

CERN-IT services for computing (introduction)


(about bytes, cores and bits)

Xavier Espinal (CERN-IT)

DUNE Collaboration Meeting
CERN & FNAL Software & Computing 31/01/2018



Who am I ?



Neutrino physicist (once upon a time) in K2K



Former member of the CERN Storage Group




IT Technical Coordinator for Experiments and Departments



WLCG: R&D on new computing models for HL-LHC



Who am I ?



Neutrino physicist (once upon a time) in K2K



Former member of the CERN Storage Group



IT Technical Coordinator for Experiments and Departments



WLCG: R&D on new computing models for HL-LHC




And more important:

Your protoDUNE - CERN/IT Liaison

All right... but what shall I use to get my analysis done ?

In *theory* all you need are **some disks**, **cpus** and **network** cables



All right... but what shall I use to get my analysis done ?

In *reality* what you need are **managed disks**, **cpus** and **network**




This is about bytes, bits and cores then...



The core:
“I couldn't *core* less
about speed”



The byte:
“byte'em and smile”



The bit:
“Biggest problem is illusion
that has taken place”

The bytes: CERN-IT Storage Services



Main Storage Platform

2.5B
250PB



Data Recording
User Analysis
Data Processing

<http://cern.ch/eos>



Cloud data storage (and much more)

Data in a shell
FUSE/batch

<https://cernbox.cern.ch>



cvmfs

Network FS based on http (distributed software delivery)

<http://cern.ch/cvmfs>

AFS

Phasing out

Network file system (home directories, works paces)

<http://cern.ch/afs>




The cores: CERN-IT Computing Services



Computing power for experiments and departments

>220k 

[How-to-batch](#)

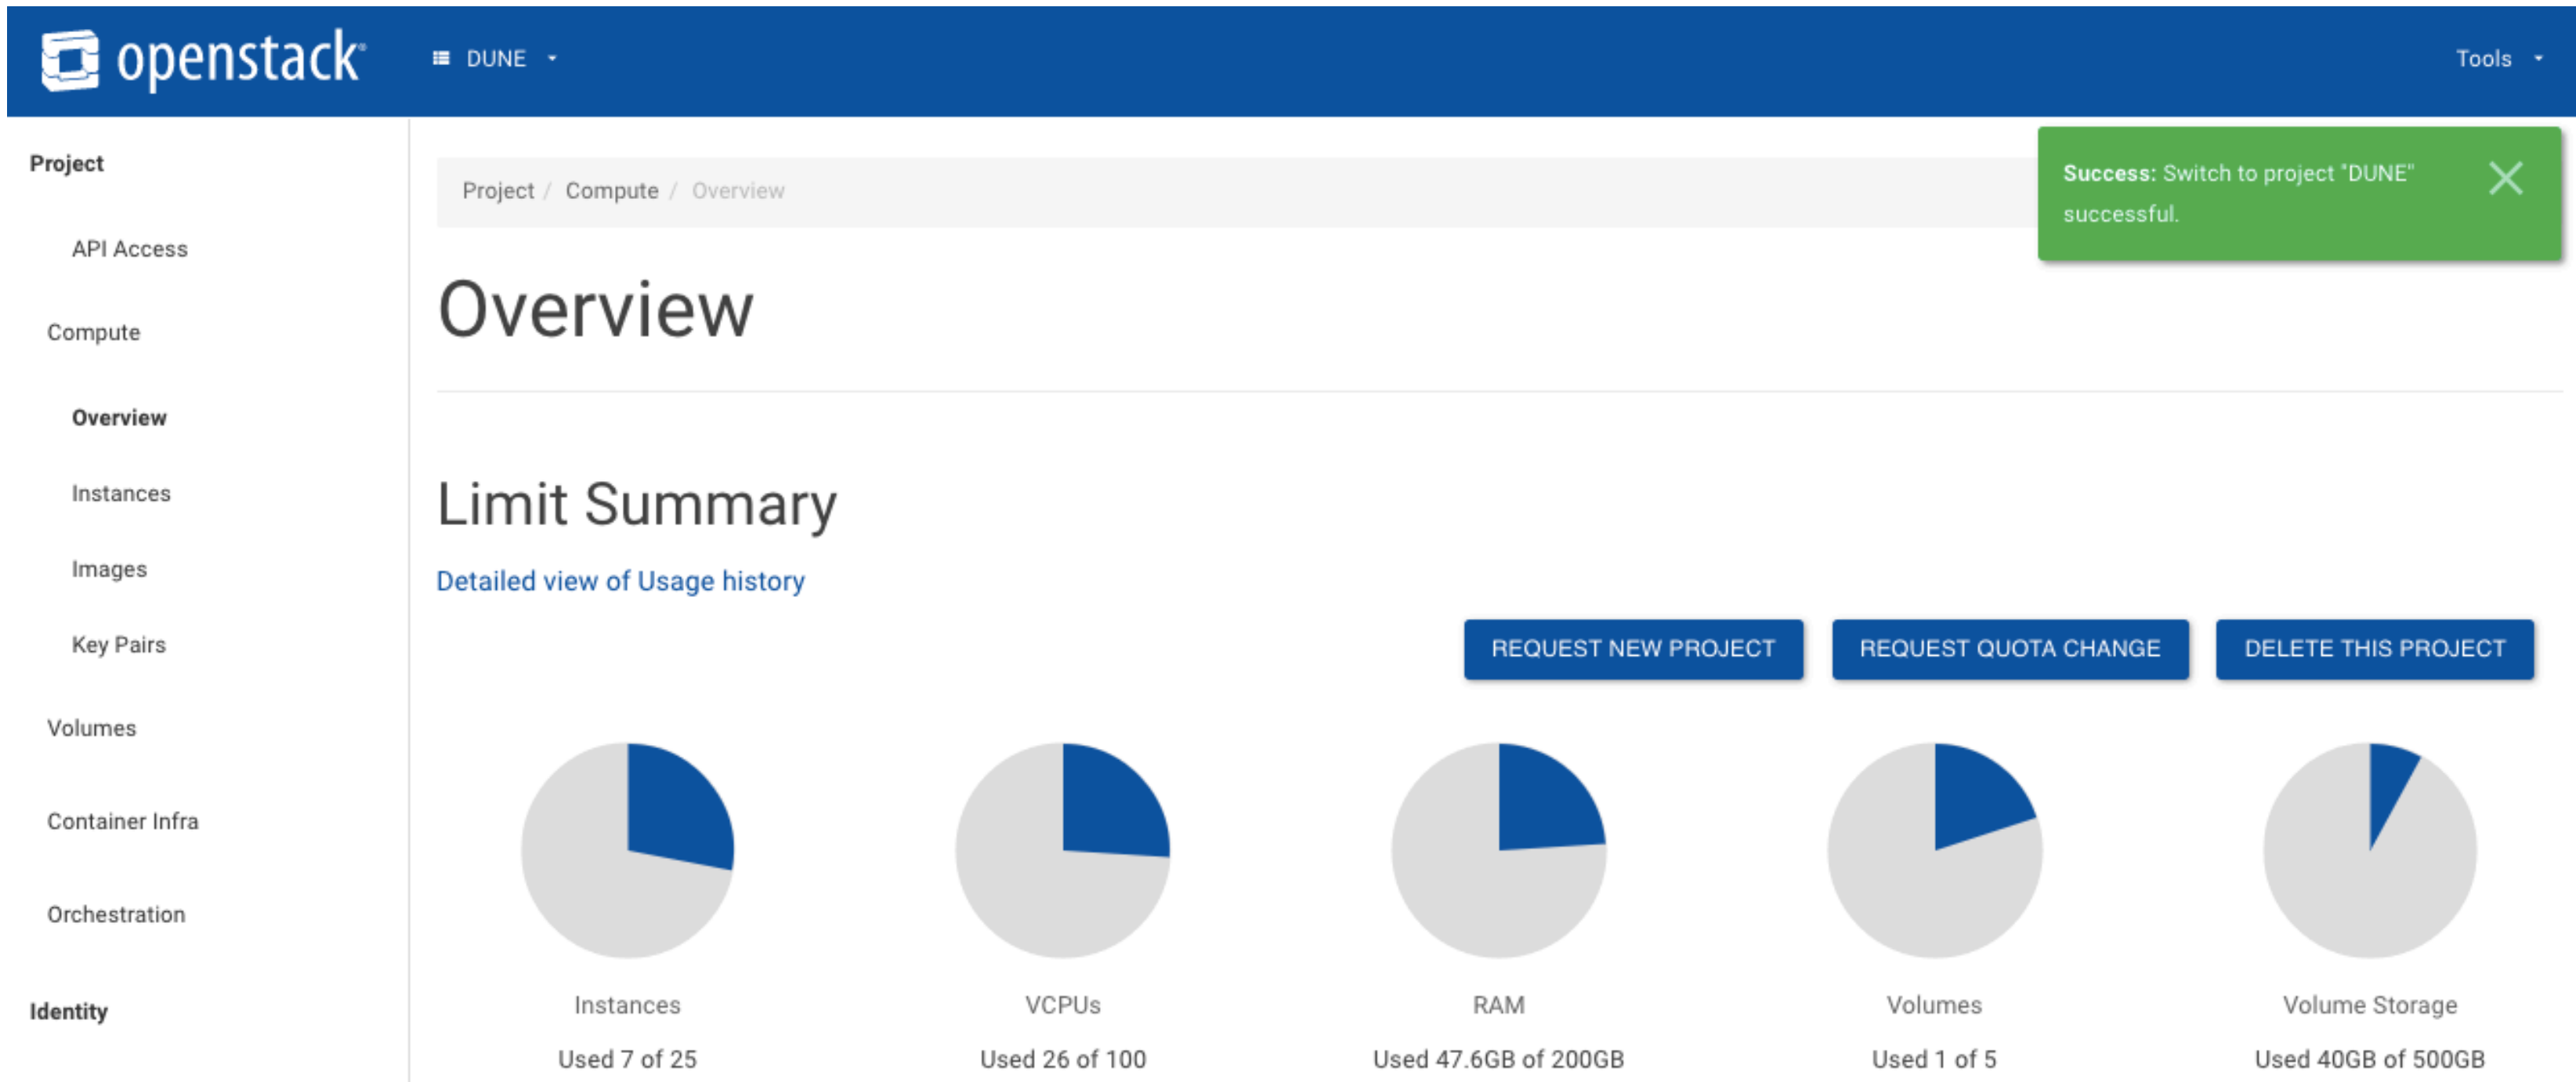


Server Provisioning Service

>9k hypervisors

Dynamic provisioning of servers for users, group, department and experiments

[how-to-cloud](#)
[Subscribe CERN Cloud Service](#)



The bits: CERN-IT Communication Services

Provides the networking resources for computing: routers, switch, cabling
...but this is *almost* invisible to you

CERN-wide in numbers

309902 devices
3832 Switches
233 Routers
667 Star points
2021 wifi access points

network.cern.ch

Register new device

Update device settings

Manage network parameters

Manage ownership

(except for VMs)

[Sign](#)

[Register](#)

[Set Mgmt](#)

[Admin Requests](#)

[News Subscribe](#)

[IP Phone](#)

Tools

[DNS Lookup](#)

[PING](#)

[MIKE](#)

[SOAP access](#)

Topology

[By Building](#)

[By StarPoint](#)

Blocked Systems

[By IP](#)

[By Hardware](#)

Apropos...

[Portables](#)

[DHCP](#)

[DNS](#)

[NTP](#)

[CERN IP](#)

[Networks](#)

More on...

[Access Rights](#)

[Computing Rules](#)

[Wireless](#)

[HELP!!!](#)

[Problems?](#)

[About](#)

[FAQ](#)

I have found the following entries where the system administrator or the main user are registered to be JAVIER ESPINAL CURULL.
Each line describes a device with its Location, Manufacturer, Type, Operating System and CERN Network Domain

- **ESPINAL CURULL JAVIER (63533 XAVIER.ESPINAL@CERN.CH) IT/ST**
- C2ADM03 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- DATATV04 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- EOSATLASHHTTP01 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- EOSLHCBHTTP01 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- EOSOPENLAB01 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- EOSPUBLICHTTP01 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- EOSPUBLICHTTP04 0000 0-0000 KVM VIRTUAL MACHINE LINUX GPN
- EOSPUBLICHTTP05 0000 0-0000 KVM VIRTUAL MACHINE LINUX GPN
- ESPINAL-WS 0031 1-0014 TTL NUC5I5MYHE LINUX GPN
- EUDAM01 0000 0-0000 KVM VIRTUAL MACHINE LINUX ITS
- GARMIN620 0031 1-0014 GARMIN 620 GARMIN
- INTELNUC 0031 1-0014 TTL UNKNOWN LINUX
- PCITFIO60 0031 1-0014 HP DC5750 MT LINUX
- PWM-ESPINAL 0031 1-0014 APPLE MACBOOK MAC-OS X
- SRM-EOSATLAS02 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- SRM-EOSCMS02 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- SRM-EOSLHCB02 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- SRM-EOSPPS01 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- SRM-EOSPPS02 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- SRM-EOSPUBLIC02 0000 0-0000 KVM VIRTUAL MACHINE LINUX LCG
- TAXEC01 0032 S-0A08 ASUS EEE PAD TRANSFORMER ANDROID
- TELEPHONE-63533 0031 1-0014 HP 4120 UNKNOWN
- WW-ESPINAL 0031 1-0014 HP 8200 ELITE WINDOWS GPN
- XAVI-ANDROID 0040 R-0008 SAMSUNG GALAXY S GT-I9000 ANDROID
- XAVI-IPAD 0031 1-0016 APPLE IPAD AIR 2 IOS
- XAVI-IPHONE 0031 1-0016 APPLE IPHONE SE IOS
- XAVIER-LGPHONE 0031 1-0014 LG OPTIMUS 4X HD ANDROID
- XAVIER-MACBOOKAIR 0031 1-0014 APPLE MAC BOOK AIR MAC OS X
- XAVIONEPLUS 0031 1-0016 ONEPLUS ONE ANDROID
- XEC-TABLET 0031 1-0014 ASUS EEE PAD TRANSFORMER ANDROID



Good, I might want to try...

(accessing Linux/lxplus, CERNBOX, Cloud Infrastructure, AFS, Windows, Web, MAC, Databases...)

Centralized management through the Resource Portal

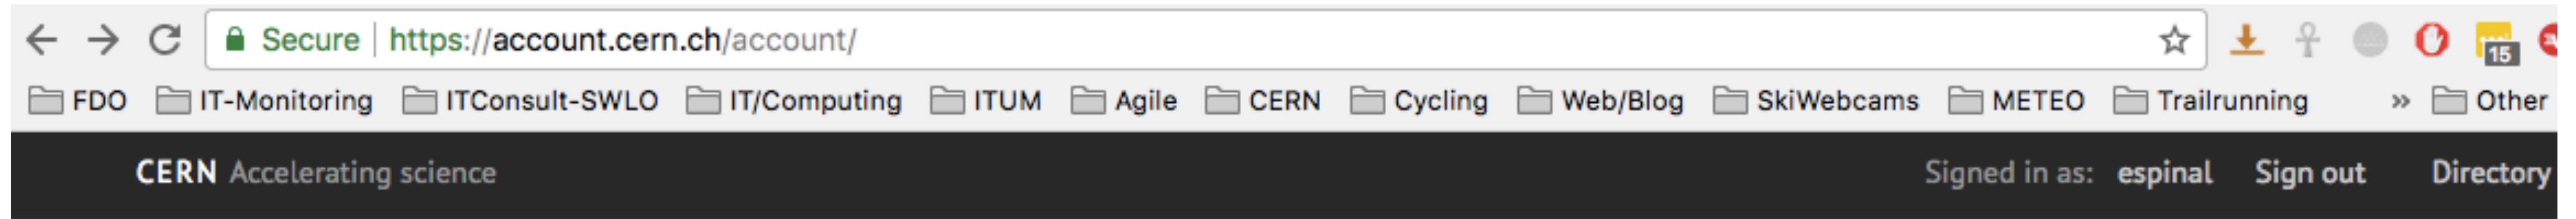
<https://resources.web.cern.ch/>

Centralized access control through the e-group portal

<https://e-groups.cern.ch/>



Accessing the Resource Portal → <https://account.cern.ch/account/>



Account Management

Manage your computing accounts and service subscriptions



Home

My Accounts


New Account

Change Password

Pending Actions

Help

Support

 Your password will expire in 17 day(s), on 09/02/2018. Please [change your password](#) before that date.

Manage Accounts



- ▶ My accounts
- ▶ Create a new account
- ▶ Personal information
- ▶ Check account status
- ▶ Sign Security & Computing Rules

View or create accounts, manage secondary and service accounts.

Manage Passwords




- ▶ CERN Account password
- ▶ Forgot your password?
- ▶ EDH Signature password

Change or reset your password, or get help setting a new password.

Authorization, Resources and CERN Applications



- ▶ Resources and Services 
- ▶ E-Groups membership

Manage your Resources and Applications subscriptions, parameters and options.

Lightweight Accounts



- ▶ Lightweight Accounts Portal
- ▶ Register an account
- ▶ Manage an account

This service allows external users who do not have a regular CERN account to get access to certain CERN applications that are open to external users.

Resource Portal



<https://resources.web.cern.ch/>

CERN Resources Portal

Manage your CERN Resources, lifecycle, settings, etc.



- Home
- List Services
- Pending Actions
- Select Account
- Help
- Support

Select Account

Select another of your accounts to manage Service subscriptions and Resources

Name, mail, login or ID

Search login only

Xavier Espinal Curull (espinal)

(no description)

Xavier Espinal Curull (xaverc3)

Created by IM portal on 02/06/2012 12:06:40

Castor Work (castorc3)

(no description)

NA62 Grid User (na62grid)

Created by IM portal on 03/07/2012 15:34:38

Xavier Espinal Curull (espinal)

Account Type: Primary
Email: xavier.espinal@cern.ch
Person ID: 548112
Disabled: **NO**
Unix User ID: 3808
Computing Group: zp (1307)



List Services

for account espinal

Create and subscribe to resources and services. Change properties, edit lifecycle, reassign to another person.

Available Services

Listing all the services, with their subscription status

[Search]

Basics

Account Management Computing Accounts

E-Mail E-Mail Services

Communication and Collaboration

Skype for Business IP enhanced telephony

Vidyo Universal videoconferencing system.

E-Groups Electronic groups for authorization and mailing lists

Unified Messaging Voice mail on Exchange

Fax Service Send and receive faxes as emails (deprecated).

Operating Systems

Windows Services Remote access to Windows computers.

Terminal

LXPLUS and Linux Linux and LxPlus services

Mac Desktops Mac Desktop Service

Windows Desktops Centrally managed Windows desktops and laptops.

Cloud Infrastructure Cloud Infrastructure Projects.

Storage

AFS Workspaces AFS File Services

DFS Workspaces Windows based file storage on highly available servers.

EOS/CERNBox EOS/CERNBox storage.

Web Services

Web Services Support for authoring and publishing content on the web.



Lxplus and Linux services

<http://information-technology.web.cern.ch/services/lxplus-service>

CERN Resources Portal

Manage your CERN Resources, lifecycle, settings, etc.

[Home](#) [List Services](#) [Pending Actions](#) [Select Account](#) [Help](#) [Support](#)

LXPLUS and Linux

Linux and LxPlus services

Service Information

Settings

Resources

Computing Groups

Linux settings

Uid: 3808
Gid: 1307 (zp)
Home Directory: /afs/cern.ch/user/e/espinal
Change Primary Group:
Unix Shell:

LXPLUS and Linux

Linux and LxPlus services

Service Information

Settings

Resources

Computing Groups

Linux and LXPLUS: This service (Public Login User Service) is the interactive logon service to Linux for all CERN users. The cluster LXPLUS consists of public machines provided by the IT Department for interactive work.

Links

- Linux (Linux @ CERN)
- LXPLUS (Interactive logon service to Linux)

Subscription is based on Computing Groups

To subscribe to LXPLUS, Linux and AFS service, your account must be added to a Computing Group, either a specific one if you participate to some specific Experiments or projects, or a generic one if you have basic needs.

You are already member of the following Computing Groups:

- it: Linux group it - IT
- np-comp: Linux group np-comp - Neutrino Platform
- zp: Linux group zp - ATLAS

Change Primary Group:

Manage your AFS subscription

Click on the button below to unsubscribe from AFS.

Specific project or experiment environment

If you need a specific environment, linked to a project or experiment, you should be added to a specific Computing Group. Select the Computing Group from the list below, and send a request to the Group Administrators, explaining why you need to join the group.

Computing Group:

Comment / reason:



Service Information

Subscribe

Resources

Accounts

Cloud Infrastructure

Cloud Infrastructure Projects.

The CERN Private Cloud provides an Infrastructure-as-a-Service solution integrated with CERN's computing facilities. Using self service portals or cloud interfaces, users can rapidly request virtual machines for production, test and development purposes. The machines can be of different capacities and run a variety of Windows or Linux operating systems.

Links

- [OpenStack Dashboard \(OpenStack Dashboard\)](#)
- [User Guide \(CERN Cloud Infrastructure User Guide\)](#)
- [Submit a ticket \(Submit a Service Now ticket to Cloud Infrastructure support.\)](#)

Service Subscription

To manage Cloud Infrastructure projects, you must be member of the cloud-infrastructure-users e-group. Click the button below to subscribe to the e-group.

Subscribe



CERN Resources Portal

Manage your CERN Resources, lifecycle, settings, etc.



Home

List Services

Pending Actions

Select Account

Help

Support

Service
Information

EOS/CERNBox

EOS/CERNBox storage.

Storage and web space for personal use, projects and experiments.

Personal space

CERNBox provides personal data storage space to anyone having a CERN account. Data can be:

- Accessed from any web browser or file explorer.
- Shared between individuals or groups of collaborators.
- Synchronised across a series of devices (smartphone, tablet, laptop, desktop).
- Accessed via the SHARED Filesystem on lxplus or via EOS.

CERNBox uses EOSUSER (an instance of EOS, the CERN disk storage system) as the storage backend.

To activate your CERNBox personal storage space (1 TB, up to 1 million files, hosted in the CERN Computer Center): login to <https://cernbox.cern.ch> (using your CERN account and password).



Projects and experiments

EOS spaces for physics data are managed by the experiments themselves. All other EOS space is managed by CERN-IT. Please consult [KB0003151](#) to request for space.

Help and support

To report a problem, contact [CERNBox Service](#).

CERN Resources Portal

Manage your CERN Resources, lifecycle, settings, etc.



Home

List Services

Pending Actions

Select Account

Help

Support

Service Information

Settings

Subscribe

AFS Workspaces

AFS File Services

Home folder

Home folder path: [/afs/cern.ch/user/e/espinal](#)

Home folder quota: **Used 0.13 / 5.00 GB**

Increase home quota to 10 GB



Workspace

No workspace is currently associated to this account.

Create AFS Workspace



Success: Switch to project "DUNE" successful.

Project / Compute / Overview

Overview

Limit Summary

[Detailed view of Usage history](#)

REQUEST NEW PROJECT

REQUEST QUOTA CHANGE

DELETE THIS PROJECT



Instances

Used 7 of 25



VCPUs

Used 26 of 100



RAM

Used 47.6GB of 200GB



Volumes

Used 1 of 5



Volume Storage

Used 40GB of 500GB

- Project
- API Access
- Compute
- Overview**
- Instances
- Images
- Key Pairs
- Volumes
- Container Infra
- Orchestration
- Identity



- Project
- API Access
- Compute
- Overview
- Instances**
- Images
- Key Pairs
- Volumes
- Container Infra
- Orchestration
- Identity


Instances

INSTANCE ID =

Displaying 7 items

<input type="checkbox"/>	Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	protodune-fts	-	137.138.76.39 2001:1458:d00:1c::39	m2.large	steve-openstack	Active	cern-geneva-a	None	Running	1 week, 6 days	<input type="button" value="CREATE SNAPSHOT"/>
<input type="checkbox"/>	dune-vm-build-03	SLC6 - x86_64 [2017-04-06]	188.185.113.233 2001:1458:d00:a::100:1e3	m2.xlarge	lxplus	Active	cern-geneva-a	None	Running	5 months, 2 weeks	<input type="button" value="CREATE SNAPSHOT"/>
<input type="checkbox"/>	dune-vm-build-02	SLC6 - x86_64 [2017-04-06]	188.185.112.32 2001:1458:d00:a::100:1a	m2.large	lxplus	Active	cern-geneva-a	None	Running	5 months, 3 weeks	<input type="button" value="CREATE SNAPSHOT"/>
<input type="checkbox"/>	np04-webgw1	CC7 - x86_64 [2017-04-06]	188.185.71.29 2001:1458:d00:7::100:317	m2.large	lxplus	Active	cern-geneva-c	None	Running	6 months, 2 weeks	<input type="button" value="CREATE SNAPSHOT"/>
<input type="checkbox"/>	p3s-content	CC7 - x86_64 [2017-04-06]	188.185.77.72 2001:1458:d00:9::100:142	m2.medium	-	Active	cern-geneva-a	None	Running	6 months, 4 weeks	<input type="button" value="CREATE SNAPSHOT"/>
<input type="checkbox"/>	p3s-db	CC7 - x86_64 [2017-04-06]	188.185.85.205 2001:1458:d00:f::100:1c7	m2.medium	-	Active	cern-geneva-a	None	Running	6 months, 4 weeks	<input type="button" value="CREATE SNAPSHOT"/>



- Project
 - API Access
 - Compute
 - Overview
 - Instances
 - Images
 - Key Pairs**
 - Volumes
 - Container Infra
 - Orchestration
- 

Project / Compute / Key Pairs

Key Pairs

Filter + CREATE KEY PAIR ↑ IMPORT KEY PAIR 🗑️ DELETE KEY PAIRS

Displaying 1 item

<input type="checkbox"/>	Key Pair Name	Fingerprint	Actions
<input type="checkbox"/>	cloud-xavier	74:51:0f:96:58:58:57:41:7f:14:e5:1c:d7:5b:bc:bc	🗑️ DELETE KEY PAIR

Displaying 1 item

Project

API Access

Compute

Volumes

Volumes

Snapshots

Container Infra

Orchestration

Identity

Project / Volumes / Volumes

Volumes

Filter

+ CREATE VOLUME

ACCEPT TRANSFER

DELETE VOLUMES

Displaying 1 item

<input type="checkbox"/>	Name	Description	Size	Status	Type	Attached To	Availability Zone	Bootable	Encrypted	Actions
<input type="checkbox"/>	ceb73204-36dd-4740-83f2-e60350f7fd09	-	40GiB	In-use	standard	/dev/vda on protodune-fts	nova	Yes	No	EDIT VOLUME

Displaying 1 item



Openstack projects

Volume Name



Description

Description:

Volumes are block devices that can be attached to instances.

Volume Type Description:

Type	standard
Usage	default
Max IOPS	100
Max Throughput	80 MB/s

Volume Limits

Total Gibibytes 40 of 500 GiB Used

Number of Volumes 1 of 5 Used

Total Gibibytes for standard (40 GiB)
500 GiB Available

Number of Volumes for standard (1) 5 Available

Volume Source

NO SOURCE, EMPTY VOLUME

Type

STANDARD

Size (GiB)

1

CANCEL

CREATE VOLUME

https://openstack.cern.ch/project/

DUNE

Project / Volumes / Volumes

Volumes

Filter




+ CREATE VOLUME

Displaying 1 item

<input type="checkbox"/>	Name	Description	Size	Status	Type	Attached To	Availability Zone
<input type="checkbox"/>	ceb73204-36dd-4740-83f2-e60350f7fd09	-	40GiB	In-use	standard	/dev/vda on protodune-fts	nova



What Lives in a Shell?



Science in a shell: `/bigdata`, `/userdata` and `/software` at the node

My code ***Htozz.root*** is
on my laptop and **synced** to
cernbox:
/eos/user/xavi/goldench/

Science in a shell: **/bigdata**, **/userdata** and **/software** at the node

My code ***Htozz.root*** is
on my laptop and **synced** to
cernbox:
/eos/user/xavi/goldench/



I'm interested in running my
analysis on the full **HtoZZ**
dataset:
/eos/atlas/phys-higgs/htozz

Science in a shell: **/bigdata**, **/userdata** and **/software** at the node

My code ***Htozz.root*** is on my laptop and **synced** to **cernbox**:
/eos/user/xavi/goldench/



I'm interested in running my analysis on the full **HtoZZ** dataset:
/eos/atlas/phys-higgs/htozz



I submit analysis jobs at the worker nodes, which **all** have:
/eos/atlas/phys-top/Htozz/*
/eos/user/xavi/*
/cvmfs/atlas/athena/*

Science in a shell: **/bigdata**, **/userdata** and **/software** at the node

My code ***Htozz.root*** is on my laptop and **synced** to **cernbox**:
/eos/user/xavi/goldench/



I'm interested in running my analysis on the full **HtoZZ** dataset:
/eos/atlas/phys-higgs/htozz



I submit analysis jobs at the worker nodes, which **all** have:
/eos/atlas/phys-top/Htozz/*
/eos/user/xavi/*
/cvmfs/atlas/athena/*

The job results aggregated on **cernbox**:

/eos/user/xavi/goldench/htozz/

And **synced** on my laptop as the jobs finished

Science in a shell: **/bigdata**, **/userdata** and **/software** at the node

My code ***Htozz.root*** is on my laptop and **synced** to **cernbox**:
/eos/user/xavi/goldench/



I'm interested in running my analysis on the full **HtoZZ** dataset:
/eos/atlas/phys-higgs/htozz



I submit analysis jobs at the worker nodes, which **all** have:
/eos/atlas/phys-top/Htozz/*
/eos/user/xavi/*
/cvmfs/atlas/athena/*

The job results aggregated on **cernbox**:

/eos/user/xavi/goldench/htozz/

And **synced** on my laptop as the jobs finished



Work on **final plots** on the **laptop** and latex-ing the paper directly on

/eos/user/xavi/goldench/htozz/paper/

Science in a shell: **/bigdata**, **/userdata** and **/software** at the node

My code ***Htozz.root*** is on my laptop and **synced** to **cernbox**:
/eos/user/xavi/goldench/



I'm interested in running my analysis on the full **HtoZZ** dataset:
/eos/atlas/phys-higgs/htozz



I submit analysis jobs at the worker nodes, which **all** have:
/eos/atlas/phys-top/Htozz/*
/eos/user/xavi/*
/cvmfs/atlas/athena/*

The job results aggregated on **cernbox**:

/eos/user/xavi/goldench/htozz/

And **synced** on my laptop as the jobs finished



Work on **final plots** on the **laptop** and latex-ing the paper directly on

/eos/user/xavi/goldench/htozz/paper/



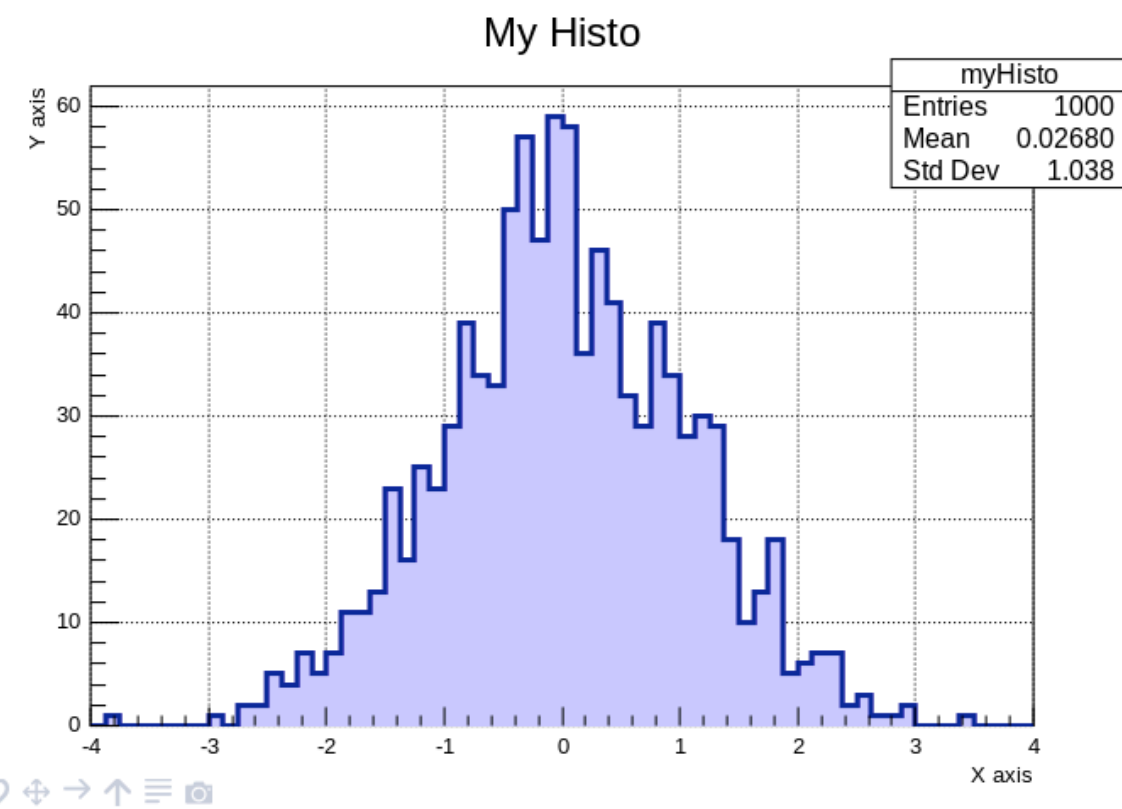
Share on-the-fly
Analysis results
n-Tuples
Plots
Publication

Storage Services: easing data access

```
c.Dra();
input_line_62:5:3: error: no member named 'Dra' in 'TCanvas'
c.Dra()
~ ^
```

Oops! We misspelled a method. Luckily ROOT informed us about the typo. Let's draw the canvas properly:

```
[11]: c.Draw();
```

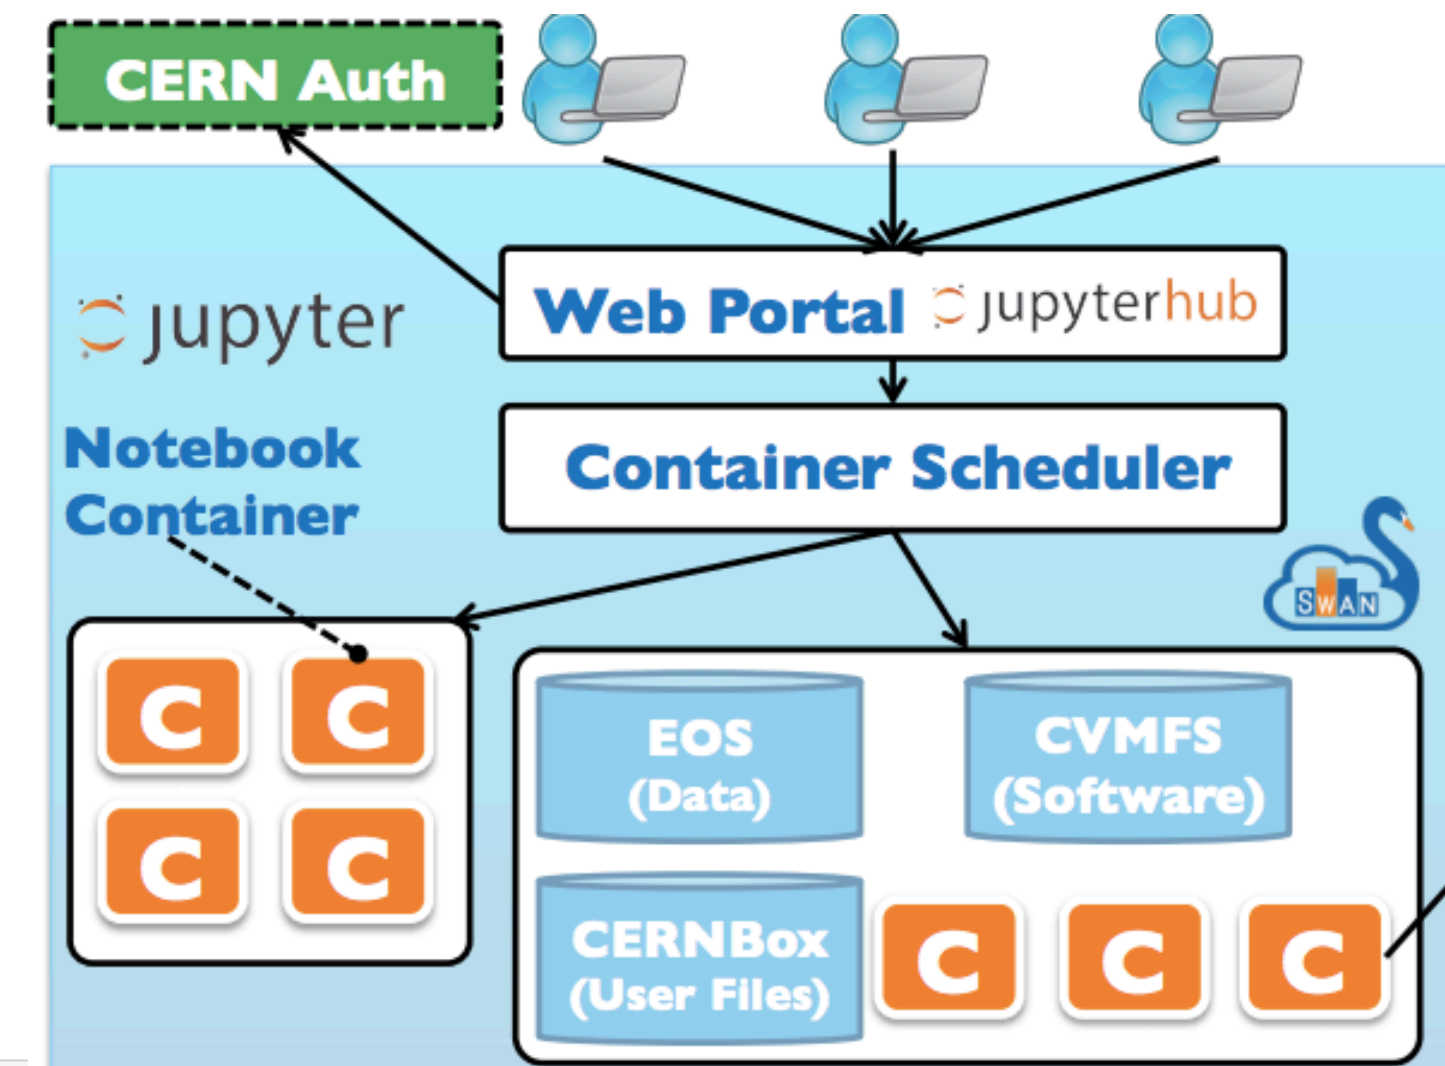
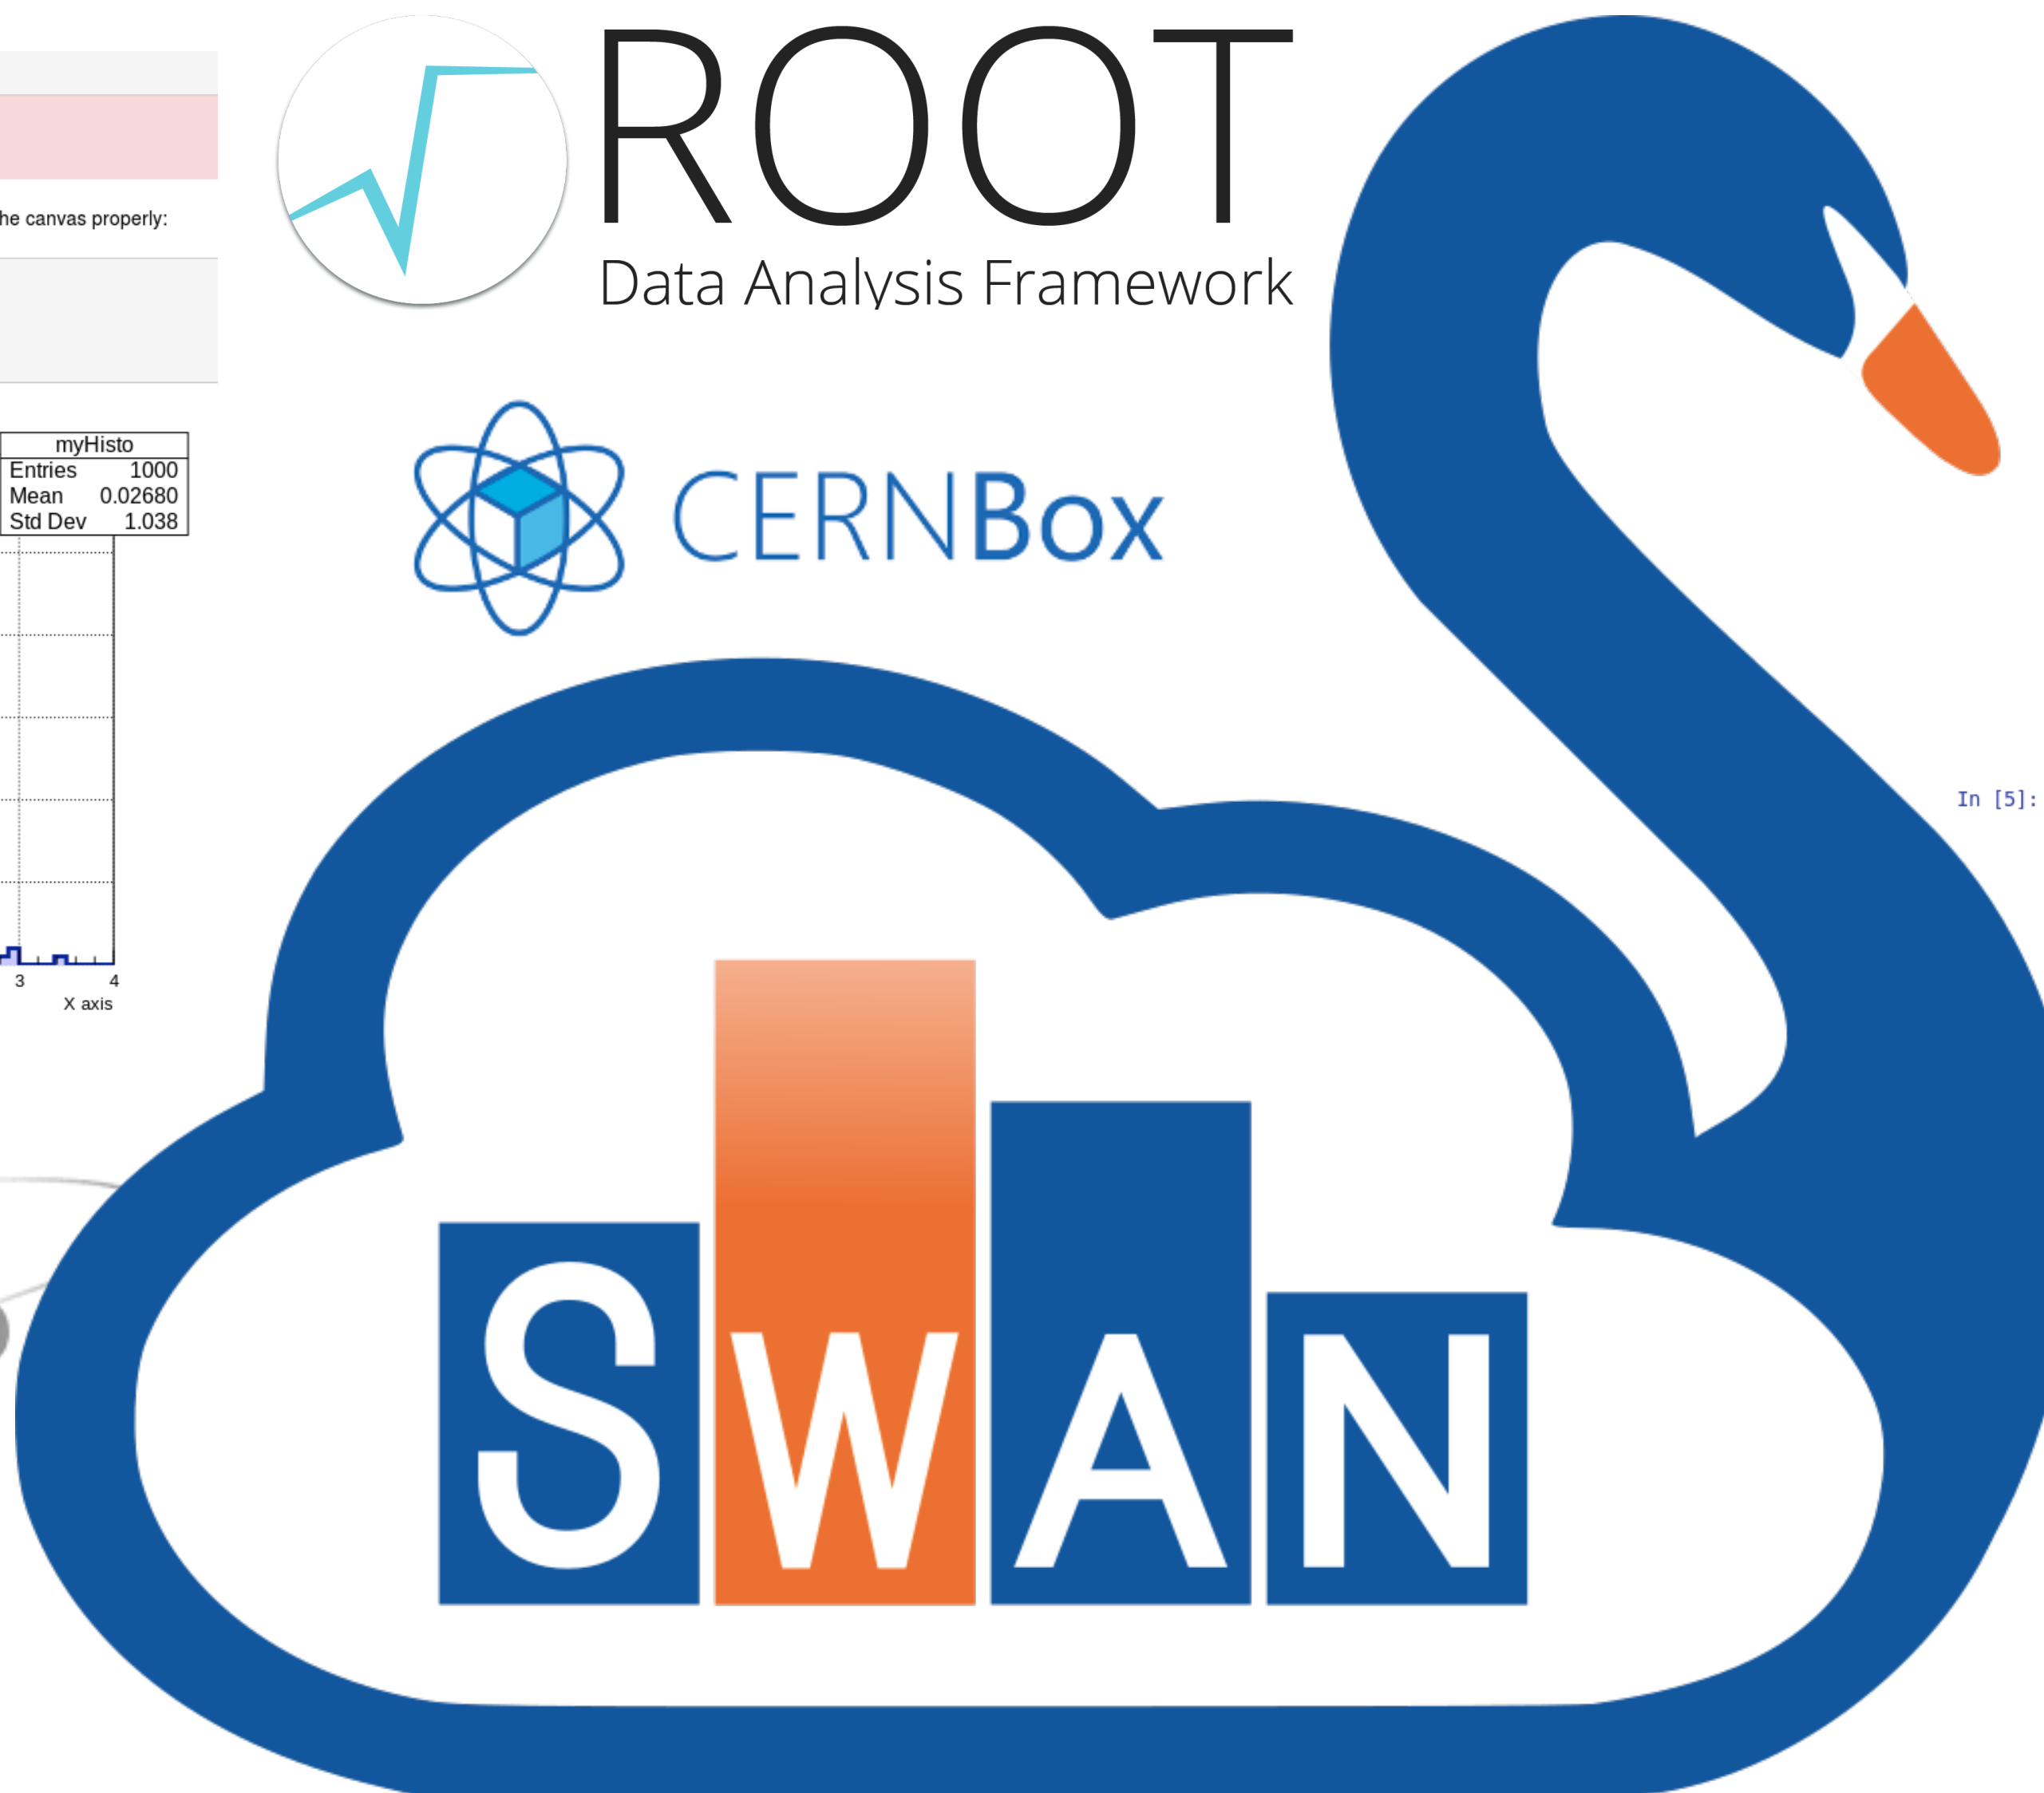


ROOT

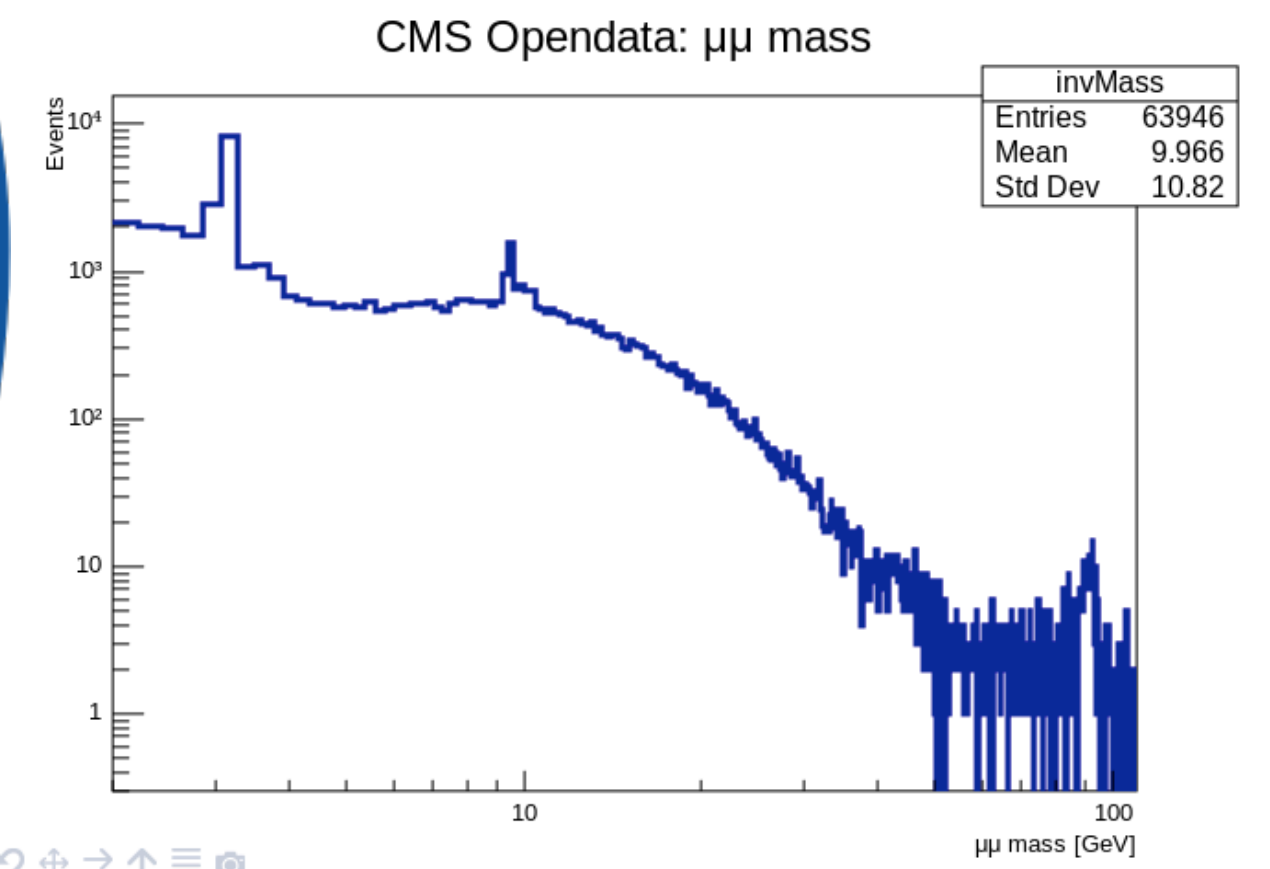
Data Analysis Framework



CERNBox



```
In [5]: invMass = ROOT.TH1F("invMass","CMS opendata: #mu#mu mass;#mu#mu mass [GeV];events",512, 2, 110)
invMassFormula = "sqrt((E1 + E2)^2 - ((px1 + px2)^2 + (py1 + py2)^2 + (pz1 + pz2)^2))"
cut = "Q1*Q2==-1"
c = ROOT.TCanvas()
dimuons.Draw(invMassFormula + " >> invMass",cut,"hist")
c.SetLogx()
c.SetLogy()
c.Draw()
```



That might have been too fast. We now make the analysis above more explicit producing a plot also for the J/Psi particle.

jupyterhub

SWAN

EOS

Web based **computing interface** combining: **data**, **code**, **equations**, text and **visualisation**

Thanks!

Your account portal

<https://account.cern.ch/account/>

Your resources portal

<https://resources.web.cern.ch/>

e-groups portal

<https://e-groups.cern.ch/>

Service status portal

<https://cern.service-now.com/service-portal/sls.do>

CERN-IT services portal

<http://information-technology.web.cern.ch/services>

