The Virtual Research Environment: towards a comprehensive analysis platform

Elena Gazzarrini[1], Enrique Garcia[1], Domenic Gosein[1,2], Alba Vendrell Moya[1], Agisilaos Kounelis[1,3], Xavier Espinal[1]

[1]European Council for Nuclear Research (CERN), 1211 Meyrin, Switzerland
[2]Mannheim University of Applied Sciences, Germany
[3]Computer Engineering and Informatics Department, University of Patras, Greece

E-mail: elena.gazzarrini@cern.ch

Abstract. The Virtual Research Environment is an analysis platform developed at CERN serving the needs of scientific communities involved in European Projects. Its scope is to facilitate the development of end-to-end physics workflows, providing researchers with access to an infrastructure and to the digital content necessary to produce and preserve a scientific result in compliance with FAIR principles. The platform's development is aimed at demonstrating how sciences spanning from High Energy Physics to Astrophysics could benefit from the usage of common technologies, initially born to satisfy CERN's exabyte-scale data management needs. The Virtual Research Environment's main components are (1) a federated distributed storage solution (the Data Lake), providing functionalities for data injection and replication through a Data Management framework (Rucio), (2) a computing cluster supplying the processing power to run full analyses with Reana, a re-analysis software, (3) a federated and reliable Authentication and Authorization layer and (4) an enhanced notebook interface with containerised environments to hide the infrastructure's complexity from the user. The deployment of the Virtual Research Environment is open-source and modular, in order to make it easily reproducible by partner institutions; it is publicly accessible and kept up to date by taking advantage of state of the art IT-infrastructure technologies.

1. Introduction

[1] Physicists working at CERN's Large Hadron Collider (LHC) experiments were historically among the first scientists to face large amounts of incoming data, and were therefore forced to find efficient, data-intensive software solutions from an early stage – the implementation of algorithms for the LHC project started back in the 80's –, much before the 'big data' trend emerged on a global scale. Nowadays, PBs of data are saved every day in the CERN Data Center; as an example, the LHCb experiment currently selects 10 GBs of the most interesting LHC collisions each second (after proccessing 4 TBs of data per second in real-time) for physics analysis [2]. CERN developers have therefore accumulated experience and expertise in engineering tools for handling, processing and analysing large data volumes. While High Energy Physics (HEP) sciences have faced these challenges for a long time, non-HEP sciences are more recently entering the exabyte-scale era [3], and therefore need the ability to efficiently track and process the generated data while meeting FAIR (Findabile, Accessibile, Interoperabile, Reusabile) data

principles [1]. However, Open Data alone is not sufficient to foster reuse and reproducibility in physics. It is also essential to capture structured information about the analysis workflows to ensure the usability and longevity of results [4, 5]. A recent literature study has unveiled that half of researchers world-wide – not only physicists, but scientists as a whole – have failed to reproduce their own results [6]; the statistics for HEP and astrophysics research should look more promising due to the historical attention that large-scale collaborations (typical of these fields) have put on reproducibility. However, preserving and reproducing results presents many challenges [7], which can be alleviated by building (re-)analysis platforms that apply logical techniques to describe, illustrate, condense and evaluate data. EU-funded H2020 projects aim to 'democratise' data-intensive technologies, allowing different sciences outside the HEP field – from High Energy Astrophysics to Gravitational Waves searches – to gain expertise on new solutions, eventually fostering cross-fertilisation of sciences. All in all, scientific collaborations are becoming more international; as a consequence, common infrastructures that allow reliable and efficient (i) Federated Data Management and Data Transfer Services, (ii) Federated Distributed Storage, (iii) Data Processing and Orchestration and (iv) Software and Analysis Reproducibility are becoming increasingly popular. The Virtual Research Environment (VRE) tries to encompass all of the above, while placing special attention on the user experience by providing the scientist with an enhanced notebook interface. The VRE's configuration can in addition be flexibly modified to access heterogeneous external resources (storage and computing) managed by EU partner institutions. The following sections will introduce the scientific value of the VRE and illustrate its main components.

2. Scientific value

The VRE concept was incubated in the European Science Cluster of Astronomy, Astroparticle and Particle Physics (ESCAPE [2]) project and is currently being developed and deployed within the EOSC (European Open Science Cloud) Future [3] project, both addressing Open Science challenges to ensure optimised access, management, organisation, processing and preservation of the enormous amount of data handled by the experiments. The tools and concepts initially developed by the ESCAPE work packages, such as the Data Lake (see next section), are hosted and implemented within the VRE, which aims at demonstrating an interdisciplinary science example from bottom-up efforts originating from different scientific domains. In fact, the experiments currently involved in the project come not only from Particle Physics (CERN), but also from High-energy Astrophysics (CTA [8], FermiLAT [9]), Neutrino Observations (KM3NET [10], Darkside [11]), Radio Astronomy (SKA [12], LOFAR [13]) and Gravitational Waves searches (LIGO [14], Virgo [15]).

Many of the aforementioned experiments tackle Dark Matter (DM) exploration from different perspectives. The problem of imposing limits on the mass of DM is a fundamental question in physics: Direct Detection methods study the interaction of particles inside underground detectors, Collider physics produces DM candidates from accelerating protons, Astrophysics observes distant phenomena in the sky and compares them with the theory to detect abnormal behaviours, while Indirect Detection methods investigate annihilating DM by looking at its decay products, such as neutrinos. Figure 1 illustrates this concept and shows how a platform such as the VRE is a useful place to collect results and to generate combined plots to impose universal DM mass limits.

[1] https://www.go-fair.org/fair-principles/
[2] https://projectescape.eu/
[3] https://eoscfuture.eu/

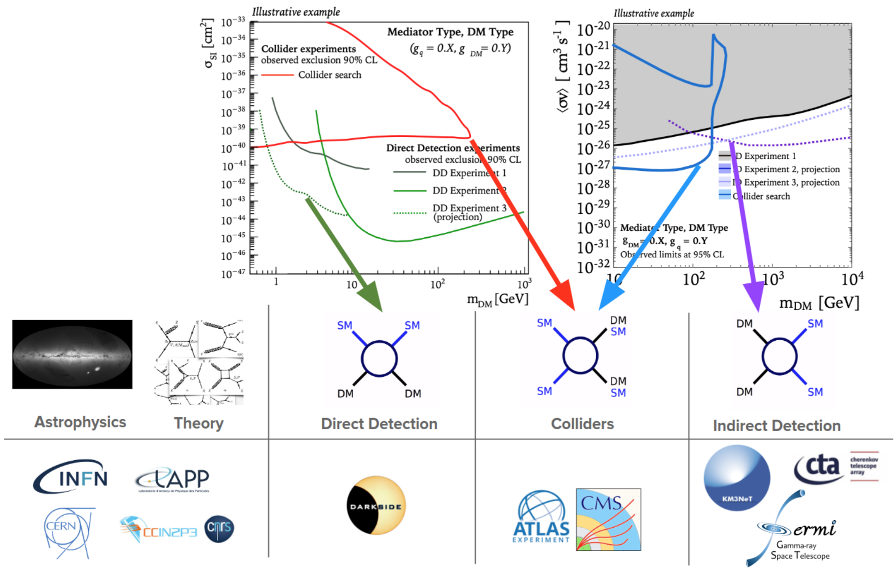

Figure 1. EOSC-Future's Dark Matter Science Project aims at bringing together different search approaches (Astrophysics, Theory, Direct Detection, Collider Physics, Indirect Detection), to ultimately investigate limits on DM mass.

3. VRE Components

In its endeavour to homogenise the technological needs of diverse scientific communities, the VRE consists of (1) a data management framework, (2) access to computing processing resources, (3) users management through a reliable Authorisation and Authentication Infrastructure (AAI) and (4) exposure to a user interface to facilitate the interaction with the underlying infrastructure. Figure 2 graphically illustrates the architecture, supported by CI/CD cycles, container orchestration and Infrastructure-as-Code (IaC) processes. The VRE's deployment is centrally managed on CERN's Cloud infrastructure (details can be found in Table 1) with Kubernetes (K8s).

vCPUs	RAM (GB)	Masters	Nodes	Remote Storage (TB)	CephFS (TB)
184	335.8	3	23	646	1.8

Table 1. The VRE technical components. The first 4 columns refer to the cloud infrastructure managed with K8s, while the last two columns refer to the total quota of the remote storage elements and of the shared object storage (CephFS) attached to the processing nodes.

3.1. Data Management, the Data Lake

The data management and storage orchestration for the VRE is largely based on a scientific software developed at CERN, Rucio [16]. Rucio is an open-source project initially developed by the ATLAS [17] experiment for managing community data. It provides services and associated libraries to manage large volumes of data spread across facilities at multiple institutions and organizations. The VRE Rucio instance is composed of (i) a cloud infrastructure described in Table 1, where Rucio servers, daemons and webUI, installed through Helm charts, manage API requests, user authentication, data upload, access, download and replication, (ii) a central relational database hosted at CERN, containing experiments' data, providing a backup copy in case of major disruptions and (iii) multiple Rucio Storage Elements (RSEs) with quotas

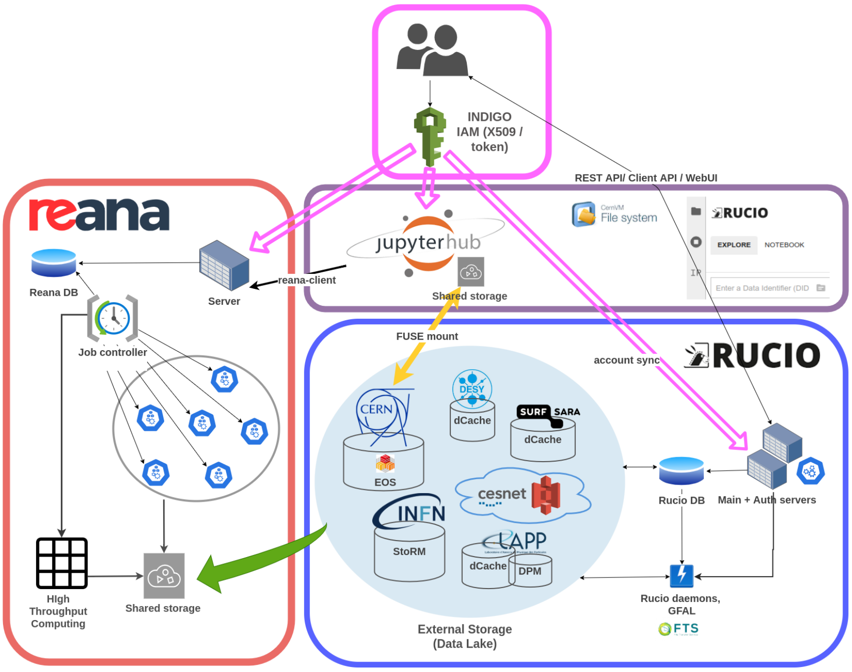

Figure 2. A graphical representation of the VRE components, i.e. (1) a federated distributed storage solution (blue), (2) a computing cluster (red), (3) a federated AAI layer (pink) and (4) an enhanced notebook interface (purple).

varying from 1 to 300 TB, hosted at each partner institution, supporting various storage technologies (EOS [18], StoRM [19], dCache [20], DPM [21], XRootD [4]) and using diverse back-ends (classic RAID systems, Ceph, and multi-replica). Data can be accessed through gridFTP, HTTP(S), XRootD and S3 protocols. Such policy-driven, reliable, distributed data infrastructure, commonly referred to as the Data Lake [22, 23], is able to deliver data on-demand at low latency to all types of processing facilities.

The main functionalities that Rucio offers are (1) data upload, download and streaming, executed by exploiting the power of GFAL [5], and (2) third-party asynchronous transfers (between RSEs) and deletion, achieved instead by CERN's File Transfer Service (FTS) [24]. The latter is granted permission to access the various RSEs by Rucio processes called daemons, which are responsible for any data management action on the infrastructure.

3.2. Computing: Reana cluster

The processing of data in the VRE is managed by an instance of CERN's reproducible analysis platform, Reana [25], which allows to run analyses on various computing backends (K8s, HTCondor, Slurm). Navigating the platform is made intuitive for the scientist, who only needs to prepare a declarative `.yaml` file containing instructions on where to find: (i) input data and parameters, (ii) code, (iii) computing environment and (iv) computational steps needed to perform a full analysis. In this way, scientists can maintain and compare lists of past runs and share the results with colleagues. Reana's workflow distribution on the cluster's virtual nodes is managed by default by K8s. However, computationally heavier analysis steps can be dispatched to High Performance Computers (HPCs) via HTCondor or Slurm, given the assumption the

[4] https://xrootd.slac.stanford.edu/
[5] https://dmc-docs.web.cern.ch/dmc-docs/gfal2/gfal2.html

user has access rights to remote resources. The independence of the Reana framework from local storage was implemented by the VRE team by adding a feature [6] that allows user authentication to the Rucio Data Lake as the first step of the Reana analysis run (represented in Figure 2 by a green thick arrow). In this way, data can be moved between the Data Lake's storage elements and the Reana shared storage, located close to the K8s processing nodes where the analysis steps are distributed. The user is then able to select the desired input data and pull it directly from the Rucio Storage Elements, without the need fo having the data locally before starting the analysis.

3.3. Authentication and Authorization Infrastructure

Given the heterogeneous composition of the VRE infrastructure, it is essential to have a single authentication and authorisation method that comprises all services and grants users the correct permissions to access them. The VRE's AAI layer is based on the INDIGO Identity and Access Management (IAM) service [26]. The VRE IAM instance (inherited by the ESCAPE one) is deployed on a K8s cluster at INFN-CNAF and supports authentication via EduGAIN, via OIDC tokens and via X.509 certificates/Virtual Organization Membership Service (VOMS [7]) attribute provisioning services. The token authentication to remote storage elements for data access and transfer – initially representing the biggest challenge – has been successfully tested on all the VRE's RSEs. The IAM authentication to the Reana instance is currently being implemented and will appear in the next software release version.

3.4. User interface: enhanced notebook service

The online entry point of the VRE is a JupyterHub interface, where scientists can run preliminary analysis. The user gets authenticated via IAM and selects the desired computational environment, automatically pulled from the VRE container registry. The software is therefore already installed in the Jupyterlab session specific to each user pod. Software distribution services such as the CERN Virtual Machine File System (CVMFS [8]) additionally allow software installation on a CephFS 800GB shared volume compatible with POSIX standards, mounted on the Jupyterhub node. In order to ensure better data security on the platform and to avoid a user filling up the shared volume (leading to an interruption of the JupyterLab session for all users), the JupyterLab interface has been enhanced with a Rucio plug-in [9] (represented in the purple box of Figure 2). This feature enables the user to browse the data in the Data Lake and make a copy of it on a CERN RSE of $\sim$0.5 T, which has been FUSE mounted on the Jupyterhub node (yellow arrow in Figure 2); the data is therefore stored close to the processing power, minimising latency. On the other hand, the Jupyterhub node consists of 14GB of RAM and its usage is limited to an exploratory analysis run; to start larger analyses, it is necessary to connect to the VRE Reana cluster via the terminal of the Jupyterhub and dispatch the computation to distributed HPCs.

4. Conclusion

The modular ecosystem of services and tools constituting the VRE represents an European attempt to demonstrate a bottom-up, FAIR approach to scientific collaboration. The weekly on-boarding of new members requesting an account to access the VRE (with a total of more than 200 users) signifies the community need of such novel infrastructure, that encompasses all the resources needed to easily run an end-to-end analysis. The project has been useful

[6] https://docs.reana.io/advanced-usage/access-control/rucio/
[7] https://github.com/italiangrid/voms
[8] https://cernvm.cern.ch/fs/
[9] https://github.com/rucio/jupyterlab-extension

in contributing to improve the software stack of consolidated technologies inside CERN, such as Rucio and Reana. The VRE's applicability to different scientific use cases has been proven successful: postdocs coming from HEP and astrophysics are already using Reana to preserve their workflows on the VRE. The deployment of the infrastructure is kept simple and is extensively documented so it can be used by other institutes as a blueprint: site administrators from collaborations such as the Einstein Telescope and the Deutsches Zentrum für Astrophysik have already demonstrated interest in emulating the VRE at their home institutions. This represents a fundamental achievement under both sociological and technological aspects for European collaborations that should address upcoming data management and computing challenges in the next decade.

Code Availability

The deployment of the VRE infrastructure is still under construction, but the code is available on the public CERN VRE Github project [10], along with the necessary documentation to reproduce it. The VRE landing page [11], provides links to the source codes and description of the various EOSC-Future Science Projects.

Acknowledgements

Authors acknowledge support from the ESCAPE and EOSC-Future projects. ESCAPE has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement 824064. EOSC-Future is co-funded by the European Union Horizon Programme call INFRAEOSC-03-2020, Grant Agreement 101017536.

References

[1] E. Gazzarrini et al. *The Virtual Research Environment: towards a comprehensive analysis platform.* 2023. arXiv: `2305.10166 [physics.comp-ph]`.

[2] C. Agrigoroae. *LHC experiments are stepping up their data processing game.* URL: `https://home.cern/news/news/computing/lhc-experiments-are-stepping-their-data-processing-game` (visited on 08/24/2022).

[3] M. Rahman et al. "Probing Diverse Phenomena through Data-Intensive Astronomy". en. In: (2019). DOI: `10.5281/ZENODO.3765389`. URL: `https://zenodo.org/record/3765389`.

[4] X. Chen et al. "Open is not enough". In: *Nature Physics* 15.2 (2019), pp. 113–119.

[5] S. Dallmeier-Tiessen and T. Šimko. *Open science: A vision for collaborative, reproducible and reusable research.* URL: `https://cerncourier.com/a/open-science-a-vision-for-collaborative-reproducible-and-reusable-research/` (visited on 08/31/2022).

[6] M. Baker. "1,500 scientists lift the lid on reproducibility". In: *Nature* 533.7604 (2016).

[7] T. Junk and L. Lyons. "Reproducibility and Replication of Experimental Particle Physics Results". In: *Harvard Data Science Review* 2.4 (Dec. 2020). https://hdsr.mitpress.mit.edu/pub/1lhu0zvn. DOI: `10.1162/99608f92.250f995b`.

[8] M. Actis et al. "Design concepts for the Cherenkov Telescope Array CTA: An advanced facility for ground-based high-energy gamma-ray astronomy". English. In: *Experimental Astronomy* 32.3 (Dec. 2011). M1 - Journal Article, pp. 193–316. ISSN: 0922-6435. DOI: `10.1007/s10686-011-9247-0`.

[9] J. M. Casandjian. *The Fermi-LAT model of interstellar emission for standard point source analysis.* 2015. DOI: `10.48550/ARXIV.1502.07210`. URL: `https://arxiv.org/abs/1502.07210`.

[10] https://github.com/cern-vre

[11] https://escape2020.pages.in2p3.fr/virtual-environment/home/

[10] A. Margiotta. “The KM3NeT deep-sea neutrino telescope”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 766 (Dec. 2014), pp. 83–87. DOI: 10.1016/j.nima.2014.05.090. URL: https://doi.org/10.1016%2Fj.nima.2014.05.090.

[11] M. Bossa. “DarkSide-50, a background free experiment for dark matter searches”. In: *Journal of Instrumentation* 9.01 (Jan. 2014), p. C01034. DOI: 10.1088/1748-0221/9/01/C01034. URL: https://dx.doi.org/10.1088/1748-0221/9/01/C01034.

[12] A. Weltman et al. “Fundamental physics with the Square Kilometre Array”. In: *Publications of the Astronomical Society of Australia* 37 (2020). DOI: 10.1017/pasa.2019.42. URL: https://doi.org/10.1017%2Fpasa.2019.42.

[13] M. P. van Haarlem et al. “LOFAR: The LOw-Frequency ARray”. In: *Astronomy & Astrophysics* 556 (July 2013), A2. DOI: 10.1051/0004-6361/201220873. URL: https://doi.org/10.1051%2F0004-6361%2F201220873.

[14] B. P. Abbott et al. “LIGO: the Laser Interferometer Gravitational-Wave Observatory”. In: *Reports on Progress in Physics* 72.7 (June 2009), p. 076901. DOI: 10.1088/0034-4885/72/7/076901. URL: https://doi.org/10.1088%2F0034-4885%2F72%2F7%2F076901.

[15] F. Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2 (Dec. 2014), p. 024001. DOI: 10.1088/0264-9381/32/2/024001. URL: https://doi.org/10.1088%2F0264-9381%2F32%2F2%2F024001.

[16] M. Barisits et al. “Rucio: Scientific Data Management”. In: *Computing and Software for Big Science* 3.1 (2019), p. 11. ISSN: 2510-2044. DOI: 10.1007/s41781-019-0026-3. URL: https://doi.org/10.1007/s41781-019-0026-3.

[17] G. Aad et al. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[18] A. J. Peters, E. A. Sindrilaru, and G. Adde. “EOS as the present and future solution for data storage at CERN”. In: *Journal of Physics: Conference Series* 664.4 (Dec. 2015), p. 042042. DOI: 10.1088/1742-6596/664/4/042042. URL: https://dx.doi.org/10.1088/1742-6596/664/4/042042.

[19] R. Zappi et al. “StoRMon: an event log analyzer for Grid Storage Element based on StoRM”. In: *Journal of Physics: Conference Series* 331.6 (Dec. 2011), p. 062047. DOI: 10.1088/1742-6596/331/6/062047. URL: https://dx.doi.org/10.1088/1742-6596/331/6/062047.

[20] P. Fuhrmann and V. Gülzow. “dCache, Storage System for the Future”. In: *Euro-Par 2006 Parallel Processing*. Ed. by Wolfgang E. Nagel, Wolfgang V. Walter, and Wolfgang Lehner. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 1106–1113. ISBN: 978-3-540-37784-9.

[21] A. Álvarez Ayllón et al. “DPM: Future Proof Storage”. In: *Journal of Physics Conference Series* 396 (Dec. 2012), pp. 2015–. DOI: 10.1088/1742-6596/396/3/032015.

[22] R. Di Maria and R. Dona. “ESCAPE Data Lake-Next-generation management of cross-discipline Exabyte-scale scientific data”. In: *EPJ Web of Conferences*. Vol. 251. EDP Sciences. 2021, p. 02056.

[23] R. Bolton et al. “ESCAPE prototypes a data infrastructure for open science”. In: *EPJ Web Conf.* 245 (2020), p. 04019. DOI: 10.1051/epjconf/202024504019. URL: https://doi.org/10.1051/epjconf/202024504019.

[24] E. Karavakis et al. “FTS improvements for LHC Run-3 and beyond”. In: *EPJ Web Conf.* 245 (2020), p. 04016. DOI: 10.1051/epjconf/202024504016. URL: https://doi.org/10.1051/epjconf/202024504016.

[25] T. Simko et al. "REANA: A system for reusable research data analyses". In: *EPJ web of conferences*. Vol. 214. EDP Sciences. 2019, p. 06034.

[26] A. Ceccanti et al. "The INDIGO-Datacloud Authentication and Authorization Infrastructure". In: *Journal of Physics: Conference Series* 898.10 (Oct. 2017), p. 102016. DOI: 10.1088/1742-6596/898/10/102016. URL: https://dx.doi.org/10.1088/1742-6596/898/10/102016.