

The VRE



The Virtual Research Environment
ET EIB workshop - Aachen March 10 2024

Today's Agenda

- VRE context
 - ESCAPE & EOSC Future project
 - Scientific context
 - Computing challenges
- VRE components
 - AAI layer
 - User Interface and Notebook service
 - Distributed storage solution: ESCAPE Data Lake infrastructure
 - Distribute storage solutions.
- Demo

Computing challenges become global

Upcoming experiments and Research Infrastructures facing a **change of scale in computing needs**, eg. the HL-LHC at CERN but also in Astronomy and Nuclear Physics.

- New generation of detectors, antennas and telescopes producing big data volumes, requiring massive processing power and support large user communities

Strong commonalities in the main distributed computing areas, eg. Data Management and Access, storage orchestration, Identity management, user analysis, etc.

European Strategy for Particle Physics

*“Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. [...] The community must vigorously pursue **common, coordinated R&D efforts in collaboration with other fields of science and industry, to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation.**”*

2020 UPDATE OF THE EUROPEAN STRATEGY FOR PARTICLE PHYSICS

4 

Other essential scientific activities for particle physics



ESCAPE project

ESFRI RIS



- Address RI's needs in Data Management, Access and Analysis for **Astro-particle, Radio-astronomy, Gravitational Waves, Cosmology and Particle Physics.**
- Provide a fully working **common data infrastructure** “The ESCAPE Data Lake” to test novel data management tools and models, giving the opportunity to influence and steer its development.
- Expand **collaborations** and foster involvement with other Scientific Communities. Maintain and strengthen collaborations with related EC initiatives and projects.
- ESCAPE finished Jan '23 and become an open collaboration [[link](#)]



Data centres

ESCAPE and the VRE

- Work done on scientific (re)analyses has been extended in the context of the EU funded **EOSC-Future project** to establish an analysis platform (a *Virtual Research Environment*) for the community.
- Being developed in collaboration and for the scientific community based on **two concrete use cases**, "Science projects":
 - Dark Matter
 - Extreme Universe
- Aim:
 - Create a **proof of concepts (PoC)** to **demonstrate** the **utility** of the **solutions and technologies developed**
 - **Integrate** them into the **EOSC landscape**, as a prototype of an ESCAPE Cell in EOSC.

Common solutions developed within ESCAPE, being integrated in the VRE:

- Data Lake infrastructure
- Software repository
- Analysis platform tools
- Integration of the IVOA
- Outreach and citizen science

EOSC

The ambition of the European Open Science Cloud (EOSC) is to provide European researchers, innovators, companies and citizens with a **federated and open multi-disciplinary environment** where they can **publish, find and reuse data**, tools and services **for research, innovation and educational purposes**.

- The five science clusters in Europe are actively collaborating and joining efforts in these common goals



VRE challenges

Can we make use of all these heterogeneous scientific solutions from a unique "entry point"?

How can we provide CERN-born and HEP-specific technologies to the wider communities?

Why do we want to do this ?

- Coordinate efforts to **share with** and **learn from** different communities
 - more people means more manpower & knowledge
- **Combine approaches** towards **common challenges**

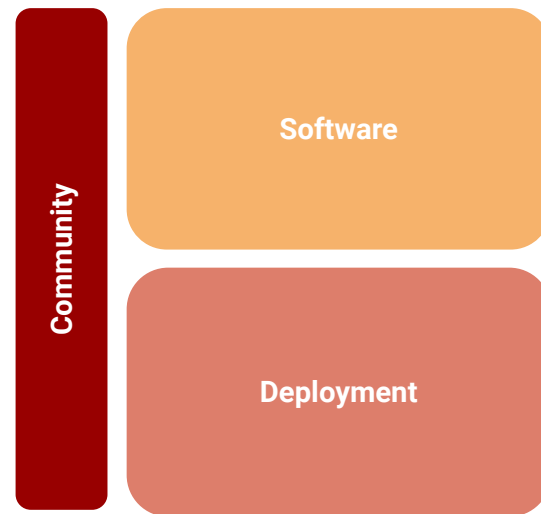
How do we do this?

- **Open Science** initiatives
- **FAIR** principles
- **Data preservation** and **reproducible research/science**

Findable
Accessible
Interoperable
Reusable

Virtual Research Environment


The VRE is an **open source analysis platform** where researchers from various scientific communities have access to all the **digital content** (software, data, computational power) needed to to develop, share and reproduce an **end-to-end scientific result** in compliance with **FAIR** principles.



Use case

Dark Matter EOSC-Future Science Project

- *Galaxy rotation curves* --> a larger amount of gravitational mass is expected to exist in the universe.
- It does not interact with the electromagnetic field and *cannot therefore be seen*.
- Many DM candidates. Many experiments target the problem. Many different research approaches.



Dark Matter EOSC-Future Science Project


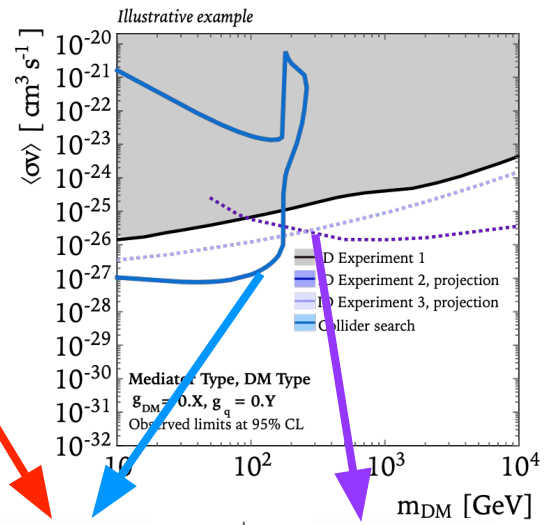
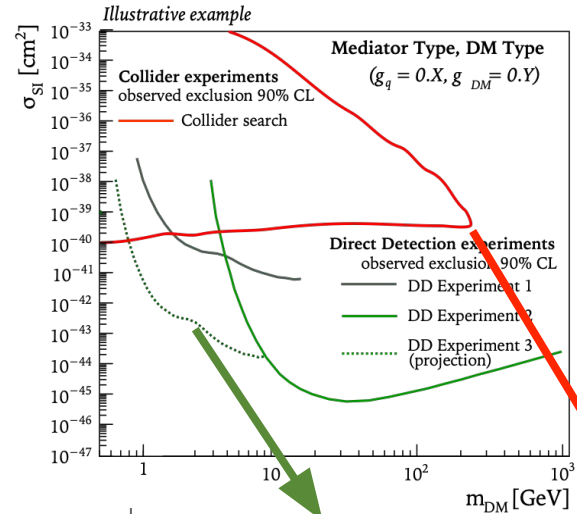
It is the perfect example for a use case, as both HEP and Astrophysics are researching limits of Dark Matter.

- They need similar infrastructures
- It would be useful to share software and expertise on how to use it


The VRE aim is to provide:

- An **infrastructure** where to run an end-to-end analysis in one place
- A solution to **preserve** the analysis steps in case of re-use
- An **interdisciplinary open science** example from a **bottom-up** effort


Combined DM plots




Astrophysics




Theory



Direct Detection



Colliders



Indirect Detection



An analysis workflow

Data Collection

Experimental data generated at site

Analysis

Run with preserved, versioned software

Result preservation

Of data, software, environment and computational steps for later reuse and share

Data processing

Data is processed for analysis

Result interpretation

Results are combined and compared across workflow

ESFRI RIs



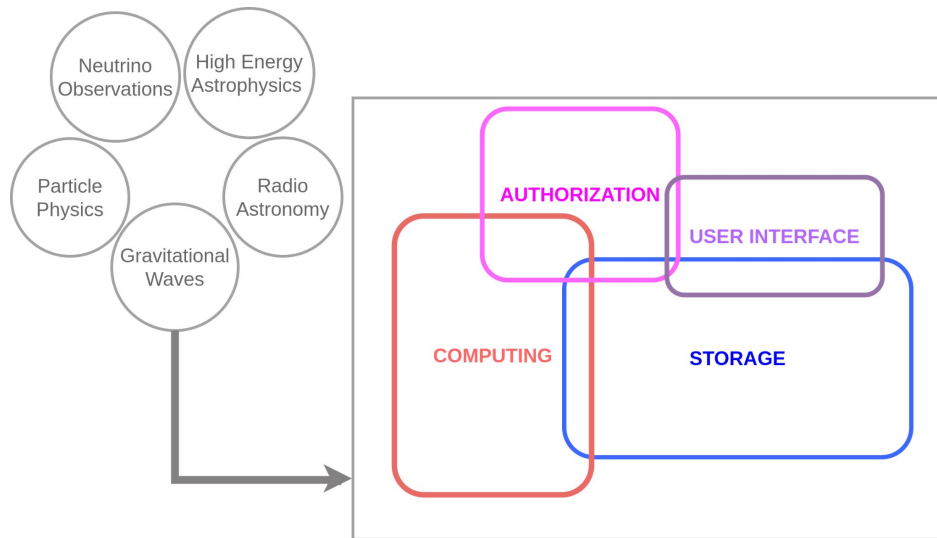
Data Centres



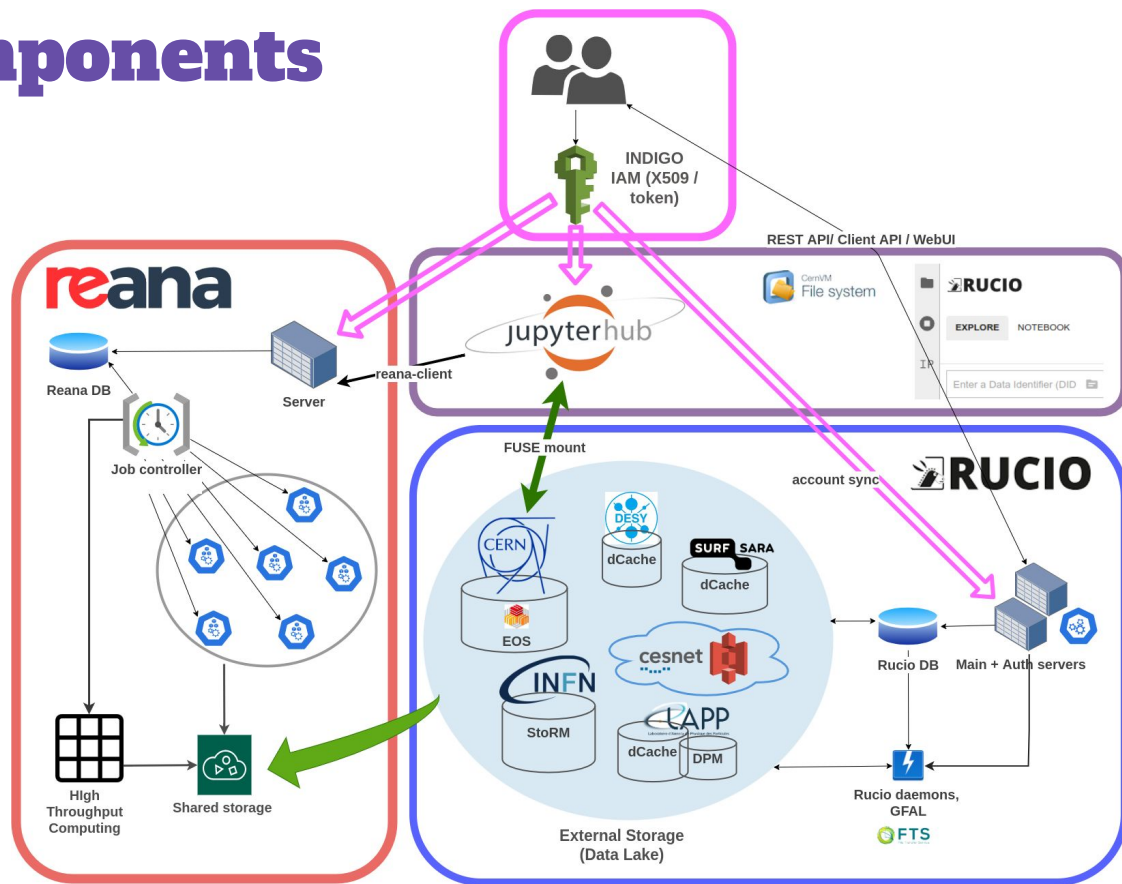
VRE components

VRE components

1. federated **Authentication and Authorization** layer
2. federated distributed **storage** solution
3. **Notebook service** with customizable containerised environments
4. **computing** cluster supplying the processing power



VRE components



1. Authentication and Authorisation

- INDIGO Identity and Access Management (IAM) service
 - adopted by WLCG for token usage
- supports authentication via
 - EduGAIN
 - OIDC tokens
 - X.509 certificates/Virtual Organization Membership Service (VOMS) attribute provisioning services
- Instance deployed on a K8s cluster at INFN-CNAF Bologna
- Rucio and Jupyterhub support IAM
 - Reana's IAM support is under active development



Welcome to **escape**

Sign in with your escape credentials

Sign in

[Forgot your password?](#)

Or sign in with

Your X.509 certificate



Not a member?

Apply for an account

[Privacy policy](#)

<https://iam-escape.cloud.cnaf.infn.it/login>

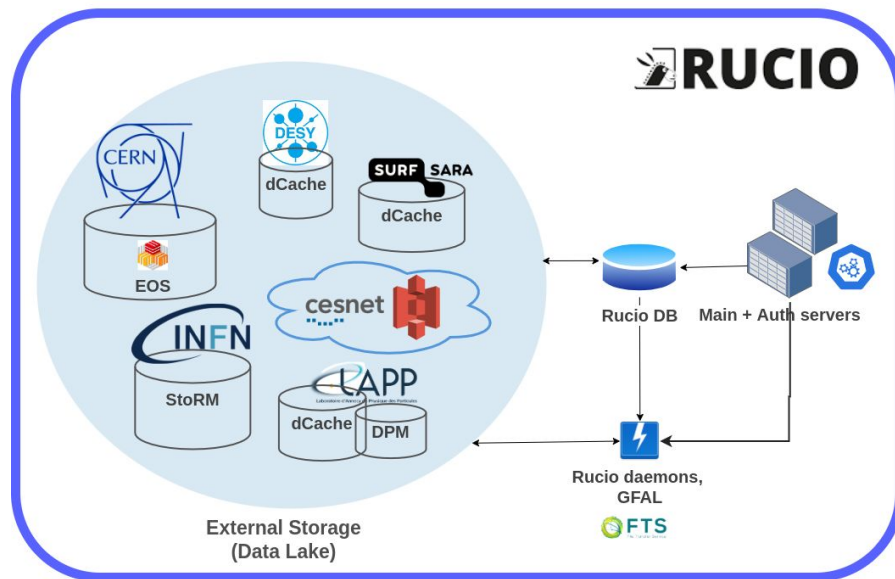
2. Storage: Rucio

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

Rucio is an **open-source data management project** initially developed by the ATLAS experiment to manage large volumes of data. It is now used by various CERN and non-CERN communities.

VRE Rucio instance:

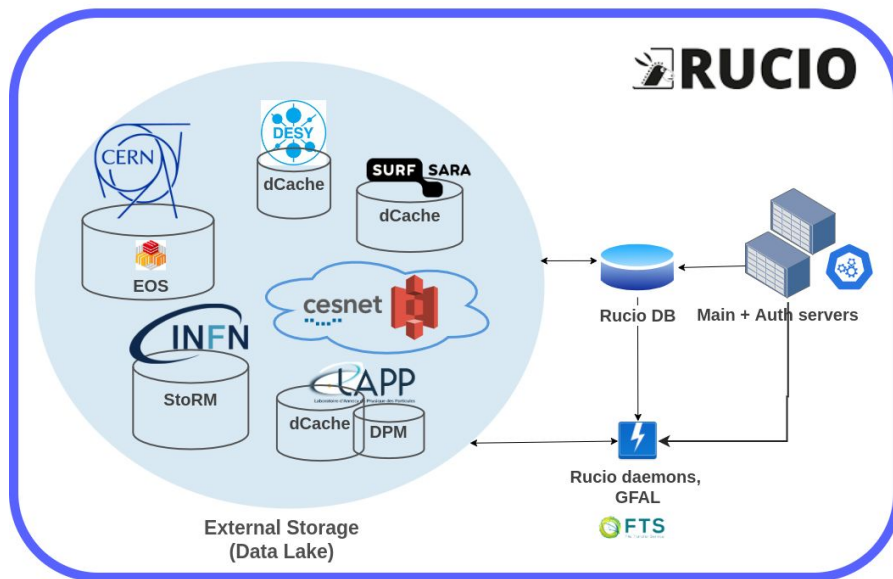
1. **cloud infrastructure**
 - a. Rucio servers, daemons and webUI
 - b. manages API requests, user authentication, data upload, access, download and replication
2. central relational **database**
 - a. backup services in case of major disruptions
3. **Rucio Storage Elements (RSEs)**
 - a. Petabyte-scale
 - b. Distributed, managed by partner institutions
 - c. Support for various storage technologies: **EOS**, **StoRM**, **dCache**, **DPM**, **XRootD**



2. Storage: Data Lake

The **Data Lake** is a **policy-driven, reliable, distributed** data infrastructure able to deliver data **on-demand** at **low latency** to all types of processing facilities.

- ensures data security, quality, access
- Distributed and heterogeneous storage
- File **transfer** with CERN's File Transfer Service (**FTS3**)
- File **upload and download** with CERN's grid file access library (**GFAL2**)
- Support for multiple protocols (**gridFTP, HTTP(S), XRoot, S3**)

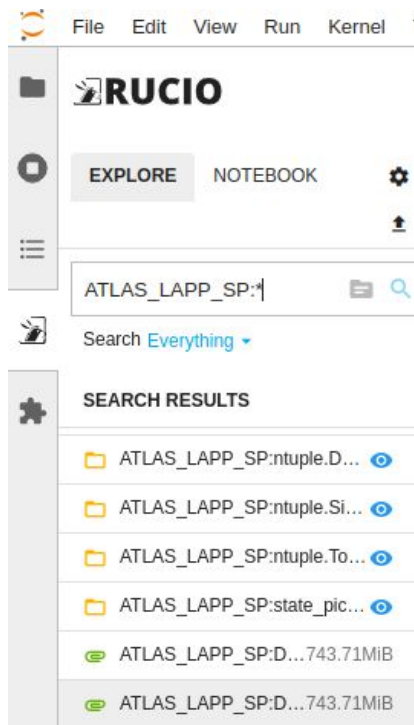


3. Notebook service

Allows users to easily interact with the underlying infrastructure by leveraging **containerised environments**. Used for **preliminary analysis**.

Rucio plug-in:

- Hides complexity of Data Lake – allows browsing Rucio database
- Performs user authentication to Rucio instance via **tokens**
- Allows **data replication** from any storage element to back-end storage element
- Allows to assign a parameter to the data and easily **import it in the notebook**



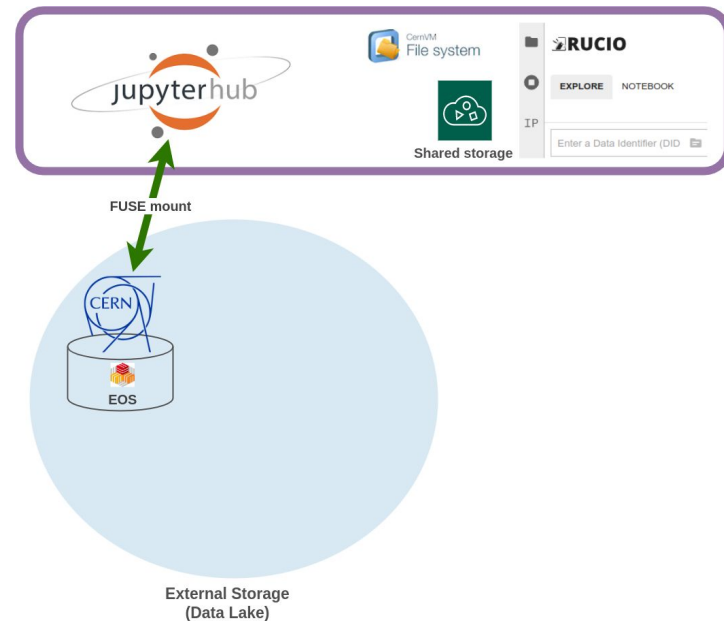
3. Notebook service

External Storage:

- EOS storage element (RSE) of half petabyte **FUSE mounted** on the JH node
- Data **replication** from the Rucio Data Lake to here when requested for analysis

Other attributes:

- 800GB **CephFS volume**, to install software, provided as shared, temporary storage solution
- Rucio and Reana libraries installed – interact as a client
- **CERN Virtual Machine FS (CVMFS)** accessible from terminal



3. User Interface: Jupyterhub

RUCIO

EXPLORE NOTEBOOK

ATLAS_LAPP_SP.*

SEARCH RESULTS

- ATLAS_LAPP_SP.DMCrossSectionGraphs_axial_e... 25.79KiB
- ATLAS_LAPP_SP.DMCrossSectionGraphs_axial_m... 25.79KiB
- Available
- ATLAS_LAPP_SP.DMsummary-dilepton-14tev-2018 39B
- ATLAS_LAPP_SP.DMsummary-dilepton.14TeV.2018 39B
- ATLAS_LAPP_SP.LimitInterpolator_CL95_14TeV.root 26.38KiB
- ATLAS_LAPP_SP.merged_DM_axial_ee_gDM_ip... 167.13MiB
- ATLAS_LAPP_SP.merged_DM_axial_muumu_gDM... 167.13MiB
- ATLAS_LAPP_SP.ntuple.data16_periodD_physics... 38.19KiB
- ATLAS_LAPP_SP.ntuple.data15_periodE_physics... 2.68MiB
- ATLAS_LAPP_SP.ntuple.data15_per... 310.72KiB
- ATLAS_LAPP_SP.ntuple.data15_periodD_physics... 4.67MiB
- ATLAS_LAPP_SP.ntuple.data15_periodE_physics... 5.05MiB
- ATLAS_LAPP_SP.ntuple.data15_period1_physics... 2.52MiB
- ATLAS_LAPP_SP.ntuple.data16_periodA_physics... 792.67KiB
- ATLAS_LAPP_SP.ntuple.data16_periodB_physics... 7MiB

Expected Limit from Dilepton Resonance Analysis

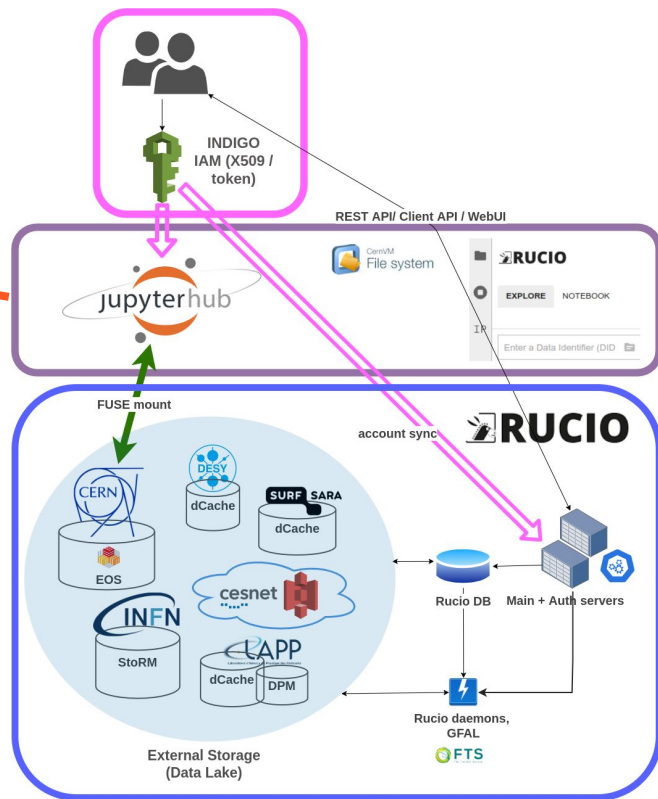
$\sigma_{\text{max}} \times \text{BR}$

Mass Z

```
[5]: cFindLimit = ROOT.TCanvas()
cFindLimit.cd()
cFindLimit.SetLogy()
explimit_ee.SetLineColor(2)

leg = ROOT.TLegend(0.35, 0.72, 0.85, 0.92, "")
leg.SetTextSize(0.04)
leg.SetFillColor(0)
leg.SetBorderSize(0)
leg.AddEntry(explimit_ee, "#font[42]{Expected e+e- limit}", "l")
leg.AddEntry(fidXsec_500_ee, "#font[42]{Vector Z' (DM) (m_{chi}=0.5 TeV)}", "l")
```

Available on FUSE mounted Rucio Storage Element

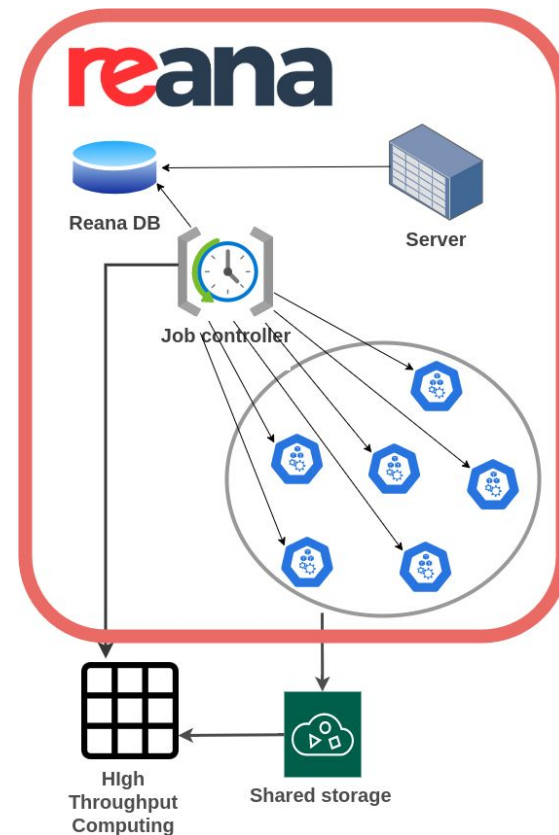


4. Computing: Reana


Reana is a **reproducible analysis** project developed at CERN, to make the preservation of an analysis seamless.

- Used for heavier analyses, it allows workflow distribution on various **computing back-ends**
 - **K8s (default)**
 - **HTCondor**
 - **Slurm**
- Intuitive declarative programming approach (`reana.yaml` file)
- Supports workflow engines (CWL, Snakemake, Yadage)
- Installed on the VRE cloud infrastructure via Helm

<https://reanahub.io/>




4. Computing: Reana



fetched locally

1. **input data** and parameters
2. code
3. computing environments
4. computational steps

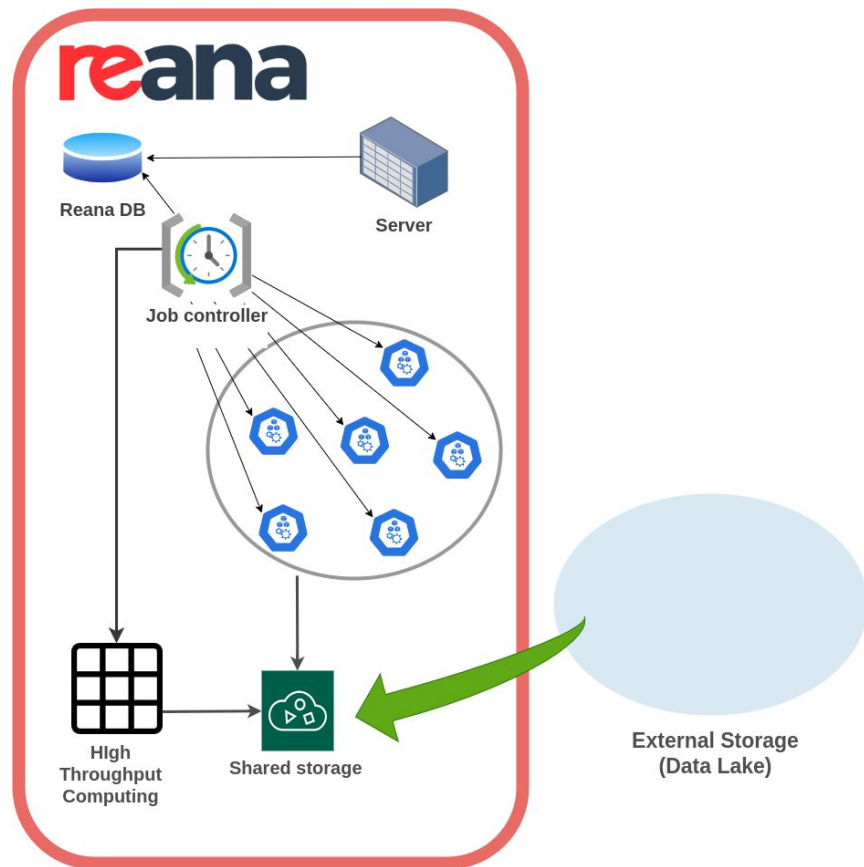


Reana + Rucio

Summer project

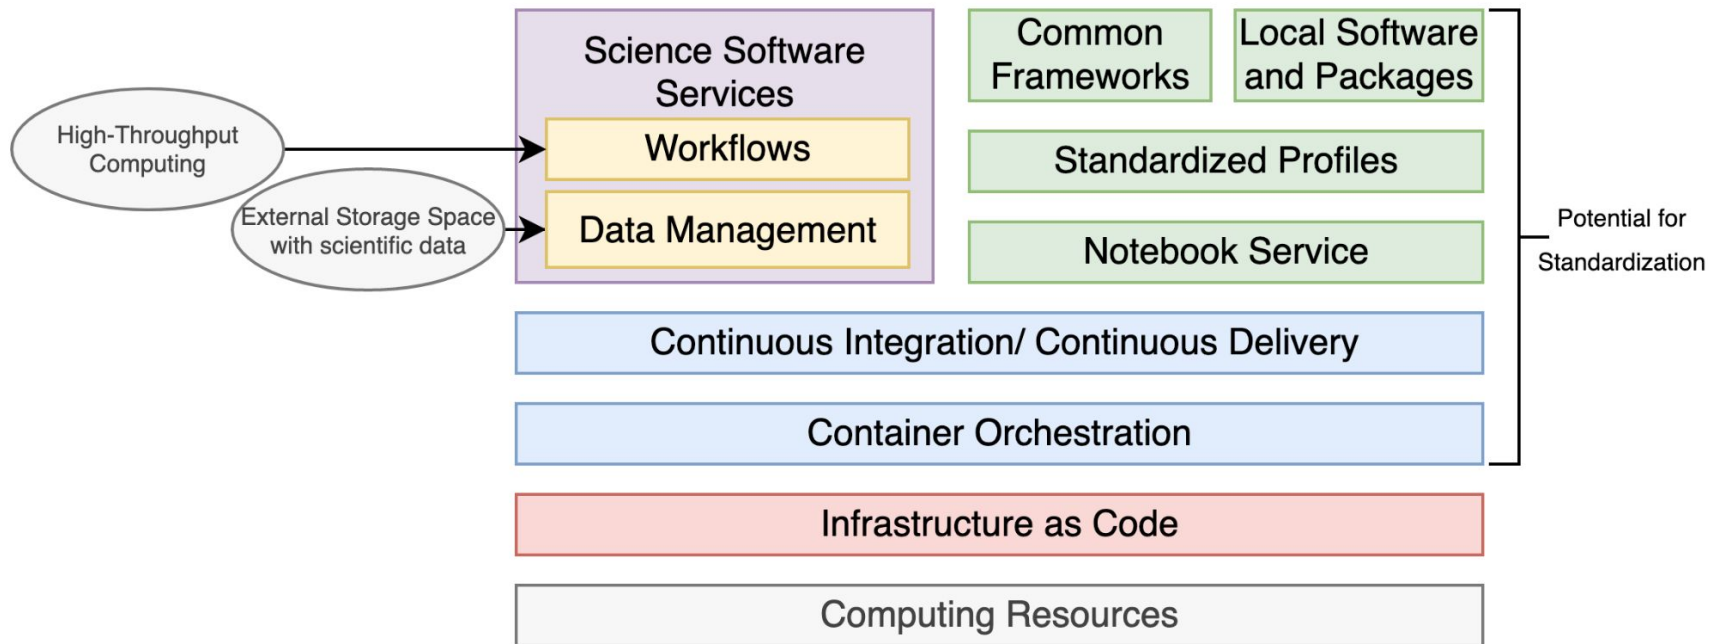
Implementation of **direct download** from Rucio Data Lake to the Reana shared storage via a side-car container that allows authentication to the VRE Rucio instance.

→ the analysis can be reproduced fully and **independently** from local storage.





Technical Aspects & Deployment

IT Architecture Overview




IT Architecture Overview

- Resources provided by CERN OpenStack
 - Computing - Container orchestrator (K8s)
 - Storage
 - Networking
- Management of Infrastructure
 - Terraform - Infrastructure as code (IaC)
 - Services installed via Helm Charts
 - Infrastructure keep updated via GitHub Actions (+ Gitlab CI/CD)
 - Migration to a single repository
 - Sealed secrets
 - CERN Data Base On Demand (DBOD)



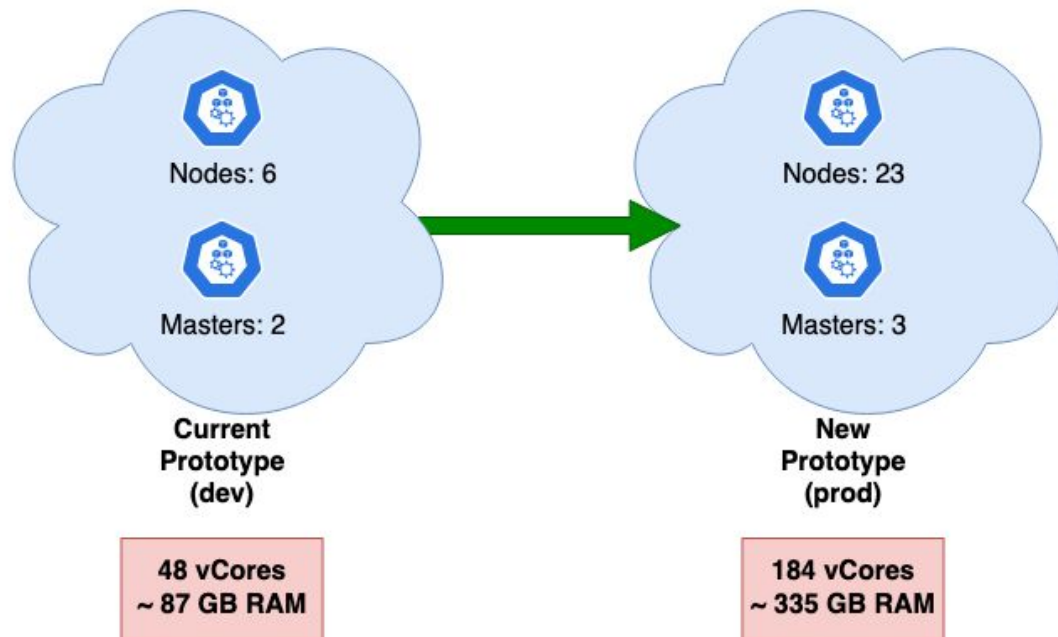
TERRAFORM CLOUD AND GITHUB ACTIONS WORKFLOW



Source: <https://developer.hashicorp.com/terraform/tutorials/automation/github-actions>

The Cluster(s)

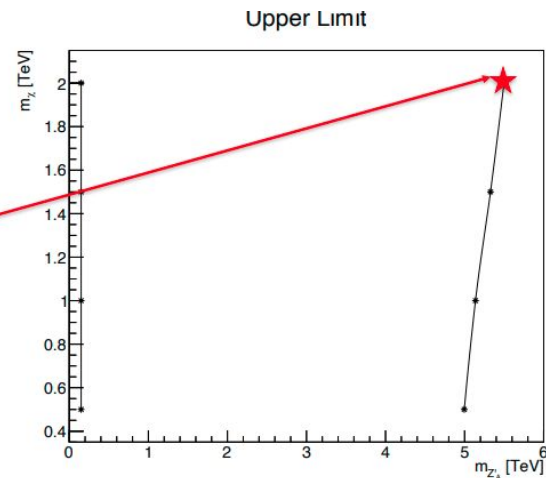
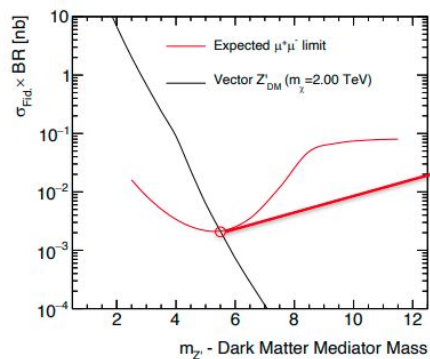

- One Prototype cluster currently in use with **~220 registered users**
- Improved and more powerful cluster under development



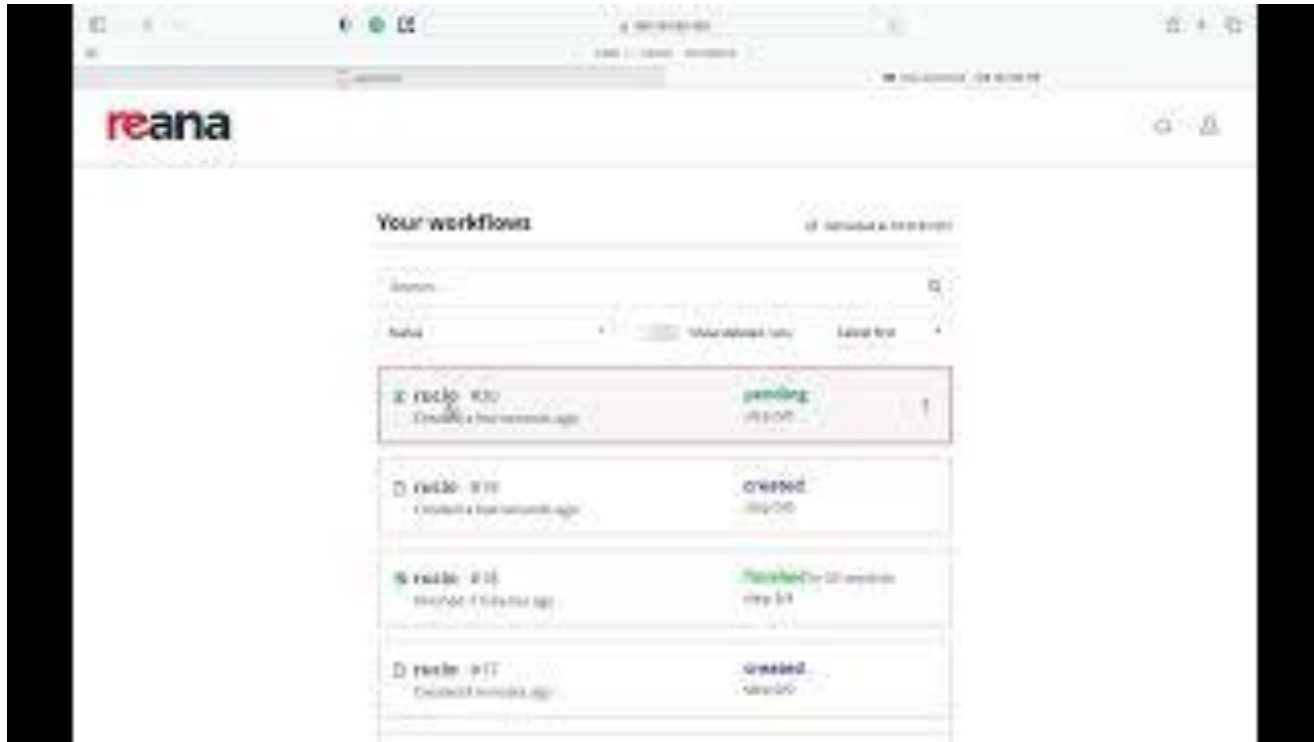
Demo

Dark Matter Dilepton Reinterpretation analysis

To set **limits** on High-Luminosity LHC constraints on $Z' \rightarrow \chi\bar{\chi}$ and project them to 14 TeV to compute the fiducial cross-sections in lower mass regions.



Demo Video



We invite you to try the analysis platform!

The current prototype is available here:

<https://escape-notebook.cern.ch>

together with the documentation here:

<https://datalake-rucio.docs.cern.ch/>

Register a new account with “*ET-EIB workshop*” in the *Notes* field.

It is still under development and subject to change!



Welcome to **escape**

Sign in with your escape credentials

[Forgot your password?](#)

Or sign in with

Not a member?

[Privacy policy](#)

Register at **escape**

This is the escape registration page.

To proceed with the registration please fill in your personal information below.

Given name

Family name

Email

Username

Notes

By submitting this registration request, you declare that you agree with the terms of this organization [Privacy policy](#).

In summary the VRE is..

.. an analysis **platform as a whole**, build from the bottom up! It **integrates important scientific software**, tools, and packages and offers a **common entry point with the same authentication** for all components. The **deployment is kept simple and is extensively documented** so it can be used by other institutes as a **blueprint**.

The platform is built **CERN independent**, with shared **European resources** and access to **data from multiple experiments also outside of HEP**. As a European effort, it is **accessible to people from other institutes** as well.

Thank you for your attention!

Where to find us: *CERN Meyrin site, room 513-1-014*

Contact: *vre-admin@cern.ch*

VRE team



Elena Gazzarrini
CERN Fellow

-
King's College London
Solid state physics
background



Domenic Gosein
Technical Student

-
Mannheim University
of Applied sciences
Computer Scientist

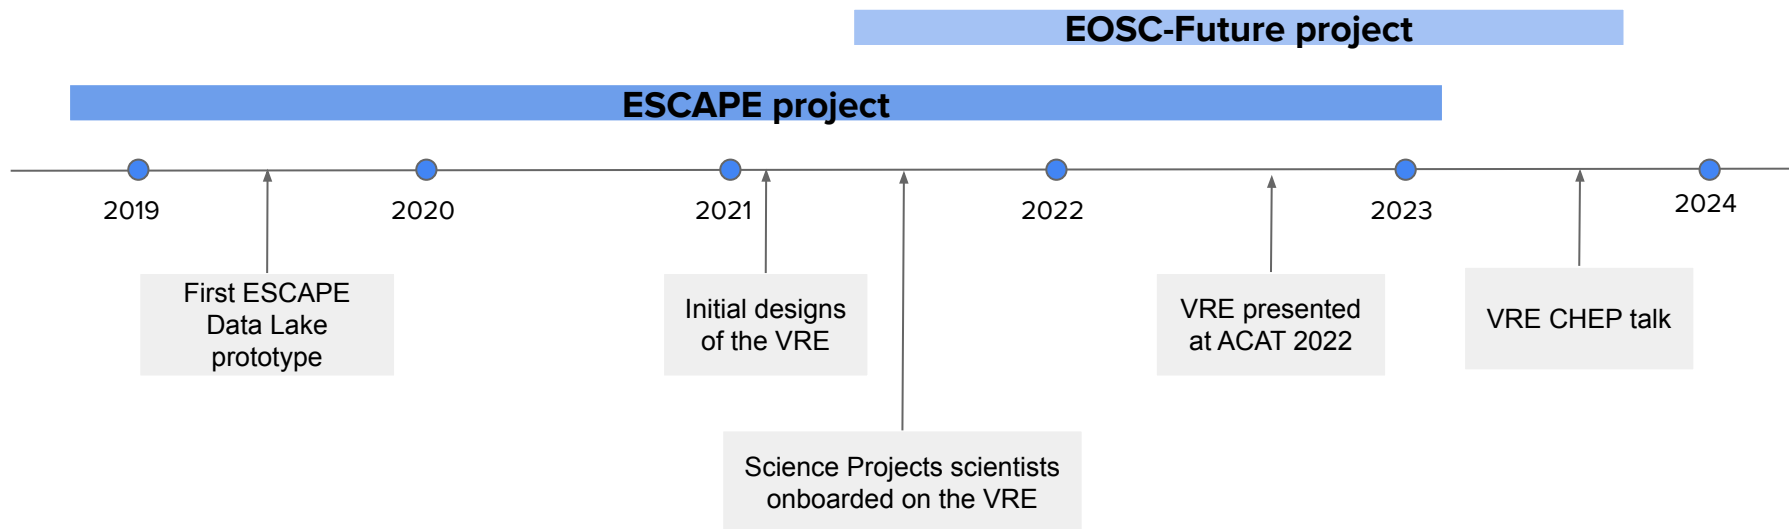


Enrique Garcia
CERN Fellow



-
Univ. Autónoma de
Madrid
Astrophysicist and
software developer
background

Back-up slides

Roadmap and timeline



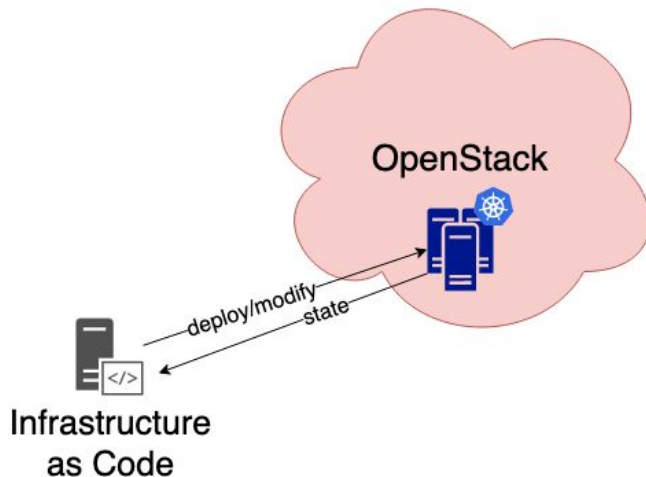
Future Development



Infrastructure as Code

Infrastructure as Code: terraform

```
19 resource "openstack_containerinfra_cluster_v1" "openstack_cluster" {
20   name           = var.cluster-name
21   cluster_template_id = data.openstack_containerinfra_clustertemplate_v1.cluster_template.id # 22b4c77f-cf63-47bb-
8006-31002375a3f3
22   master_count   = 3
23   node_count     = 23
24   keypair        = var.cluster-keypair-name
25   merge_labels   = true
26   flavor         = "m2.xlarge"
27   master_flavor  = "m2.large"
28   labels = {
29     cern_enabled      = "true"
30     cvmfs_enabled     = "true"
31     cvmfs_storage_driver = "true"
32     eos_enabled       = "true"
33     monitoring_enabled = "true"
34     metrics_server_enabled = "true"
35     ingress_controller = "nginx"
36     logging_producer   = var.logging-producer
37     logging_installer  = "helm"
38     logging_include_internal = "true"
39     grafana_admin_passwd = "admin"
40     keystone_auth_enabled = "true"
41     auto_scaling_enabled = "true"
42     min_node_count     = "4"
43     max_node_count     = "23"
44   }
45   provisioner "local-exec" {
46     command = "sh ../../scripts/post_cluster_setup.sh"
47     environment = {
48       cluster = var.cluster-name
49     }
50   }
51 }
52
53 resource "openstack_sharedfilesystem_share_v2" "share_1_reana" {
54   name           = var.reana-share-name
55   description    = "Share for reana"
56   share_proto    = "CEPHFS"
57   size           = 1000
58   share_type     = var.cephfs-type
59 }
```



```
7   required_providers {
8     openstack = {
9       source = "terraform-provider-openstack/openstack"
10      version = "1.49.0"
11    }
12    kubernetes = {
13      source = "hashicorp/kubernetes"
14      version = "2.16.0"
15    }
16    helm = {
17      source = "hashicorp/helm"
18      version = "2.8.0"
19    }
20  }
21 }
```

Cluster Modification with Scripts

```
1 #!/bin/bash
2
3 # post setup cluster configuration --> alternatively also use tf with k8s provider to label resources and get nodes through kubernetes_resource data
4
5 # get cluster config for kubectl
6 $(openstack coe cluster config $cluster)
7 export KUBECONFIG="config"
8
9 # extract node prefix
10 NODE_PREFIX=$(kubectl get nodes -l magnum.openstack.org/role=worker --sort-by .metadata.name -o name | head -n 1)
11 NODE_PREFIX=${NODE_PREFIX%-0}
12 echo $NODE_PREFIX
13
14 openstack loadbalancer set --description "vre-rucio.cern.ch" $LB_ID_MAIN
15 openstack loadbalancer set --description "vre-rucio-auth.cern.ch" $LB_ID_AUTH
16
17 ## set reana HA node labels
18 kubectl label "${NODE_PREFIX}-3" reana.io/system=infrastructure
19 kubectl label "${NODE_PREFIX}-4" reana.io/system=runtimebatch
20 kubectl label "${NODE_PREFIX}-5" reana.io/system=runtimejobs
21 kubectl label "${NODE_PREFIX}-6" reana.io/system=runtimesessions
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45     provisioner "local-exec" {
46         command = "sh ../../scripts/post_cluster_setup.sh"
47         environment = {
48             cluster = var.cluster-name
49         }
50     }
```


Terraform with Helm and Kubernetes


```
73 resource "helm_release" "jupyterhub-chart" {
74   name      = "jhub-${var.resource-suffix}"
75   repository = "https://jupyterhub.github.io/helm-chart/"
76   chart     = "jupyterhub"
77   version   = "2.0.0"
78   namespace = var.ns-jupyterhub
79
80   values = [
81     "${file("jhub/config.yaml")}"
82   ]
83
84   set {
85     name = "hub.db.url"
86     value = data.kubernetes_secret_v1.jhub_db_secret.data.dbconnectstring
87   }
88   set {
89     name = "hub.config.GenericOAuthenticator.client_id"
90     value = data.kubernetes_secret_v1.jhub_iam_secret.data.client_id
91   }
92   set {
93     name = "hub.config.GenericOAuthenticator.client_secret"
94     value = data.kubernetes_secret_v1.jhub_iam_secret.data.client_secret
95   }
96 }
```

```
68 resource "kubernetes_storage_class_v1" "sc_manila-meyrin-cephfs" {
69   metadata {
70     name = "manila-meyrin-cephfs" # ref.: https://kubernetes.docs.cern.ch/docs/storage/fileshares/
71   }
72   storage_provisioner = "cephfs.manila.csi.openstack.org"
73   reclaim_policy      = "Delete"
74   allow_volume_expansion = true
75   parameters = {
76     type = "Meyrin CephFS" # ref.: https://clouddocs.web.cern.ch
77     /file_shares/share_types.html
78     "csi.storage.k8s.io/provisioner-secret-name" = "os-trustee"
79     "csi.storage.k8s.io/provisioner-secret-namespace" = "kube-system"
80     "csi.storage.k8s.io/controller-expand-secret-name" = "os-trustee"
81     "csi.storage.k8s.io/controller-expand-secret-namespace" = "kube-system"
82     "csi.storage.k8s.io/node-stage-secret-name" = "os-trustee"
83     "csi.storage.k8s.io/node-stage-secret-namespace" = "kube-system"
84     "csi.storage.k8s.io/node-publish-secret-name" = "os-trustee"
85     "csi.storage.k8s.io/node-publish-secret-namespace" = "kube-system"
86   }
}
```

Terraform Automation with GH Actions

```
10 terraform:
11   name: "Terraform"
12   runs-on: ubuntu-latest
13   permissions:
14     pull-requests: write
15   steps:
16     - name: Checkout
17       uses: actions/checkout@v3
18
19     - name: Setup Terraform
20       uses: hashicorp/setup-terraform@v1
21       with:
22         # terraform_version: 0.13.0:
23         cli_config_credentials_token: ${ secrets.TF_API_TOKEN }}
24
25     - name: Terraform Format
26       id: fmt
27       run: terraform fmt -check
28
29     - name: Terraform Init
30       id: init
31       run: terraform init
32
33     - name: Terraform Validate
34       id: validate
35       run: terraform validate -no-color
36
37     - name: Terraform Plan
38       id: plan
39       if: github.event_name == 'pull_request'
40       run: terraform plan -no-color -input=false
41       continue-on-error: true
```

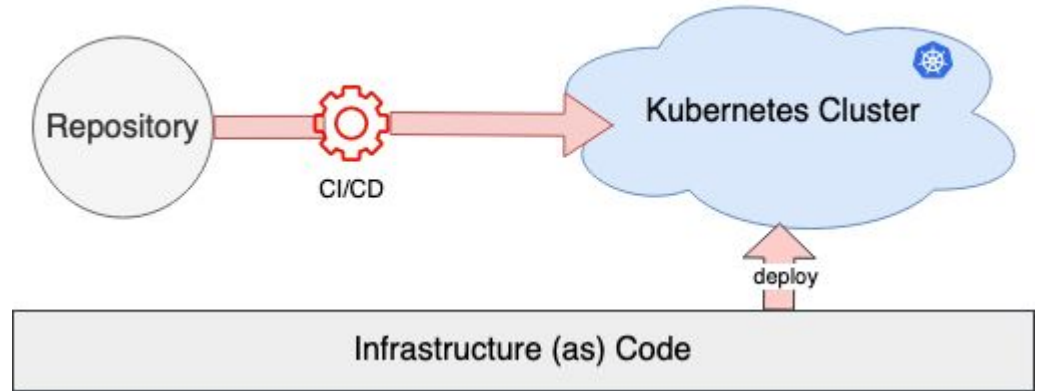
TERRAFORM CLOUD AND GITHUB ACTIONS WORKFLOW



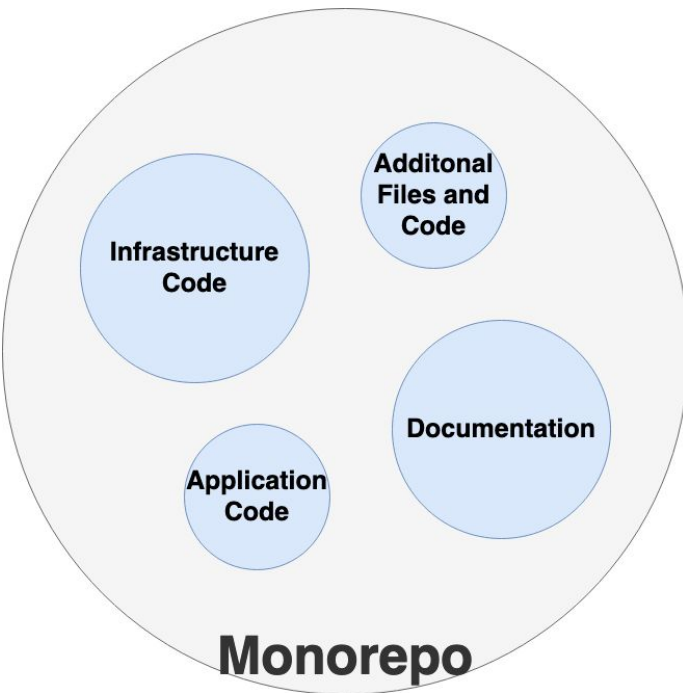
Source: <https://developer.hashicorp.com/terraform/tutorials/automation/github-actions>

Additional CI/CD

- IaC for cloud resources
- GitOps Tools additionally for K8s Manifests and Applications
- GitHub Actions for other workflows like the image registry



Code Management and Documentation



main 3 branches 0 tags Go to file Add file Code

egazzarr Merge pull request #58 from cern-vre/ruccio-dev 23:4428 2 weeks ago 191 commits

- .github/workflows Update merge-check.yml 2 weeks ago
- iac necessary 'main' of https://github.com/cern-vre/cern-vre into ruccio-... 2 weeks ago
- .gitignore change secrets dir, mod .gitignore and add tmp dir last month
- README.md Update README.md 2 weeks ago
- cern-vre-logo.png update logo and rm (#28) last month

README.md

This repository is still under construction

Merge Check passing

CERN VRE

Designed with [freelogodesign](#)

The CERN Virtual Research Environment.

Please find the technical documentation for this repository here [CERN VRE Wiki](#).

Home egazzarr edited this page last week - 14 revisions Edit New page

CERN VRE Technical Documentation

The following is the technical documentation of the CERN Virtual Research Environment, including CERN-specific matters, which are marked accordingly.

TL;TR

In order to create or interact with this cluster you'll need: `kubectl`, `terraform`, `helm` and `kubeseal` installed locally. Further, you'll need the OpenStack RC configuration and the kubeconfig (once the cluster has been created).

Repository Structure

The repo is designed to be a Monorepo, containing all relevant files to this project:

- .github: Workflows
- add: additional configuration files
- iac: Infrastructure as Code
 - scripts: Scripts used for the infrastructure
 - secrets: sealed secrets
 - tmp_local_secrets: not tracked dir for local secrets to work with the scripts
 - tf: Terraform files
 - cluster: Cluster main files
 - modules: Terraform modules used in the cluster

CERN VRE Technical Documentation ©CERN 2023

Page 1 of 1

- Home
- CERN VRE Technical Documentation
- TL;TR
- Repository Structure
- Components
- Developer Setup
- Infrastructure

Clone this wiki locally <https://github.com/cern-vre/cer>

Secrets Management: Sealed Secrets

- Controller installed in the cluster
 - Has the key
 - Creates the K8s secret
- Client side utility to encrypt values and create the custom resource “SealedSecret”

```
1  apiVersion: bitnami.com/v1alpha1
2  kind: SealedSecret
3  metadata:
4    creationTimestamp: null
5    name: jhub-cvre-iam-secrets
6    namespace: jhub
7  spec:
8    encryptedData:
9      client_id: AgB40MgVJ6LSkkXsFBSQ9P1e5tiyXVqUu76sc39bEzwUXvhjZzrkdvLW0X8bzMaX66mKwvesla0CNLj00rjLBdZ8cAgVBX/Kpxo7in/DPfRJKI
/RA0o88X8C2GtcQzV6vKXTgRM+Z0adx+crJMLAoGmg5jSaYmR7NlCy532pEtJ3FWZ81mqQCAF9+S1o0LukGdwF+7ihz/bVEUoPuf0kjAP9oDdL9PBXvL9lHe6Udj5
/q3cEbotoCuVvW0JH1+ygfWeDlps3rc
/ah613FZYjfhHmek0fiP0fg1xXzxHrq8dvSnFkLaVe3YNyn15CJVutkooGjxZs9n6Ucw4XufxGFSf3Ho6mGpkrSmmX9XLYoao4A1ItLA6c46YFcv2nLaDpRTAfgzH
/pNYyNrPERYNgFXoUyC/+aTgcfzLiY75f0F0Hkp+hxuvqkcZhiFesE+LxgxLtThLjJdLpIwCMMUhiDcbf9vaCwe6cQjQEF8B6GSWZ+2dxXM8T1Yzk2cWs/XqJlp
/e049smycj\lHsYkebMB0K0Wpx0pSXhjAm8BLgL7vaPKgbkZfTE19Ks9xZnCZPeiIvXfuKmwA0Tbvboe8GILasSBjJq/Y8Yd+78TgSsxDH8e0IQF/JCdZUDemXZ
10  client_secret: AgBmvm2fe9eJ79Pn1+NEVadUhjAxdVy2TuPVk25Tz9vrK1uuQ/euMCmhoZjzsE5LYWc46AXv
/d05aE052RB8ZEIH4JEJzEK2c+bLb3gu04Nneh7wkG1hsVYHxiTKZomM71cJE3vud7fR1+YoDec0kdZCIHI8Ygz+g6Y7yQvzAZSahc75h0GakJQNrksmzLP92vs
/v4TdploHWSFZ6PWjYJOXTPREBQ0GFjaywKAowrR+iqm8pE0Q4NTFXHXFY5eptVuu4Mv74d8vc6wXATHkre00c5pw8nVorrxh9g0njXXkpwcdL0PKgxQ9uBg1Ny6j
/kQL1jo7psyz21UX5PcCbVewLvR6zjGYj8VP3syd0XuvqY4N5tkaqF3US8F9MAQZzLvhLneeopRwbCMCQVP+e3mP05kLij5Uf1HlSzfj7ekRjL0MgmhMAiF9iunr
/31nHTo1FDyuXV6ZmYHfvRzjP2cj5jeRvyo2GW+ChFCK5jnW6JdFanUoH9h5UK3L0le04wfh73cBUTt6sA4V52DCPPFMJ8YtYL0DNXcIEp7LR+dJM4
/V3VMAeIadZrr4ymA0nA5boU3t23gCVLq310DR6ia8NtCi+DvEtCX/S8H/MN0WN8ch6o7pZfV/BEc9XINrPjvpSDPC/SjNf27zqzIVR
/DwrwhPA6uTMkurhrkFPU5+AFcbQ3FuMycJCB9aqWvVzhn1Cg/EQoxM0iHMPmUmeEpk9vu6T7nbEzTVwRIYYiPnKs0CoG07kkjTVmXpyQVD+Dvpo9PdWnZn1r9YN2
11  template:
12    metadata:
13      creationTimestamp: null
14      name: jhub-cvre-iam-secrets
15      namespace: jhub
16
```

Database

- CERN Database on Demand Instance
- Postgres

The screenshot displays the DBOD (Database on Demand) interface. At the top, there is a navigation bar with the DBOD logo, a 'Light theme' toggle, a '+ REQUEST NEW INSTANCE' button, and a user profile section indicating 'Signed in as Domenic Frency Gosein from CERN'. Below the navigation bar, the instance details for 'vre' (owned by IT-GOV Domenic Frency Gosein) are shown. A green button allows to 'Change owner, admin group or delete instance'. The instance description is 'Replacement of the falsely named vre_rucio instance.' Below this, a table lists instance parameters:

dogosein	escape-cern-ops	Project EOSC Future	Type PG	Version 14.6
Category PROD	Charge Group none	Port 6600	dbod-vre.cern.ch	Creation date 16/12/2022

An 'Expiry date' field is also present. Below the instance details, there are tabs for 'Jobs', 'Logs', 'File Editor', 'Backup and Restore', and 'Clones'. The 'Jobs' tab is active, showing a table of job execution results:

Status	Starting Date	Ending Date	Description
Succeeded	2/13/23, 4:01 PM	2/13/23, 4:02 PM	Submit File Pg_hba.conf requested by Domenic Frency Gosein

Core Notebook Service (JupyterHub)

docker-images
Project ID: 15418

99 Commits 1 Branch 0 Tags 20.1 MB Project Storage

Docker images used for the ESCAPE VRE (Virtual Research Project)

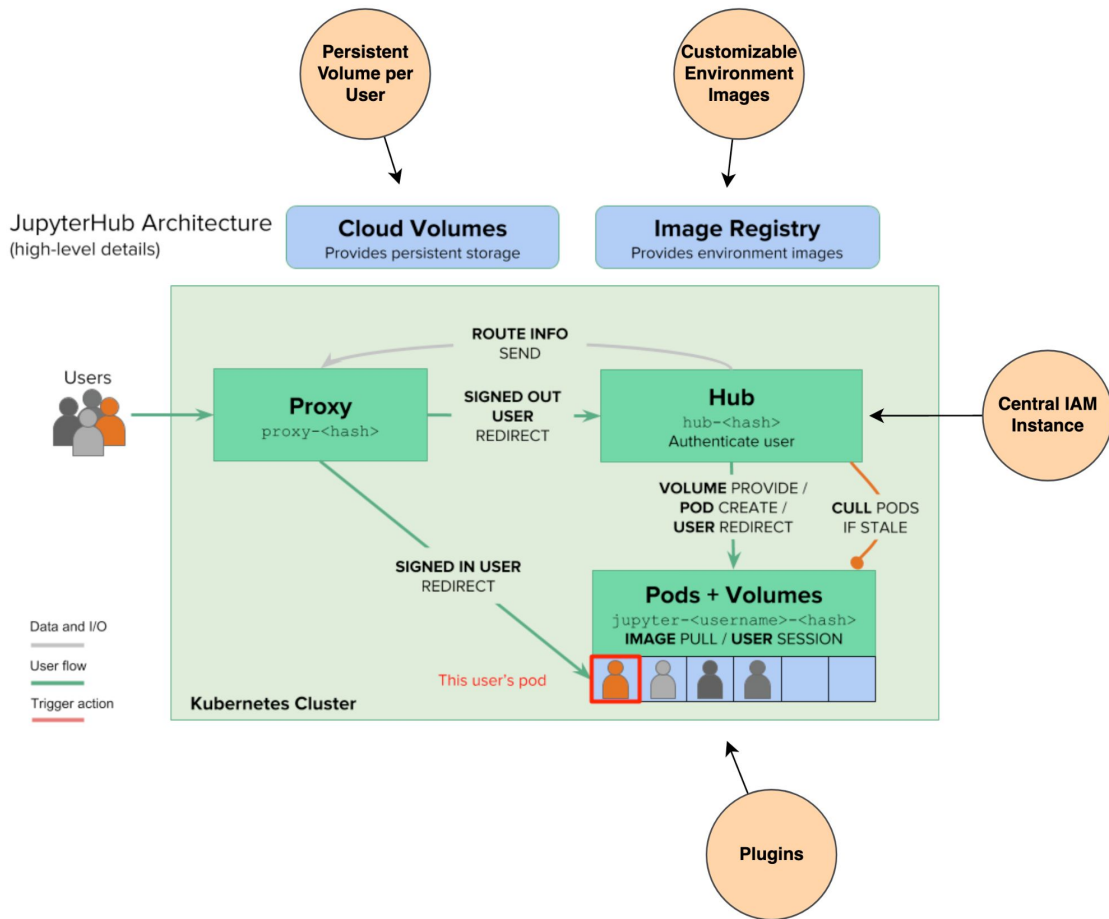
Merge branch 'rebuild_wdf' into 'master'
Enrique Garcia authored 1 week ago

5662237e

master docker-images Find file Clone

README GNU GPLv3

Name	Last commit	Last update
SP-images	update env with gammapy>1.0 and rem...	2 months ago
datalake-singleuser-compact...	latest singleuser	4 months ago
datalake-singleuser-gammato...	Update DLaaS envs	8 months ago
datalake-singleuser-km3net_lrfs	update env with gammapy>1.0 and rem...	2 months ago
datalake-singleuser-micromega	fix micromega env	3 months ago
datalake-singleuser-root-xcac...	Update dockerfiles baselayer	9 months ago
datalake-singleuser-root	Update root env with latest datalake-sing...	4 months ago
datalake-singleuser-virtualLob...	Fix VO project CI	4 weeks ago
datalake-singleuser-wdf	trigger build new WDF env	1 week ago



Source: <https://z2jh.jupyter.org/en/latest/images/architecture.png> (modified)

Community

VRE@CERN
 Technologies developed by CERN within the EOSC Future project to promote open science and collaboration between astroparticle physics communities. Unfollow

3 followers Switzerland https://eoscfuture.eu @EOSCFuture vre-admin@cern.ch

Overview Repositories 10 Discussions Projects 1 Packages Teams 1 People 5 Settings

README.md

CERN Virtual Research Environment

EOSC Future is an EU-funded H2020 project that is implementing the European Open Science Cloud (EOSC). EOSC will give European researchers access to a wide web of FAIR data and related services.

Our team at CERN is developing and contribution to the infrastructure code base of EOSC.

More information can be found on our [website](#).

View as: **Public** ▾
 You are viewing the README and pinned repositories as a public user.

[Get started with tasks](#) that most successful organizations complete.

Top discussions this past month

Discussions are for sharing announcements, creating conversation in your community, answering questions, and more.

[Start a new discussion](#)

Pinned

VRE@CERN

cern-vre Public CERN Virtual Research Env Shell ☆ 3

Overview Repositories 10 Discussions Projects 1 Packages Teams 1 People 5 Settings

Search all discussions Sort by: Latest activity ▾ Label ▾ Filter ▾ New discussion

profiles Public JupyterHub Profiles

Categories

- View all discussions
- Announcements
- General
- Ideas
- Polls
- Q&A
- Show and tell

Discussions


Welcome to Organization Discussions!

Organization discussions are to broadcast news, create conversation in your community, answer questions, and share ideas. To get started, you can [create a new discussion](#).

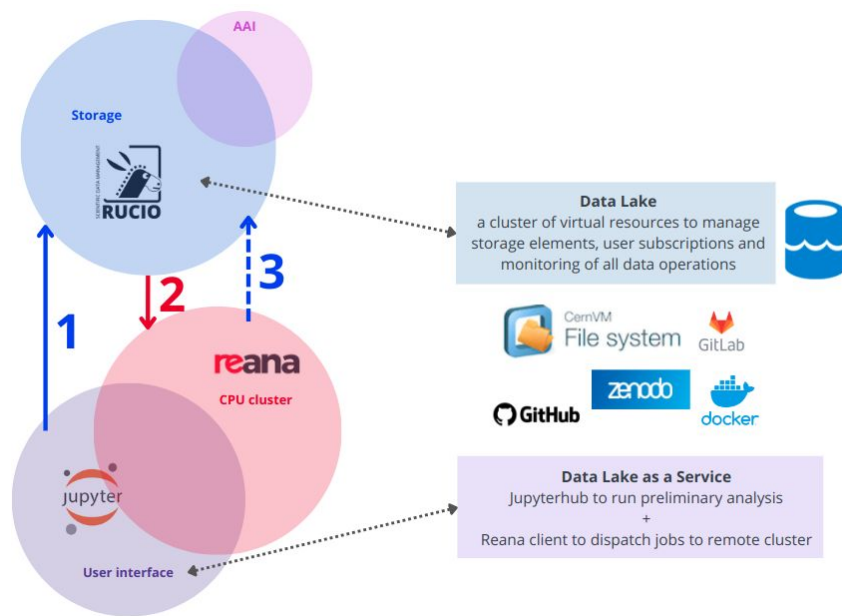

Most helpful

Be sure to mark someone's comment as an answer if it helps you resolve your question — they deserve the credit! 💖

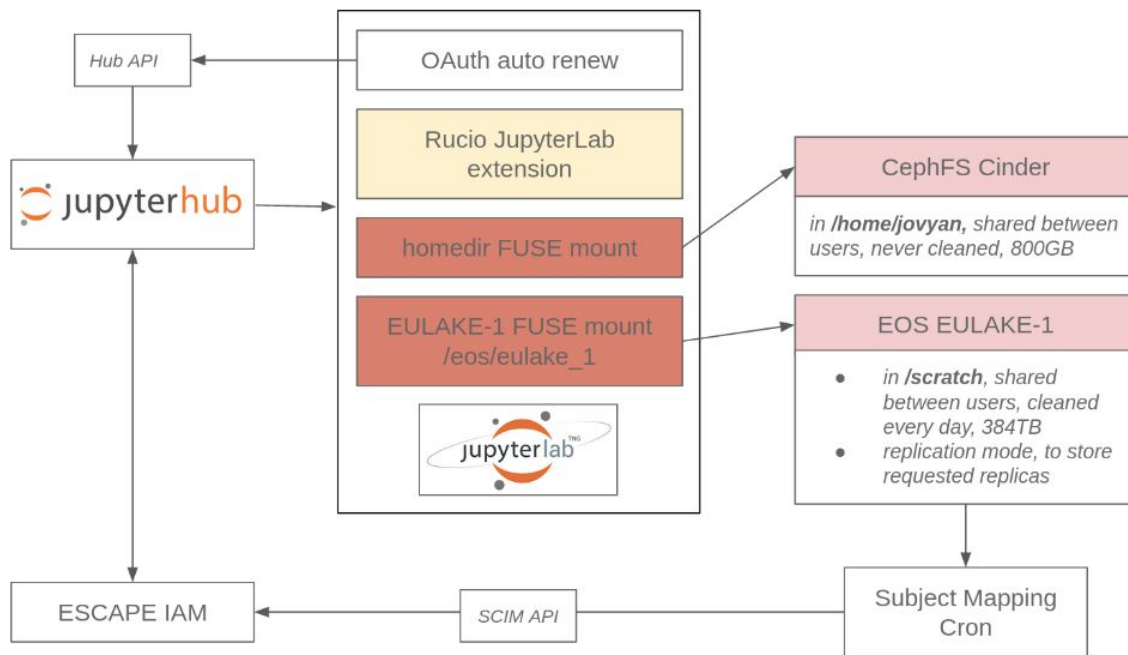
- Community guidelines
- Community insights



Rucio-Reana: summer student project



Jupyterhub current architecture



Technologies

