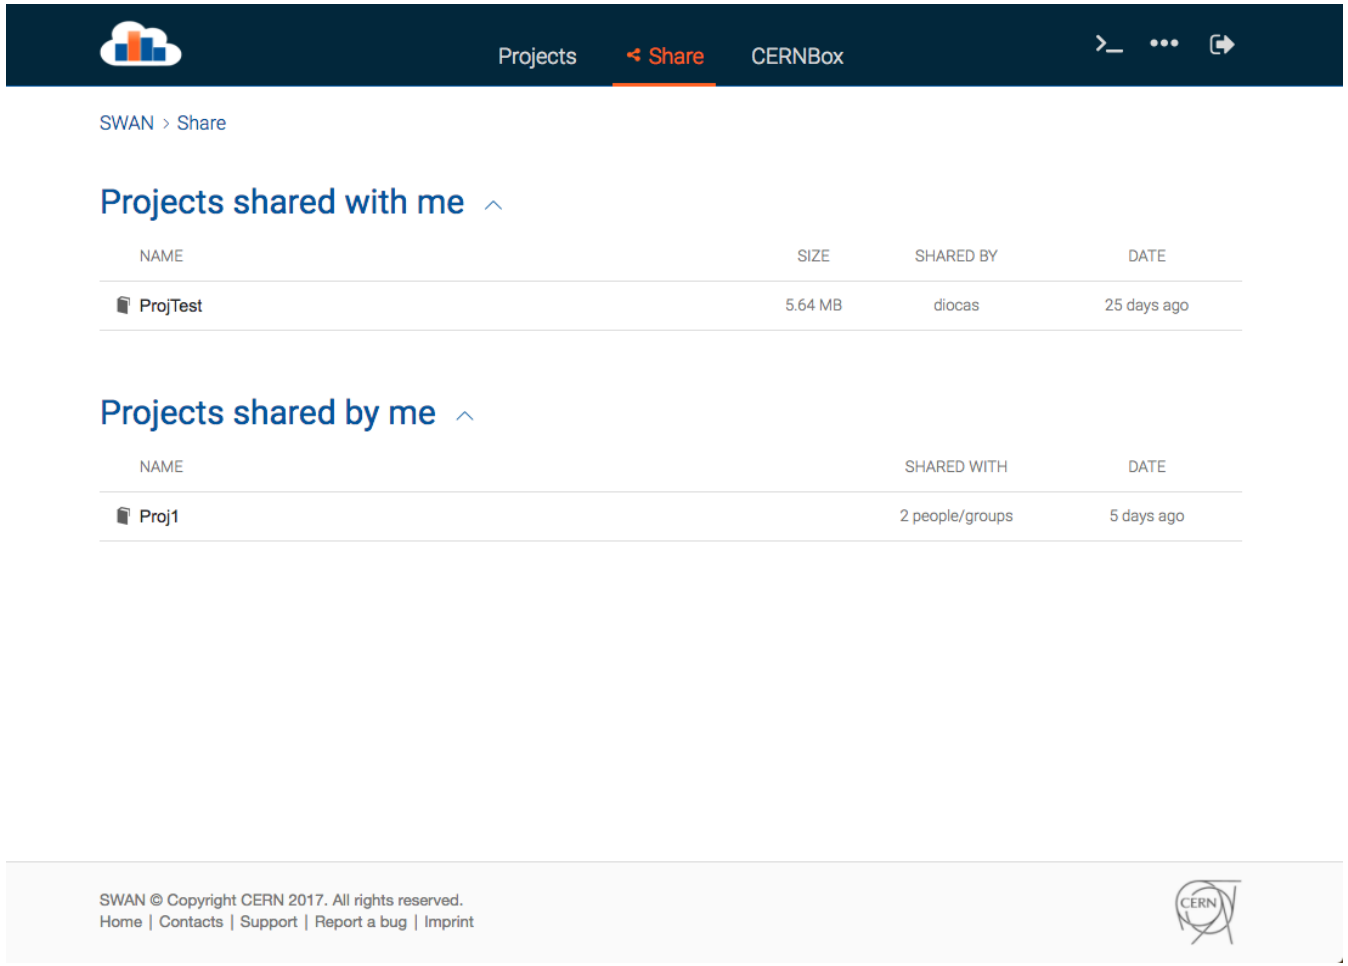

Inspecting a Project

- By clicking on a shared project, a user can inspect its contents
  - Browsing of the project files
  - Static rendering of notebooks
- Useful to decide whether to accept or not the shared project

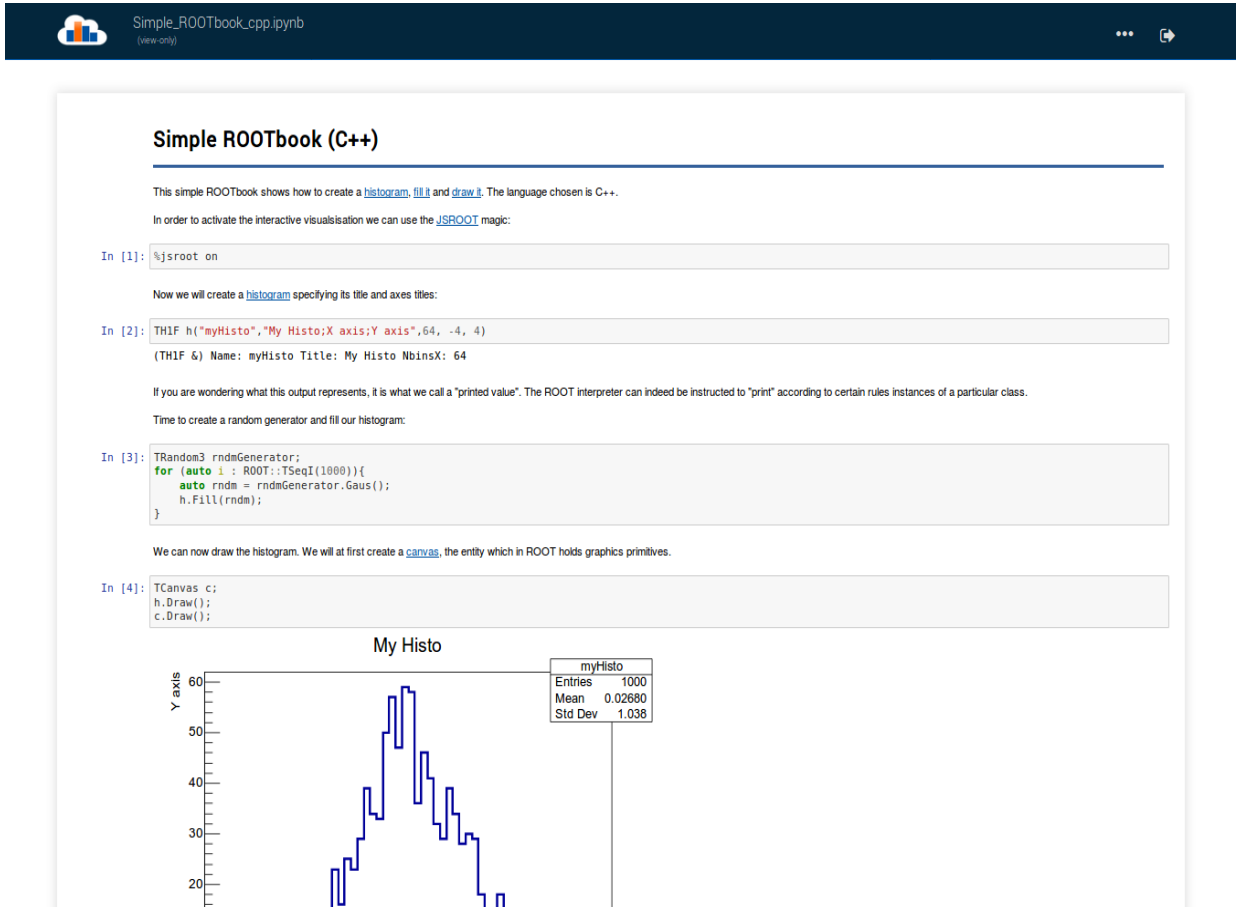

Accepting a Shared Project

- When accepting a shared project, its contents are cloned to the receiver's CERNBox
  - The receiver will work on their own copy
- Concurrent editing not supported by Jupyter
  - Safer to clone

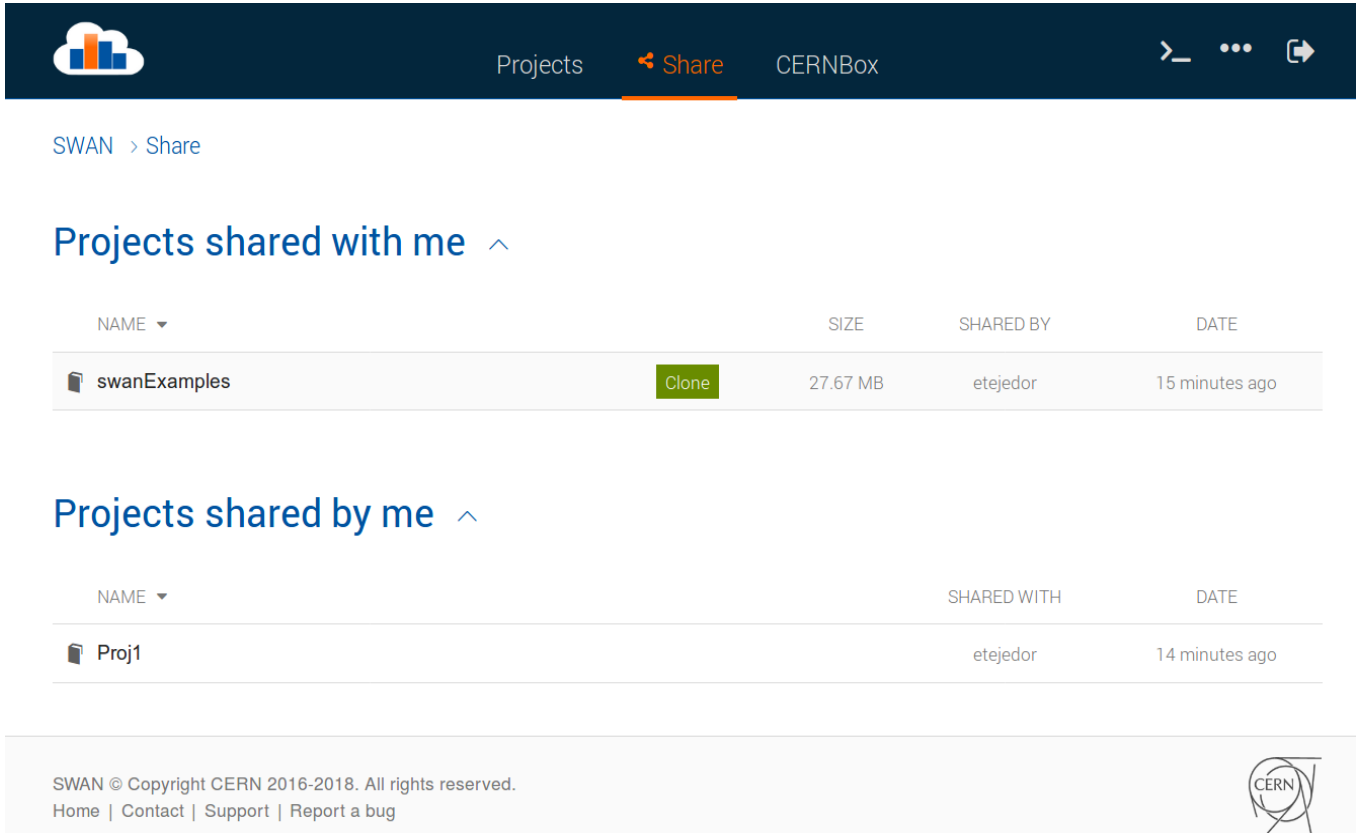

Sharing Spark Projects

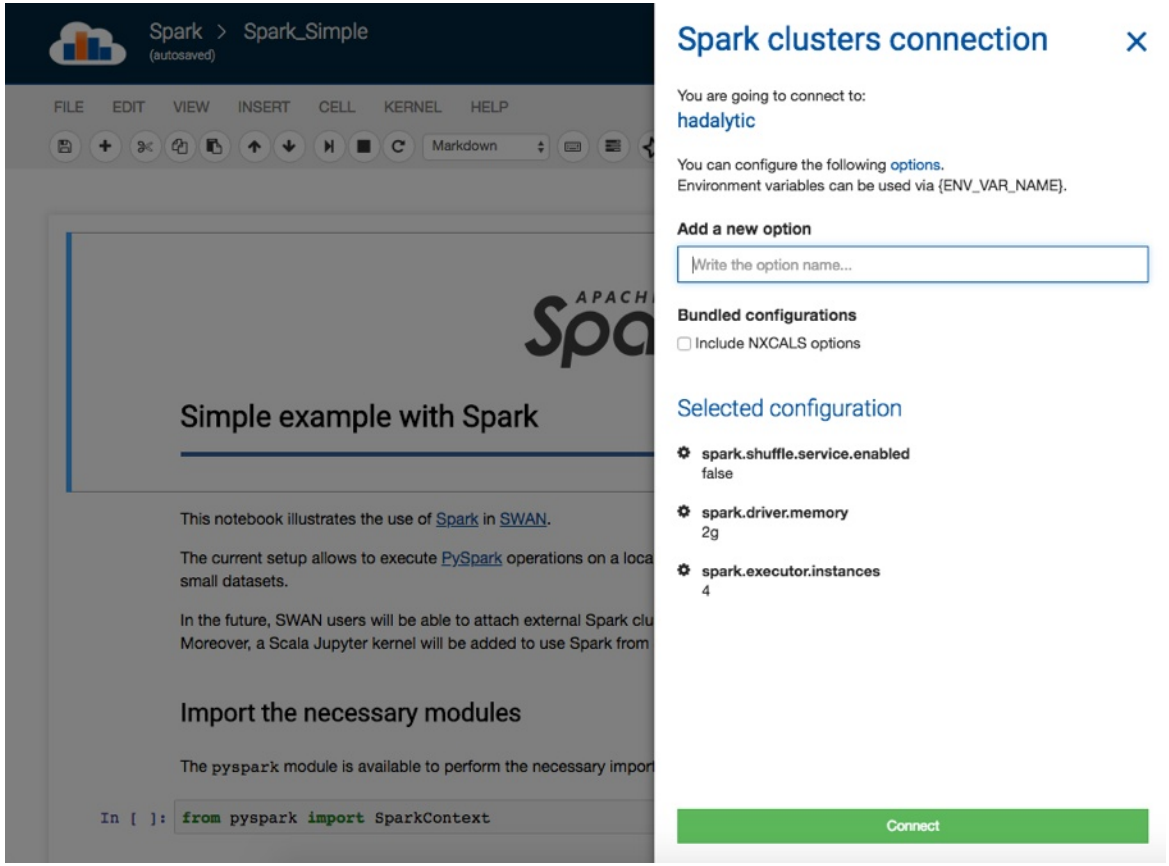

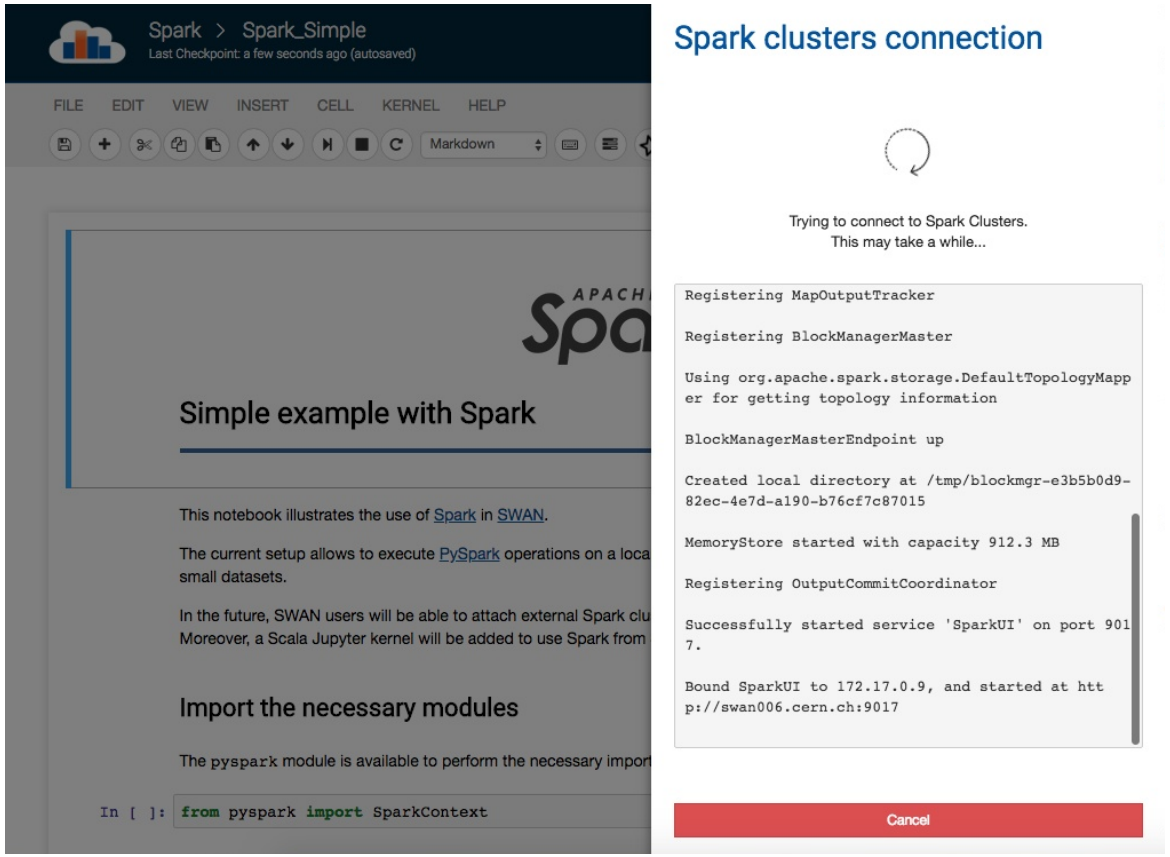

[Talk: Apache Spark usage and deployment. P. Kothuri, E. Tejedor et al]

SWAN for Education

- UP2University European Project
  - Bridge the gap between secondary schools, higher education and the research domain
- SWAN used by students to learn physics and other sciences
  - Let them use the very same tools & services used by scientists at CERN
- SWAN Boxed: distribution easily deployable on premises
  - https://github.com/cernbox/uboxed

Poster: The EU Up to University Project, E.Bocchi, J. Moscicki

Conclusion

Summary

- The interface of SWAN has been completely redesigned to foster collaboration and sharing of results among scientists
- Sharing is now fully integrated in SWAN, where users can share their work in the form of Projects (notebooks + data)
- New functionality to share a project, list shared projects, inspect and clone shared projects into users CERNBox
- Concurrent edition of notebooks in Jupyter would give another dimension to sharing

Facilitating Collaborative Analysis in SWAN

Thank you

Backup Slides

Integrating Jupyter

- Configurable environments through user defined scripts
- Jupyterhub to allow multiple Jupyter instances
  - Single instance of Jupyter per user
- User sessions spawned as Docker containers
  - Enforces resource limits per user
  - To isolate users work

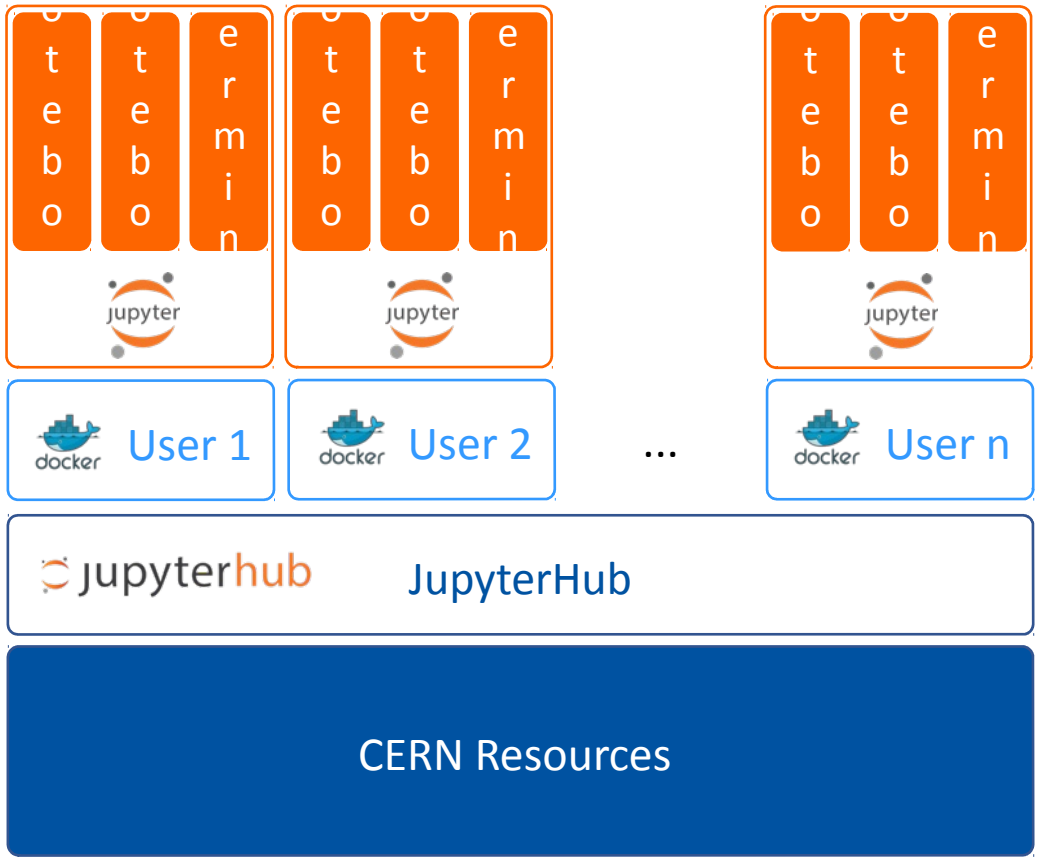

Software

- Software distributed through CVMFS
  - LCG Releases: distribute a series of compatible packages
  - Software used by researchers is available
- Possibility to install other libraries in user storage (CERNBox)

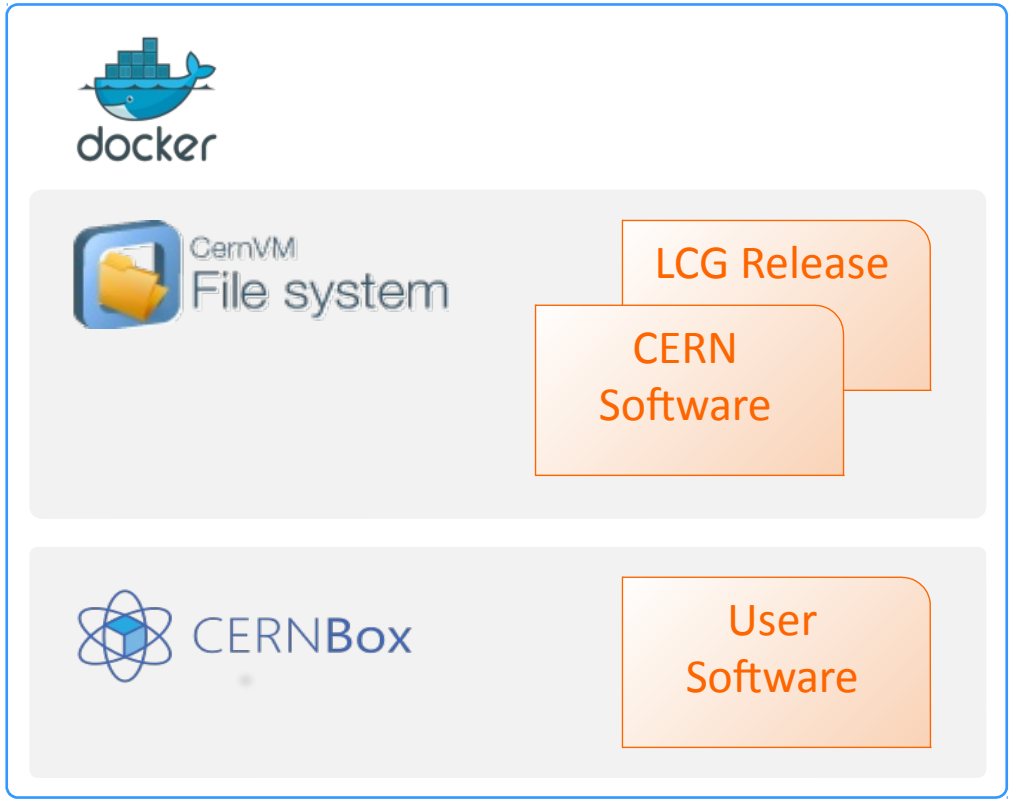

Architecture

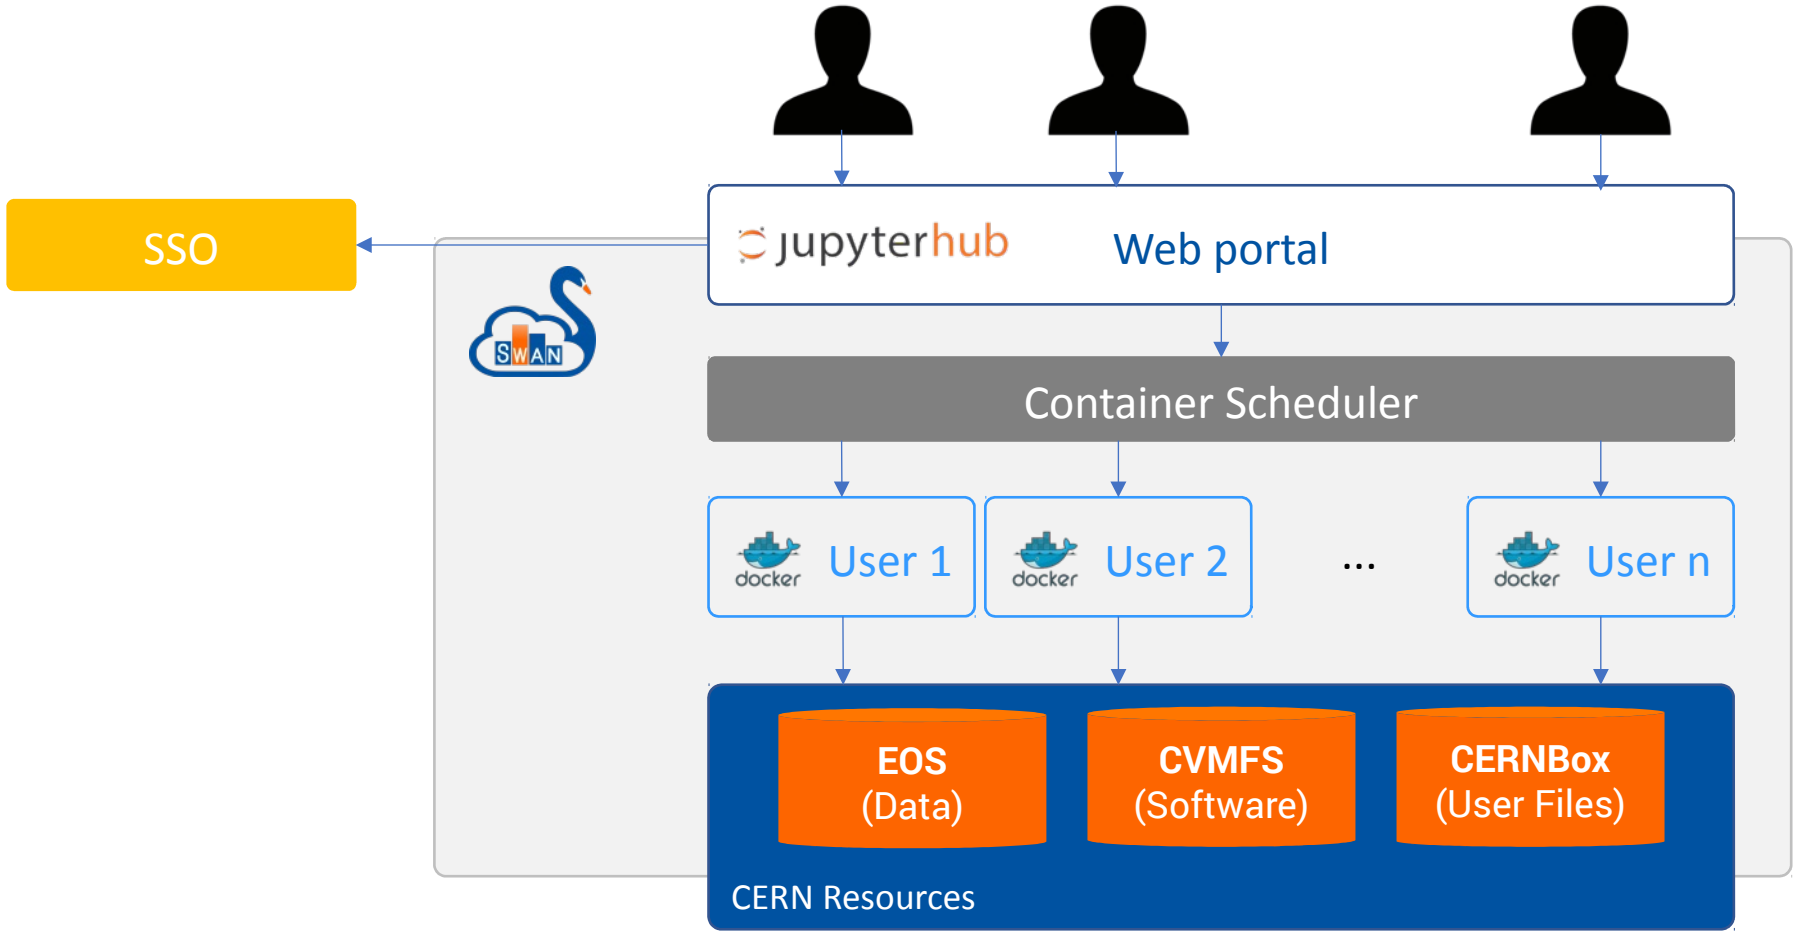

The Notebook as Interface

- A web-based interactive interface and platform that combines code, equations, text and visualisations
  - Ideal for sharing/collaboration
  - A “shell opened within the browser”
- Many supported languages (kernels)
  - In SWAN: Python, ROOT C++, R and Octave
- Interactive, usually lightweight computations
- Very useful for some use cases
  - Final steps of an analysis, exploration, teaching, documentation and reproducibility

New User Interface

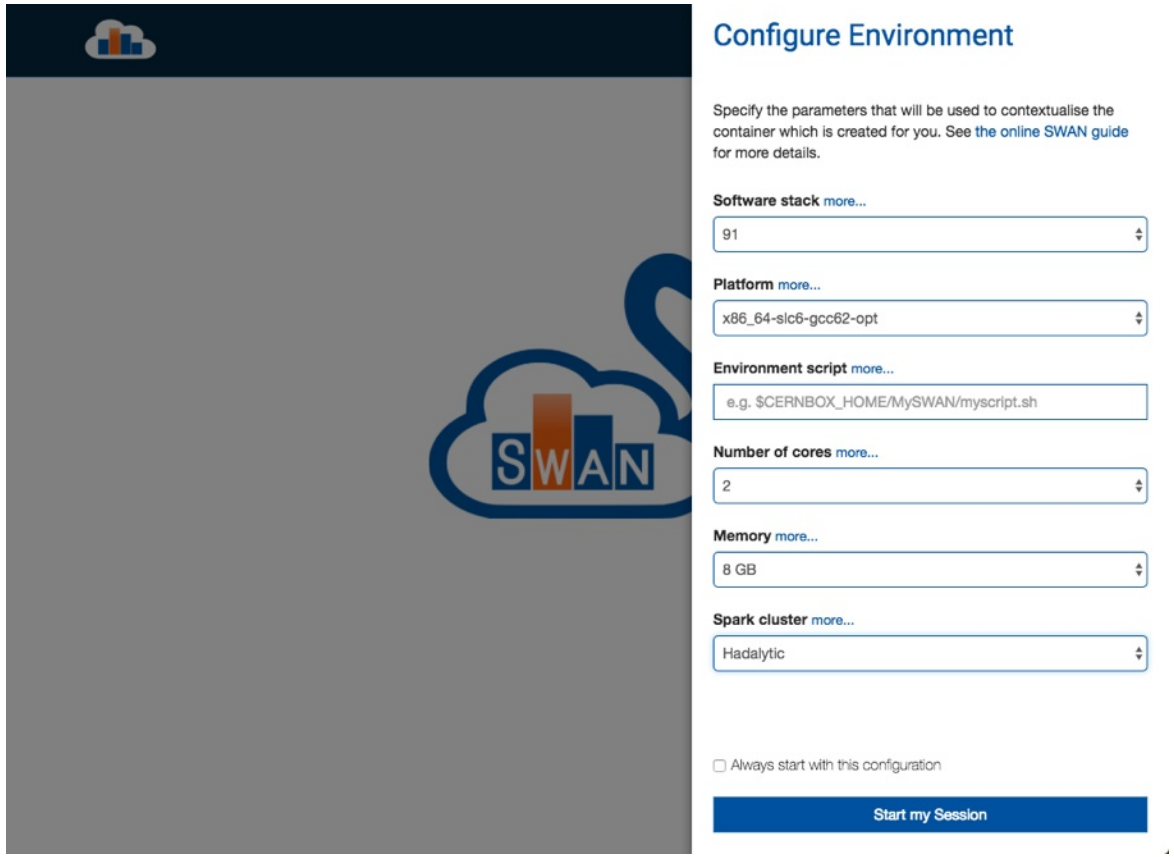

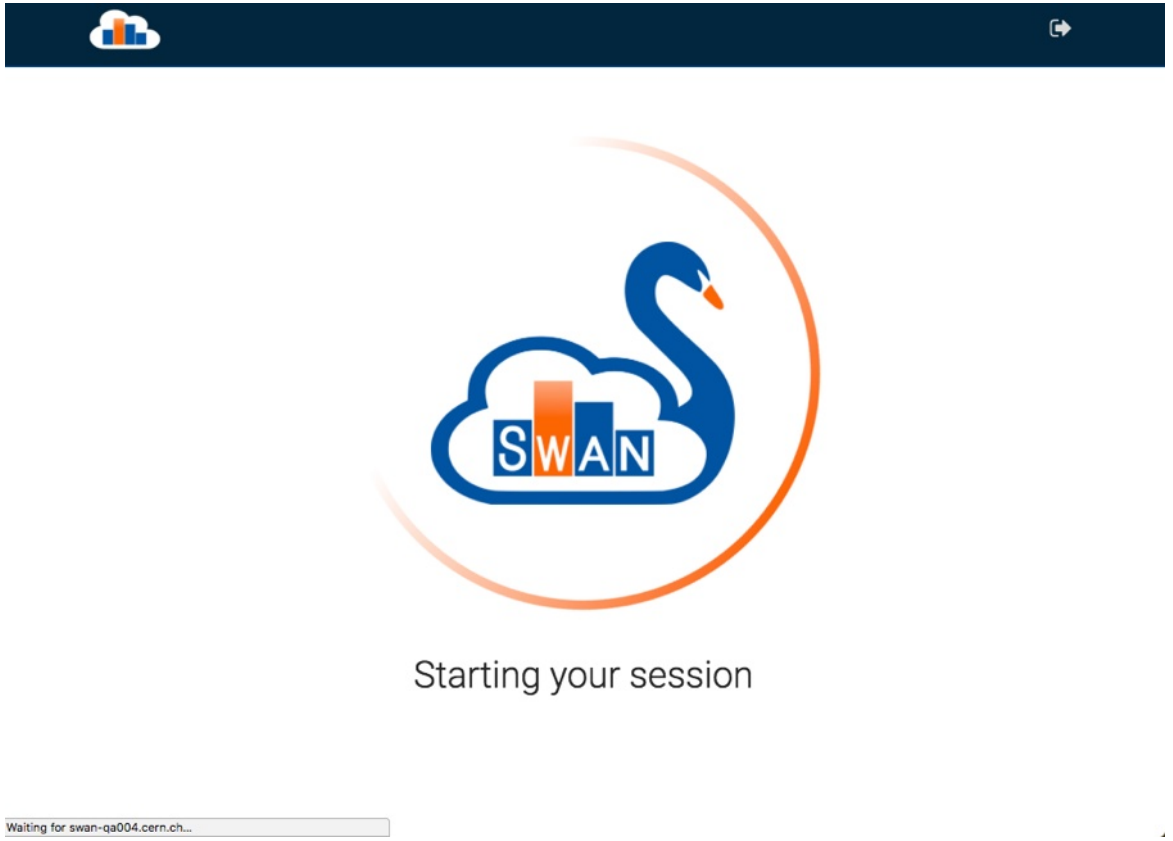

Boxed

- Containerized version of all the infrastructure
  - Includes EOS, CERNBox, CVMFS and all Swan services (Jupyter Docker image, JupyterHub)
  - Available in https://github.com/cernbox/uboxed
- Easily deployable on premises
  - Installable in Linux systems
  - Based on Docker Compose

Integration with Spark

- Connection to CERN Spark Clusters
- User data accessed through EOS
- Graphical Jupyter extensions developed
  - Spark Connector
  - Spark Monitor

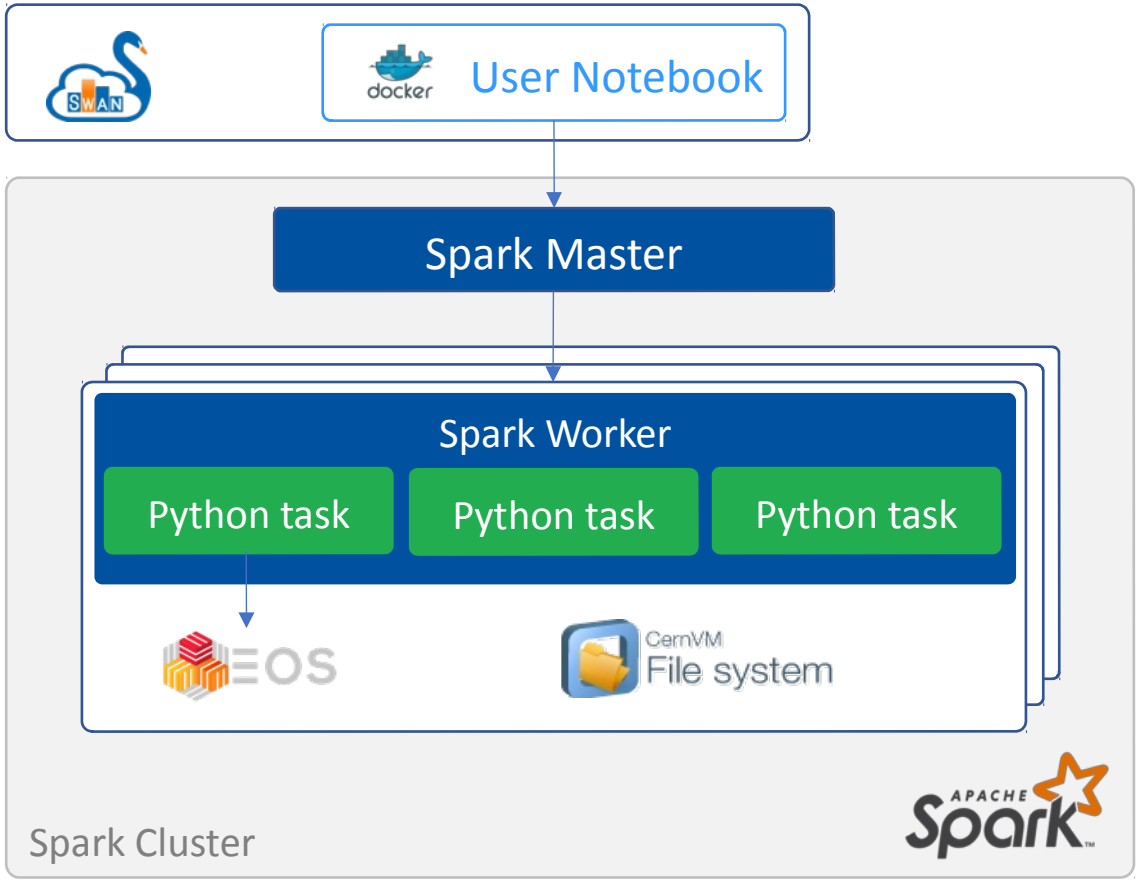

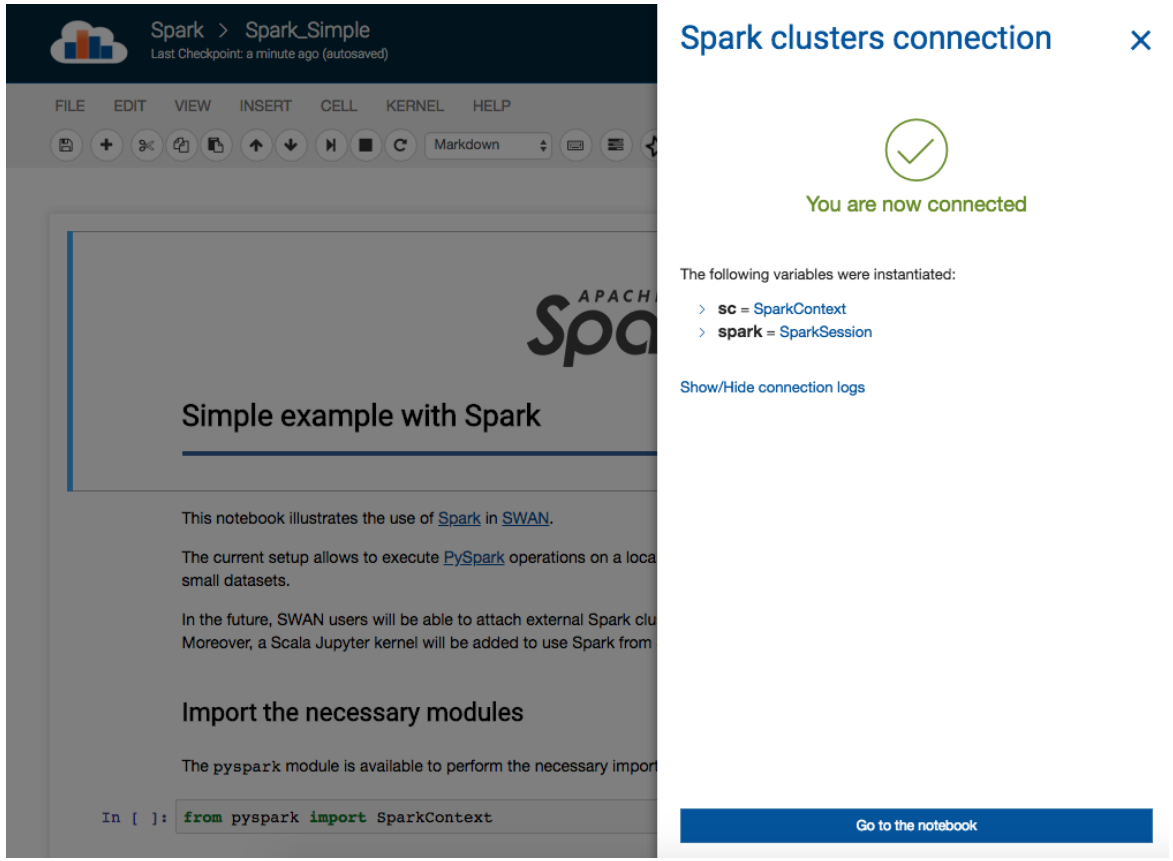
Spark > Spark_Simple
Last Checkpoint: a minute ago (autosaved)
FILE EDIT VIEW INSERT CELL KERNEL HELP
Markdown
Simple example with Spark
This notebook illustrates the use of Spark in SWAN.
The current setup allows to execute PySpark operations on a loca
small datasets.
In the future, SWAN users will be able to attach external Spark clu
Moreover, a Scala Jupyter kernel will be added to use Spark from
Import the necessary modules
The pyspark module is available to perform the necessary import
In [ ]: from pyspark import SparkContext

Spark clusters connection
You are now connected
The following variables were instantiated:
sc = SparkContext
spark = SparkSession
Show/Hide connection logs
Go to the notebook

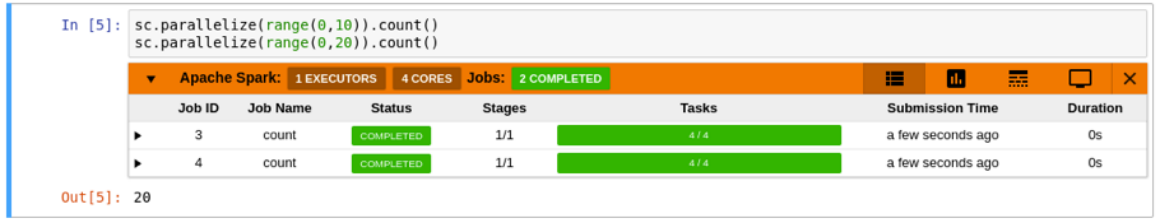
In [5]: sc.parallelize(range(0,10)).count()
sc.parallelize(range(0,20)).count()
Apache Spark: 1 EXECUTORS 4 CORES Jobs: 2 COMPLETED
Job ID Job Name Status Stages Tasks Submission Time Duration
3 count COMPLETED 1/1 4 / 4 a few seconds ago 0s
4 count COMPLETED 1/1 4 / 4 a few seconds ago 0s
Out[5]: 20

HEP User Community

- SWAN development is guided by our user community
  - New features (libs, kernels, ...) are requested by users from their real usage needs
- Gallery of examples
  - Made in collaboration with our users
  - Almost 50 notebooks in 7 categories

Example notebooks at **swan.web.cern.ch**

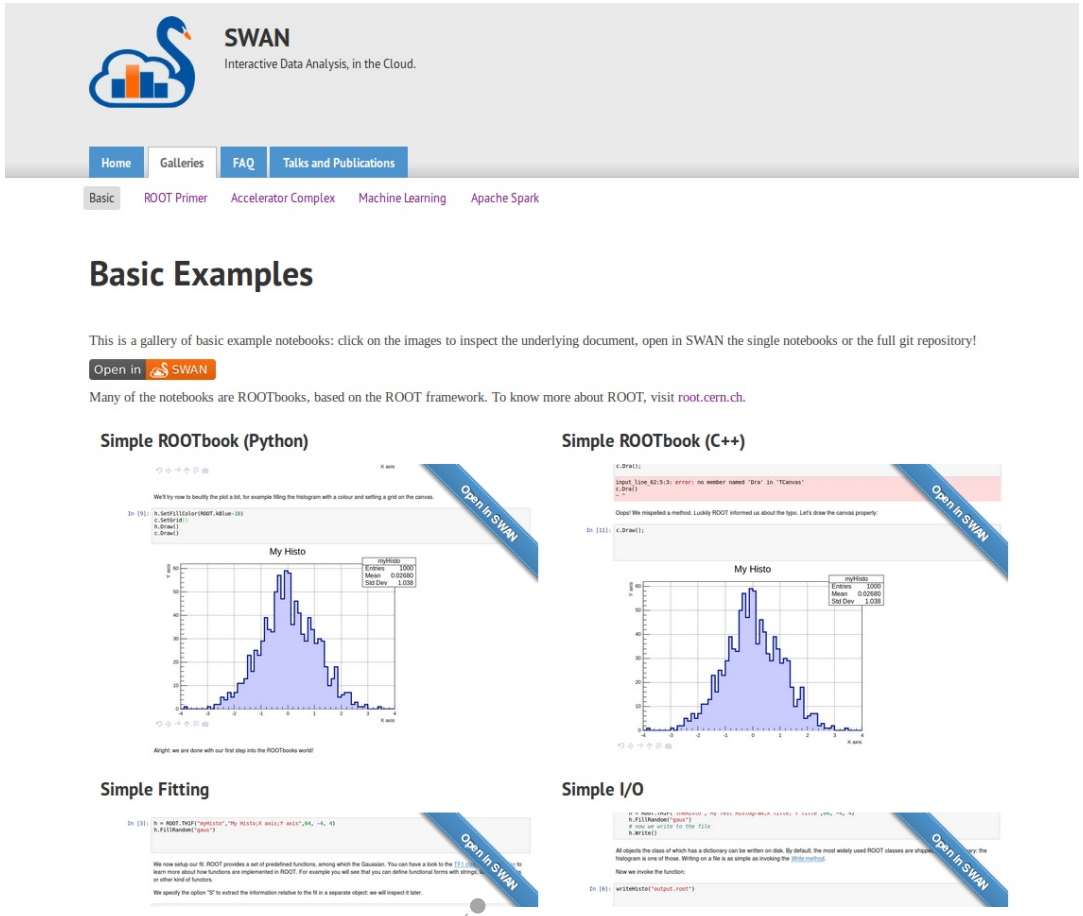

Access with only a click