

CS3 2023 - Cloud Storage Synchronization and Sharing

Monday, 6 March 2023 - Wednesday, 8 March 2023



Book of Abstracts

Contents

Introduction	1
Welcome	1
Technical and scientific challenges in biomedicine; an European perspective	1
Evolving SWAN towards an Analysis Facility system	1
Providing on-demand interactive notebook-based environments with transparent access to cloud storage and specialised hardware through the INFN Cloud Platform	1
VOIS library: pushing data science dashboards to the limits	2
Seafile, what's new in the year 2022	4
Nextcloud. State of the nation	4
ownCloud Infinite Scale	4
Site Report Summary	5
Standing innovation for 10 years: CERNBox	5
Sync & Share self service provisioning on INFN Cloud	5
Research Drive, a platform for active research data management	6
Sunet Drive - An Academic EFSS Packaged for EOSC	7
Sciebo Site Report: Challenges in supporting more than just the core project	8
The GÉANT Special Interest Group on Cloud Interoperable Software Stacks	8
Tripleo: exploit ansible to customize both the Undercloud and the Overcloud	8
Sovereign Cloud Stack - A common Open-Stack for digital sovereignty	9
LIQO – Building and Orchestrating your Service Continuum	9
Learning Management System based on sync & share storage at ETH Zürich	9
ownCloud Infinite Scale: New scenarios for Research and Higher Education	10
Stop data sharing.	10
Sciebo RDS - reducing friction of FAIR data handling for researchers	10

Keeper: certifying research data combining Seafire and blockchain technology	11
HIFIS & Nextcloud: VO Federation for EFSS	12
Cubbit Next Generation Cloud	12
AAI-powered Federated Groups for SURF Research Drive	13
ScienceMesh: an interoperable federation of EFSS services	13
OCM Test Suite current status	14
A View on the CS3 APIs	14
OCM and CS3 discussion	14
ScienceMesh Demos Introduction	15
Connecting Nextcloud & OC-10 to the ScienceMesh	15
Invitation Workflow and Federated Sharing: NC, OC10	15
User-friendly OCM Invitation and Sharing: OCIS, CERNBox	15
Data Science Environments: JupyterLab sharing and collaborative editing	15
Data Transfers: Connecting Science Mesh and ESCAPE Data Lakes	16
Applications integration beyond local clouds with OCM	16
OFORMs for document automation and collaboration	16
Collabora Online: Better, Digitally Sovereign Document Collaboration	17
Status Update of the no-code platform SeaTable	17
An unhappy end for S3	18
From 1VM+1LUN to k8s+Ceph - the uneasy way...	18
Comparison between CephFS, CFFS (Comtrade FastFS), HDFS (Apache Hadoop), GPFS (IBM Spectrum Scale), Lustre	19
C(ERN) BACK(UP): consolidated multi-petabyte backup solution for heterogenous storage and filesystems	19
Data at scale: from Storage to Data Governance	20
Protocol Plumbing: Presenting iRODS as WebDAV, FUSE, REST, NFS, SFTP, K8s CSI, and S3	20
Secure Zones for Sunet Drive	21
On-demand cloud-based secure environments for analysing personal and health data	22
OIDC, it seems to be hip	22
Driving the ScienceBox package into the future	23

An overview of Nextcloud monitoring 23

Modern notification system for Sync and Share using NATS 23

Visualize big data and create individual customer frontends 24

FaaS Data Processing with Onedata 24

Data Science environments in ScienceMesh 25

Summary and Conclusions 26

Science Mesh demos 26

Introduction & Welcome / 122

Introduction

Corresponding Author: jakub.moscicki@cern.ch

Introduction & Welcome / 121

Welcome

Corresponding Author: jonathan.wareham@esade.edu

Keynote / 114

Technical and scientific challenges in biomedicine; an European perspective

Corresponding Author: alfonso.valencia@bsc.es

Collaborative Data Science and Visualisation / 76

Evolving SWAN towards an Analysis Facility system

Author: Diogo Castro¹

¹ CERN

Corresponding Author: diogo.castro@cern.ch

With the HL-LHC, the HEP community will experience orders of magnitude more data at a multi-exabyte scale. To prepare for such unprecedented scientific data collection, the different research sites are combining their diverse resources into integrated Analysis Facilities' systems.

SWAN, CERN's Service for Web based ANalysis, is following this approach, evolving from a plain notebook-based service into a fully fledged Analysis Facility. This allows it to serve as a simple to use, single entry point to the multiple and heterogeneous storage, software and computational resources provided to CERN's research community.

This presentation will focus on the challenges of integrating such different resources and the future direction of the service, including the increasing role of sync&share in the scientific workflows.

Collaborative Data Science and Visualisation / 69

Providing on-demand interactive notebook-based environments with transparent access to cloud storage and specialised hardware through the INFN Cloud Platform

Authors: Marica Antonacci¹; Ahmad Alkhansa^{None}; Luca Giovanni Carbone^{None}; Daniele Cesini²; Diego Ciangottini³; Vincenzo Ciaschini^{None}; Alessandro Costantini⁴; Giacinto Donvito⁵; Doina Duma¹; Federica Fanzago⁶; Nadina Foggetti^{None}; Emidio Giorgio⁷; Alessandro Italiano^{None}; DIEGO MICHELOTTO⁵; Barbara Martelli^{None}; Gianluca

Peco²; Carmelo Pellegrino^{None}; Davide Salomoni²; Massimo Sgaravatto⁶; Francesco Sinisi^{None}; Daniele Spiga⁸; Vincenzo Spinoso¹; Stefano Stalio^{None}; Lucio Strizzolo⁵; Sergio Traldi^{None}; Marco Verlatto⁶; Enrico Vianello^{None}

¹ INFN

² *Universita e INFN, Bologna (IT)*

³ *INFN, Perugia (IT)*

⁴ *INFN-CNAF*

⁵ *INFN - National Institute for Nuclear Physics*

⁶ *Universita e INFN, Padova (IT)*

⁷ *Istituto Nazionale Fisica Nucleare (IT)*

⁸ *Universita e INFN, Perugia (IT)*

Corresponding Author: diego.ciangottini@cern.ch

The Italian National Institute for Nuclear Physics (INFN) has a long history of designing and implementing large-scale computing infrastructures and applications.

INFN has spent the past ten years heavily investing in developing solutions to enable, optimise and simplify transparent access to a multi-site federated Cloud infrastructure. A primary goal of this effort is to provide a generic model that allows INFN and other users to access resources in a fair and simple manner, regardless of the complexity of their requirements, of their proximity to a powerful computing centre, or their ability to administer advanced resources such as those offering GPUs. The ultimate objective is to shorten both the “time-to-market” and the learning curve for deploying, managing, and utilising computing services on a federated cloud system.

For this purpose, INFN Cloud provides a rich set of compute and storage services that can be automatically deployed on geographically distributed sites in an easy and transparent way.

One of the most frequently requested services by members of different scientific communities is based on jupyter notebooks. Therefore, we have been adapting the standard JupyterHub setup to provide a flexible and extensible multi-user service with some key integrations. First of all, the authentication mechanism is based on OpenID-Connect, while the authorization is based on OAuth attributes (like the user’s subject and groups) to grant admin or regular permissions. JupyterLab instances are spawned in containers which may start from custom images that encapsulate the needed libraries, depending on users’ needs (i.e. experiment software, big data analytics tools, etc.). All the containers mount two different types of storage space: a local area, where data is stored on the node filesystem; and a remote storage area, that allows to access the INFN Cloud Object Storage via posix. Files (notebooks, data, etc.) saved in the local storage area can persist until the node hosting the notebook servers is up and running, whereas data saved in the cloud area can be accessed at any time either through the notebook, or through the web interface of the INFN Cloud Object Storage service.

The usage of GPUs is also supported for running compute-intensive workloads. The automated configuration has also been tested with partitioned A100 GPUs: in this case, each notebook container gets an available partition of the GPU.

This contribution will provide details about the implementation of the service and some example use-cases running on INFN Cloud.

Collaborative Data Science and Visualisation / 91

VOIS library: pushing data science dashboards to the limits

Author: Davide De Marchi¹

Co-authors: Armin Burger¹; Franck Eyraud²; Pierre Soille¹

¹ *European Commission - Joint Research Centre*

² *Freelance consultant*

Corresponding Author: davide.de-marchi@ec.europa.eu

The Joint Research Centre (JRC) of the European Commission has set up the JRC Big Data Analytics Platform (BDAP) as a multi-petabyte scale infrastructure to enable EC researchers to process and analyse big data in support to EU policy needs [1]. One of the service layers of the platform is based on Jupyter notebooks and the Python programming language to enable exploratory visualization and interactive analysis of big geospatial and non-geospatial datasets [2]. In this context, we have gained a lot of expertise in the design, development and deploy in production of many complex Voilà dashboards [3] that enable JRC scientists and research groups to achieve a better communication of their scientific results and policy relevant insights to a non-technical audience as well as the public.

Although Voilà Jupyter plugin [4] automatically transforms a notebook into a dashboard, creating an impactful dashboard is still a hard task. Beyond the classical communication issues (regarding, for instance, the story to tell, the message to convey, the graphic elements to use), from a pure web-development point of view, the designers and developers have to clearly define the single-page or multi-page style of the application they want, how to position elements on the page and how to intercept user inputs. One can tackle all these aspects using the standard tools available in the Jupyter world, as for instance the ipywidgets library [5] that provides a basic set of input widget and that is widely used in the data science community. Nevertheless, when a dashboard intends to create a strong impact, we found that it needs to exploit more advanced UI (User-Interface) frameworks.

One open-source library recently gained a lot of interest for creating rich and engaging user experiences: ipyvuetify [6]. It is a widget library based on Vuetify/Vue JavaScript library [7] for making modern looking GUI's in Jupyter notebooks and dashboards. It implements the Google material design philosophy [8] best known from the Android user interface and provides a large set of widgets with multiple variants, all highly customizable. The usage of ipyvuetify is not easy at all given that, for creating non-trivial components, one needs to dig into the details of the Vuetify widgets and the JavaScript API syntax. For this reason, we started to develop a library with the aim to simplify the complex tasks involved in the creation of modern dashboards and to provide easy-to-use and reusable components: the VOIS library [9]. This pure Python library provides many ready-to-use widgets and exposes an "app" class that can serve as the base for creating the dashboard structure. With few lines of code, the "app" can be customised using styles, colours, fonts, images, and all the graphic elements that contribute to its unicity. The VOIS library has many functions for the easy creation of complex geospatial visualisations (like bi-variate and tri-variate choropleth maps for vector data, or fast display of multi-terabytes raster datasets). It contains several custom-made SVG interactive charts that allow for modern user interaction and widgets for the display of hierarchical and tabular data.

Among the functions of the VOIS library, we can cite: file uploads (to enable users to send their local input data to the web application); file downloads (to have a local copy of charts, tables, CSVs, reports in PDF or Word .docx format); management of parameters passed into the dashboard URL (in reading and writing mode); easy creation of dialog-boxes; support for responsive application development; etc.

This presentation will illustrate the concepts that are at the base of the development of the VOIS library, as well as demonstrate some of the main dashboards that we created using the library for impactful policy support.

It is interesting to note that the recently introduced VaaS service (Voilà as a service) enables BDAP users to autonomously create and deploy their dashboard in production, through an automated procedure based on GitLab repositories. This new service, together with the intensive training on the usage of the VOIS library, is contributing to the spreading of Voilà dashboards usage by many research groups in the JRC.

The VOIS library will soon be available as a fully open source project at <https://code.europa.eu/>, the main repository of OSS for the European Commission, with the aim to create a community of users and, possibly, of interested contributors.

These developments were partially funded by the H2020 project CS3MESH4EOSC, led by CERN and to which JRC participates providing support in the Earth Observation use case.

The JRC Big Data Analytics Platform is a living demonstration of a complex ecosystem of cloud applications and services that allows data scientists' navigation inside a multi-petabyte scale world. In particular, the exploratory visualization and interactive analysis tools and the Voilà/VaaS services are widely used to create GUI applications to communicate scientific research results to end-users ranging from policy makers to citizens.

- [1] P. Soille, A. Burger, D. De Marchi, P. Kempeneers, D. Rodriguez, V. Syrris, and V. Vasilev. “A Versatile Data-Intensive Computing Platform for Information Retrieval from Big Geospatial Data”. *Future Generation Computer Systems* 81.4 (Apr. 2018), pp. 30-40.
<https://doi.org/10.1016/j.future.2017.11.007>.
- [2] D. De Marchi, A. Burger, P. Kempeneers, and P. Soille. “Interactive visualisation and analysis of geospatial data with Jupyter”. In: *Proc. of the BiDS'17*. 2017, pp. 71-74.
<https://zenodo.org/record/3248741>
- [3] <https://jeodpp.jrc.ec.europa.eu/bdap/voila/>
- [4] <https://github.com/voila-dashboards/voila>
- [5] <https://ipywidgets.readthedocs.io/en/latest/>
- [6] <https://ipyvuetify.readthedocs.io/en/latest/>
- [7] <https://vuetifyjs.com/en/>
- [8] <https://m2.material.io/>
- [9] https://jeodpp.jrc.ec.europa.eu/services/shared/vois/1_intro.html

EFSS Products / 85

Seafile, what's new in the year 2022

Author: Jonathan Xu^{None}

Corresponding Author: xjqkilling@gmail.com

Seafile is a popular open-source file sync and share solution. The focus of Seafile is reliability, security and performance. It's widely used by many large educational institutes in Europe.

This talk will review the new features and improvements that we have made to Seafile in 2022. Topics include:

1. Performance improvements to the server
2. Improved user experience
3. New drive client for macOS
4. Progress on OCM support

EFSS Products / 66

Nextcloud. State of the nation

Author: Frank Karlitschek^{None}

Corresponding Author: frank.karlitschek@nextcloud.com

This talk will give an overview of the Nextcloud developments and improvements in the last 12 month. Several noteworthy things happened in the last Nextcloud releases. From architectural improvements to changes on APIs and the sync engine, to useability and functionality. This Talk will give a full overview.

EFSS Products / 103

ownCloud Infinite Scale

Author: Holger Dyroff^{None}

Co-author: Klaas Freitag

Corresponding Author: hd@owncloud.com

The first supported version of ownCloud Infinite Scale was released in November 2022 and was received very well in the community and by the first customers of ownCloud.

This presentation gives a brief overview about the status of the stable and released product, before we start to look at the current “hot” new features and developments. That will include features and new APIs, deployment models and possibly new integrations as examples.

An outlook on the roadmap as foreseeable will round the presentation up.

Site Reports / 108

Site Report Summary

Corresponding Author: steiger@id.ethz.ch

CS3 community summary

Site Reports / 98

Standing innovation for 10 years: CERNBox

Author: Hugo Gonzalez Labrador¹

¹ CERN

Corresponding Author: hugo.gonzalez.labrador@cern.ch

CERNBox is key enabler service for users at CERN and beyond. The service is used by more than 37K users and stores over 15PB of data, representing all the user communities at the laboratory.

In this talk we will explain the current status of the service, the challenges we faced in 2022 and we look into the future: CERNBox as the gateway for heterogeneous storage spaces at CERN and beyond.

Site Reports / 67

Sync & Share self service provisioning on INFN Cloud

Authors: Ahmad Alkhansa^{None}; Alessandro Costantini¹; Alessandro Italiano^{None}; Barbara Martelli^{None}; Carmelo Pellegrino^{None}; DIEGO MICHELOTTO²; Daniele Cesini³; Daniele Spiga^{None}; Davide Salomoni³; Diego Ciangottini⁴; Doina Cristina Duma⁵; Emidio Maria Giorgio⁶; Federica Fanzago⁷; Francesco Sinisi^{None}; Giacinto Donvito⁸; Gianluca Peco³; Luca Giovanni Carbone^{None}; Lucio Strizzolo²; Marco Verlati⁷; Marica Antonacci⁵; Massimo Sgaravatto⁷; Nadina Foggetti^{None}; Sergio Traldi^{None}; Stefano Stalio^{None}; Vincenzo Ciaschini^{None}; Vincenzo Spinoso⁵

¹ INFN-CNAF

² INFN - National Institute for Nuclear Physics

³ Universita e INFN, Bologna (IT)

⁴ INFN, Perugia (IT)

⁵ INFN

⁶ INFN LNS

⁷ Università e INFN, Padova (IT)

⁸ INFN-Bari

Corresponding Author: stefano.stalio@lngs.infn.it

Following up on 20 years of successful development and operation of the largest Italian research e-infrastructure through the Grid, the Italian National Institute for Nuclear Physics (INFN) has been running for the past three years INFN Cloud, a production-level, integrated and comprehensive cloud-based set of solutions, delivered through distributed and federated infrastructures.

INFN Cloud provides a large and customizable set of services, ranging from simple IaaS to specialized SaaS solutions, centered through a PaaS layer built upon flexible authentication and authorization services offered via INDIGO-IAM, and optimized resources and services orchestration.

Sync & Share instances based on ownCloud or NextCloud are among the several applications and services that users can self-deploy via the INFN Cloud dashboard.

Besides giving a general overview of INFN Cloud and its federated model, this talk will describe how the deployment of Sync & Share services for small to medium sized communities is fully automated. We will show how added-value features are integrated: a geographically distributed S3 storage backend, automated database and configuration backup, dedicated resource monitoring, secure connections and centralized authentication/authorization. We will also describe how INFN Cloud may provide a dedicated solution for supporting sensitive data privacy that exploits end-user level encryption of the storage block devices, used by ownCloud or NextCloud to store user data.

Site Reports / 77

Research Drive, a platform for active research data management

Author: Narges Zarrabi^{None}

Corresponding Author: narges.zarrabi@surf.nl

One of the challenges of research institutes is finding a proper tooling and platforms to practice and support research data management. SURF has more than 100 institutes in different sectors and has various user types from researchers and data stewards to librarians and research supporters. This translates to a wide range of requirements for research data management in our user community. To address those requirements, we have developed a national service called Research Drive, which accommodates different roles mapping to the roles and responsibilities involved in research and research support within an institute.

A regular user can store data, have access to sync and share functionalities, and collaborate on data as primary needs of researchers. An owner of a project folder is a regular user with extra privileges on project level to make project folder structure and invite and give permissions to others to use the data. To accommodate power users, i.e., data stewards, research supporters or IT staff at institutes, we have developed integrated apps such as the Dashboard app for managing users, projects and storage quotas and the Reporting app for an overview of all shares, permissions and other service-related information such as service availability. We setup Research Drive as dedicated instances per institute which makes customizations possible. The RDSettings app enables the power users to enforce policies institute wide, such as enforcing policies for sharing, passwords, and multi-factor authentication (MFA).

We are continuously working to improve the service performance and user experience and address the feature requests of our institutes. To improve the performance we are going to replace the storage backend of Research Drive from Object Store to Ceph. We are going to use Keycloak as Identity and Access Management provider. This will be run on a kubernetes cluster and integrated with LDAP to support OpenIDconnect and eduID. One of the main feature requests we have is the need for archiving and publishing data from Research Drive. For this we are integrating the Sciebo RDS app to facilitate data archiving and publication workflows. A strict auditing feature is also being

implemented to be able to closely monitor activities. This feature is of interest for hosting sensitive data and can be configured per project bases. In this talk we present the latest developments and plans for the Research Drive service.

Site Reports / 73

Sunet Drive - An Academic EFSS Packaged for EOSC

Authors: Gabriel Paues¹; Richard Freitag²

Co-author: Guido Aben²

¹ *Safespring*

² *SUNET*

Corresponding Authors: gabriel.paues@safespring.com, freitag@sunet.se

Sunet Drive is a federated and scalable Enterprise File Sync and Share solution, that has been developed, deployed, and packaged as part of the European Open Science Cloud and can be transparently extended to new participating organizations. The two main building blocks of Sunet Drive are Nodes and Buckets, both elements designed to promote data sovereignty and FAIR principles. Participating organizations co-manage their Sunet Drive node as part of a global scale setup, meaning that every node is governed by the operating organization, while being able to collaborate and share data with users within the federation, but also external partners that support the open cloud mesh protocol (OCM), such as the ScienceMesh. New organizations have been and can be onboarded by migrating existing provisioned users to a full node associated to an organization or institution. Buckets, specifically S3-compatible buckets, are used as logical storage entities that can be assigned for different purposes: research projects, institutions, laboratories. They are technically independent from the EFSS layer and their lifecycle can therefore be managed beyond the lifetime of the selected EFSS software, an important step towards long-term sustainability for FAIR handling of data.

The infrastructure stack is implemented in collaboration with the commercial actor Safespring, and data generally resides in at least two different data centers. This ensures a scalable stack built on best practice open source components, together with experience in running large scale deployments. Certification standards such as ISO 27001 guarantee a mature handling of the infrastructure and data in the solution.

By having chosen a state-of-the-art EFSS solution for Sunet Drive, researchers, scientists, and their collaborators can align the requirements of their funding body and associated data management plans with their primary data sources by using modern like synchronization clients and mobile applications. On the other hand, integrated and connected services ensure that scientists will be able to collaboratively work on their projects without having to leave the ecosystem.

Collaboration is encouraged by allowing any eduGAIN connected identity provider to provision user accounts, and subsequently accept documents, shares, and data from their collaboration partners. The lack of support for a discovery service on the EFSS side has been solved by using a global site selector through a SaToSa proxy that delegates users to their respective Sunet Drive Node. External collaboration is enabled via Euid.se.

During the runtime of a research project, research data can be curated and prepared for publication. The integration of Research Data Services, RDS, enables the preparation and publication of datasets directly from the EFSS solution. This includes external services like InvenioRDM (e.g., Zenodo), Harvard Dataverse, or Doris from the Swedish National Dataservice, SND. This includes domain-specific customizations. While research object crates (RO-Crate) are used as an intermediate lightweight package for the data, and respective metadata, connectors will ensure compliance with each publication service. While data is being actively pushed to InvenioRDM, the Doris connector uses a more lightweight approach where the metadata is pushed to Doris, while the data storage remains under the sovereignty of the publishing institution.

Having data stored in S3-compatible buckets associated to a federated EFSS node managed by specific organizations ensures data sovereignty and helps to ensure compliance with local, national, and

international guidelines for storing of research data, including FAIR principles. After a project has finished, ownership of the data can be transferred during the data retention period, and for long-term archival purposes.

Site Reports / 94

Sciebo Site Report: Challenges in supporting more than just the core project

Author: Marcel Wunderlich^{None}

Co-author: Holger Angenent¹

¹ *University of Münster*

Corresponding Author: wunderlich@uni-muenster.de

Toil is the enemy of any admin and SRE.

Automation of our processes has made operations quite smooth, allowing us to spend more time in supporting and coevolving some long awaited projects (overleaf, scuot_ng, rds...)

The technical side of things almost seems trivial compared to keeping an overview of the todos across ticket systems or resolving dependencies of people waiting for each other.

Restructuring workflows is challenging, but can really help removing unnecessary roadblocks.

Automating project management in gitlab and github allows to have many smaller repos, which reduces merge backlogs and allows for a finer granularity of code ownership.

Sometimes some vertical integration is needed to further automate, which brings its own challenges.

Interoperable Cloud Infrastructure Stacks / 112

The GÉANT Special Interest Group on Cloud Interoperable Software Stacks

Author: Mario Reale^{None}

Corresponding Author: mario.reale@geant.org

This short (5 minutes) presentation will summarise the goals and the activities of the Special Interest Group on Cloudy Interoperable Software Stacks of the GÉANT Community Programme. The SIG-CISS will meet on Wednesday, March 8, in the afternoon, co-located with the CS3 Conference in Barcelona.

Interoperable Cloud Infrastructure Stacks / 113

Tripleo: exploit ansible to customize both the Undercloud and the Overcloud

Authors: Andrea Dell'Amico¹; Franca Debole¹

¹ *CNR-ISTI*

Corresponding Author: andrea.dellamico+cern@isti.cnr.it

TripleO, <https://docs.openstack.org/tripleo-docs/latest/>, is a set of tools for the deployment and management of OpenStack. Its strategy consists in using a underlying OpenStack installation (undercloud) to install and manage the main one (overcloud).

It's the installation method used by RDO, <https://www.rdoproject.org/>.

In our project to deploy a HyperConverged (HCI) OpenStack cloud infrastructure, we wrote some ansible roles to prepare the undercloud and setup the templates required to configure the overcloud. We also added some additional steps, to fully support OIDC authentication in keystone and to add external prometheus metrics to the physical nodes.

Interoperable Cloud Infrastructure Stacks / 107

Sovereign Cloud Stack - A common Open-Stack for digital sovereignty

Author: Christian Schmitz^{None}

Corresponding Author: cschmitz@wwjkd.net

The talk will provide an overview of the Sovereign Cloud Stack, an open-stack release partially funded by the German government and German open-source companies, enabling governments, organizations and companies to deploy and manage their own public/private/hybrid clouds based on a common stack. This stack is designed to address the growing concern of data sovereignty, which is the ability to control and protect data within a specific geographic region and software supply chains. The Sovereign Cloud Stack provides a set of tools for creating and managing virtual infrastructure, as well as for deploying and managing applications. The talk will also discuss the key features of the stack, such as its support for advanced deployments, its ability to handle high availability, and its integration with popular open-source tools. Additionally, the talk will provide a demonstration of how to set up and use the stack, and will also cover best practices for using it in a production environment, for example in combination with Kubernetes.

Interoperable Cloud Infrastructure Stacks / 111

LIQO – Building and Orchestrating your Service Continuum

Author: Fulvio Giovanni Risso¹

¹ *Politecnico di Torino*

Corresponding Author: alessandro.olivero@polito.it

LIQO (<https://liqo.io>) is an open-source multi-cluster orchestrator that enables the creation of “virtual Kubernetes clusters” spanning across an arbitrary number of real clusters, even crossing multiple administrative boundaries.

Liqo enables the sharing of resources (e.g., CPU, memory, GPUs) and services (e.g., an existing cloud-native service) among different clusters, and facilitates the portability of workloads among different cloud providers, hence reducing the lock-in of the infrastructure provider.

This talk presents the idea behind liqo, it summarizes its main architectural pillars, and it shows how this framework can be used to facilitate the creation of infrastructure-independent cloud native services.

Education and Research Use-cases / 106**Learning Management System based on sync & share storage at ETH Zürich**

Author: Tilo Steiger^{None}

Co-authors: Gianluca Caratsch ; Melanie Walter ; Tony Moser

Corresponding Author: steiger@ethz.ch

Today requirements in teaching, learning and ultimately also examination at universities make more and are more digital alignments and resilient Learning IT Management Systems necessary.

Within this contribution we want to show the system components and their interaction. We will show what added value the use of sync & share storage provides.

Education and Research Use-cases / 105**ownCloud Infinite Scale: New scenarios for Research and Higher Education**

Author: Reinhard Schüller^{None}

Corresponding Author: rschueller@owncloud.com

In this talk we'd like to share recent experiences with customer deployments.

The product oCIS has a new architecture that allows new approaches to practical requirements of data management and its corresponding processes.

We'd discuss customer scenarios with their key challenges. This will include a setup with a few 1000 tenants to share some specific administrative and provisioning processes. The architecture and deployment will be discussed.

Keynote / 115**Stop data sharing.**

Corresponding Author: barend.mons@go-fair.org

FAIR Data Management / 71**Sciebo RDS - reducing friction of FAIR data handling for researchers**

Authors: Juri Hößelbarth¹; Lennart Hofeditz²; Richard Freitag³

¹ *University of Münster*

² *University of Duisburg-Essen*

³ *SUNET*

Corresponding Authors: hoesselbarth@uni-muenster.de, freitag@sunet.se

Research Data Services (RDS) is a self-hosted cross-platform interoperability layer which allows research data to be curated, prepared and published directly from an EFSS solution such as Sciebo (own-Cloud) or Sunet Drive (Nextcloud). It provides modular interoperability to external data repositories like the Open Science Framework (OSF), InvenioRDM (e.g., Zenodo), Harvard Dataverse, or Doris from the Swedish National Dataservice, each including domain-specific customizations.

Publishing data sets, corresponding metadata and persisting the information with a digital object identifier (DOI) is not only increasingly required by project funding entities such as the European Union or scientific journals, but also positively contributes to researchers' credibility and visibility. However, publishing research data often requires a specific data repository. This challenge is addressed by using Sciebo Research Data Services (RDS) as an interface between the enterprise file sync and share solution (EFSS) and the data repository. While research object crates (RO-Crate) are used as an intermediate lightweight package for the data and respective metadata, individual connector microservices ensure compliance with each publication service and easily developed support for additional data repositories. Metadata is added through Describo and based on schema.org annotations in JSON-LD and aims to make best-practice in formal metadata description accessible and practical for use in a wider variety of situations, from an individual researcher working with a folder of data, to large data-intensive computational research environments.

An integral part of the research data lifecycle process is the publication of research data, either by itself or as part of the supplementary information (SI) of a published article. Contributing to the Open Science movement, research data should be FAIR: findable, accessible, interoperable, and reusable. However, not only researchers' fears (e.g., fear of misuse of the data or fear of errors in the data or analysis) but also technical barriers (such as a lack of functionality or too long of a tool chain) are causing friction, often stopping researchers from adequately managing and sharing their research data.

Originally funded by the Deutsche Forschungsgemeinschaft (DFG) and implemented by WWU Münster in collaboration with the University of Duisburg-Essen, Sciebo RDS has sustainably evolved into a cross-platform solution that is used by several institutions and NRENs, notably SURF and Sunet. Its EFSS application has been ported from ownCloud to Nextcloud, thus also enabling cross-platform interoperability with the ScienceMesh.

To ensure that Sciebo RDS offers the highest possible benefit and ease of use for researchers, the development is accompanied by an iterative scientific evaluation process. This involves an extensive qualitative requirements analysis with researchers from various disciplines and qualitative as well as quantitative usability studies at different phases of prototyping.

To use Sciebo RDS, researchers can log on to their respective EFSS system and will find Sciebo RDS directly in the EFSS main menu. After connecting Sciebo RDS to appropriate data repositories (e.g. OSF or Zenodo), users are guided through a four-step data publication process, including:

1. the configuration of a research data project,
2. the collection and management of the data,
3. adding and editing the metadata, and
4. the transfer and publication of the data to external services.

Sciebo RDS integrates into academic enterprise file sync and share solutions and lets researchers collect data, collaborate on documents, and publish valuable scientific results directly from one simple solution.

Keeper: certifying research data combining Seafile and blockchain technology

Author: René Ranger¹

Co-author: Vladislav Makarenko ¹

¹ *Max Planck Digital Library*

Corresponding Author: ranger@mpdl.mpg.de

The Max Planck Society runs a customized installation of Seafile called KEEPER (<https://keeper.mpdl.mpg.de/>) for its scientists which offers the possibility to certify research data with or without metadata by leveraging on blockchain technology. Snapshot data and a certificate representing the data on the blockchain are stored on application side and presented to the user.

This implementation is tailored to our scientists and stores transaction data via a publicly available API on bloxberg, the blockchain run by a association of worldwide research institutions (<https://bloxberg.org/>). We would like to showcase the implementation and its potential benefits for the CS3 community.

Federated Infrastructures & Clouds / 75

HIFIS & Nextcloud: VO Federation for EFSS

Authors: Andreas Klotz^{None}; Björn Schießle¹

¹ *Nextcloud GmbH*

Corresponding Authors: andreas.klotz@helmholtz-berlin.de, bjoern@nextcloud.com

After two years of planning for Virtual Organisations (VO; Community AAI based group of any size) as the basis for a new kind of EFSS Federation [2,3] by HIFIS in coordination with the CS3 community, the development of this new feature for the Nextcloud software has been completed, thanks to the strong support of Nextcloud and their subcontractor publicplan.

Admins of Nextcloud instances now have the option of joining a federation in the context of any Community AAI which supports the necessary interfaces [4]. This will enable users to share with VOs which they are a member of within these Community AAIs, while every other VO member who has connected their Community AAI account will receive the share, no matter on which Nextcloud instance within the federation they work on.

This first approach on federated group sharing builds upon the OCM protocol [5], adding a new share type to it. We want to propose this new extension of OCM to the CS3 and OCM Community and discuss it in depth, to achieve a level of standardisation which will allow for further software vendors to implement similar features and for the research community to span VO based federations across providers of heterogeneous EFSS solutions.

1 AARC Blueprint for Community AAIs: <https://aarc-project.eu/architecture/>

[2] CS3 contribution 2021: <https://indico.cern.ch/event/970232/contributions/4157924/>

[3] CS3 contribution 2022: <https://indico.cern.ch/event/1075584/contributions/4658939/>

[4] SCIM rfc: <https://www.rfc-editor.org/rfc/rfc7644.html>

[5] OCM protocol: <https://wiki.geant.org/display/OCM/Open+Cloud+Mesh>

Federated Infrastructures & Clouds / 81

Cubbit Next Generation Cloud

Author: Stefano Baldi¹

Co-author: Alessandro Cillario ²

¹ *Cubbit srl*

² *Cubbit*

Corresponding Author: alessio.paccoia@cubbit.io

In the last year Cubbit has delivered many Cubbit cells in Italy. A cubbit cell is a very simple device that provides encrypted block storage service. These cubbit cells connect to each other from different datacenters of different Italian companies. Each cell relies on a different data link and is powered by a different power line. Even cubbit cells hosted in a specific company do not contain data blocks of that company but mainly data blocks of others and vice versa, data blocks of a specific company are mostly stored in cubbit cells hosted elsewhere.

This design combines hardware redundancy and hyper-distribution of risk and thus ensures that data stored in a Cubbit swarm (i.e. the p2p network made up of all Cubbit cells) is strongly defended against accidental and catastrophic loss, at a fraction of the cost that companies typically spend for that level of service.

But this, good as it is, is just the basic idea of Cubbit. To convince a BTB customer to join the Cubbit swarm, Cubbit had to change some of the basic assumptions that made up its initial proposal to the customer: today Cubbit is no longer a Sync and Share solution, but an S3 compatible Object Storage service. This means that the Customer can continue to store data via the Cubbit web interface and can continue to share it with other Cubbit users, but can also activate the versioning service on the files or prevent them from being modified for a certain period of time. And it can do much more: it can connect Cubbit to a third-party solution designed for the S3 protocol. Data uploaded from widely used clients (Cyberduck, Nextcloud, CloudBerry, Veeam...) are visible to each other and also from the Cubbit web interface, giving the client complete control over the data flow and a variety of data cases . use, which was previously not possible. By adopting a de facto industry standard such as the AWS S3 protocol, Cubbit can synchronize its cloud object storage with third-party cloud object storage and allows the customer, for example, to synchronize data from Azure to Cubbit with minimal effort .

Adding more protocols (as we did with S3) will multiply the use cases, but a key step in the integration will be making the Cubbit cell available in the form of a virtual machine or container, allowing the customer to implement their own portion of a cubbit swarm in a virtualized data center or even create a private cubbit swarm from scratch, and deploy it in a private wide area network.

Federated Infrastructures & Clouds / 78

AAI-powered Federated Groups for SURF Research Drive

Author: Michiel de Jong^{None}

Corresponding Author: michiel@unhosted.org

For SURF (the Dutch NREN), Ponder Source are proud to build a connection between ResearchDrive (based on ownCloud) and SRAM (their AAI solution).

As part of this process, we also implement Federated Groups in OC-10, and improve OC-10's Open Cloud Mesh implementation to allow not only OCM-sharing to a user, but also to a (local) group.

In this presentation we will briefly present the design issues we encountered, demo the results, and discuss the exciting similarities and synergies with the related Federated Groups ("Virtual Organizations") project of Nextcloud / Helmholtz, and we would like to discuss the implications for the Open Cloud Mesh protocol.

Federated Infrastructures & Clouds / 109

ScienceMesh: an interoperable federation of EFSS services

Corresponding Authors: pedro.ferreira@cern.ch, jakub.moscicki@cern.ch

ScienceMesh is an interoperable research platform developed for the European Open Science Cloud (EOSC), in the context of the CS3MESH4EOSC project.

It is designed to be an interoperable research platform for seamless sharing and collaboration on data across different EFSS systems, including major open-source platforms such as Owncloud, Nextcloud and others.

ScienceMesh builds on the common experience and contributions from the CS3 community, ranging from interoperable protocols and APIs (OCM, CS3APIs) to service integration for collaborative research: access to large scale storage, Jupyter environments, digital repositories and files transfers.

During this presentation, we will provide an update on the progress of the project and outline our plans for 2023 and beyond.

For more information visit sciencemesh.io.

Discussion / 80

OCM Test Suite current status

Author: Michiel de Jong^{None}

Corresponding Author: michiel@unhosted.org

This short presentation will give a brief overview of the current status of the OCM test suite, and present our latest knowledge of the compatibility matrix, showing how various EFSS systems (now also including Reva) can act as an OCM sender or an OCM receiver (with / without the new invite flow that is used on the ScienceMesh).

Discussion / 104

A View on the CS3 APIs

Author: Klaas Freitag^{None}

Corresponding Author: kfreitag@owncloud.com

The CS3 APIs are the most important technical building block of the CS3 community. ownCloud Infinite Scale is implementing these APIs and has added some new functions to continue to drive innovation.

In this talk we will present some of the most important changes to the CS3 APIs that come with the so called “edge” branch, and how that is beneficial to the community.

Furthermore, the CS3 API protocol GRPC has turned out to be very valuable for internal server communication, but has challenges for communication with clients. For them, HTTP based APIs have proven to be more suitable.

oCIS is using some new HTTP based APIs that will be explained. There will be a fly over the new functionality and some opinions on how the APIs could be added to the next version of the CS3 APIs.

Discussion / 116

OCM and CS3 discussion

ScienceMesh workshop / 117

ScienceMesh Demos Introduction

Corresponding Author: pedro.ferreira@cern.ch

ScienceMesh workshop / 79

Connecting Nextcloud & OC-10 to the ScienceMesh

Author: Michiel de Jong^{None}

Corresponding Author: michiel@unhosted.org

We built a two-way connection between Nextcloud / OC-10 and Reva, which is deployed at many of the sites that are currently connected to the ScienceMesh testnet.

In this short presentation we'll explain how the connection between the EFSS GUI and Reva works in different scenarios.

ScienceMesh workshop / 118

Invitation Workflow and Federated Sharing: NC, OC10

Corresponding Author: milan.danecek@cesnet.cz

ScienceMesh workshop / 90

User-friendly OCM Invitation and Sharing: OCIS, CERNBox

Authors: Elizaveta Ragozina¹; Gianmaria Del Monte¹; Samuel Alfageme Sainz¹

¹ CERN

Corresponding Author: elizaveta.ragozina@cern.ch

The Open Cloud Mesh (OCM) protocol has been adopted as part of the CS3MESH4EOSC project to enable interactive and agile collaboration across various file synchronization and sharing providers at a pan-European level.

We present CERNBox production service that enables the management of invitations and shares via OCM Invite and Share APIs and integrates with ownCloud and nextCloud. After the initial step of user discovery with the "ScienceMesh" extension, the inviters are supported by the UI in the seamless sharing workflow while the invitees get an overview of shared with them contents. The prototype aims the ease and intuitiveness of use, as well as the consistency with the existing ownCloud based UI.

ScienceMesh workshop / 119

Data Science Environments: JupyterLab sharing and collaborative editing

Corresponding Author: marcin.sieprawski@softwaremind.com

ScienceMesh workshop / 120

Data Transfers: Connecting Science Mesh and ESCAPE Data Lakes

Corresponding Author: ron.trompert@surf.nl

Collaboration Products / 86

Applications integration beyond local clouds with OCM

Authors: Gianmaria Del Monte¹; Giuseppe Lo Presti¹; Michiel De Jong²

¹ CERN

² Pondersource

Corresponding Author: giuseppe.lopresti@cern.ch

The Open Cloud Mesh (OCM) protocol has been adopted as part of the ScienceMesh infrastructure to enable interactive and agile collaboration across various file synchronization and sharing providers at a pan-European level.

With such infrastructure emerging as the collaboration space across institutions, an important added value is the ability to open document files and work with other collaborators, by exploiting concurrent editing capabilities of common online editors such as Collabora and CodiMD.

In this presentation, we demonstrate the work done towards integrating applications on top of OCM-shared files, both in terms of OCM extensions and web UI, and how this enables organizations that are part of ScienceMesh to share their application engines. We discuss licensing issues related to such cross-sites use cases, and we conclude with future prospects in the landscape of collaborative apps for cloud storage systems.

Collaboration Products / 70

OFORMs for document automation and collaboration

Author: Mikhail Korotaev^{None}

Co-author: Oleksii Ivanov

Corresponding Author: info@onlyoffice.com

Aimed at improving the automation of document creation and co-authoring tasks, OFORMs allow building, exporting, and sharing fillable forms for standardized paperwork. The technology behind combines flexibility of working with fields, properties, and protection tools with complete editing, formatting, and collaboration instruments of ONLYOFFICE Docs.

ONLYOFFICE continues expanding the functionality of form creator and introducing new methods of OFORM application – integration with document and content management systems, implementation in web interfaces, and use in native ONLYOFFICE Workspace and ONLYOFFICE DocSpace environments.

With introduction of encryption mechanisms and recipient roles in forms, it is possible to incorporate personalized digital signature requests for select fields and assume the use of OFORMs in more complex legal scenarios.

This presentation will include:

- Overview of ONLYOFFICE forms (OFORMs);
- Models of implementation: file sharing environments, web applications, ONLYOFFICE DocSpace;
- New form fields and field parameters;
- Recipient roles in form filling;
- ONLYOFFICE Docs 2022 overview;
- 2023 roadmap.

Collaboration Products / 68

Collabora Online: Better, Digitally Sovereign Document Collaboration

Author: Michael Meeks^{None}

Corresponding Author: michael.meeks@collabora.com

Come and see how Collabora Online (COOL) can be a pleasure to deploy and integrate into your File Sync & Share or LMS provision. Hear about the work we've done in the last year to make it even better. From dynamic load balancing in Kubernetes, to accelerated compression of tiles reducing both CPU and bandwidth use and improving interactivity.

Admins should know about our new Grammar checking server integration with LanguageTool & DudenCorrector, as well as new easy font management APIs & tooling to improve interoperability.

Adding lots of core features to the underlying LibreOffice technology we have brought SparkLines, writer table change-tracking, new form building content-controls support, efficient 16k column spreadsheets, chart data tables, color theming and much more to COOL.

We've also improved accessibility and tooling for the impaired adding a document accessibility checker, and PDF & EPUB export as well as PDF form creation.

Come and see how we've made things beautiful from UX improvement & polish for users, to expanded Prometheus metrics for admins.

See how we can deliver scalable, secure, on-premise editing of your documents with a simple, easy to deploy office for the free world.

Collaboration Products / 83

Status Update of the no-code platform SeaTable

Author: Christoph Dyllick-Brenzinger^{None}

Corresponding Author: cdb@datamate.org

SeaTable is like a lego kit for IT. It enables you to develop and build efficient business process in the shortest possible time. You can easily design your database structure, store any kind of data, define access rights for your team or externals and visualize your data with various charts. Automations help to streamline your work.

In this presentation, I will give an overview of the improvements that happened in SeaTable in the last year:

- **Common datasets** allow synchronization of data between bases. This is useful if you want to use parts of a central user database in other bases, for example.
- **External App** is an interface designer to create individual views for the different stakeholders.
- A variety of new **automations** and **visual evaluations** have been added

Scalable Storage Backends / 92

An unhappy end for S3

Author: Tom Wezepoel¹

¹ SURF

Corresponding Author: tom.wezepoel@surf.nl

Research Drive, the Dutch Sync & Share service based on ownCloud, uses OpenStack Swift S3 as its storage backend. Where the integration of S3 within the software is not that good, we will migrate back to a posix compliant file system, namely CephFS. But how to migrate almost 2 PB of data without too much downtime...

Scalable Storage Backends / 82

From 1VM+1LUN to k8s+Ceph - the uneasy way...

Authors: Adam Prycki¹; Eugeniusz Pokora¹; Filip Blicharczyk¹; Jacek Kochan¹; Jan Brózdziak¹; Krzysztof Dudek¹; Krzysztof Wadówka¹; Maciej Brzeźniak¹; Mateusz Piątkowski¹; Norbert Meyer¹; Rafał Mikołajczak¹; Sebastian Petruczynik¹

¹ PSNC

Corresponding Author: kwadowka@man.poznan.pl

From 1VM+LUN to k8s+S3 - an uneasy way...

Since 2015 PSNC has provided a sync & share service for science and academia in Poland, based on the Seafile software. We started small by running a 1VM+1LUN setup and community version of the software, integrated with the local PSNC's LDAP. In 2016 we began to build a fully-fledged setup based on a cluster of application servers, background jobs servers, DB servers, and a dedicated 2-servers GPFS cluster as a storage backend, that become operational in 2017 and serves most of our users until today. We are currently migrating our service towards the most modern and fancy setup based on k8s and Ceph/S3.

In our presentation, we discuss the experiences and observations on the impact that changing cloud and storage technologies and infrastructure compute/storage infrastructures' features have on your system, services, data, and users, while you are trying to follow the trends in the system architectures, services deployment approaches, management practices, etc. We also discuss the pros and cons of

the simplified 1-VM setup vs the bare-metal multiserver infrastructure vs the fully containerized setup with lots of automation.

Surprisingly (or not) we faced the unobvious and uneasy-to-accept fact, that the ‘ancient’ simplistic setup of our sync & share system required the least effort to keep it up, and caused almost no operational issues, faults & failures over 8+ years of operation, while the complexity of the management processes such as data and users migration, systems and application upgrades grows non-linearly with growing the level of ‘fanciness’ and ‘intelligence’ of the infrastructure and application setup. Obviously, it would not be fair to say that the capabilities of a 1VM+1LUN platform intended to serve ~500 users with <40TB data and ~8 million files are comparable to a fully-fledged clustered setup for 1000s of users with ~1PB storage and quarter billion files.

Therefore we will explain the reasoning behind the design decisions made since the start of the minimalistic service, where the HA features were based on a rock-solid 1VM + hypervisor + orchestration platform :), through its extension to a full bare metal setup comprising almost a whole rack of servers and disk arrays junk :), until the cutting-edge setup based on top-down software-defined infrastructure including k8s-fuelled containers, software-defined storage back-end (Ceph with S3 gateways) and SDN user for network mgmt.

We will also discuss the impact that particular decisions had on the complexity of the system management with a special focus on the sync & share application and underlying operating systems and platforms (DB engines) upgrades. For this purpose we will provide a deep dive into the process of upgrading Seafile 7.x to Seafile 9 along with the required operating systems updates and users’ and users’ data migration from GPFS-based POSIX-speaking backend to a Ceph-based S3/object storage backend.

We will also overview our efforts on the preparation of a fully containerized setup allowing us to deploy, manage and maintain an arbitrary number of testing, development, and production Seafile instances, ensuring full coverage of system manageability, high availability, scalability, high performance, and data security.

Scalable Storage Backends / 102

Comparison between CephFS, CFFS (Comtrade FastFS), HDFS (Apache Hadoop), GPFS (IBM Spectrum Scale), Lustre

Author: Gregor Molan¹

¹ *Comtrade 360's AI Lab*

Corresponding Author: gregor.molan@cern.ch

Different high-performance, high-available file systems can store big data (hundreds of PB) and provide high data throughput (hundreds of TB per second). Each of these solutions highlights its advantages, and it is challenging to compare them.

Based on 30 years of storage development experience, Comtrade provided test scenarios to compare these file systems. On the appropriate high-performance hardware, here are the results from January 2023 that helps companies to choose appropriate high-performance file systems that fulfil given requirements. The limitation of this comparison is limited only to performance, while other essential features, such as software maintenance and upgradeability, are not covered in this comparison.

Scalable Storage Backends / 89

C(ERN) BACK(UP): consolidated multi-petabyte backup solution for heterogenous storage and filesystems

Author: Gianmaria Del Monte¹

Co-authors: Elizaveta Ragozina¹; Roberto Valverde Cameselle¹

¹ *CERN*

Corresponding Author: gianmaria.del.monte@cern.ch

The IT storage group at CERN is responsible to ensure integrity and security of all the stored data for physics and general computing services. In the last years a backup orchestrator, cback, has been developed based on the open source backup software restic. Cback is able to backup EOS, CephFS and any local mountable file system, like NFS or DFS. cback is currently used to daily backup CERNBox data (2.5 billion of files and 18PB), including experiment project spaces and user home directories, CephFS Manila shares, the CVMFS home folders and the CERN gitlab instance.

The data copy is stored in a disk-based S3 cluster in another geographical location in the CERN campus 4km away from the main data center (protecting against natural disasters). The usage of restic allows us to reduce the storage costs thanks to the deduplication of the data. In the last months, the cback portal server has been implemented, exposing a set of REST APIs to allow the integration with end-user backup utilities to navigate snapshots and restore data.

In this presentation, we will describe the architecture and the implementation of cback, the integration with CERN services and the future integration with tape archive for long term data preservation.

Scalable Storage Backends / 96

Data at scale: from Storage to Data Governance

Author: Jean-Thomas Acquaviva¹

¹ *DDN Storage*

Corresponding Author: jtacquaviva@ddn.com

Data are said to live forever, however their life is a complex journey. Initiated at acquisition or production date, data start a whole life cycle. During the different epochs of this life cycle, data will be moved, processed, compressed, shipped, archived.

To ease the management of this data orchestration, modern storage systems provide powerful tools. The foundation of these tools remains the ability to describes data with metadata

Metadata can be simple file information (date, size, format) or more complex, defining a structure including discipline-specific schema (or ontologies) used to address specific elements needed by a discipline.

In this talk we will present the layered approach of file systems, notably Lustre, to help end-users to implement a data governance solution.

Scalable Storage Backends / 100

Protocol Plumbing: Presenting iRODS as WebDAV, FUSE, REST, NFS, SFTP, K8s CSI, and S3

Authors: Kory Draughn^{None}; Terrell Russell^{None}

Corresponding Author: korydraughn@renci.org

The open source iRODS (Integrated Rule-Oriented Data System) data management platform presents a virtual filesystem, metadata catalog, and policy engine designed to give organizations maximum control and flexibility over their data management practices and enforcement. Since iRODS defines its own RPC API and protocol, interoperability with other software has always lagged new server features and functionality. In the last few years, the iRODS Community has been working on multiple ways to present the iRODS protocol more conveniently to other software.

This talk covers the efforts to present iRODS as WebDAV, FUSE, REST, NFS, SFTP, K8s CSI, and S3.

The iRODS Consortium was started as an open-source software development organization in 2013 by members of the research and storage communities. The technology has roots from an earlier project started in 1995. The Consortium was launched as a response to a major scale increase in management and storage needs driven by the advent of “big data”. The member community is now comprised of over 30 members and spans the globe from the Australia to Japan and much of the EU.

Security and Authentication / 72

Secure Zones for Sunet Drive

Authors: Michiel de Jong¹; Micke Nordin²

Co-authors: Leif Johansson²; Richard Freitag²

¹ *Ponder Source*

² *SUNET*

Corresponding Authors: michiel@pondersource.com, kano@sunet.se

Enterprise File Sync and Share (EFSS) systems have become an integral part of every researcher’s life, handling an abundance of scientific data for multiple projects. Those projects generally span multiple collaborators and can extend over a significant geographic area. However, there is an inherent conflict when handling research data, between the researcher’s need to collaborate and share data with each other and the sensitive nature that that data can sometimes have. Secure Zones for Sunet Drive is a technical implementation of protected data zones in the EFSS system, guarded by step-up authentication. The idea is to only give access to protected data to users that have been properly identified and help those users when handling the data so that they do not give further access to someone they should not, by mistake.

The complexity of multi factor authentication (MFA) can be understood, when one considers all parameters involved in its implementation. Multiple technologies like SMS, TOTP, or FIDO2 devices can be implemented either by the identity provider (IdP), the service provider (SP), or potentially even both. Among other things MFA also requires administration for lost or stolen devices. Identity providers must implement MFA individually and different technologies can be used for different IdPs.

Secure Zones for Sunet Drive have been developed in collaboration with Ponder Source and they implement MFA on the service provider side, with hooks being built into the EFSS solution such that a seamless transition between the general use of data, and corresponding secure zones can be done almost seamlessly. Since many EFSS systems have support for single-sign on via SAML, but no support for Discovery Services (i.e., aggregators of SAML/SSO-logins), Sunet Drive uses SaToSa, a configurable proxy for translating between different authentication protocols and providers. Users can opt to log on to the EFSS directly via their identity provider, with or without MFA and then step-up with MFA at a later point if necessary. The EFSS is made aware whether a user has logged on using MFA and if certain data storage areas of the EFSS should be accessible or not. Users can also control whether certain files or folders will require access via step-up authentication.

Secure Zones are an important technical tool that can be used by organizations and research groups to be compliant with the handling of sensitive data.

Security and Authentication / 74**On-demand cloud-based secure environments for analysing personal and health data**

Authors: Federico Zambelli¹; Giacinto Donvito²; Marco Antonio Tangaro³; Marica Antonacci⁴; Nadina Foggetti^{None}

¹ *Department of Biosciences, University of Milan*

² *INFN - National Institute for Nuclear Physics*

³ *Institute of Biomembranes, Bioenergetics and Molecular Biotechnologies*

⁴ *INFN*

Corresponding Author: ma.tangaro@ibiom.cnr.it

Galaxy is the de facto standard workflow manager for bioinformatics providing a complete collaborative platform for researchers. Even though several Galaxy public servers are currently available, there are some situations where users would benefit more from having full administrative control over a private Galaxy instance. These situations include, but are not limited to, worries about data privacy, the need for customization, the need to prioritise particular job types, the development of tools, and training activities.

The Laniakea 1 software platform facilitates the provisioning of on-demand Galaxy instances over heterogeneous Cloud infrastructures, by leveraging on the open source INDIGO-DataCloud cloud stack [2], which aims to make cloud infrastructures more accessible by scientific communities.

End users interact with Laniakea through a web front-end that allows a general setup of the Galaxy instance. The deployment of the virtual hardware and of the Galaxy software ecosystem is subsequently performed by the INDIGO Platform as a Service layer. At the end of the process, the user gains access to a private, production-grade, fully customizable, Galaxy virtual instance. Laniakea features the deployment of stand-alone or cluster backed Galaxy instances, shared reference data volumes, and rapid development of novel Galaxy flavours for specific tasks.

Moreover, to extend the usage of this platform in clinical scenarios, where the analysis of sensitive data, in compliance with the GDPR, requires strong countermeasures to grant data privacy and security, Laniakea guarantees the creation of isolated and secure environments, exploiting storage encryption and access control to Galaxy through VPN, in order to carry out data analysis.

Laniakea allows the on-demand encryption of the entire storage volume attached to the virtual machine, using the Linux kernel encryption module. The level of disk encryption is completely transparent to software applications, in this case Galaxy: data are encrypted and decrypted on-the-fly when writing and reading, respectively. The procedure has been completely automated through the web Dashboard of the PaaS orchestration service [3], taking advantage of Hashicorp Vault for storing user passphrases.

We have implemented a robust mechanism to create secure encryption keys and prevent user credentials or the encryption passphrase from being transmitted unencrypted to the virtual infrastructure, compromising its security.

The oral contribution will provide details about the platform architecture and the service implementation strategy.

References

- 1 Tangaro et al. , Laniakea: an open solution to provide Galaxy “on-demand” instances over heterogeneous cloud infrastructures, *GigaScience*, Volume 9, Issue 4, April 2020, g1aa033, <https://doi.org/10.1093/gigascience/giaa033>
- [2] Salomoni, D., Campos, I., Gaido, L. et al. INDIGO-DataCloud: a Platform to Facilitate Seamless Access to E-Infrastructures. *J Grid Computing* 16, 381–408 (2018). <https://doi.org/10.1007/s10723-018-9453-3>
- [3] <https://github.com/indigo-dc/orchestrator>

Security and Authentication / 93**OIDC, it seems to be hip**

Author: Tom Wezepoel¹

¹ SURF

Corresponding Author: tom.wezepoel@surf.nl

How to change the login method for almost 90000 users from 5 different login scenarios and different backends to 1 method with OIDC. Welcome in the world of flows with Keycloak. What could possibly go wrong?

Technology Bricks: Virtualisation, Monitoring and Notifications / 87

Driving the ScienceBox package into the future

Authors: Enrico Bocchi¹; Jimil Desai²; Samuel Alfageme Sainz¹

¹ CERN

² Ahmedabad University (IN)

Corresponding Author: samuel.alfageme.sainz@cern.ch

In this talk we describe the 2022 reboot of the **ScienceBox project**, the demonstrator package for some of CERN's storage and analysis services. We evolved the original implementation to make use of **Helm charts** across the entire dependency stack.

We've also incorporated the major architectural update to **CERNBox, replacing the previous PHP backend with a catalog of distributed microservices based on Reva**. Besides enhancing our existing use cases, the new CERNBox implementation enables and streamlines interoperability with additional applications and sites deployed under the CS3 APIs or compatible with them.

We present this update as a **self-contained and easy-to-use package with minimal dependencies** and with the same goals as the original ScienceBox: Provide a sandbox to evaluate the storage, sharing, and analysis services we run at CERN on external premises to non-CERN users. We believe there is not only a great value in releasing and contributing back to the open source projects that sustain these services, but also in **describing the configuration and artifacts that make operating such complex software systems at scale possible**.

Technology Bricks: Virtualisation, Monitoring and Notifications / 99

An overview of Nextcloud monitoring

Author: Pietro Marini¹

¹ Nextcloud GmbH

Corresponding Author: pietro.marini@nextcloud.com

Deploying Nextcloud at scale implies a close monitoring of critical software and infrastructure components. In enterprise environments, Nextcloud is typically run in a clustered setup and it requires both infrastructure and application monitoring. In this talk we are going to discuss the basic elements of monitoring, with a focus on understanding why some metrics are important to monitor to judge the global healthiness of the application. The content will be based on a blend of theoretical product features and the experience gathered in years of production operations of our customers instances.

Technology Bricks: Virtualisation, Monitoring and Notifications / 88**Modern notification system for Sync and Share using NATS****Author:** Javier Ferrer¹**Co-author:** Elizaveta Ragozina ¹¹ CERN**Corresponding Author:** javier.ferrer@cern.ch

The new CERNBox platform was successfully released for CERN-Wide usage on October 24. It has been quickly adopted by the whole community of about 27.000 users. The platform is comprised of a new Web User Interface based on modern web framework technologies, and a scalable distributed micro-service backend architecture based on Reva.

One of the most prevalent feature requests received after the release was to provide users with the means to generate notifications for the different workflows on the platform (e.g., sharing a folder with another user or group, or uploading a file to a shared space).

A notification system has been implemented with different types of media (email, web); a templating system for e-mails, a hook system which enables extensibility to all parts of the platform, and a notification center for the web. The notifications service leverages the NATS open-source communication system to enable fast and scalable communication.

In this talk we describe the architecture of the new service, the hooks provided to enable developers to add new notifications, and the planned future extensions.

Collaborative Data Science and Visualisation / 84**Visualize big data and create individual customer frontends****Author:** Christoph Dyllick-Brenzinger^{None}**Corresponding Author:** cdb@datamate.org

One of the main challenges in dealing with large amounts of data is to find a **suitable presentation for the different target groups**. With the new module **External App**, SeaTable allows you to build individual frontends for the different stakeholders and process participants in no time.

In this way, processes can be streamlined and the transfer of information can be made more efficient.

In this presentation, I will introduce and explain the functions of the interface designer.

Collaborative Data Science and Visualisation / 101**FaaS Data Processing with Onedata****Authors:** Bartosz Kryza¹; Michał Orzechowski²; Łukasz Dutka^{None}¹ ACC Cyfronet-AGH² AGH University of Science and Technology, Academic Computer Centre Cyfronet AGH, Krakow, Poland**Corresponding Author:** orzechowski.michal@gmail.com

Onedata is a distributed, global, high-performance data management system, which provides transparent and unified access to globally distributed storage resources and supports a wide range of use cases from personal data management to data-intensive scientific computations. Due to its fully

distributed architecture, Onedata allows for the creation of complex hybrid-cloud infrastructure deployments, with private and commercial cloud resources. It allows users to share, collaborate and publish data and perform high-performance computations on distributed data. Onedata enables users to collaborate, share and perform computations on data using applications relying on POSIX-compliant data access.

Onedata has recently been enhanced with a powerful workflow execution engine powered by OpenFaaS. This allows for the creation of complex data processing pipelines that have transparent access to distributed data provisioned by Onedata. The workflow functionality can be used for embedded data processing and includes a library of ready-to-use functionalities such as metadata extraction and format conversion. Custom functions can also be easily added and shared among user groups. The solution has been thoroughly tested on auto-scalable Kubernetes clusters. In addition to the transparent access to distributed data, the use of a Function as a Service (FaaS) platform for data processing offers flexibility and innovation for contemporary data tasks. The use of FaaS allows for the creation of custom, modular functions that can be combined and executed in a variety of ways to meet the specific needs of a given data processing task. This modular approach makes it easy to scale and update individual functions as needed, making FaaS well-suited for the dynamic and constantly evolving nature of modern data processing. The combination of FaaS and Kubernetes further enhances the flexibility and scalability of this data processing solution. By running the FaaS platform on top of Kubernetes, it is possible to easily scale the number of functions being executed and the resources allocated to them based on the needs of the specific task at hand. This allows for efficient and effective use of resources, making it possible to tackle even the most demanding data processing tasks. The use of Kubernetes also enables seamless integration with other tools and technologies, further expanding the capabilities of the FaaS platform. Overall, the application of FaaS on top of Kubernetes makes this data processing solution highly flexible and well-suited for a wide range of contemporary data tasks.

Currently, Onedata is used in the European EGI-ACE PRACE-6IP, and FINDR project, where it provides a data transparency layer for computation, and data processing automation deployed on dynamically hybrid clouds containerised environments.

Acknowledgements. This work was supported in part by 2018-2020's research funds in the scope of the co-financed international projects framework (project no. 5145/H2020/2020/2).

1. Onedata project website. <https://onedata.org>.
2. OpenFaaS - Serverless Functions Made Simple. <https://www.openfaas.com/>.
3. David Giaretta, CCSDS Group, and CCSDS Panel. Reference model for an Open Archival Information System (OAIS). 06 2012.
4. EGI-ACE: Advanced Computing for EOSC. <https://www.egi.eu/projects/egi-ace/>.
5. Partnership for Advanced Computing in Europe - Sixth Implementation Phase. <http://www.prace-ri.eu>.
6. FINDR: Fast and Intuitive Data Retrieval for Earth Observation

Collaborative Data Science and Visualisation / 95

Data Science environments in ScienceMesh

Author: Marcin Sieprawski^{None}

Co-author: Diogo Castro¹

¹ CERN

Corresponding Author: marcin.sieprawski@softwaremind.com

Data Science is a complex field that requires a high level of expertise and collaboration among teams of experts. With the rise of big data, it has become increasingly important to create collaborative workflows that enable data scientists to combine their skills and knowledge to create better results. This, however, can be a challenge in an environment of heterogeneous cloud and storage systems.

ScienceMesh, developed in CS3MESH4EOSC project, creates the Federated Scientific Mesh providing federated sharing of data across different sync-and-share services, federated use of applications (such as collaborative document editing, data archiving, and data publishing), fast transfer of large datasets and remote data analysis (Data Science environments).

For ScienceMesh distributed data science environments we developed a JupyterLab extension, integrating with ScienceMesh – so that we can provide file browsing and additional share and collaboration functionalities for notebooks and resources.

This talk will present the development of Data Science environments in ScienceMesh and demonstrate how it supports collaborative workflows in federated sync-and-share environment.

123

Summary and Conclusions

Corresponding Author: jakub.moscicki@cern.ch

ScienceMesh workshop / 110

Science Mesh demos