


Cloud Services for Synchronisation and Sharing (CS3)

Book of Abstracts

27 - 29 January 2020

Copenhagen, Denmark



<http://www.cs3community.org>

Introduction

This is the "Book of Abstracts" of CS3 2020 -- the 6th Cloud Synchronization&Sharing Conference (Copenhagen, 27-29 January 2020).

A lot happened since the last meeting in Rome. The European Open Science Cloud (EOSC) is in full swing, digital sovereignty became one of the critical topics for European public institutions and literally tens of organization are looking for opportunities to offer services based on open-source software, on-premise data storage and collaborative work environments cross-institutionally. This opens up new opportunities for the CS3 community.

Last year the CS3 participants prepared the CS3MESH¹ project proposal which was approved as a 3-year EU project starting January 2020. This new project has been as a jump-start for the entire CS3 community to become a permanent and integral part of the newly shaping landscape of European cloud computing for research and education.

This edition of the annual CS3 conference is immediately preceding the CS3MESH project kick-off meeting. This was a deliberate choice: we believe that CS3MESH is a tool and an opportunity for all the actors of the CS3 community. The project aims to deliver the initial and coherent implementation of the long-term CS3 collaborative infrastructure. The CS3 2020 conference is an important opportunity to give input and help co-design the common future.

We expect the project to deliver the initial implementation of the CS3 infrastructure, hosting all sites interested to use it and to contribute to its evolution. We placed a dedicated session "MeetCS3MESH" in the CS3 2020 programme and we encourage input on specific use cases: requirements ("my users desperately want this"), success stories ("this is the way we operate our system and we believe it might be of interest of others") or other experience ("it looked a good idea but eventually we abandoned it"). So, if you want something delivered from CS3MESH, or you are prepared to co-design it or you are offering experience or resources, this is your session. We expect that some contributions might be called to give more technical details in one of the following tracks while concentrating here on the high-level messages.

We believe that original formula focusing on innovative storage systems and their integration with user environments remains the key to enable progress in data sciences at all levels: local laboratory, regional collaborations and global science. Therefore CS3 will continue to follow applications ranging from innovative big-data analysis to science outreach and education promoting exchange of experience, innovative ideas and collaboration across all actors in this industry. In short: CS3 will continue to drive the evolution in the field via an inclusive community effort and CS3MESH project will help to structure this effort.

Flexible CS3MESH federation across installations will promote global scientific collaboration and integration, avoiding disconnection of infrastructures. The main strength of CS3MESH is the CS3 community and the project will be listening to all proposals!

The CS3 community certainly has the momentum to transform the activities of these past years into the sustainable fabric of cloud storage of the future and establish it as a European mainstream infrastructure for research, science and public institutions.

Massimo Lamanna and Jakub T. Mościcki

¹ Full project acronym: CS3MESH4EOSC, Fact Sheet: <https://cordis.europa.eu/project/id/863353>



<http://www.cs3community.org>

Contents

May we please store your personal data?	1
CS3MESH Introduction	1
Interoperable Research Workspaces and the Mesh	1
The next steps for Nextcloud	2
File Cloud Innovation at ownCloud	2
Seafile development in year 2019	3
ownCloud infinite scale - Concepts, Protocols and Extensions	3
Nextcloud architecture update	3
CS3 Community Summary with the CS3MAP – Animated MAP of the CS3 community	4
bwSync&Share at KIT	4
A SOLID Distributed Architecture for Sciebo Research Data Services	4
Supporting Keycloak in iRODS systems with OpenID authentication	5
Preparing EOS for Enterprise Users	6
UNCloud, a cloud storage and collaborative work platform in a french university.	7
Backup and DR with Large Datasets	8
Breaking Cloudstor since 2019, so our users won't	8
Large scale Home Directories migration: from DFS to CERNBox	9
BNLBox: A New Cloud Storage Service at the BNL SDCC	9
Evolving the Nextcloud ecosystem - during spare time	9
Powerful document collaboration: Online, Desktop, Mobile	10
Digital sovereignty in communication & collaboration	10
Collabora Online: secure, on-premise collaborative editing	11
Sovereignty or Surrender: Europe has a choice	11

The importance of open source software in delivering large storage services	12
Digital Sovereignty can only be achieved through pure Open Source Software	12
From instrument to publication: A First Attempt at an Integrated Cloud for X-ray Facilities	12
Sharing your scientific data with Rucio	13
iRODS and Sync&Share: Getting the best out of both worlds	14
Jupyter on Earth: how SWAN is powering CERN use cases	15
ELIXIR Cloud & AAI: Standardised and Interoperable Services for Human Data Communi- ties	15
From EFSS to Smart Workspace	16
Interoperability Protocols: OCM, CS3APIs, ... and more – Campfire Discussion	16
CS3 APIs: from the community to the community	16
CS3 APIs for data science environments	17
WOPI and Custom Apps in the CS3APIs	17
The HIFIS Cloud Competence Cluster	18
FAIR Implementation Profiles: Driving Convergence onto an Internet of FAIR Digital Ob- jects	18
ESCAPE ESFRI Science Analysis Platform	19
Oracle for Data Science – A Vision for CS3MESH Connector	20
Interoperability as the new frontier	20
Onedata Jupyter Integration Using OnedataFS	21
Dismantling the Monolith	22
Sharing data between HPC, hadoop and POSIX with EFSS	22
iRODS - Beyond Discoverability: Metadata to drive your data management	23
Building a scalable and fully automated public cloud platform based on Seafle and Ceph	23
The growth of the Nextcloud apps ecosystem	24
Owncloud Infinite Scale Extensions	24
Migrating CERN e-mail system to open source	24
Real-time communication and collaboration in a file sync & share environment	25
ScienceBox as Deployment Model for CS3MESH	25
Elettra Drive(s) data far and FAIR	26

Drive RENATER	26
OSF and CS3MESH	27
Ceph in the real world - what have we learned	27
File and Object synchronization in the Hybrid Multi-Cloud world. (NetApp CloudSync) .	28
Scalable and innovative high-performance HPC solution with BeeGFS backed by NetApp E-Series storage	28
ARCHIVER - Data Archiving and Preservation for Research Environments	29
Seafile - great but vendor lock-in? Evaluating Seafile openness in practice	29
From Jupyter notebooks to web dashboards for big geospatial data analysis	30
The Cern VM File System as global software delivery fabric	32
Reva: IOP (Inter-Operability Platform)	32
Just-in-time data ingestion for distributed metagenomics assembly in multiple clouds with OneData	33
Improving user and customer experience in CS3 services	33

Welcome and Keynote / 117**May we please store your personal data?****Corresponding Author:**

Michiel de Jong will talk about his work on Terms of Service; Didn't Read, Unhosted web apps, and more recently, Solid.

Meet CS3MESH / 111**CS3MESH Introduction**

Author: Jakub Moscicki¹

¹ *CERN*

Corresponding Author: jakub.moscicki@cern.ch

Cloud storage services for synchronisation and sharing are an indispensable element of the daily workflow routine, allowing research groups, scientists and engineers to share, transfer and synchronise data in simple but powerful ways. The services are operated and funded by major e-infrastructure providers such as the National Research and Education Networks and major research institutions. However, these services remain largely disconnected from each other. The EU-funded CS3MESH4EOSC project will integrate the existing application and storage ecosystem by promoting vendor-neutral application programming interfaces and protocols. State-of-the-art connected open-source infrastructure will provide researchers with a broader access to services and boost collaborative research.

CS3MESH4EOSC will implement a service for the European Open Science Cloud (EOSC) with a built-in sustainability model using the on-premise service delivery by utilizing existing key technology enablers: Open Cloud Mesh (OCM) standardized protocol and EduGAIN service. It will consolidate and integrate the existing application ecosystem following the open-source strategy for delivering services.

CS3MESH4EOSC will empower service providers in delivering state-of-the-art, connected infrastructure to boost effective scientific collaboration across the entire federation and data sharing according to FAIR principles. It will also help strengthening the role of European industry in delivering competitive cloud solutions on global scale. Finally, the project will deliver the core of a scientific and educational infrastructure for cloud storage services in Europe for research, education and public institutions.

Meet CS3MESH / 71**Interoperable Research Workspaces and the Mesh**

Authors: Gavin Kennedy¹ ; Peter Sefton² ; Jason Nairnsey³ ; Andrew Brazzatti⁴ ; Carina Kemp¹

¹ *AARNet*

² *University of Technology Sydney*

³ *University of Wollongong*

⁴ *Queensland Cyber Infrastructure Foundation*

Corresponding Author: gavin.kennedy@aarnet.edu.au

In our collective efforts to deliver data services to the research community the separation between data and service is often blurred. Research data services are often either highly specific to the data or generalised across all possible domains. Features and functions are overlapping or redundant, and research integrity is at risk from data stored across multiple systems (“Research Workspaces”), where data becomes hard to manage, publish and archive. Such Research Workspaces include public and private cloud storage and compute services and applications such as data repositories, code repositories, electronic notebooks, online surveys and instrument data applications.

To address these issues an Australian consortium including the University of Technology Sydney, the University of Wollongong, Queensland Cyber Infrastructure Foundation and AARNet are developing a sustainable approach to managing Research Workspaces and the interchange of data and metadata.

The approach is to extend research data management platforms, such as the Australian ReDBox platform, to provision Research Workspaces in platforms such as AARNet’s CloudStor. This provisioning will happen during the creation of research Data Management Plans (DMPs) or Data Management Records (DMRs). DMPs/DMRs are then linked to one or multiple Research Workspaces and surfaces metadata to assist in the Findability and Accessibility of data (the “F” and “A” in FAIR[1]). As data is generated and managed in these Research Workspaces, the metadata is subsequently enriched.

This approach will allow institutional, national and international linking of a wide variety of Research Workspace applications. Combined into CS3Mesh enabled environments researchers can work in a variety of applications, move data between them and be assured that they will always be able to export their data into a neutral archival/preservation package. Currently in development for ReDBox is the utilisation of the emerging specifications of OCFL[2] and ROCrate[3], this approach ensures data are not locked-in to particular platforms or applications - making data Interoperable and maximising the potential for Reuse (the “I” and “R” in FAIR).

[1] FAIR Data - Findable, Accessible, Interoperable, and Reusable, <http://wilkinsonlab.info/node/FAIR>

[2] OCFL: Oxford Common File Layout, <https://ocfl.io/>

[3] ROCrate: Research Object Crate, <https://researchobject.github.io/ro-crate/>

File Sync&Share Products for Home, Lab and Enterprise / 90

The next steps for Nextcloud

Author: Frank Karlitschek¹

¹ *Nextcloud*

Corresponding Author: frank@nextcloud.com

Nextcloud started as an Enterprise File Sync and Share solution. Nowadays Nextcloud turned into a full Content Developer Platform and an ecosystem for developers and 3rd party applications. 2019 was a particular interesting year with 3 major releases from Nextcloud. This talk will give an overview over the changes in 2019 and what it means for the future. A special focus will be on the latest release which is a very big step forward and a change in strategy. This talk will also discuss the roadmap and what it means for the education and research community. The talk will cover several concrete unique features with live demonstrations.

File Sync&Share Products for Home, Lab and Enterprise / 105

File Cloud Innovation at ownCloud

Author: Felix Boehm¹

¹ *ownCloud GmbH*

Corresponding Author: felix@owncloud.com

ownCloud will launch a rewrite of ownCloud's server and front-end in 2020. A peak preview will show some highlights. How can you benefit even with ownCloud 10? And what comes down the roadmap?

File Sync&Share Products for Home, Lab and Enterprise / 61

Seafile development in year 2019

Author: JiaQiang Xu^{None}

Corresponding Author: xjqkilling@gmail.com

In this presentation we'll provide a review of the development of Seafile project in the year 2019 and what's coming in the near future. Major efforts include:

- Implementing OCM protocol
- A major rewrite of the Drive client to use Windows Cloud Files API
- Python 3 support

File Sync&Share Products for Home, Lab and Enterprise / 85

ownCloud infinite scale - Concepts, Protocols and Extensions

Author: Jörn Dreyer¹

¹ *ownCloud GmbH*

Corresponding Author: jfd@owncloud.com

Do you know what makes a reliable sync an share platform? At ownCloud we have been able to iterate on your feedback and rethink some of the underlying concepts of our architecture. In this session we will examine the most important paradigm shifts and explain how they led us to the different open protocols that we are building the new architecture upon. Finally, we are going to give an overview of how these protocols allow building extensions in multiple languages and even allow end users to run custom code without affecting the core stability of the platform.

File Sync&Share Products for Home, Lab and Enterprise / 91

Nextcloud architecture update

Author: Frank Karlitschek¹

¹ *Nextcloud*

Corresponding Author: frank@nextcloud.com

Nextcloud has a well tested and established server side architecture. Unlike other product the Nextcloud strategy is continuous improvement without rewrites or feature regressions. This talk will give an overview over the current architecture with concrete examples how to run it in high performance environments. Additionally it will discuss the latest improvements of the unique Global Scale

architecture that allows to run Nextcloud in a globally distributed environment for tens of millions of users. The talk will discuss ways to scale the Nextcloud push notification system, the Nextcloud Talk Chat and Video system and other components. Additionally it will cover some roadmap features that will reduce the server load by several magnitudes.

Site reports / 116

CS3 Community Summary with the CS3MAP – Animated MAP of the CS3 community

Authors: Dr. Tilo Steiger¹ ; Stefano Stalio²

¹ *ETH Zuerich*

² *INFN*

Corresponding Authors: steiger@id.ethz.ch, stefano.stalio@lngs.infn.it

We will present a summary of the results of the Site Reports survey as well as a short, helpful and graphical overview of the installations within the CS3 Community: <https://cs3map.ethz.ch>.

Site reports / 97

bwSync&Share at KIT

Author: Klaus Scheibenberger¹

¹ *KIT*

Corresponding Author: klaus.scheibenberger@kit.edu

At the SCC (Steinbuch Centre for Computing) - the scientific data centre of the KIT - we operate a country-wide Sync&Share Service called "bwSync&Share" for Baden-Wuerttemberg since 2014. This service is embedded in the country-wide Identity federation called "bwIDM" which is operational since 2013.

Both projects were sponsored by the Ministry for Science, Research and Arts in Baden-Württemberg.

Beneath the pure functionality provided by the software the service evolved with processes closely coupled with the bwIDM federation, e.g. provisioning/deprovisioning of accounts, registration based on entitlements, privacy.

Currently 35 Institutions are using this service with about ca. 40.000 accounts.

In 2020 we'll introduce a cost allocation for the institutions using bwSync&Share so the service becomes financially more self-reliant. Thus legal aspects and contracts become very important and we started working on this topic already in 2019. At the moment we also develop the migration of the service to a new software basis.

I may report about some of our experiences operating this service and the legal, financial aspects generated with the transformation to a country-wide service financed by the institutions using bwSync&Share.

Site reports / 67

A SOLID Distributed Architecture for Sciebo Research Data Services

Authors: Peter Heiss¹ ; Jens Stegmann¹ ; Holger Angenent¹

¹ *University of Muenster*

Corresponding Authors: peter.heiss@uni-muenster.de, jstegmann@uni-muenster.de, holger.angenent@uni-muenster.de

The project sciebo Research Data Services (sciebo RDS) [1] of the University of Münster and the University of Duisburg-Essen aims to extend the functionality of the ownCloud-based sciebo service, in particular by integrating existing services (e.g. Zenodo, DataCite, EPIC, CLARIN Webservices), in order to provide one-click solutions and easily reusable workflows for integrated research data management (RDM). The RDM services are supposed to enable the users in their day-to-day research work via sciebo.

Currently, sciebo [2] has more than 100.000 users with 35.000 of them being scientists, i.e. they are employed at participating universities in the German federal state of North Rhine-Westphalia. Hence, infrastructure and architecture have to support scalability. In order to make the integration, the development of new services and the management of the sciebo RDS as easy as possible, a modular, flexible and sustainable architecture is required. A suitable architecture should follow clear principles concerning both structure and dependencies, such that developers can contribute and reuse easily without knowledge of all implementation details.

The architecture to be presented is based on the Clean Architecture approach of Robert C. Martin. It follows the so called SOLID principles (Single responsibility, Open-closed, Liskov substitution, Interface segregation, Dependency inversion) [4], which simplify and clarify the design, the implementation and the maintenance of the whole ecosystem. Some of these principles were postulated as early as the 1980s. They are grounded in observations about different kinds of programming paradigms (structured, object-oriented, functional). With these principles in place, we strive to enable and sustain reusability of the architecture and its components.

The talk will introduce the applied architectural features and discuss aspects of their implementation in depth, i.e., we intend to show how specific parts of a scalable ecosystem based on microservices and containers on top of Kubernetes can integrate and function for this purpose. Concerning the interfaces of the microservices under discussion, we showcase the usage of the OpenAPI v3 specification, which additionally fosters reusability.

Sciebo RDS is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - project number 403637381.

[1] <https://www.research-data-services.org>

[2] <https://www.sciebo.de>

[3] <https://blog.cleancoder.com/uncle-bob/2012/08/13/the-clean-architecture.html>

[4] <https://en.wikipedia.org/wiki/SOLID>

Site reports / 74

Supporting Keycloak in iRODS systems with OpenID authentication

Authors: Rubén Jesús García-Hernández¹ ; Martin Golasowski²

¹ *Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities*

² *VSB Technical University Ostrava*

Corresponding Author: garcia@lrz.de

The Integrated Rule-Oriented Data System (iRODS) is an open source data management software aimed at deployment in mission critical environments, which virtualizes data storage resources. iRODS can take advantage of OpenID login by using an authentication plugin.

The approach works by letting iRODS authenticate using tokens provided by an OpenID provider, which are verified by the plugin (e.g. [1]).

However, the token is transferred between the iRODS client and the iRODS server by using the user-name field,

which has a maximum length of 1024+64 bytes. This is plenty for most OpenID implementations; the only exception the authors are aware of is Keycloak, since its JWT tokens contain extensive information about the user,

their identity and their permissions. Some changes were required in the different components of the authentication plugin

to avoid issues with such tokens.

Additionally, when such a token is sent from the iRODS client to the iRODS server, an iRODS USER_PACKSTRUCT_INPUT_ERR error

is produced, and the iRODS server prints a URL to enable re-authentication.

Once this URL is used by the user to authenticate, the iRODS command can complete.

While this is enough in some situations, some workflows cannot use this approach.

In particular, the iinit command will not work, and web-based applications will require modifications.

There are plans within iRODS to more tightly integrate OpenID in the future, avoiding these issues and simplifying implantation.

However, this solution will very probably be provided outside of our LEXIS project timeframe (2019-2021).

We implemented two solutions to the issue, taking care to maintain the overall security level:

a) for web-based applications interfacing directly with the user, a helper thread or process is launched to perform a query for (meta-data)

to the iRODS system and later (after successful authentication) store the result. The main thread gets the authentication URL from the helper

thread (while the helper is waiting for results from the iRODS) and redirects the user to it. The helper thread obtains the result data from

the iRODS server as soon as the user authenticated. After the user is eventually redirected to the web portal, the web application can retrieve

this data. No token ever traverses the iRODS client-server interface, so no error is produced.

b) for back-end applications, the solution above is not applicable. We therefore modified the iRODS authentication microservice [2] to implement

opaque tokens. The microservice now sends a hash of the token to the user (which is small enough to traverse the iRODS interface). Upon receiving

a hash, the microservice database is used to retrieve the token and perform the verification activities.

The project received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 825532 with LEXIS - Large-scale EXecution for Industry & Society.

We thank the ITS Group of Leibniz Supercomputing Centre (LRZ, Garching) of the BAdW for supporting our research with its Compute-Cloud infrastructure.

[1] https://iRODS.org/uploads/2019/Cacciari-CINECA-OpenID_Connect-paper.pdf

[2] https://github.com/heliumdatacommons/auth_microservice/

Preparing EOS for Enterprise Users

Author: Molan Gregor¹

¹ *Comtrade*

Corresponding Author: gregor@comtrade.com

EOS has been developed to the stage where it became potentially interesting for enterprise users. Being developed for research use however, means that EOS had to be modified for commercial use - that is where Comtrade with its vast experience with enterprise software development came in.

The first step in adoption of EOS in enterprise environment is development of robust installation and detailed documentation. Functionalities of EOS are summarized into 8 topics (work packages). For each work package we examined:

- Real use cases of functionality
- Source code corresponding to specific work package
- Feedback from end users
- Internal documentation regarding the functionality prepared by developers
- Verification of documentation prototype at the live EOS storage system at CERN
- Set up of EOS corner case environments on virtual hosts (Geo-scheduling, Balancing, Upgrades)

Based on this examination we prepared prototype of new documentation. This prototype was validated both by domain experts (developers) and end users. Feedback was implemented into version that was edited by professional proof readers. Iterations of revisions between proof readers and experts was repeated several times. The final version was also edited by graphic design team.

Some work packages were merged into documents covering additional functionalities like:

- Best practice guide
- Installation guide

We believe that work presented here serves as the first step towards establishing EOS as a viable enterprise storage product.

Site reports / 68

UNCloud, a cloud storage and collaborative work platform in a french university.

Authors: Matthieu Le Corre^{None} ; Arnaud Abélard^{None}

Corresponding Authors: matthieu.lecorre@univ-nantes.fr, arnaud.abelard@univ-nantes.fr

In 2017, The University of Nantes, France, launched UNCloud, a web service project aiming to facilitate interaction and collaboration between the faculty, staff and students as well as external collaborators. Both a storage service and a collaboration platform, offering 100 GB to their 70,000 faculty and staff members and students, UNCloud has become, with its 10,000 users in only a few months, a cornerstone of the establishment.

To begin, we'll discuss the politics that led to the birth of the project in a context of mistrust towards free consumers services in an unclear legal landscape. After studying several alternatives, the opensource solution Nextcloud was selected.

We will consider its implementation, both technically and organizationally. Special attention will be given to the test phase and its management. We will explore the technical challenges that led us to building a fully redundant infrastructure and we'll explain its deployment in detail.

Finally, we'll draw on our experience and the lessons we have learned. Especially in user support and technical optimization..

Now in production for almost two years, a new era starts : transforming a sharing platform into a whole digital collaboration environment by integrating the existing collaboration services in use at the university : email, calendar, learning management systems, etc.

Site reports / 76

Backup and DR with Large Datasets

Author: Fergus Kerins^{None}

Corresponding Author: fergus.kerins@switch.ch

SWITCHdrive is the File and Sync Solution from SWITCH. We are running on our internal SWITCHengines (Openstack) IaaS with CEPH storage. With the success of the service we are managing a relative large dataset.

We currently use CEPH rbd snapshots for backup, but with the growth of our service and growing Customer demands for disaster recovery and regulatory requirements we are looking at options to provide offsite backups

Due to our current setup of multiple NFS drive we are currently evaluating solutions for an offsite Disaster Recovery solution and backup.

This talk we will look at using backy2 a deduplicating block based backup to backup rbd volumes to S3 object storage. We will also look at other options and compare our results.

We would also like to reach out to the community to share experience and hear how other people are solving this issue.

Site reports / 70

Breaking Cloudstor since 2019, so our users won't

Author: Brad Marshall¹

Co-author: Gavin Kennedy²

¹ AARNet Pty Ltd

² AARNet

Corresponding Author: brad.marshall@aarnet.edu.au

AARNet's CloudStor is one of the most successful collaboration tools across the Australian research data ecosystem. We operate a significant software stack across multiple geographically dispersed nodes. Our service end points are exponentially increasing, along with the complexity of the interacting components, and an ever increasing number of end users.

As the AARNet Cloud Services team have been busily scaling our systems up, we've noticed that its been harder and harder to test our changes in a way that accurately reflects our production environments. Complexity introduces challenges with scale, and we have to scale up some of our processes to cope with these challenges.

This talk will cover the road we have started to walk down for quality control, specifically in how we are testing changes, and a mindshift in how we are looking at monitoring. This will include testing environments, chaos monkey environments, infernal testing environments and client focused monitoring and eventually full continuous integration.

The aim with our QA platform is to provide an environment we can rapidly spin up full deployments using Openstack, Kubernetes, Terraform and Ansible. The combination of this rapid deployment, improved monitoring and ability to test multiple complex environments will facilitate the addition of new features and capabilities to our systems. This agility will benefit our interoperability testing with the CS3Mesh, and our deployment testing for the new ownCloud Infinite Scalability deployment.

Site reports / 99

Large scale Home Directories migration: from DFS to CERNBox

Author: Apostolos Smyrnakis¹

Co-authors: Vincent Nicolas Bippus¹; Michal Kwiatek¹; Sebastian Bukowiec¹; Sebastien Dellabella¹; Joel Alcala Perez²; Hugo Gonzalez Labrador¹; Luca Mascetti¹; Diogo Castro¹; Jakub Moscicki¹

¹ CERN

² Universidad de Oviedo (ES)

Corresponding Author: apostolos.smyrnakis@cern.ch

Moving users' data is never an easy task. When it also involves changing the way they work and collaborate, then it can become even more complex. In this presentation you will discover some of the different challenges we faced when designing and deploying a data migration automation that moved the Home directory for more than 15,000 Windows users.

Site reports / 107

BNLBox: A New Cloud Storage Service at the BNL SDCC

Authors: Hironori Ito¹; Ofer Rind^{None}

¹ Brookhaven National Laboratory (US)

Corresponding Author: hito@bnl.gov

The Brookhaven National Laboratory (BNL) is a multi-disciplinary US Department of Energy (DOE) lab, supporting a wide range of scientific research. The Scientific Data and Computing Center (SDCC) at BNL provides computational resources to the scientists and engineers engaged in this research. In order to support this wide variety of science, the SDCC supports various types of data storage services: very fast data caches; large, multi-PB distributed stores; network file systems; parallel file systems; object storage, etc. BNLBox, which is based on NextCloud software, is one of the newer storage services at BNL. It is a cloud storage service, allowing users easy access to their storage wherever they are on the network. The storage is built on top of a robust and reliable Lustre File system. The HPSS API in LustreFS is being used to provide the users with true archive storage. With redundant and backed-up services wherever possible, BNLBox provides a stable and easy-to-use storage service for all users at BNL. This presentation will describe the detailed setup and current features of the BNLBox service at BNL.

Site reports / 110**Evolving the Nextcloud ecosystem - during spare time****Author:** Marcel Scherello^{None}**Corresponding Author:**

Providing meaningful apps can be challenge when not being a professional developer.

I created the Audio Player for Nextcloud which is one of the top active NC apps. Following a roadmap between issues,feature requests and version updates required coordination and prioritisation - within the development cycles as well as in private time planning.

With the new "Data Analytics" app for Nextcloud, the focus is shifting from leisure usage to more professional use cases.

I want to outline the major challenges and achievements and motivate others to participate in open source. Even on small scale!

Collaborative applications / 60**Powerful document collaboration: Online, Desktop, Mobile****Author:** Mikhail Korotaev^{None}**Corresponding Author:** mikhail.korotaev@onlyoffice.com

ONLYOFFICE by Ascensio System SIA is an HTML5-based online office for editing text documents, spreadsheets and presentations online and collaborating on them in real time. Relying on three core principles - innovative technology, high security standards and smart service architecture - ONLYOFFICE reached wide applicability in multiple spheres of business and public structures worldwide.

With an established technological ground ensuring browser-agnostic content display, maximum OOXML compatibility and seamless collaborative data transfer, ONLYOFFICE focused on enhancing the workflow flexibility. Rich cloud and local editing experience is equally accessible from web, desktop and mobile clients, with security aspects and third-party integration strategy in mind.

Having a strong focus on integration into Sync&Share platforms, ONLYOFFICE angled its development strategy via technological partnerships with solution developers such as Nextcloud, ownCloud, Seafile, Pydio and others.

This year's presentation will cover:

- Technological ground of ONLYOFFICE solutions;
- Recent versions of ONLYOFFICE online editors;
- Roundup of the recent integration app releases;
- Mobile editing and integration of mobile apps for iOS, iPadOS and Android;
- ONLYOFFICE Desktop Editors integration;
- Data protection and access rights management;
- End-to-end encryption of documents at rest and in editing; ONLYOFFICE roadmap.

Collaborative applications / 69**Digital sovereignty in communication & collaboration****Author:** Andreas Rösler¹

¹ *Kopano GmbH*

Corresponding Author: a.roesler@kopano.com

Keeping control over one's own data means first and foremost, knowing who owns the results of our work - that is the files containing the content generated by us. It is not primarily a question of who the content belongs to. Above all, it is about whether someone has the possibility to keep me away from my data. This conference has long provided answers to the question of how to keep control over your own files.

But what about communication? Email is commodity and has been taking place for years in the clouds of Microsoft or Google - or should, according to the will of the vendors. In addition to communication via e-mail, teams organize themselves in shared calendars, tasks, online meetings or chats. Is it possible to do this in a digital sovereign way?

Solutions are being evaluated in the sciebo project and other scientific institutions. This lecture will show why one should think about digital sovereign communication, what possibilities Kopano, for example, can offer and what learnings out of the ongoing evaluations are, so far.

Collaborative applications / 78

Collabora Online: secure, on-premise collaborative editing

Author: Michael Meeks¹

¹ *Collabora*

Corresponding Author: michael.meeks@collabora.com

Come and hear how Collabora Online can give all the security benefits of much more complex architectures, while allowing de-centralized, on-premise hosting.

Hear about the significant improvements in functionality in the last year everywhere from a myriad of interoperability improvements with UX and Mobile wins, as well as Global Scale integration, and an initial WOPI-like locking implementation. Hear some thoughts on how our simple architecture allows easy deployment, simple scaling, high availability and more for your EFSS.

Digital Sovereignty / 93

Sovereignty or Surrender: Europe has a choice

Author: Christian Schmitz¹

¹ *ownCloud GmbH*

Corresponding Author: owncloud@wwjkd.net

The transformation into the information age is seeing social, industrial and economical changes of gigantic magnitude. The hard reality that quickly manifests itself is that only China and the USA are able to keep up with the transformational pace of the information age, leading to a bi-polar world dominated by the USA and China.

Given that China has just announced (6th Dec 2019) to remove all foreign hardware and software components from its governmental and public computing infrastructure and services, the question is open if a digital cold war is arriving at an ever accelerating pace.

Europe has to make hard choices in all areas if it wants to maintain its sovereignty to be able to orient, decide and act in the information age. The presentation will elaborate developments that led to the current situation and explore options and consequences that could be taken.

A special focus will be taken towards how a CS3 Stack as precursor could lead to a “European Standard Stack” that could technologically lead to a next generation technology leap akin and in the tradition of the www.

Digital Sovereignty / 102

The importance of open source software in delivering large storage services

Author: Alberto Pace¹

¹ CERN

Corresponding Author: alberto.pace@cern.ch

This (short) presentation will address the aspect of on on-premises versus cloud storage and the importance of using open source software in maintaining data sovereignty while delivering large storage services.

Clearly commercially licensed software can also be used as a part of a general complex architecture, but the presentation will discuss the checklist to be validated to avoid vendor lock-in or uncontrolled growing infrastructure costs.

Digital Sovereignty / 124

Digital Sovereignty can only be achieved through pure Open Source Software

Corresponding Author: frank@nextcloud.com

Fabric and platforms for Global Science / 108

From instrument to publication: A First Attempt at an Integrated Cloud for X-ray Facilities

Authors: Rasmus Munk¹ ; Brian Vinter¹ ; Zdenek Matej² ; Artur Barczyk²

¹ Niels Bohr Institute

² MAX IV Laboratory

Corresponding Authors: rasmus.munk@nbi.ku.dk, vinter@nbi.ku.dk

Large scientific X-ray instruments such as MAX IV [1] or XFEL [2] are massive producers of annual data collections from experiments such as imaging sample materials. MAX IV for instance has 16 fully funded beamlines, where 6 of which can produce up to 40 Gbps of experimental data during a typical 5 to 8 hour time-slot, resulting in up to 90 to 144 TBs for a particular beamline experiment.

Scenarios like this calls for solutions that can manage petabytes of datasets in an efficient manner, while enabling scientists with a path of least resistance to define on the fly and subsequent batch processing that often seeks to find needle answers in the data haystack. General outlier detection, pattern recognition and basic statistics just as bin counting are some of the typically tasks conducted during the post analysis phase. To enable scientist with such capabilities, the current challenges calls for an integrated solution that is both able to scale horizontal in terms of available storage, but also be able to make on the fly informed decisions that could potentially either reduce the experimental data stream before it is persistently stored, or enable feedback mechanisms to the instrument itself about which data is of interest to the scientists and that which has no or little value.

The continuous collaboration between the eScience group [3] at the Niels Bohr Institute and the MAX IV facility through their Data Storage and Management Project (DataSTaMP) [4] and European Open Science Cloud (EOSC) [5] participation aims to provide just such an integrated cloud solution to elevate the combined data services available to researchers in general.

The architecture design to enable this is made of two distinct services. The High Throughput Storage System (HISS) and the Electronic Research Data Archive (ERDA) [6]. HISS is a developing distributed system that is designed as a high speed I/O gateway of storage nodes for stream oriented data collections. The system does this through temporary buffering during storage and retrieval of high bandwidth streams, acting in a sense as a front proxy to a subsequent persistent storage location such as a PFS or tape archive system. In addition to being a mere set of buffer nodes that allows for temporal storage reservations, the system is also being designed to allow for an on the fly scheduling of operations to be conducted during the I/O of datasets by scheduling preprocessing tasks on an FPGA accelerator. This enables for both in situ decisions about particular data points mid stream or general data reduction/prefiltering as specified by a user defined kernel, that may also introduce feedback streams to the data provider itself. A provider in this instance could be a beamline instrument at MAX IV.

The system enables access through a REST API that is inspired by and aims to be compatible with the AWS S3 [7] service and commandline tools. To define the computational kernels that are targeted for accelerated computation the proposal is to transpile python kernels into VHDL through an eScience developed toolchain that consists of Bohrium [8], SMEIL [9] and SME [10].

The target of the HISS offloading in our instance, is the ERDA system that is subsequently responsible for retaining the incoming collections, that is either stored GPFS or tape archives. The archive then on top of this provides a rich set of features for both managing and post-processing of data upon being stored. This includes Dropbox like sharing and synchronization in addition to efficient data access to home and collaborative datasets through standard secure protocols like WebDAV over SSL/TLS (WebDAVS), FTPS and SFTP. For processing the service enables processing of existing datasets through a JupyterHub [11] environment with container based JupyterLab [12] sessions for interactive executions of personal or collaborative resources.

It is the aim of this integrated cloud solution to enable both the receipt of instrumentation data streams directly from the source, while allowing user defined decision making to take place before the data is persistently stored. For instance, the user could specify a reduction or statistics kernel that would alleviate the need to schedule such processing upon finishing the experimental phase. Enabling them to immediately interpret the results generated from the computed metadata.

[1] <https://www.maxiv.lu.se>

[2] <https://www.xfel.eu>

[3] <https://www.nbi.ku.dk/Forskning/escience/>

[4] <https://www.maxiv.lu.se/accelerators-beamlines/technology/kits-projects/datastamp/>

[5] <https://ec.europa.eu/research/openscience/index.cfm?pg=open-science-cloud>

[6] <http://www.erda.dk>

[7] <https://docs.aws.amazon.com/cli/latest/userguide/cli-services-s3.html>

[8] <https://bohrium.readthedocs.io>

[9] <https://github.com/truls/lib sme>

[10] <https://github.com/kenkendk/sme>

[11] <https://jupyterhub.readthedocs.io/en/stable/>

[12] <https://jupyterlab.readthedocs.io/en/stable/>

Fabric and platforms for Global Science / 86**Sharing your scientific data with Rucio****Authors:** Mario Lassnig¹ ; Martin Barisits¹ ; Thomas Beermann² ; Cedric Serfon³¹ CERN² Bergische Universitaet Wuppertal (DE)³ Brookhaven National Laboratory (US)**Corresponding Authors:** mario.lassnig@cern.ch, martin.barisits@cern.ch, thomas.beermann@cern.ch, cedric.serfon@cern.ch

Rucio is an open-source software framework that provides scientific collaborations the functionality to organize, manage, monitor, and access their distributed data across heterogeneous infrastructures. Rucio was originally developed to meet the requirements of the high-energy physics experiment ATLAS, and is continuously extended to serve a diverse set of scientific communities, from agricultural to radioastronomy. In 2019, Rucio orchestrated more than an Exabyte of data across a billion files on 130+ data centres.

In this contribution we want to address potential future improvements to scientific data managed with Rucio: (1) transparent provisioning of data for interactive analyses, (2) publishing and annotation of data according to FAIR principles, and (3) selective synchronisation of data for users and desktop applications. A special focus across all three topics will be dynamic adaptation of dataflows to protect global system performance.

Fabric and platforms for Global Science / 84**iRODS and Sync&Share: Getting the best out of both worlds****Author:** Stefan Wolfsheimer¹**Co-author:** Ron Trompert¹¹ SURFsara**Corresponding Author:** ron.trompert@surfsara.nl

Within the Netherlands, iRODS is gaining substantial traction with universities and other research institutes as a tool to help manage large amounts of heterogeneous research data. In this context iRODS is usually used as middleware, providing value through data virtualization, metadata management and/or rule-driven workflows. This is then typically combined with other tools and technology to fully support the diverse needs of researchers, data stewards, IT managers, etc.

While integrations with other RDM tools are facilitated by iRODS' flexibility, a significant amount of work is usually still required to develop and test them with users in their specific context. For this reason SURF – as the collaborative ICT organisation for Dutch education and research – sees a role for itself to spearhead the development of such integrations as that effectively means pooling of resources which lowers the collective development cost and accelerates the pace of adoption.

In this contribution, we will focus on a recent project undertaken by SURF to explore the integration between OwnCloud and iRODS. OwnCloud is an open-source, “sync and share” solution to manage data as an individual or as a research team. OwnCloud is the technology behind two successful existing SURF products: SURFdrive and Research Drive. Offering a GUI, versioning, off-line sync and link-based sharing, it's functionality is in many ways complementary to iRODS. This makes integrating the two technologies attractive, yet there are several challenges in terms of file inventory synchronization, metadata management, and access control.

As an outlook into future work, this integration could be extended to support seamless publication of research data in trusted, long-term data repositories. Existing data publication workflows have many common tasks, but also significant variance in the “details” of how these tasks are strung together and how they need to be operationalized. To address this balance, we are exploring an approach that essentially abstracts data publication tasks into an overarching workflow framework, so as to allow for flexibility yet also benefit from standards and common patterns.

Fabric and platforms for Global Science / 87

Jupyter on Earth: how SWAN is powering CERN use cases

Authors: Diogo Castro¹ ; Jakub Moscicki¹ ; Enrico Bocchi¹ ; Enric Tejedor Saavedra¹ ; Prasanth Kothuri¹ ; Piotr Mrowczynski¹

¹ CERN

Corresponding Author: diogo.castro@cern.ch

SWAN (Service for Web-based ANalysis) is CERN’s general purpose Jupyter notebook service. It offers a preconfigured, fully fledged and easy to use environment, integrates CERN’s storage, compute and analytics services and is available at a simple mouse click.

Due to this simplicity, the Jupyter usage via SWAN has been steadily increasing at CERN in the last years (more than 2000 unique users in last 6 months).

The 1st SWAN User Workshop - held on October 11, 2019 - was the first opportunity to get the overview of the CERN’s user community, to discover typical and unexpected use-cases, discuss user feedback, new features and further evolution of Jupyter Notebooks at CERN. Overall, 19 presentations from the users covered a rich set of topics: Data Analysis for the LHC Experiments, Beams operations, Engineering applications, Education & Outreach use-cases and others. SWAN is a member of the CS3 community, has been deployed outside of High Energy Physics and is being integrated into the new, pan-European research cloud service.

In this presentation we will give a summary of present experience and future evolution of SWAN both at CERN and in a larger context.

Fabric and platforms for Global Science / 103

ELIXIR Cloud & AAI: Standardised and Interoperable Services for Human Data Communities

Authors: Susheel Varma¹ ; Jonathan Tedds² ; Shubham Kapoor³ ; Alexander Kanitz⁴

Co-authors: Steven Newhouse¹ ; Tommi Nyrönen⁵ ; Ludek Matyska⁶

¹ EMBL-EBI

² ELIXIR

³ CSC

⁴ University of Basel

⁵ C

⁶ Unknown

Corresponding Author: jonathan.tedds@elixir-europe.org

Currently, patient data are geographically dispersed, difficult to access, and often, patient data are stored in siloed project-specific databases preventing large-scale data aggregation, standardisation, integration/harmonisation and advanced disease modelling. The ELIXIR Cloud and Authentication

& Authorisation Infrastructure (AAI) for Human Data Communities project aim to leverage a coordinated network of ELIXIR Nodes to deliver a Global Alliance for Genomic Health (GA4GH) standards-compliant federated environment to enable population-scale genomic and phenotypic data analysis across international boundaries and a potential infrastructure to enable 1M Genome analysis across ELIXIR Nodes (member states).

The ELIXIR Cloud & AAI project will lay the groundwork to deliver the foundational capability of “federation” of identities, sensitive data access, trusted hybrid cloud providers and sensitive data analysis services across ELIXIR Nodes by underpinning the bi-directional conversation between partners with the GA4GH standards and specifications and ELIXIR trans-national expertise. The project is also developing a framework for secure access and analysis of sensitive human data based on national federations and standardised discovery protocols. The secure authentication and authorisation process alongside guidelines and compliance processes is essential to enable the community to use these data without compromising privacy and informed consent.

The project therefore, provides mechanisms to enable a globally available curated repository to store bioinformatics software containers and workflows (Biocontainers - GA4GH TRS), a service to discover and resolve the locations of datasets (RDSDS - GA4GH DRS) and distributed workflow and task execution service (WES-ELIXIR/TESK - GA4GH WES/TES) to leverage the federated life-science infrastructure of ELIXIR.

The ambition of the project is to provide a global ecosystem of joint sensitive data access and analysis services where federated resources for life science data are used by national and international projects across all life science disciplines, with widespread support for standard components securing their long-term sustainability. Connecting distributed datasets via common standards will allow researchers unprecedented opportunities to detect rare signals in complex datasets and lay the ground for the widespread application of advanced data analysis methods in the life sciences.

Collaborative applications / 120

From EFSS to Smart Workspace

Author: Martin Meusbürger^{None}

Corresponding Author: martinm@dropbox.com

600 million people and 400000 organizations around the world use Dropbox to collaborate and have been relying on our Cloud Sync and Share solution since 2007. But what does the future hold for Dropbox and its customers? In this presentation, we will give you a peek into the announcements made at our customer conference in September 2019 and the technical challenges we are solving in order to build a smart workspace for individuals and organizations across the world.

How do we incorporate cloud documents such as Google Docs or Dropbox Paper files into the sync and share ecosystem to break down silos, leverage productivity focused machine intelligence to prioritize the workspace, or build data governance capabilities for millions of users. This session will give you a look behind the scene of what is happening at Dropbox today.

Interoperability & APIs / 121

Interoperability Protocols: OCM, CS3APIs, ... and more – Campfire Discussion

Interoperability & APIs / 104

CS3 APIs: from the community to the community

Authors: Hugo Gonzalez Labrador¹ ; Jakub Moscicki¹ ; Alberto Pace¹ ; Massimo Lamanna¹ ; Giuseppe Lo Presti¹ ; Diogo Castro¹

¹ CERN

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

The CS3 APIs is an initiative created in the core of the CS3 community to promote inter-operability between a vast ecosystem of Sync and Share Services and Application Providers to reduce the friction of collaboration across cloud boundaries.

The project is an umbrella to:

1. Promote the usage of a common set of well-known APIs between the CS3 community (OCM, WOPI, ...) and foster the adoption among different stakeholders.
2. Design new APIs to address the challenges faced by many European institutions with the aim at creating decentralized storage and application APIs with state-of-the-art technology to improve data governance, increase user freedom and benefit from the economy of scale.

In this contribution we give an overview of the project: we explain its origins, the current status, the role in the CS3Mesh H2020 project and the future ahead of us to conceive a friction-less decentralized cloud collaboration network.

Interoperability & APIs / 88

CS3 APIs for data science environments

Author: Diogo Castro¹

¹ CERN

Corresponding Author: diogo.castro@cern.ch

CS3 APIs are at the core of the CS3 community and are the reason why multiple platforms will be able to easily connect to different storage vendors and services.

Jupyter, the de facto analysis platform for science, is used at CERN under the name SWAN, integrating CERNBox share capabilities in its UI. With CERNBox's adoption of the CS3 APIs, SWAN will be able to go deeper in its integration with the sync&share platform while, at the same time, becoming compatible with many others.

With this presentation we discuss the possibilities open by creating a CS3APIs storage layer for Jupyter and how this would benefit the larger community.

This presentation can be done as part of a technical discussion/debate.

Interoperability & APIs / 96

WOPI and Custom Apps in the CS3APIs

Authors: Giuseppe Lo Presti¹ ; Hugo Gonzalez Labrador¹

¹ CERN

Corresponding Author: giuseppe.lopresti@cern.ch

WOPI is a well-known standard for Office apps interoperability, used by Microsoft Office 365 and Collabora Online.

In this contribution, we will look into the integration of WOPI-compliant applications in the ecosystem foreseen by the CS3 APIs. In particular, what is the most effective way forward to adopt WOPI? Are vendors ready to interface their Office-like applications via WOPI to cloud storages, as opposed to using custom connectors? Shall a reference implementation of a WOPI connector to the CS3 APIs be provided to ease further integrations?

We aim at a fruitful discussion with the CS3 community to further define the CS3 APIs in this respect.

Meet CS3MESH / 62

The HIFIS Cloud Competence Cluster

Authors: Andreas Klotz¹; Patrick Fuhrmann^{None}; Laura Schollmaier¹

¹ *Helmholtz Berlin*

Corresponding Authors: andreas.klotz@helmholtz-berlin.de, patrick.fuhrmann@desy.de

The HIFIS Cloud Cluster aims to bring existing, outstanding Helmholtz IT Services into a federated cloud environment and to make them available to the whole Helmholtz Community.

This includes many different scenarios, ranging from the Helmholtz-wide accessibility of HPC and HTC Computing clusters to a federation of different existing Next- and Owncloud instances. To achieve the latter, we are interested in finding solutions to mesh Sync&Share solutions from different vendors in order to provide a consistent HIFIS view.

Special attention goes to the creation of possible synergies with other Helmholtz Incubator clusters such as HAICU[1] and with already up and running initiatives like HDF[2].

To get the HIFIS Cloud up and running, we started an initial Helmholtz-wide service survey aiming to create an objective overview of demand and supply. This survey is now used to create an initial Service Portfolio which will then be brought into operation. Furthermore, an organizational rule-set including e.g. processes for the integration of new services as well as definitions for consistent service reviews will be set up. The extension of the initial Service Portfolio will be planned using the experiences made in operation.

To access the services, the HIFIS Backbone Cluster will provide an AAI which will be used in an access layer platform to authenticate and authorize users on many different channels, in a similar way as can currently be seen at EGI or EUDAT.

The workload within the Cloud Cluster is distributed over several Helmholtz Centers.

Together with the other HIFIS Clusters, Backbone and Software, HIFIS will amplify the network of Helmholtz scientists and administration.

HIFIS Website: <https://www.hifis.net>

1: <https://www.haicu.de/>

2: <https://www.helmholtz.de/en/research/information-data-science/helmholtz-data-federation-hdf/>

Meet CS3MESH / 64

FAIR Implementation Profiles: Driving Convergence onto an Internet of FAIR Digital Objects

Author: Erik Erik Schultes¹

¹ *GO FAIR International Support and Coordination Office*

Corresponding Author: erik.schultes@go-fair.org

For communities striving to adhere to FAIR Principles, the collective list of FAIR implementation choices compose the FAIR Implementation Profile (FIP) for that community. The FIPs of numerous communities can be systematically acquired from the FAIR Convergence Matrix, which is an online platform that compiles for any community of practice (“columns” in the Matrix), an inventory of their FAIR implementation choices and challenges (“rows” in the Matrix)

[https://www.mitpressjournals.org/doi/abs/10.1162/dint_a_00038].

In the Convergence Matrix environment, the Community Data Steward is prompted to systematically declare the implementation choices for each of the FAIR Principles. The FIPs are then themselves exposed as FAIR and Open data. Taken together, the accumulated FIPs from the global community give a bird’s eye view of the technology landscape supporting FAIR Digital Objects [<https://github.com/GEDE-RDA-Europe/GEDE/tree/master/FAIR%20Digital%20Objects>]. Based on patterns of use and reuse among the FIPs, strategies for optimal coordination on standards and technologies can be formulated (e.g., maximizing the reuse of existing resources or maximizing interoperability within or between domains). Ready-made and well-tested FIPs created by trusted community-authorized representatives could find widespread reuse and thus vastly accelerate well-informed implementation of FAIR Digital Objects.

Meet CS3MESH / 73

ESCAPE ESFRI Science Analysis Platform

Authors: Zheng Meyer-Zhao¹ ; Michiel van Haarlem¹ ; Yan Grange¹

¹ *ASTRON*

Corresponding Author: meyer@astron.nl

ESCAPE (European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures) addresses the Open Science challenges shared by ESFRI facilities (SKA, CTA, KM3NeT, EST, ELT, HL-LHC, FAIR) as well as other pan-European research infrastructures (CERN, ESO, JIVE) in astronomy and particle physics. ESCAPE actions are focused on developing solutions for the large data sets handled by the ESFRI facilities. These solutions shall: i) connect ESFRI projects to EOSC ensuring integration of data and tools; ii) foster common approaches to implement open-data stewardship; iii) establish interoperability within EOSC as an integrated multi-messenger facility for fundamental science.

To accomplish these objectives, ESCAPE aims to unite astrophysics and particle physics communities with proven expertise in computing and data management by setting up a data infrastructure beyond the current state-of-the-art in support of the FAIR principles. These joint efforts are expected result into a data-lake infrastructure as cloud open-science analysis facility linked with the EOSC. ESCAPE supports already existing infrastructure such as astronomy Virtual Observatory to connect with the EOSC. With the commitment from various ESFRI projects in the cluster, ESCAPE will develop and integrate the EOSC catalogue with a dedicated catalogue of open-source analysis software. This catalogue will provide researchers across the disciplines with new software tools and services developed by astronomy and particle physics community.

The main objectives of ESCAPE Work Package 5, ESAP - ESFRI Science Analysis Platform, are to define and implement a flexible science platform for the analysis of open access data available through the EOSC environment that will allow EOSC researchers to identify and stage existing data collections for analysis, tap into a wide-range of software tools and packages developed by the ESFRIs, bring their own custom workflows to the platform, and take advantage of the underlying HPC and HTC computing infrastructure to execute those workflows.

Our approach is to provide a set of functionalities from which various communities and ESFRIs can assemble a science analysis platform geared to their specific needs, rather than to attempt providing a single, integrated platform to which all researchers must adapt. Deploying an EOSC-based science

platform provides a natural opportunity to integrate with the data and computing fabric this environment encompasses while simultaneously accessing the tools, techniques, and expertise other research domains bring to that environment.

The ESFRI Science Analysis Platform (ESAP) developed through ESCAPE WP5 will provide a flexible and expandable analysis environment for the astronomy and physics community and constitute an absolutely essential resource for the big data challenges of the next generation of ESF/RIs.

Meet CS3MESH / 98

Oracle for Data Science – A Vision for CS3MESH Connector

Author: Peter Szegedi¹

¹ Oracle

Corresponding Author: peter.szegedi@oracle.com

A year ago it was noted “Oracle is now keen on collaborating with the CS3 community as part of its open research engagement campaign...” Indeed, our recently announced Oracle for Research program intends to collaborate with academic researchers closer than ever. It primarily offers researchers, scientists and university-associated innovators access to Oracle Cloud technology and also a global community working to address complex problems and drive meaningful change in the world.

Oracle HPC Cloud and Data Science Platform

Oracle Cloud Infrastructure offers exceptional performance, security, and control for today’s most demanding high-performance computing (HPC) research workloads. Oracle Data Science Cloud – recently acquired – is a collaborative platform for data scientists to build and manage ML models. Oracle supports both Jupyter and Zeppelin notebooks for real-time collaborative research cases, but how about sharing files and folders with co-workers?

Box Connector for Oracle Integration Cloud

At OOW’19 Oracle announced the collaboration with Box [3] that will allow customers to connect their cloud and on-premises Oracle and third-party applications with Box via Oracle Integration. Through this integration, enterprise customers will be able to seamlessly connect applications with Box as their unified cloud content management layer to power secure collaboration and workflows around their most valuable content in the cloud. This is all about the Business User, but how about the Researcher?

Let’s work towards a CS3 connector for Oracle Infrastructure Cloud

The CS3 Community’s vision is that a flexible CS3MESH federation across installations will promote global scientific collaboration and integration, avoiding disconnection of infrastructures. The project, called CS3MESH, aims to build a global interoperable mesh of synchronization and sharing cloud services as part of the European Open Science Cloud by federating cloud storage sites and software providers around the world [4].

Some of the research data though is gradually moving towards public cloud services where certain research workflows can be executed at an attractive price-performance level. Oracle and CERN have been collaborating for more than 15 years in the context of the CERN Openlab initiative in order assess public cloud solutions.

The conceptual diagram above (Fig. 3) depicts a possible scenario where a specific CS3 connector can be built for the CS3MESH project and implemented in the Oracle Integration Cloud could provide a bridge for scientific workloads being executed in a hybrid community/public cloud ecosystem.

Oracle is keen on investigating this deployment scenario and the potential research collaboration opportunity with CS3 further.

Scalable Storage Backends for Cloud, HPC and Global Science / 101

Interoperability as the new frontier

Author: Jean-Thomas Acquaviva¹

¹ *DDN Storage*

Corresponding Author: jtacquaviva@ddn.com

Parallel file systems have reach new height in performance and scale. Storage systems delivering in the 1+TB/sec at a 100+PB scale are now available in several HPC environments including enterprise ones.

It went with efforts, sweat if not tears, but with little surprise since the performance community has such a long track record of success in challenging the new order of magnitude. Recently a strong push has been made on the I/O patterns. Nowadays, storage systems can address sequential and none sequential I/O without major performance impact. This first part of this talk will briefly discussed this recent achievement using IO500 data.

However, it seems that in our data world, the extreme efficiency of the data center is not the right scale of thinking: interoperability is key, data life cycles and flow are the new paradigms.

As organizations are moving toward the multicould environment with complex interactions, the real challenge is to leverage the ultra fast data center in the most efficient way.

The second part of this talk will be focused on the efforts made by HPC players in order to open-up the box and bring inter-operations from external stakeholders in the very core of the data factory.

Scalable Storage Backends for Cloud, HPC and Global Science / 81

Onedata Jupyter Integration Using OnedataFS

Authors: Lukasz Dutka¹ ; Michał Orzechowski¹ ; Bartosz Kryza¹

¹ *ACC Cyfronet-AGH*

Corresponding Author: lukasz.dutka@cyfronet.pl

Onedata is a global high-performance data management system that unifies data access across globally distributed environments and multiple types of underlying storages, such as NFS, Lustre, GPFS, Amazon S3, CEPH, as well as other POSIX-compliant file systems. It allows users to share, collaborate and perform computations on their data. Due to its fully distributed architecture, Onedata enables the creation of complex hybrid-cloud infrastructure deployments, including private and commercial cloud resources. It allows users to share, collaborate and publish data as well as perform high-performance computations on distributed data.

Globally Onedata [1] comprises of Onezones, distributed metadata management and authorisation components that provide entry points for users to access Onedata; and Oneproviders, that expose storage systems to Onedata and provide actual storage to the users. Oneprovider instances can be deployed, as a single node or an HPC cluster, on top of high-performance parallel storage solutions with the ability to serve petabytes of data with GB/s throughput.

Onedata introduces the concept of Space, a virtual directory, owned by one or more users. The Spaces are accessible to users via an intuitive web interface, which allows for Dropbox-like file management and file sharing, Fuse-based client that can be mounted as a virtual POSIX file system, or REST and CDMI standardized APIs. Onedata does not provide users with any physical storage, and each Space has to be supported with a dedicated amount of storage by one or more providers, who are running Oneprovider component. The newly released python library - OnedataFS [2] - allows for even faster access to that data located in Onedata spaces. Thanks to integrating OnedataFS with Jupyter Content Manager API [3], one can not only access the data when using OnedataFS library inside the Notebook but also store the Jupyter Notebooks in Onedata Space.

Currently, Onedata is used in European Open Science Cloud Hub [4], eXtreme DataCloud [5], PRACE-5IP [6], and EOSC Synergy [7], where it provides data transparency layer for computation deployed on hybrid-clouds.

1. Onedata project website. <http://onedata.org>.

2. OnedataFS - PyFilesystem Interface to Onedata Virtual File System. <https://github.com/onedata/fs-onedatafs>.
3. Implementation of the Jupyter Content Manager API, for running Jupyter Notebooks on top of Onedataspaces. <https://github.com/onedata/onedatafs-jupyter>.
4. European Open Science Cloud Hub (Bringing together multiple service providers to create a single contact point for European researchers and innovators.). <https://www.eosc-hub.eu>.
5. eXtreme DataCloud (Developing scalable technologies for federating storage resources). <http://www.extreme-datacloud.eu>.
6. Partnership for Advanced Computing in Europe - Fifth Implementation Phase. <http://www.prace-ri.eu>.
7. European Open Science Cloud - Expanding Capacities by building Capabilities. <https://www.eosc-synergy.eu>.

Scalable Storage Backends for Cloud, HPC and Global Science / 72

Dismantling the Monolith

Author: Crystal Chua¹

Co-authors: David Jericho¹; Michael Usher¹

¹ AARNet Pty Ltd

Corresponding Author: crystal.chua@aarnet.edu.au

Every year at CS3 we all come together to talk about the things we've built and how they've grown - more users, more files, more shares, more storage used than in past years, more features we've added. Last year, we introduced a particularly interesting feature to the AARNet CloudStor ecosystem: S3 gateways as a means of convenient, high-speed data transfer directly to our backend storage. This year, we're going to talk about how that effort led us to a turning point in CloudStor's history, and where we're going from here.

Our earlier (per 2019) deployments of the S3 gateways revealed a critical issue; resource contention between different access pathways to the backend storage resulted in multiple outages across the entire CloudStor ecosystem. This experience highlighted a different, perhaps more pressing problem - that we'd inadvertently built a monolith, and that one component of the system could take out everything else.

Through the next few months, we addressed this issue by splitting worked to split out the S3 backend storage from the CloudStor Prime backend storage, and going one step further, we decided to shard the new environment we were building. In the new model, institutions or groups of institutions are allocated to a separate storage shard, greatly reducing the blast radius of an outage - even if one shard is experiencing issues, customers on the other shards remain unaffected. Additionally, leveraging both Kubernetes as well as the new QuarkDB namespace for EOS, we've managed to cut outage/upgrade downtime from approximately an hour down to a matter of seconds.

This new model has worked so well that we're looking to apply it to the rest of CloudStor, which will be a significant amount of work, but worth the effort. It's been a great run, but perhaps it's time to dismantle the monolith that CloudStor has become, and transform it into something more robust, scaleable, and modular.

Scalable Storage Backends for Cloud, HPC and Global Science / 79

Sharing data between HPC, hadoop and POSIX with EFSS

Author: Olaf Weiser^{None}

Corresponding Author: olaf.weiser@de.ibm.com

SpectrumScale is a software defined parallel file system, which can scale over multiple nodes, networks and block storage types. SpectrumScale R5.x supports Watchfolder, which is somehow comparable to linux inotify, but WF supports to be used over multiple directories and sub-trees of a file system and even over the complete name space recursively.

Based on Watchfolder, NEXTCLOUD and IBM created a combined solution. This integration enables Nextcloud to get notified and so share data, which is ingested or changed by any source connected or directly running on SpectrumScale. So, directly to the file system ingested data from any other application gets fully automated shared or can be managed then by NEXTCLOUD.

Scalable Storage Backends for Cloud, HPC and Global Science / 59

iRODS - Beyond Discoverability: Metadata to drive your data management

Authors: Terrell Russell¹ ; Jason Coposky^{None}

¹ *iRODS Consortium / RENC I*

Corresponding Author: tgr@renci.org

As commercial, governmental, and research organizations continue to move from manual pipelines to automated processing of their vast and growing datasets, they are struggling to find meaning in their repositories.

Many products and approaches now provide data discoverability through indexing and aggregate counts, but few also provide the level of confidence needed for making strong assertions about data provenance. For that, a system needs policy to be enforced; a model for data governance that provides understanding about what is in the system and how it came to be.

With an open, policy-based platform, metadata can be elevated beyond assisting in just search and discoverability. Metadata can associate datasets, help build cohorts for analysis, coordinate data movement and scheduling, and drive the very policy that provides the data governance.

Data management should be data centric, and metadata driven.

Scalable Storage Backends for Cloud, HPC and Global Science / 123

Building a scalable and fully automated public cloud platform based on Seafile and Ceph

Authors: Luc Mader¹ ; Schwabe Danilo¹ ; Verges Martin²

¹ *luckycloud GmbH*

² *Croit GmbH*

Corresponding Authors: luc.mader@luckycloud.de, danilo.schwabe@luckycloud.de, martin.verges@croit.io

Building a scalable public cloud platform for hundred of thousands from scratch can be a difficult and challenging task.

This session will cover a short introduction of luckycloud and how we integrated and fully automated the deployment of Seafile clusters with highly scalable multi petabyte storage backends. We will also

show how we build a reliable and powerful storage backend for our Seafile clusters based on Ceph with our partner Croit.

We will show how to use Croit to easily setup a Ceph based S3 object storage and optimize it for Seafile. Furthermore we show different functions how Seafile and Ceph work together best and speak about the challenges that have to be overcome.

We will also show how we configured the core components to make the setup efficient and stable. This will not focus on automation itself but the automation of Seafile deployment in combination with Ceph and S3 as storage backend.

Collaborative applications / 100

The growth of the Nextcloud apps ecosystem

Author: Olivier Paroz¹

¹ *Nextcloud GmbH*

Corresponding Author: olivier.paroz@nextcloud.com

Nextcloud is designed as a platform, empowering organisations to meet their needs through a large ecosystem of apps covering various enterprise capabilities such as collaboration, office productivity, research, authentication, reporting and more.

This year I will be focusing on the evolution of the app store over the past 12 months, from the perspective of both users and developers.

Collaborative applications / 106

Owncloud Infinite Scale Extensions

Author: Alex Unger¹

¹ *Owncloud GmbH*

Corresponding Author: hello@zyxan.io

How taking advantage of the micro-service architecture enables us at OwnCloud implementing a robust extension system and wrap it all up under a single convenient binary.

Collaborative applications / 92

Migrating CERN e-mail system to open source

Author: Pawel Grzywaczewski¹

¹ *CERN*

Corresponding Author: pawel.grzywaczewski@cern.ch

E-mail service is considered as a critical collaboration system. I will share our experience regarding technical and organizational challenges when migrating 40 000 mailboxes from Microsoft Exchange to free and open source software solution: Kopano.

Collaborative applications / 94**Real-time communication and collaboration in a file sync & share environment****Author:** Björn Schießle¹¹ *Nextcloud GmbH***Corresponding Author:** bjoern@nextcloud.com

What has real-time communication to do with enterprise file sync & share? Already for many years, Nextcloud is way more than just a file sync & share solutions. It is a collaboration platform centered around your data. In order to move the collaboration aspect to the next level, real-time communication was introduced almost two years ago. Nextcloud Talk provides a complete collaboration platform, enabling both real-time and asynchronous communication based on WebRTC. Like always, we build our solution on Open Standards and Open Source software. Real-time communication comes with some unique challenges like instant notifications on events. In order to tackle this challenges Nextcloud has introduced some additional components like a push proxy and a high-performance back-end to handle large groups. Nextcloud Talk improved a lot over the last twelve months and the integration become even deeper, reaching a level which is unique compared to all other solutions. This talk will show you how the real-time communication and collaboration platform Nextcloud can move your productivity to the next level and give you some technical insides how the different components work together.

Meet CS3MESH / 89**ScienceBox as Deployment Model for CS3MESH****Authors:** Enrico Bocchi¹ ; Jakub Moscicki¹¹ *CERN***Corresponding Author:** enrico.bocchi@cern.ch

Docker containers are the de-facto standard to package, distribute and deploy applications on cloud-based infrastructures. Commercial providers and private clouds expand their offer with container orchestration engines (e.g., Kubernetes, Docker Swarm, Apache Mesos), making the management of cloud resources and container-based applications tightly integrated.

A key feature of container orchestration consists in decoupling the container images from the runtime configuration. This simplifies the release management of containerized software (i.e., developers provide a single, immutable image that is uniquely identified by a tag) and also the customization of services to the specific deployment context (i.e., local administrators only maintain the configuration parameters), ultimately enabling the re-usage of one container image in different scenarios.

In this context, CERN Storage provides ScienceBox: An integrated software bundle with distributed storage and computing services for general purposes and scientific use. ScienceBox features 1. EOS, the CERN storage technology for physics data and users' files, 2. CERNBox, the cloud synchronization and sharing platform for science, 3. SWAN, the Jupyter notebook service at CERN, and 4. CVMFS, the software distribution service used by the worldwide computing grid. ScienceBox can run on a single machine with Docker Compose or scale-out across multiple hosts when used jointly with Kubernetes.

ScienceBox is evolving into a modular and fully customizable bundle where each service component can be deployed through Helm charts. This provides all-round configuration flexibility and allows each site part of the CS3MESH to install on-premise the complete stack of ScienceBox services or only a subset to be further integrated with pre-existing services. The ease of deployment provided

by container technologies and the modular architecture of ScienceBox aim at fostering the distribution of open-source scientific software across multiple institutions to increase the interoperability beyond the borders of single clouds and support a collaborative work environment for scientific research.

Meet CS3MESH / 65

Elettra Drive(s) data far and FAIR

Author: Ivan Andrian¹

¹ *Elettra - Sincrotrone Trieste SCpA*

Corresponding Author: ivan.andrian@elettra.eu

As a natural evolution of the internal CIFS-based shares, in 2015 the *Seafile*-based **Elettra Drive** Sync and Share was launched at Elettra Sincrotrone Trieste.

The features that such a product offer enabled *delegation* of the *authorisations* for data access directly to the users and easily enabled cross-area network access and sharing.

During the past four years its use grew and spread through scientists in particular for synchronisation of projects (coding, data, etc.) with several workstations present in different labs or areas (beamlines, offices, ...).

Elettra Drive evolved during the years from a cluster of physical servers sharing a Gluster filesystem to a single virtual machine in a KVM cluster on top of a CEPH storage system.

Future enhancements of this infrastructure will be discussed, in particular its integration with other systems (from general Office-based documents management to collaborative platforms - Zimbra -, to massive FAIR scientific data hosted on the same CEPH cluster and now accessible only via dedicated shares and custom web tools).

Use of the Sync and Share paradigm for Remote Data Analysis As a Service (DAAS), now under exploitation in many EU-funded projects like CALIPSOplus and PANOSC, will also be studied and tested in the coming months; the goal of CS3MESH about federated systems will shorten the distances between data and users, and Elettra Drive wants to play a significant role in this collaboration.

Meet CS3MESH / 77

Drive RENATER

Author: Alexandre Salvat¹

¹ *RENATER*

Corresponding Author: alexandre.salvat@renater.fr

Over the past decade, various Internet players have been increasing cloud data storage offerings with, in some cases, additional features. However, the equilibrium of the economic model is often ensured on the one hand by a usage that becomes time-consuming or depending on the use and, on the other hand, by the exploitation that can be made of data and metadata. To address these issues, many institutions have implemented their own solution, thus constituting a rich functional and application ecosystem.

However, several issues remain:

- How can users of different community platforms share data in an authenticated and trusted way?
- What kind of architecture to meet the needs of hundreds of thousands of users?
- What mechanisms to allow geographic distribution this type of service?
- How to guarantee the minimum levels of security, in particular on the control of access to the service and stored information?

Following the evaluation of several free solutions likely to provide a “drive” type service, GIP RENATER has started a process of building a highly scalable solution in terms of access control, capacity (users, volumes, etc.), distributed deployment and possibly interoperable with other similar services in the community.

The implementation challenges are multiple and concern as much the choice of the solution as the design of the associated technical architecture as well as taking into account the changes and the organization of the MCO.

Meet CS3MESH / 115

OSF and CS3MESH

Authors: Nici Pfeiffer¹ ; Nicole Pfeiffer¹

¹ *Center for Open Science*

Corresponding Author: nici.pfeiffer1@gmail.com

COS builds and maintains open source infrastructure, OSF (<https://osf.io/>) for researchers to manage their research, collaborate on projects, and share their outcomes. As part of the COS mission to increase the openness, integrity and reproducibility of research there is great benefit from a connection with the CS3MESH as a member of the community. Participation in the track would allow COS to learn from the MESH API architects on the API design, from other members of the community of their use cases and platform API architecture, and to share with other members of the community the OSF and its value in the ecosystem.

Research producers and consumers would have maximum benefit if the cloud service tools they use could interoperate seamlessly to share data and metadata across platforms efficiently with little to no effort by the uploader/updater. These workflows should transcend research institutions, storage providers, and geographic locations to not put unnecessary barriers between research collaborations, coordination of research activities, and sharing of research artifacts and outcomes. The way to move forward is with an open-source ecosystem supported by infrastructure with neutral, agnostic APIs and standard protocols. As a tool for providing many of these services to researchers, OSF’s public API using standard schemas belongs in the ecosystem of CS3MESH.

Leveraging existing interfaces like OSF to populate the ecosystem with useful services that are being used by researchers to generate a proof of concept demonstrating the API capability to interoperate between services and demonstrate value to the research lifecycle bringing efficiency gains to the research community. Early adopters will be able to populate the proof of concept with use cases and further build out robust workflow support capacity in the CS3MESH API.

As part of the Meet CS3MESH track, COS can share details on the OSF, the OSF’s API and the possible interoperability with the CS3MESH API. We can demonstrate the many workflows possible, the API needs for bidirectional interoperability, and FAIR metadata standards. Drilling in on the benefits of one streamlined connection with the CS3MESH to connect OSF with institutional repositories, storage locations and data stores and how the full ecosystem can support delivery of this goal.

Scalable Storage Backends for Cloud, HPC and Global Science / 66

Ceph in the real world - what have we learned

Author: Anders Bruvik¹

¹ *Safespring*

Corresponding Author: anders.bruvik@safespring.com

We are a Nordic cloud provider that have been operating Storage as a Service running on ceph clusters for the last three years - providing storage service to the academic sector in Sweden and Norway. In this session I will share some of our experience. I will not go deeply into technical details, but I will rather share some lessons we have learnt about how to build a good team, how we organise development and operations, share a little about some incidents we have had. I will talk about how we do long term planning, and what difference it makes once we start to operate as a service – to be accessed through an API. And last but not least - I will talk about some of the business models, give examples of some customer stories and why it is important to have independent local cloud providers.

Scalable Storage Backends for Cloud, HPC and Global Science / 82

File and Object synchronization in the Hybrid Multi-Cloud world. (NetApp CloudSync)

Author: Mathieu Van Schaik^{None}

Corresponding Author: mathieu.vanschaik@netapp.com

Consider a 100 TB NFS data-set on your on-premises file server that you need to import in Azure Blob storage for further processing using Azure Machine Learning Studio and you need the data there fast.

Also consider having this repeated several times with slightly changed data-sets.

Some might consider this to be a challenge.

With Cloud Sync NetApp offers:

- a fully managed easy-to-use, fast and versatile SAAS-service
- for securely transferring, migrating, replicating and or synchronization data-sets
- between on-premises and cloud environments and vice versa, between hyperscaler cloud providers (AWS, Azure, GCP), and or between different on-premises environments
- and even between different formats (SMB, NFS, Object).

In this session we'll cover the architecture, what works with what, use cases and we'll demonstrate the intuitive web-based user interface.

Scalable Storage Backends for Cloud, HPC and Global Science / 83

Scalable and innovative high-performance HPC solution with BeeGFS backed by NetApp E-Series storage

Corresponding Author:

As technologies continue to evolve, the size and amount of data that your organization must work with is growing exponentially.

Keeping ahead of this data growth requires a scalable and innovative high-performance solution with a lightning-fast, highly reliable IT infrastructure to process, store, and analyse your data. However, the cost and complexity of deploying and operating an HPC infrastructure to manage this critical data can be daunting. Whether you're looking for the origins of the universe, the next big oil reserve, a fool proof way to predict financial markets, or a cure for cancer, ThinkParQ and NetApp can help. The award winning BeeGFS parallel cluster file system backed by NetApp E-Series storage is a proven, integrated solution with a simple, reliable, scalable, and cost-effective HPC infrastructure that keeps pace with your most extreme workloads.

Together, BeeGFS and the E-Series Storage are the optimal combination no matter what size your organization is, and whether you are an experienced HPC guru or are taking your first steps into HPC. The combined storage solution enables clients and researches to easily analyse, discover, share and store data much faster whilst lowering operating costs.

This session will cover in detail how BeeGFS and NetApp E-Series can further accelerate and scale customers storage backends, along with an overview of how Simula Research Laboratory are paving the way for exascale computing in Norway with BeeGFS and NetApp E-Series storage.

Scalable Storage Backends for Cloud, HPC and Global Science / 63

ARCHIVER - Data Archiving and Preservation for Research Environments

Authors: João Fernandes¹ ; Jamie Shiers¹ ; Bob Jones¹

¹ CERN

Corresponding Author: joao.fernandes@cern.ch

Commercial services for Digital Preservation that are currently available have not been proven to scale to the “petabyte region and beyond”, not address the complex data types, often domain-specific, that are needed by many scientific disciplines. In-house services, where they exist have often not acquired the degree of “trustworthiness” verified through certification schemes.

Using a Pre-Commercial Procurement instrument, the ARCHIVER project will introduce significant improvements in the area of archiving and digital preservation services, thus closing critical gaps between what is increasingly required by funding agencies, requested by data creators and eventual (re-)users and what is currently commercially available. ARCHIVER will combine multiple ICT technologies, including extreme data-scaling, network connectivity, federated authentication, service interoperability and business models adapted to the research community, in a hybrid environment to deliver end-to-end archival and preservation services that cover the full research lifecycle. By acting as a collective of procurers, the consortium will create an eco-system for specialist ICT companies active in digital archiving, who would like to introduce new services capable of supporting the expanding needs of research communities but are currently prevented from doing so because there is no common procurement activity for the advanced stewardship of publicly funded research data in Europe. ARCHIVER’s final goal is to allow research group to retain responsibility ownership of their data whilst leveraging best practices, standards and economies of scale.

Scalable Storage Backends for Cloud, HPC and Global Science / 113

Seafile - great but vendor lock-in? Evaluating Seafile openness in practice

Authors: Maciej Brzezniak¹ ; Krzysztof Wadówka² ; Eugeniusz Pokora² ; Norbert Meyer³

¹ PSNC Poznan Poland

² PSNC

³ Unknown

Corresponding Author: maciekb@man.poznan.pl

Sync&share systems are widely used at universities and commercial institutions in order to address data storage and sharing as well as data synchronisation needs. Academic users mostly use open source solutions, while companies, especially SMEs prefer commercial products with paid support. PSNC decided to use Seafile, a scalable, purpose-made, reliable and performant sync&share system. The main motivation for choosing Seafile was its high performance, low overheads and known reliability. PSNC has built the local pilot service in 2015 based on the community version of the software and

expanded it in 2016 to the country-wide production system, box.pionier.net.pl, using Seafile Pro, deployed in a fully redundant setup with two application servers, database cluster and cluster file

system.

PSNC made the decision with awareness that using Seafile in an academic context may bring also challenges such as possible vendor lock-in including difficulties while opting-out from the paid version, obscurity of the code and more complicated integration with systems and applications around etc.

In our presentation we analyse the openness of Seafile in the context of data and meta-data migration. We will also shortly discuss Seafile server API, that can be used to integrate Seafile into software stacks. PSNC maintained two instances of the sync&share service since 2015. Over 2018 we made preparatory efforts in order to integrate both instances, which included migrating user's data (organised in 'libraries') from the old system, based on community version to the new version based on Seafile Pro. While Seafile provides basic tools for exporting data from the system, exporting meta-data such as public share links to data objects and information on sharing data among named users is not supported by this tool.

In our presentation we will discuss features of the more advanced data migration tool we have developed at PSNC and tested on the demonstration and production instances of the services. We will also discuss the experience related to the data migration process specific to Seafile as well as share general comments on the massive data migration. We will also show that the Seafile's internal data and

meta-data organisation enables exploring relationships among data objects and users, including data sharing and other user-level meta-data as well as extract this important information for use of our Seafile.

While the migration tool is still a work in progress (e.g. for now we can only migrate external sharing links, user-level shares are not yet supported), the fact that we can export both data and important meta-data demonstrates that internal architecture is transparent enough to enable Seafile to be used as a technical component of the long-term service. By developing the tool for comprehensive data and meta-data migration we further decreased the lock-in risk related to vendor-specific data organisation.

Another possible approach to data migration in large scale is to use the Seafile API, that is documented and exposes complete functionality needed to build Seafile clients and user interfaces. It is in fact used to implement all Seafile clients including web interface, GUI, CLI and mobile clients and virtual drives. In

particular it could be used to access data objects and explore sharing information. However developing migration solution based on API requires more effort and decision on taking this approach requires more analysis and goes beyond the capabilities of the system administrator team.

Fabric and platforms for Global Science / 75

From Jupyter notebooks to web dashboards for big geospatial data analysis

Authors: Armin Burger¹ ; Davide De Marchi¹ ; Pierre Soille¹

¹ *European Commission - Joint Research Centre*

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

From Jupyter notebooks to web dashboards for big geospatial data analysis

The Joint Research Centre (JRC) of the European Commission has set up the JRC Big Data Platform (JEODPP) as a petabyte scale infrastructure to enable EC researchers to process and analyse big geospatial data in support to EU policy needs[1]. One of the service layer of the platform is the JEO-lab environment[2] that is based on Jupyter notebooks and the Python programming language to enable exploratory visualization and interactive analysis of big geospatial datasets. JEO-lab is set-up with deferred processing, using multiple service nodes to execute the Jupyter client processing workflow starting from data stored in the CERN EOS distributed file system deployed on the JEODPP. In this context, many new applications and services were recently added in order to expand the platform attractiveness towards data scientists and researchers. The presentation will make a tour of the many new features added to the JEO-lab, providing use cases and demos that will include topics like:

• **Sentinel2explorer: an advanced remote sensing application that fully exploits the Jupyter widgets**

It allows users to browse, search and display the full set of Copernicus Sentinel-2 images stored in the JEODPP platform. Selecting any band combination, calculating vegetation or water indexes, creating videos, animations or other types of exports, drawing vector features on top of the displayed images, extracting the full story of the images covering a polygonal feature are among the many functions available, which were created by using at their maximum extent the ipywidgets[3] collection of standard GUI elements as well as some other Jupyter widgets[4]. The outcome is an application that helps end-users to easily navigate inside the many petabytes of Sentinel-2 images available in the JEODPP platform.

• **From interactive to distributed computing of land parcel signatures using HTcondor**

A demonstration of an integrated solution, which comprises interactive and heavily parallel batch processing, to support the new CAP (Common Agricultural Policy) in the monitoring of agricultural parcels at regional or national scale. Using HTcondor orchestrator, a batch extraction of yearly vegetation profiles over millions of polygons is launched and the results are visualized and assessed in the JEO-lab interactive environment. The users can easily view the full story of any by accessing a single, indexed, multi GBytes binary file containing all the results of the batch extraction.

• **ML classification inside a Jupyter notebook using server-side injection of custom Python code**

An example of interactive training for a Symbolic Machine Learning (SML)[5] algorithm inside a Jupyter notebook, exploiting the capability of JEO-lab to execute any custom python code inside the server-side processing chain, via the on-the-fly creation of a Python interpreter inside the server C++ tile engine. Users can profit from the pyjeo[6] EO Python library to execute complex tasks, like image classification or segmentation, thus greatly expanding the analytic capabilities of JEO-lab.

• **Dynamic API to browse and display the full catalogue of Sentinel-2 data in geo-spatial web portals**

The Biodiversity and Protected Areas Management (BIOPAMA) Programme[7] assists the African, Caribbean and Pacific countries to address their priorities for improved management and governance of biodiversity and natural resources. BIOPAMA provides a variety of tools, services and funding to conservation actors in the African, Caribbean and Pacific (ACP) countries. Inside its web portal, a new service provided by JEODPP is implemented: web maps can now show the full JEODPP Sentinel-2 catalogue by using a dedicated REST API that provides discovery, query and fast display capabilities. Inside a Mapbox client, the JEO-lab tile engine dynamically serves TMS layers.

• **Porting notebooks and applications to Voilà to grant access without authentication**

Voilà[8] turns Jupyter notebooks into standalone web-dashboard applications; it supports Jupyter interactive widgets, while not permitting arbitrary code execution, thus posing less security threats. Many applications developed inside the JEO-lab environment are going to be brought into the Voilà world, where they will be accessible without the need for user authentication, and thus greatly expanding the impact of the JEODPP platform and providing an easy way to publish complex interactive visualization environments.

JEODPP platform is a living demonstration of a complex ecosystem of cloud applications and services that allows data scientists' navigation inside a petabyte scale world. In particular, the exploratory visualization and interactive analysis tools in the JEO-lab component can run custom code to prototype the generation of scientific evidence as well as create GUI applications that can be used by 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#.XeDvSuhKg2w>.

[3] <https://ipywidgets.readthedocs.io/en/latest/>

[4] <https://github.com/quantopian/qgrid>

[5] M. Pesaresi, V. Syrris and A. Julea. "A New Method for Earth Observation Data Analytics Based on Symbolic Machine Learning". *Remote Sens.* 2016, 8(5), 399; <https://doi.org/10.3390/rs8050399>

[6] P. Kempeneers, O. Pesek, D. De Marchi, P. Soille. “pyjeo: A Python Package for the Analysis of Geospatial Data” ISPRS International Journal of Geo-Information, Volume 8, Issue 10, October 2019. <https://doi.org/10.3390/ijgi8100461>

[7] <https://www.biopama.org/>

[8] <https://blog.jupyter.org/and-voil%C3%A0-f6a2c08a4a93>, <https://github.com/voila-dashboards/voila>

Fabric and platforms for Global Science / 95

The Cern VM File System as global software delivery fabric

Author: Enrico Bocchi¹

¹ CERN

Corresponding Author: enrico.bocchi@cern.ch

The Cern VM File System (CVMFS) is a service for fast and reliable software distribution on a global scale. It is capable of delivering scientific software onto physical nodes, virtual machines, and HPC clusters by providing POSIX read-only file system access. Files and metadata are downloaded on demand by means of HTTP requests and take advantage of aggressive caching on the client and at intermediate caches. The choice of the HTTP protocol enables the exploitation of standard web servers and web caches, including commercially-provided content delivery networks.

CVMFS was developed to assist the High Energy Physics (HEP) community to run data processing applications on the Worldwide LHC Computing Grid (WLCG). The scale of the deployment for HEP counts more than 1 billion files accessed by 100,000 nodes and cached on 5 replica servers and 400 web caches.

Potential applications of CVMFS, however, are not confined to the HEP world. The recent addition of S3 as data storage backend for CVMFS makes it readily deployable on Amazon Web Services and compatible with the Ceph-provided S3 API. In addition, the specialized DUCC (Daemon that Unpacks Container images into CVMFS) component supports the publication of container images in their extracted form into CVMFS. Such functionality replaces and goes beyond the service provided by container registries (e.g., Docker Hub) as published images are usable by container daemons (e.g., Docker, Singularity) without the need of pulling and unpacking them first.

Fabric and platforms for Global Science / 109

Reva: IOP (Inter-Operability Platform)

Authors: Hugo Gonzalez Labrador¹ ; Jakub Moscicki¹ ; Diogo Castro¹ ; Giuseppe Lo Presti¹

¹ CERN

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

The Reva project is dedicated to create a platform to bridge the gap between Cloud Storages and Application Providers by making them talk to each other in an inter-operable fashion by leveraging on the community-driven CS3 APIs. For this reason, the goal of the project is not to recreate other services but to offer a straightforward way to connect existing services in a simple, portable and scalable way.

Besides that, Reva is also the reference implementation of the CS3 APIs and a proven-to-scale software platform running on CERN premises since 2018 to power the core of the CERNbox service, a massive collaborative cloud platform used by more than 8,000 users accounting for 7 petabytes of

user data. The successful operational experience of running it at CERN and the desire to open its development to the community led to CERN releasing the software alongside the CS3 APIs under a permissive license under the CS3 Virtual organisation in 2019.

In this contribution we explain the roots of the project, the importance of having a reference implementation for the CS3 community APIs and what can Reva bring to your service by connecting to it.

Fabric and platforms for Global Science / 112

Just-in-time data ingestion for distributed metagenomics assembly in multiple clouds with OneData

Author: C.D. Tiwari¹

¹ *EMBL-EBI*

Corresponding Author: cdtiwari@ebi.ac.uk

- Accessing large amounts of data in the cloud poses several problems:
 - Many bioinformatics applications require POSIX access, which does not scale well. Rewriting the application is not always an option.
 - Data sitting in the cloud costs money, whether it's being used or not.
- An ideal solution in many cases would be to provide federated data access to data stored on-premises, with caching in the cloud to reduce latency and use of network bandwidth.
- We have been using OneData to provide federated, secure, scalable access from the cloud to our on-premises storage. OneData offers tuneable caching, block-level access, and many other features that make it very attractive for our use-cases
- I present the tests we have done with OneData, including performance measurements from our first production use-case, distributed assembly of marine metagenomes in multiple clouds. We discuss our plans for rolling it out as a production service in the coming year.

Site reports / 80

Improving user and customer experience in CS3 services

Author: Ricardo Makino¹

Co-author: Rodrigo Azevedo²

¹ *RNP - Rede Nacional de Ensino e Pesquisa*

² *Rede Nacional de Ensino e Pesquisa*

Corresponding Author: ricardo.makino@rnp.br

The National Education and Research Network (RNP) is an organization that plans, designs, implements and operates the national network infrastructure under contract with the Ministry of Science, Technology, Innovation and Communications (MCTIC). A current government program includes five ministries - MCTI, Education (MEC), Culture (MinC), Health (MS) and Defense (MD), and annually define the objectives of the contract and its plan.

The increasing production of scientific data (eg, environmental monitoring, biodiversity databases, a variety of simulation and visualization systems such as climate forecasting, high physical energy data

collection, astronomy and cosmology), cultural datasets and others, brings the need for a scalable, sustainable and high availability IT infrastructure to support these demands, and these facilities must be located in a distributed manner and locations offering telecommunications, energy and safety services, as well as appropriate physical space and infrastructure.

In addition to the question of where these data are stored and processed, there is a need to create different institutions in research programs, such as the Large-scale Biosphere-Atmosphere Experiment in the Amazon (LBA), which is composed of 280 institutional and international, with about 1400 Brazilian scientists, 900 researchers from Amazonian countries and 8 European nations and from American institutions, aiming to study and understand the climate and environmental changes in the Amazon. In this type of community sharing information securely and in compliance with legislation, whether from data collected from sensors, or in the production of articles or articles, is vital, and a cloud file synchronization and sharing solution meet this demand.

edudrive@RNP is cloud file synchronization and sharing offered by RNP to your community, it allows users to sync your data between your desktops, notebooks, smartphones, and tablets and share this data with others users of the service or not, supporting researchers, teachers, and students in research projects and during his academic studies.

The service is developed by RNP, in partnership with Anolis IT and is based on Owncloud software, which acts as the frontend of the service offering a web portal, desktop, and mobile clients that users use to synchronize and share their files. Besides, Openstack Swift is also used as a multi-tenant object storage backend, which provides high scalability, cost savings, and significant resiliency, as a Software-Defined Storage (SDS). Additionally, one of the key service requirements it's the integration with the Shibboleth-based SAML federation for user authentication and authorization.

Although EFSS services are very simple to use, several design gaps have been identified in the service's construction journey on the open-source platforms that provide user interfaces, especially for those non-science and technology users. Also, a lack of visibility was identified from the point of view of who contracts the service, and in some cases was very difficult to report or justify the use of that service.

To mitigate these problems, strategies, and features based on surveys focused on customer experience and user experience were implemented so that their adherence and satisfaction were improved. The purpose of this presentation is to demonstrate the advances made in Owncloud usability, considering a service provider role, from both, user and customer perspective.

User Experience Improvements:

1. First Access / Use Tutorial
2. User Feedback
3. Account Closure
4. Login Notifications
5. User's Guide
6. Integrations (Draw.io, Collabora, etc)
7. Shared Folder Uploader Identification
8. Shared Link View Counter
9. Password Sharing Identification in Web Interface

Also, the enhancements developed to improve the customer experience will be presented:

1. Customer Status Reports
2. Management Portal
3. Management API (under development)
4. Privacy and usage policy management (under approval)

