

Cloud Services for Synchronisation and Sharing (CS3)

Book of Abstracts

28 - 30 January 2019

Roma, Italy

Workshop on Cloud Services for Synchronisation and Sharing

28 - 30 January 2019, Roma, Italy



 cs3.infn.it

International Advisor Committee

- Guido Aben (AARNET)
- Lukasz Dutka (Cyfronet)
- Fabio Farina (GARR)
- Agostino Funel (ENEA)
- Massimo Lamanna (CERN)
- Jakub Moscicki (CERN)
- Tilo Steiger (ETH)
- Stefano Stalio (INFN)
- Ron Trompert (SURFsara)

Organising Committee

- Lukasz Dutka (Cyfronet)
- Francesco Failla (INFN Roma)
- Piero Gatta (INFN Roma)
- Claudio Grandi (INFN Bologna)
- Massimo Lamanna (CERN)
- Mauro Mancini (INFN Roma)
- Jakub Moscicki (CERN)
- Enrico Pasqualucci (INFN Roma)
- Stefano Stalio (INFN LNGS)



DOI 10.5281/zenodo.2545482

<https://doi.org/10.5281/zenodo.2545482>

Cover page photo by DAVID ILIFF. License: CC-BY-SA 3.0

Editor: Belinda Chan (CERN)
Publication date: 2019-01-24

Introduction

This is the "Book of Abstracts" of the 5th CS3 (Cloud Services for Synchronisation and Sharing) conference (28-30 January 2019, Rome).

The 5th edition of this conference marks an important milestone and there are some reflections that we would like to share with you.

Five years ago, the CS3 community simply did not exist. Cloud storage technologies were a convenient extension of traditional storage technologies. Universities, National Research and Education Networks, and Research Centres were all exploring this area by adopting solutions proposed by a few innovative emerging companies, many of them from Europe.

Cloud storage was immediately a success: all the installations have grown by leaps and bounds as the number of active users and number of files per user increased. More importantly though, cloud storage has become a vital part of our day-by-day activities. This success triggered the need for reflection on scalability, data durability and overall sustainability – all of these areas are still progressing, as illustrated by the abstracts collected at this conference.

At the time of the first conference, the actual potential was probably not entirely clear but the conference became the forum to exchange experiences and ideas, and to advance together in this new and exciting domain. Five years hence, and the understanding of cloud storage technology, its maturity and the extension of its usage is continuously increasing.

The sharing potential of these technologies gradually became clearer, surpassing the initial emphasis on synchronisation technologies. Several institutes and some companies (notably OwnCloud, PyDIO, Seafiler, and then NextCloud and PowerFolder) who presented their work to the CS3 community proposed, prototyped and subsequently defined a standard to federate heterogeneous installations. This standard, Open Cloud Mesh, is currently one of the potential cornerstones of a new pervasive infrastructure for collaboration and sharing in European research. Such an infrastructure, however, is yet to emerge.

As well as innovative approaches from the academia and the SME world, CS3 has attracted established companies providing storage solutions. Larger and more efficient data stores mean that the "content" of these repositories is changing as is their usage. Used initially for personal data convenience and backups, more and more users now actually store and share scientific data along with the algorithms for data analysis. These sharing capabilities make cloud storage the natural environment for scientific data, opening the door to new challenges, namely on how to "connect" the data with analytics procedures whilst avoiding unnecessary data replication and transfers.

The emerging vision is that CS3 services will become the natural home for all the computing activities of our scientific communities. This includes the "daily-usage" services, starting from office applications (such as OnlyOffice and CollaboraOnline), to data processing and scientific data analysis (as reported by CERN, JRC and AARNET).

The CS3 community is now a solid, pro-active community. It continues to attract new users: for example, leading HPC facilities are looking to CS3 for possible solutions to make it easier for their users to get the data in and out from the HPC systems. At the same time, traditional on-premise sync&share services are becoming complete environments with advanced collaborative services side-by-side with the original data functionality.

What is next? The network of trust within CS3 is the main asset of our community. The accumulated know-how may allow a new class of federated services to emerge. Such a federation could provide a useful service for current CS3 users and beyond. We believe CS3 is the environment where more and more companies will grow their products and invent new ones.

This is probably the moment to capitalize on the remarkable work done in several regions in Europe, Australia and North America. This community certainly has the momentum to transform the activities of these past years into the sustainable fabric of cloud storage of the future.

Massimo Lamanna and Jakub T. Moscicki



Contents

Conference Keynotes	1
Open Deep Learning and Data Management of large datasets in hybrid Clouds: a practical view	3
Building an EOSC in practice	4
The Magic Pocket - How Dropbox is handling large data sets	5
Scaling Tightly Coupled Algorithms on AWS	6
Sharing and Collaborative Platforms	7
CERNBox as the CERN Apps Hub	9
Sharing the Success, Pain and remedies of a Sync-Software-Solution	10
The strengths of the Nextcloud apps ecosystem	11
Extensive & Intensive Integration of Collaborative Editing in Sync&Share Platforms	12
Collabora Online: Reduced Latency and Other News	13
Building the Layers of File, Editing, and Co-editing Security	14
Kopano and Ceph - High available - High Scalable groupware & data platform	15
A side serve of metadata: evolving AARNet's tools for data packaging	16
Site Reports	17
CS3 Community Report Summary	19
Connecting Clouds	20
NWU: Expanding Nextcloud to supply an ever-growing demand	21
DevOps with sync and share	22
Synchronization and Sharing Services in IHEP	23
SWITCHdrive: Concerns and expectations of our S&S service	24
GARRbox: status and future directions	25
Administering Users, Guests and Projects in our sync & share service	26
Sync & Share on the INFN Corporate Cloud	27
File Sync&Share Products for Home, Lab and Enterprise	29
Nextcloud: From EFSS to CCP	31
Seafile Development and Outlook	32

Cubbit: the distributed cloud.....	33
Is “zero-knowledge” privacy achievable in the distributed cloud? An in-depth analysis of encryption architecture of Cubbit’s sync&share service.	34
Syncthing - Introduction, Technology and Use Cases	35
Introducing Pydio Cells	36
Enterprise File Sync & Share by Oracle	37
Data Science: Applications and Infrastructure	39
Exploring Callysto	41
Image and data analytic pipeline for studying single-cell signaling dynamics in cancer	42
Advanced geo-spatial data analysis with Jupyter	43
DLCM: A Swiss solution for preserving research data on the long term	45
SURF Reseach Drive - JupyterHub integration within OwnCloud	46
SWAN and its analysis ecosystem	47
Application store for data analytics	48
CloudStor SWAN: Data Processing and Analysis Challenges in the Cloud	49
CloudStor Minio: Improving S3 performance in CloudStor	50
Open Data Ecosystem and CS3	51
Open, FAIR, and GO FAIR	53
Building a Scalable, Flexible, Interoperable Federated Cloud to Power Australian Research	54
Open Data analysis with EOSC-hub services.....	55
Hybrid Cloud storage deployment models	56
Connecting and Supporting the Research Workflow with OSF	57
Sync/share Technology&Research	59
Designing and Implementing for Scale, Distribution & Async.....	61
On the way to OCM 1.0, a first stable version ready for mass adoption	62
A journey to the mesh : microservices-based sync and share using gRPC and Protobuf	63
Scalable Storage Backends for Cloud, HPC and Global Science	65
Anomaly detection of large scale distributed storage system based on machine learning	67
From performance monitoring to predictive maintenance.	68
Converge heterogeneous legacy filesystems at the Pawsey Super-computing Centre	69

dCache - active storage for science and cloud	70
Cynny Space Cloud Object Storage integrated with Network Intelligence platform Vegoia.....	71
Creating a load balancer for Ceph backed S3 service using open-source tools.....	72
iRODS: Managing Data from the Edge to HPC	73
Onedata - Managing Data and Metadata in Hybrid Clouds	74

Conference Keynotes

Open Deep Learning and Data Management of large datasets in hybrid Clouds: a practical view

Davide Salomoni¹

¹ *INFN*

Corresponding Author(s): davide.salomoni@cnaf.infn.it

Machine learning methods are routinely used to explore large data sets, handle web searches, recommend products, perform voice recognition, recognize features in pictures and videos, and in countless other applications. Neural networks have been around for a long time. But in recent times the pervasive availability of huge amount of data, associated to large numbers of hidden parameters and to advancements in computing, storage and network technologies, has led to a rapid development of methods to train, optimize and exploit machine and deep learning techniques.

In scientific domains, resources such as data sets and computing or storage facilities are typically distributed across multiple providers, often located in different countries. The resources might be dedicated or opportunistic, and could be offered through worldwide computing Grids, via private or public Cloud infrastructures, via single data centers, or through a composition of the above.

This plethora of opportunities in accessing and exploiting the resources needed by deep learning methods brings about a number of issues. In this keynote, Davide Salomoni from the Italian National Institute for Nuclear Physics (INFN) will discuss how a “Deep learning as a Service” spanning heterogeneous, public or private infrastructures, could be designed, drawing on the experience gained by several international projects and initiatives. In particular, the keynote will cover topics such as orchestration of data and computing resources, data placement, automated instantiation of deep learning solutions and publication of reusable, trained data sets into catalogues.

Building an EOSC in practice

Isabel Campos Plasencia¹

¹ *Consejo Superior de Investigaciones Científicas (CSIC) (ES)*

Corresponding Author(s): isabel.campos@csic.es

The European Open Science Cloud (EOSC) is a coordinated effort to provide an European, open environment for scientific data and related processing that promotes convergence of infrastructures and thematic services provided at national or European level. Realizing this level of support requires coordinated actions to provide services to access, process and compute on research data in a scalable way at the European level, ensuring research data is broadly exploitable as a public good.

For the EOSC implementation to succeed in providing such level of integration across research infrastructures, a number of technical integration areas need to be addressed. In particular a higher level of harmonization is needed in policies for data access, policies for resource allocation but also those for enacting policies for authentication, authorisation, and security incident response.

In this presentation we will make a comparative analysis between existing policies and recommendations from EOSC and specifically the results of H2020 projects such as EOSC-hub, AARC and INDIGO-Datacloud. In particular it will address harmonization of resource allocation in the EOSC context, organization of data access scalable in a federated environment, enabling security among collaborating infrastructures and Authorization and Authentication of users in a consistent and scalable way.

Conference Keynotes

The Magic Pocket - How Dropbox is handling large data sets

Said Babayev¹ ; Martin Meusburger¹

¹*Dropbox International*

Corresponding Author(s): said@dropbox.com, mmeusburger@dropbox.com

600 million people around the world use Dropbox to collaborate the way they want, on any device, wherever they go. With 300,000 businesses on Dropbox Business, we are transforming everyday workflows. In this keynote session, Said Babayev, Site Reliability Engineer, from Dropbox's storage team and Martin Meusburger the Head of Central and Northern Europe will give us a peek of the challenges they are facing to handle more than 500 PB of data stored in Dropbox and how collaboration is transforming entire industries. The team has to tackle the challenges around durability, availability and cost in Dropbox's in-house multi-exabyte storage system, making Dropbox one of the 10 largest cloud providers globally.

Scaling Tightly Coupled Algorithms on AWS

Scott Eberhardt¹

¹ *Amazon*

Corresponding Author(s): hpcscott@amazon.com

Many HPC, tightly-coupled workloads scale linearly on AWS in the Gustafson-Barsis sense. This presentation will discuss some large CFD calculations (400M cells and more). The ability to tailor the architecture to the application allows these applications to run at speeds similar, and sometimes faster, to on-prem facilities. In particular scaling of the openFOAM (simpleFoam) will be explored, including where various limits are reached for several architectures. Strategies that allowed us to work around these limits will be presented.

Sharing and Collaborative Platforms

CERNBox as the CERN Apps Hub

Giuseppe Lo Presti¹ ; Michal Kwiatek¹

¹*CERN*

Corresponding Author(s): lopresti@cern.ch

This contribution will present the work being done at CERN to enable running multiple diverse applications on top of CERNBox, the CERN cloud storage solution.

Powered by EOS and ownCloud, CERNBox has evolved in the past few years as the central storage platform for the CERN users community, with an ever-growing user base that is now beyond 15K users. The type of data stored and shared within CERNBox spans from office-like documents to High-Energy Physics specific formats like ROOT.

On the other hand, CERN is reviewing its portfolio of applications offered to the user community, with the aim to incorporate open-source and - wherever possible - web-based solutions, which are gaining more popularity. CERNBox is the ideal hub where a number of online applications have already been integrated, from simple viewers to collaborative editors for office formats (Microsoft, ONLYOFFICE, Draw.io), to more sophisticated apps for data analysis (SWAN, Spark).

While each of those apps required specific tweaks to make them work with CERNBox, we will present the vision for a generic application layer, which will enable diverse apps contributed by different user communities to be seamlessly integrated in CERNBox.

Sharing the Success, Pain and remedies of a Sync-Software-Solution

Christian Schmitz¹

¹*ownCloud*

Corresponding Author(s): cschmitz@owncloud.com

This talk will have a very concise train of thought and revolve around the problems of success in a non technical manner focused around the owncloud ecosystem and its many shortcomings both technical and non-technical.

The key to this is not, as may be expected, a oc vs nc talk, but actually a focus on how success creates problems, more success creates even more problems and higher expectations to the point where seemingly everything fails. Because that is the moment, the darkest of the night, where it only goes uphill, and how the next day can and will be one brighter than the previous one, IF one grabs the opportunity of reshaping the days schedule. This is where we will dive into aspects of future usecases and how ecosystems can scale beyond initially anticipated and failures can be turned into successes.

I will adress the talk with technical “bits” to appease the entire crowd, the focus here is that the success of a service delivery is a big one, and that high level retrospective view is an important one for success, not only the bits and bytes.

The strengths of the Nextcloud apps ecosystem

Olivier Paroz¹

¹*Nextcloud GmbH*

Corresponding Author(s): 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 talk presents the most popular applications in each area, picked from the about 200 applications available from the Nextcloud app store.

Extensive & Intensive Integration of Collaborative Editing in Sync&Share Platforms

Author(s): Mikhail Korotaev¹

Co-author(s): Olga Golubeva ¹

¹ *Ascensio System SIA*

Corresponding Author(s): mikhail.korotaev@onlyoffice.com

ONLYOFFICE by Ascensio System SIA is an HTML5-based online suite for collaborating on text documents, spreadsheets and presentations online. Here, maximizing format compatibility, establishing browser-agnostic content display and optimizing real-time data transfer in co-authoring are the important core principles in building applicable editing software.

But once the solution is tailored to the industry standards internally, deeper integration in popular Sync&Share solutions becomes a reasoned next stage. That strategy is defined by both end user-side demand and the results the 'division of labor' practice has delivered in cloud software development so far.

However, the design lays not only in own on-demand development of the integration apps, but also in closer strategic collaboration between the vendors themselves, creation of the accessible documentation and joint support, and encouraging the developer community to take part in the efforts.

This year's presentation will cover:

- Review of the technical basis of ONLYOFFICE solutions;
- Providing open API to allow creating third-party connectors;
- Recent achievements in integrating the suite with Sync&Share platforms (e.g. Nextcloud and ownCloud);
- Connecting desktop environments to cloud sharing platforms;
- Security technologies in collaborative editing;
- ONLYOFFICE integration roadmap.

Collabora Online: Reduced Latency and Other News

Jan Holesovsky¹

¹ *Collabora Productivity*

Corresponding Author(s): kendy@collabora.com

The last year I have introduced the Collabora Online and how it integrates into many File Sync&Share Products to create a powerful, secure, real-time document editing experience.

This year I would like to build on top of that presentation and talk about the recent improvements: How we have reduced the typing latency considerably, or what we have implemented in the protocol for even tighter integration into the File Sync&Share solutions.

Building the Layers of File, Editing, and Co-editing Security

Author(s): Korotaev Mikhail¹

Co-author(s): Olga Golubeva ¹

¹ *Ascensio System SIA*

Corresponding Author(s): mikhail.korotaev@onlyoffice.com

When document processing and storage is carried out in private or public cloud environments, the data is more vulnerable to external access than in a desktop system. Storage protection instruments are not always a sufficient measure as security of this information can still be breached by insiders or simply mishandled.

A possible and yet the only solution is client-side protection of the information itself, which is the data transfer in online editing, including collaboration between multiple users.

ONLYOFFICE project has developed the next generation of its desktop client that is not only able to encrypt documents in OOXML and OpenDocument formats end-to-end, but also connects to the cloud and performs immediate encryption actions in online editing and collaboration, protecting the data in transfer. The mechanics of this technology are based on blockchain.

This technology, combined with access regulation backed by JSON Web Token, adds an extra layer of data security which compliments the efforts in storage protection and monitoring.

This talk will cover the following topics:

- end-to-end encryption of documents,
- reliable password storing and transferring using blockchain technology with asymmetric encryption,
- secure file sharing and encrypted real-time co-editing within ONLYOFFICE collaboration platform,
- prospects of technology integration with third-party cloud platforms.

Kopano and Ceph - High available - High Scalable groupware & data platform

Michael Kromer¹

¹ *Kopano*

Corresponding Author(s): m.kromer@kopano.com

Mike will show how to build a cloud service with **Ceph**, **MySQL** group replication and the corresponding **Kopano** functionality that provides 100.000 users with files, calendar data, contact data and email.

800 TB of data distributed on a multiserver architecture as well as being redundantly mirrored between two data centers. Using S3 protocol with rados as backend for attachments and file storage at the same time as the environment providing true scalability by containers shows a dynamic and highly scalable setup ready for some serious work. Providing high availability and scalability with extra components such as **haproxy** and **Pacemaker** are included as well.

Technologically Kopano works here together with **SUSE** and **owncloud**.

Michael Kromer, VP Technology and Architecture Kopano

A side serve of metadata: evolving AARNet's tools for data packaging

Gavin Kennedy¹ ; Ingrid Mason¹ ; Peter Sefton²

¹ AARNet

² University of Technology Sydney

Corresponding Author(s): gavin.kennedy@aarnet.edu.au

A data package is simply the combination of a dataset with metadata that describes the dataset. The purpose of the data package is to provide sufficient contextual data to make the dataset usable to others. It therefore becomes the basis of a loosely coupled data management platform in which the information is in the dataset, not in the platform. A receiving platform just needs to be able to access and interpret the package.

But of course this is where the complexity, and our story, begins. At AARNet we partnered with Intersect Australia and Western Sydney University to implement a packaging plug-in tool called Collections [1], an extension of a tool called cr8it [2], an ownCloud plugin, which in turn was developed on top of the Baglt [3] file manifest specification. Our goal was simple, give the researcher a tool to support data sharing and repository deposits, basically to assist with data publication (including open data publication). So we made it available through our CloudStor [4] console, where the data is visible to the user.

This talk will discuss the next steps in research data packaging. As our CloudStor service has become ubiquitous amongst Australian researchers, the need has grown for a tool that supports not only sharing, but open data publishing and archiving of data, placing data management and curation at the front of the Research Data Management Lifecycle.

For our next iteration of the Collections tool we are considering the work of two other Australian initiatives. The first is the State Library of New South Wales use of Baglt in their PanDA [5] development for ingest as part of the archiving workflow. The second is the University of Technology Sydney on a new data packaging specification that also builds on the Baglt packaging specification and makes data easier to disseminate and consume. DataCrate [6] formats its machine-readable metadata in JSON-LD and follows the schema.org vocabulary, making data packages instantly consumable for semantic driven workflows and instantly consumable and indexable for discovery platforms. DataCrate also creates rich human-readable metadata in the form of web pages that also describe the “who, what, where” metadata that helps to make user understanding of the data and its provenance.

By refining our packaging plugin AARNet can make our CloudStor service interoperable with institutional repositories and digital archives, or maybe, just maybe, make CloudStor the repository and the archive.

[1] <https://support.aarnet.edu.au/hc/en-us/sections/115000264274-CloudStor-Collections>

[2] https://github.com/IntersectAustralia/cr8it_doc

[3] <https://en.wikipedia.org/wiki/Baglt>

[4] <https://support.aarnet.edu.au/hc/en-us/categories/200217608-CloudStor>

[5] <http://www.sl.nsw.gov.au/blogs/panda-digital-asset-ingestion-scale>

[6] <https://github.com/UTS-eResearch/datacrate>

Site Reports

Site Reports

CS3 Community Report Summary

Dr. Tilo Steiger¹ ; Stefano Stalio²

¹*ETH Zuerich*

²*INFN*

Corresponding Author(s): stefano.stalio@lngs.infn.it

Site Reports

Connecting Clouds

Crystal Michelle Chua

Corresponding Author(s): crystal.chua@aarnet.edu.au

At the CS3 conference in 2018, a panel of very wise people gravely informed us that sync and share is dead.

Moving beyond simple sync and share, over this past year the focus of AARNet's CloudStor team has been towards building a platform that further enables our users to do more with their data without leaving the system. From collaborative document editing to data analytics with Jupyter Notebooks, we're now investigating connecting CloudStor with the Australian NeCTAR Research Cloud, combining compute and storage to provide researchers the best of both worlds.

This talk covers the development process of an API that builds on top of our existing tenant portal, which is designed to allow institutions access to manage their users and provide self service functionality. Of particular relevance to this use case is the ability to create and manage group drives - shared project spaces for any given group of researchers, with their own quota allocations.

By exposing this functionality via an API, we aim to enable users to provision virtual machines on NeCTAR's OpenStack cluster with automatically attached storage from CloudStor via the WebDAV interface. This allows researchers to access data in a group drive directly from a file system, instead of having to download, process and upload. In cases where a group drive does not already exist within CloudStor, one can be automatically provisioned for the corresponding NeCTAR project, further facilitating the objective of data life cycle management along the entire arc of research – from data creation to eventual publication.

Site Reports

NWU: Expanding Nextcloud to supply an ever-growing demand

Hans Erasmus

Corresponding Author(s): hans.erasmus@nwu.ac.za

A well-known but underrated fact is the growth-rate of research data. Any enterprise can be caught off guard if they do not keep their finger on the pulse. But data on hard drives is useless, if it is not accessible. With global collaboration in research, a demand for anywhere-availability, security and legal compliance, many enterprises and academic institutions are left with a major headache to supply on all of these demands.

I hope to provide a short overview of the road up to now, but more importantly, the road ahead with components like Global Scale Architecture and OnlyOffice, to name but a few.

Site Reports

DevOps with sync and share

Fergus Kerins¹

¹ *SWITCH*

Corresponding Author(s): fergus.kerins@switch.ch

SWITCH as the national research and educational network runs an ownCloud service called SWITCHdrive for all the institutions of higher education. In order to run our service in a structured and time efficient way we have adopted a DevOps approach to operations, based on tools such as ELK, Grafana, Docker, Ansible and Python.

DevOps pillars we use:

1. Accept failure as normal mode of operation: failure in our system happens, we try to deal with it gracefully.
2. Implement gradual changes: controlled rollout of features and bug fixes allow to identify potential issue source.
3. Leverage tooling and automation: in order to facilitate change and recovery and operation tasks we will look at some tooling and automation solutions
4. Measure everything: most importantly of all is the monitoring which allows us visibility to see how the system is performing and plan for the future

In this talk we will look at how this approach has worked and what challenges we face.

Site Reports

Synchronization and Sharing Services in IHEP

Qi mengyao¹ ; Tao Lin¹

¹*Institute of High Energy Physics, CAS*

Corresponding Author(s): qmy@ihep.ac.cn

In order to provide a secure personal storage space for IHEP users, IHEPBox is built with the open source ownCloud. The platform integrates an LDAP database (AD domain) ,and we use EOS to storage users data, deploy the onlyoffice to provide office online services , and setup jupyter services to provide users with data analysis based on WEB, which will meet the needs of interactive data analysis for high-energy physics experiments.

Site Reports

SWITCHdrive: Concerns and expectations of our S&S service

Urs Schmid¹

¹ SWITCH

Corresponding Author(s): urs.schmid@switch.ch

SWITCHdrive is the sync & share service for universities and other institutions of higher education in Switzerland. We are using a branded version of ownCloud V10 which runs on top of our IAAS offering (SWITCHengines) based on OpenStack/Ceph and are serving nearly 30'000 users at 43 institutions.

I will present how the service evolved since its start 5 years ago and point out, what have been (and mostly still are!) the biggest concerns of our customers and users when they use SWITCHdrive. In fact, our customers - which are the IT managers of the Swiss universities - have a completely different view about these concerns than our users - which are the members of the Swiss universities, so mostly employees and students.

I will show our plan how to solve this dilemma and will give a short summary of our cloud survey from fall 2018 that involved all Swiss universities and which produced some surprising results.

GARRbox: status and future directions

Fabio Farina¹ ; Pasquale Mandato¹ ; Paolo Velati¹

¹GARR

Corresponding Author(s): fabio.farina@garr.it

GARRbox is the synch and share storage service built and operated by Consortium GARR, the Italian National Research and Education Network. GARRbox, built on top of OwnCloud, has been designed in 2016 as the GARR response to a specific commitment from the Italian Ministry of Health for supporting the needs of the biomedical research community. The service main focus has been since the liaison between the ease of use and security and resiliency.

At a later stage Universities, Research Institutions and Collaboration have been allowed to access the service.

GARRbox developed a rich authorization framework to control resource access policies through domain specific language ACLs and authority delegation. These features let principal investigators and local sysadmins select access criteria and quota assignments.

In this talk we will give a quick overview on the service status both from a technological perspective showing how Ansible, VMware, OpenStack and Docker are used during the service delivery chain, and from the management point of view, by discussing the challenges we face in supporting the growing user community. Then we will discuss the evolution of the service, focusing on the authentication and authorization features for the future deployments: an improved security SAML-only registration chain, a richer authorization automation language supporting also dynamic differentiated attributes in case of federated account linkings, and a better telemetry approach to the service monitoring.

Administering Users, Guests and Projects in our sync & share service

Renato Furter¹

¹SWITCH

Corresponding Author(s): renato.furter@switch.ch

The past

SWITCH as the national research and educational network runs an ownCloud service called SWITCHdrive for all the institutions of higher education. As we have all user on the same instance we needed an administrator role on tenant level. As we already have an external website where all the users need to register for the service, we could build on top of that an administrator interface where privileged users were able to give away vouchers to invite external people or upper quota for a single user.

The present

Now as we are now running this service for about 4 years it turned out that having an application user database and the ownCloud database is sometimes a little hard to synchronise. Also is the login process with the registration on an external site not the most convenient for the user. So we might change the login to our standard federated login, which is based on Shibboleth and as the Shibboleth integration of ownCloud got also better in the last 4 years this seems to be a better solution. Also user are asking for project based disk-space instead of person-bound data shared with others.

The future

We plan to have the whole user administration portal revamped and will consider whether we want to integrate that in SWITCHdrive as an app. Also add the group/project based folder possibility to that portal.

The audience will see what we have developed and why and how. There will probably a small demo of the administrator functions of the current state at that time.

Site Reports

Sync & Share on the INFN Corporate Cloud

Author(s): Doina Cristina Duma¹ ; Giacinto Donvito² ; Marica Antonacci¹ ; Stefano Nicotri² ; Stefano Stalio¹

Co-author(s): Alessandro Costantini³ ; Alessandro Italiano² ; Antonio Budano⁴ ; Claudio Grandi⁵ ; Davide Salomoni⁶ ; Diego Michelotto² ; Domenico Diacono¹ ; Emidio Maria Giorgio² ; Federico Zani⁷ ; Lorenzo Chiarelli⁸ ; Lucio Strizzolo⁹ ; Massimo Pistoni¹⁰ ; Riccardo Gargana¹⁰ ; Roberto Valentini¹¹ ; Stefano Zani¹² ; Tomaso Tonto¹⁰ ; Vincenzo Spinoso¹

¹ INFN

² INFN - National Institute for Nuclear Physics

³ INFN - CNAF

⁴ Roma Tre Università Degli Studi (IT)

⁵ INFN - Bologna

⁶ Università e INFN, Bologna (IT)

⁷ INFN - Sezione di Roma Tor Vergata

⁸ GARR & INFN

⁹ INFN - Sezione di Trieste

¹⁰ INFN - Laboratori Nazionali di Frascati

¹¹ INFN - Sezione di Bari

¹² INFN - CNAF

Corresponding Author(s): donvito@infn.it, marica.antonacci@ba.infn.it, nicotri@infn.it

INFN Corporate Cloud (INFN-CC) is a geographically distributed private cloud infrastructure, based on OpenStack, that has recently been deployed in three of the major INFN data centres in Italy. INFN-CC has a twofold purpose: on one hand, its fully redundant architecture and resiliency make it the perfect environment for providing critical network services for the INFN community; on the other hand, the fact that it is hosted in modern and large data centres makes it the platform of choice for a number of scientific computing use cases.

Sync and share services on INFN-CC exploit the geo-distributed OpenStack Swift setup, the cross-domain DNS load balancing, and remote data synchronization technologies for resilient installations that can span across multiple sites.

Also, templates for deploying natively encrypted sync & share services are available for immediate use for scientific collaborations and communities.

We present the INFN-CC hardware and software architecture, the challenges and advantages of this solution, together with the first sync & share use cases that we have already implemented.

File Sync&Share Products for Home, Lab and Enterprise

Nextcloud: From EFSS to CCP

Frank Karlitschek¹

¹*Nextcloud*

Corresponding Author(s): frank@nextcloud.com

Nextcloud was started as a classic enterprise file sync and share solution. It's an 100% open source software that can be used for on-prem file sync and share services. But nowadays the enterprise file sync and share market is growing into the content collaboration platform market. Nextcloud is a leading solution in this area. This talk covers the latest innovations and discusses the interesting new features of the latest releases.

Seafile Development and Outlook

JiaQiang Xu

Corresponding Author(s): xjqkilling@gmail.com

In the year of 2018, many exiting new features have been added into Seafile. In this presentation, we'll go through new features such as department hierarchy, group-owned libraries, multiple storage backends, finer-grained share permissions.

We'll also present the prototype of Seafile Docs, an online document collaboration product based on Markdown and rich-text format. This product combines an easy-to-use Markdown/rich-text editor with many efficient collaboration work flows, such as versioning, diff, review. The final goal is to implement an user experience similar to Dropbox Paper, while being more focused on professional collaboration.

Cubbit: the distributed cloud

Lorenzo Posani¹ ; Marco Moschettini¹ ; Alessio Paccoia¹

¹ *Cubbit*

Corresponding Author(s): lorenzo.posani@cubbit.io

Since its debut decades ago, “the Cloud” has been stagnating in the same centralized architecture: one single webfarm serving millions of users. Parallely, the advent of peer-to-peer (p2p) solutions, from file sharing to VPN, showed that a distributed architecture can outperform a centralized one, but the lack of control often led to other inefficiencies.

Cubbit is a hybrid cloud architecture comprised of a network of p2p devices (the “Swarm”) coordinated by a central optimization server. By intelligent management of bandwidth, CPU time, and storage contributed by the Swarm, Cubbit recycles these unused raw resources to offer performant, encrypted, and green cloud services. Operating costs are strongly reduced with respect to the centralized paradigm, allowing for Cloud services that are sustainable without the need for monthly fees from the consumer.

Here, we will present the architectural principles behind Cubbit “sync&share” cloud service. A particular focus will be dedicated to the environmental impact of the Cloud, which is today estimated to consume as much as the whole Brasil. Cubbit reduces the impact per Gigabyte of a factor 10 compared to standard webfarm hardware, providing an example of how a radical shift of paradigm can benefit both the final consumer and society as a whole.

Is “zero-knowledge” privacy achievable in the distributed cloud? An in-depth analysis of encryption architecture of Cubbit’s sync&share service.

Marco Moschettini¹ ; Alessio Paccioia¹ ; Lorenzo Posani¹

¹ *Cubbit*

Corresponding Author(s): marco.moschettini@cubbit.io

Cubbit is a hybrid distributed cloud where a central “coordinator” server organizes the resources of a “swarm” of peer-to-peer interacting devices to offer performant and encrypted cloud services. One of Cubbit’s core principles is “privacy by design”, meaning that no third party, including the coordinator server, can access the content of the users’ files.

The usual paradigm of encrypted cloud solutions is to securely store the encryption keys on the server to facilitate operations such as file sharing and synchronization on different devices. However, this solution allows the cloud provider to potentially access the user’s data, thus not guaranteeing total privacy. Implementing a performant sync&share service where the provider itself is unable to access the contents of the files poses several technological as well as architectural challenges.

Here we will present the encryption architecture implemented by Cubbit to allow performant and seamless sync&share while at the same time guaranteeing zero-knowledge privacy. We will focus on key management protocols as well as upload, download, and file-share procedures.

Syncthing - Introduction, Technology and Use Cases

Jakob Borg¹

¹ *Kastelo Inc.*

Corresponding Author(s): jakob@kastelo.net

The talk introduces Syncthing, describes the operation and underlying technology, and goes through a couple of common use cases. Emphasis is on how Syncthing works and what it can accomplish, how the user is in control of the flow of information (synced data), privacy, and ease of deployment. The described use cases show how Syncthing scales from a small office deployment to datacenter replication with terabytes of data. We round off by mentioning the availability of commercial support and turn-key deployment.

Operation

- Omni directional sync
- No central authority
- Peer to peer (non-cloud)
- Discovery and relays

Technology

- Encryption and security
- Restricted operation modes (read only, etc)
- Single binary deployment
- Cross platform (25 supported platforms)
- GUI or REST API

Use cases

- Multi site replication
- Office file sharing

Introducing Pydio Cells

Charles du Jeu

Corresponding Author(s): charles@pydio.com

Started as Ajaxplorer ten years ago, Pydio has been a major player in open source sync-and-share for enterprise since then. After hitting the limitations of PHP stack, the core team decided to rewrite their software from ground up using the Go language: reinventing both the user experience and the underlying technical stack, the team unveiled a brand new product called Pydio Cells in may 2018.

This presentation will go through Cells features and internals, and show how the new architecture helps both quick deployments for non-technical people as well as massive scalability for critical installations.

Subjects such as micro-services, JWT & authentication, datasources synchronisations and object-storage protocols will be covered.

Enterprise File Sync & Share by Oracle

Peter Szegedi¹

¹ *Oracle*

Corresponding Author(s): peter.szegedi@oracle.com

Abstract – The Enterprise File Sync and Share (EFSS) service function is part of the broader Enterprise Content Management (ECM) portfolio of Oracle. It is available on-premises as licensed software, in the cloud in form of a Platform as a Service (PaaS), and also in any hybrid deployment models, soon to be part of the Cloud at Customer offering too. It's important to emphasize the "Enterprise" nature of the solution that leaves the full control of the service at the enterprises, in contrary to any public Consumer-focused FSS services such as Dropbox, Google Drive or Microsoft OneDrive where users put unnecessary risks to their organizational IT processes. Oracle's EFSS solution is targeting a large spectrum of information workers, especially in the area of document management and content publishing use cases, not necessarily in the research and education context. In this paper, however, we explore how the Oracle EFSS might be repositioned to be in line with the needs of the CS3 community.

Data Science: Applications and Infrastructure

Exploring Callysto

Joe Topjian

Corresponding Author(s): joe.topjian@cybera.ca

In 2017, Cybera began a joint collaboration with the Pacific Institute for the Mathematical Sciences (PIMS) called Callysto: a free, multimodal learning program for grades 5-12 students in Canada. The goal of this project is to use open source tools to enhance the computational thinking, coding, and data skills of teachers and students. At its core, Callysto uses JupyterHub to provide easy access to interactive notebooks that have been specially created for younger students, and cover a broad range of subjects, from math to history.

This talk will cover Callysto's JupyterHub technical infrastructure in detail, including:

- The tools and processes we use to deploy and manage JupyterHub
- Scaling JupyterHub
- Managing personally identifiable information
- And since this is CS3, there will be a good amount of storage discussion, too:
- Creating custom storage drivers for Jupyter
- Storage considerations when scaling JupyterHub

Image and data analytic pipeline for studying single-cell signaling dynamics in cancer

Maciej Dobrzynski¹

¹ *University of Bern*

Corresponding Author(s): maciej.dobrzynski@izb.unibe.ch

Signaling pathways are an important part of cellular machinery responsible for processing biochemical signals and translating them into fate decisions such as proliferation, migration, or programmed cell death. Aberrant signal processing due to mutations leads to errors in fate determination, such as an increased rate of proliferation which is one of the hallmarks of cancer. The process of cell fate determination depends on the dynamics rather than the steady state of signaling components, thus to dissect network wiring that orchestrates fates, it is crucial to monitor signaling over time and subject the network to genetic and chemical perturbations.

In the Cellular Dynamics Lab at the University of Bern, we have established a number of experimental techniques based on time-lapse microscopy to measure and perturb the dynamics of signaling pathways in individual cells over time. Such experiments generate vast amount of images, a problem exacerbated by a recent shift towards the study of 3D organoids that mimic physiological conditions better than conventional 2D cell cultures. Storage and the analysis of images, extraction of relevant features from time-series data, visualization of multidimensional feature spaces, and fitting of mechanistic models to data all pose a significant challenge to the field.

Our overarching aim is to understand oncogenic signaling and sources of drug resistance, and pave the way for personalized approaches to cancer treatment. In my talk, I will present a pipeline realized in our group and roadblocks we have stumbled upon during the implementation. I will demonstrate how efficient data analytic approaches coupled to modern visualization techniques allow us to quickly gain insight into the measurements and plan the next-best experiment. I will focus on a suite of homegrown tools for spectral analysis of time series data as well as approaches based on machine learning.

Advanced geo-spatial data analysis with Jupyter

Armin Burger¹ ; Paul Hasenohr¹ ; Pierre Soille¹ ; Davide De Marchi¹

¹ *European Commission - Joint Research Centre*

Corresponding Author(s): paul.hasenohr@ec.europa.eu

The JRC Earth Observation Data and Processing Platform (JEODPP) is serving JRC projects and their partners for big data applications with emphasis on geo-spatial data. It has evolved into a multi-petabyte scale platform, offering advanced Web-enabled services for container-based batch processing, remote desktop, as well as interactive analysis and visualization through the JEO-lab service.

The JEO-lab service is providing a powerful and flexible Web-based environment to both data science specialists and less experienced occasional users for interactively analyzing and visualizing geo-spatial data. JEO-lab is based on Jupyter notebooks using a Python kernel and an API based on C++ libraries exposed to Python via the SWIG interface.

The design of the JEO-lab service separates the coding inside notebooks from the actual data processing that is deferred and executed by a series of back-end service nodes. The processing is initiated via an interactive map, requesting map tiles with the processing results from the back-end service nodes acting as tile engine. All processing chains from the notebooks are encoded into JSON objects and stored in a REDIS key-value store. The back-end tile engine nodes are retrieving and decoding the JSON objects, applying the processing chains, and sending the results back to the notebook clients.

The move to the JupyterLab environment largely improved the usability of JEO-lab. It allows to better manage coding and side-by-side display of the results on an interactive map. Split-maps allow for a convenient comparison of different analysis workflows.

A set of widgets provided by various Jupyter extensions enables the creation of advanced user interfaces for data analysis, in a style of desktop tools where the parameters of the underlying python functions are interactively controlled by appropriate widgets. This way, powerful analysis tools serve the needs of both specialists and desk officers without requiring any programming knowledge from the end-user. A series of customized thematic processing interfaces based on Jupyter note-books has been developed to support JRC projects in various data analysis fields and visualization modes.

The separation of coding (notebook) and processing (tile engine) nodes improves security and scalability, but makes it difficult for users to extend the existing API. In order to overcome this limitation, a mechanism has been implemented to embed Python code provided by the user into modules and functions in the JSON objects of the processing chain. The code is then executed by the tile engine nodes where a Python on-the-fly interpreter is instantiated by the C++ libraries.

Various export functions allow to retrieve the results of the processing in various formats for further analysis, reporting, or distribution. The output files can then be retrieved through the JEODPP NextCloud instance.

./..

In addition to the interactive processing environment, two new data processing tools based on Jupyter notebooks for large-scale data processing are currently in a prototype phase:

- The DASK environment (<https://dask.org>) provides an interface to a Kubernetes cluster with DASK workers and allows to launch parallel Python processing like NumPy or machine learning algorithms over multiple nodes, transparently for the user;
- Through the integration of Kubernetes with HTCondor, users have the possibility to submit jobs to HTCondor (the main batch processing environment of the JEODPP) through notebooks via the Kubernetes cluster. Kubernetes acting as meta-scheduler allows to use the same computing resources in an harmonious way leading to a more efficient use of the infrastructure.

It is expected that those new services will allow the users to be more autonomous to fully exploit the processing capabilities offered by the JEODPP.

References:

<https://doi.org/10.1016/j.future.2017.11.007>

DLCM: A Swiss solution for preserving research data on the long term

Pierre-Yves Burgi¹ ; Hugues Cazeaux¹

¹ *Université de Genève*

Corresponding Author(s): pierre-yves.burgi@unige.ch

Developed in the context of the 2017-2020 Swiss national program “Scientific information: Access, processing and safeguarding”, the DLCM solution (dlcm.ch) consists of an open and modular architecture for long-term preservation of research data, compliant with the OAIS standard (ISO 14721). The independent modules, once deployed in the cloud, offer a range of services that allow researchers to prepare their data for preservation, namely: to submit them with a pre-ingest step followed by ingest (submission package), to store them physically (archival package), to index the metadata (METS container with PREMIS and DataCite fields) and to be able to access them according to specific rights (distribution package). This set of services, available via RESTful APIs, guarantees on the one hand the implementation of best practices in the domain: virus detection, format detection, checksum calculation, integrity check, replication, etc., and on the other hand a tight integration to Laboratory Information Management Systems (LIMS). A preservation planning module (preservation-centric workflows) handles replication and synchronization of data, stored in different data centres, with the possibility to be interfaced to the LOCKSS technology, which is based on the Byzantine Generals’ Problem, with $3n+1$ geographically-distant storage nodes, allowing to tolerate up to n faults.

SURF Research Drive - JupyterHub integration within OwnCloud

Tom Wezepoel¹

¹ *SURFsara*

Corresponding Author(s): tom.wezepoel@surfsara.nl

In addition to our SURFdrive ownCloud environment, where we support today more than 40,000 users from Dutch higher education and science with their need for personal data storage. We have set up last year SURF Research Drive, where researchers are able to share more and larger datasets with a group of researchers from home or abroad and third party.

SURF Research Drive is an ownCloud based solution running on Docker, in combination with Swift S3 Object Storage. Within this environment we have made an integration with Jupyter Hub, where researchers are able to create and share their notebooks and also be able to use their datasets immediately.

In this presentation we want to tell you something about our setup and the challenges we have encountered.

SWAN and its analysis ecosystem

Author(s): Diogo Castro¹ ; Jakub Moscicki¹ ; Massimo Lamanna¹

Co-author(s): Enrico Bocchi¹ ; Enric Tejedor Saavedra¹ ; Danilo Piparo¹ ; Prasanth Kothuri¹

¹CERN

Corresponding Author(s): diogo.castro@cern.ch

CERN, and High Energy Physics (HEP) in general, face unprecedented challenges in data storage, processing and analysis. With the planned improvements to the Large Hadron Collider (LHC), including the High-Luminosity LHC, there is an expected increase of data in one order of magnitude. After processing and filtering these data, new tools and solutions, capable of dealing with such large datasets, are particularly important for the last phases of analysis.

SWAN (Service for Web-based ANalysis) is a service that provides an interactive interface to access data analysis tools from the web, allowing users to perform their work in a simpler way and with much larger datasets. Its integration with CERN's infrastructure, more precisely with users synchronized storage and syncing capabilities (via CERNBox), computing resources, experiments data (via EOS) and software, allows a seamless experience across all of our user scenarios and devices.

But even more data means even more resources and different approaches. SWAN has recently been integrated with CERN's Spark Clusters and we are already working to provide access to our Worldwide LHC Computing Grid (WLCG) batch service. We're also experimenting with the integration outside of CERN's infrastructure, using external cloud vendors to install and deploy our Science Box service (which bundles SWAN, CERNBox, EOS and CVMFS). These last experiments are also important in our pursuit of bringing SWAN to new education scenarios, namely via the UP2University European Project.

Application store for data analytics

Gurvinder Singh Dahiya

Corresponding Author(s): gurvinder.singh@uninett.no

At Uninett, we have chosen Kubernetes as a platform to provide various services. We have currently two production clusters serving different communities. In addition we have used public cloud resources through kubernetes for some use cases as well.

To make use of platform easier by end users, we have made an application store where users have commonly used tools such as Jupyter Notebook, Jupyter Hub, Rstudio, Spark to process their stored data. Users can customize these tools with the packages they need and run it with the same ease. We also provide Minio as sync & share solution due to its simplicity to manage and operate it.

Deep learning is another area where our users have shown interests and to address those needs, we provide an application to enables use of most common frameworks e.g. Tensorflow, Pytorch, Keras etc on a GPU enabled infrastructure. User can request multiple GPUs to scale their workflow as needed.

This presentation will go through our platform and application store. I will also share experience from our users using these application for various use cases.

CloudStor SWAN: Data Processing and Analysis Challenges in the Cloud

Michael D'Silva¹

¹ *AARNet*

Corresponding Author(s): michael.dsilva@aarnet.edu.au

CloudStor SWAN (Service for Web based ANalysis) is AARNet's first attempt at providing data processing and analysis in the cloud to the research community in Australia. This presentation will discuss AARNet's experiences, challenges and tools used to provide research data computing in the cloud.

SWAN (Service for Web based ANalysis) helps users run scientific data processing and data analysis in the AARNET cloud quickly. One of the problems we are faced with is that researchers upload data to CloudStor and then need to download the data in order to do any processing on it. This is undesired by users as it causes issues and interrupts the natural workflow. For this reason, we have developed and deployed a modified version of the SWAN service (<https://swan.web.cern.ch>) developed by CERN and presented during earlier CS3 conferences which is an extended implementation of Jupyter Notebooks that also integrates directly into ownCloud.

Out of the box, CERN's SWAN requires CERNBox and EOS to provide authentication and storage whereas AARNet's CloudStor SWAN has been modified to interact with our ownCloud instance directly. When a user requests a SWAN notebook, the ownCloud SwanViewer App generates an ownCloud App Password and passes it onto SWAN for authentication. Once authenticated, the ownCloud App Password is invalidated and removed from ownCloud. For users who access SWAN directly, they simply use ownCloud sync client credentials.

In order to use CloudStor storage in SWAN a WebDAV connection is made to ownCloud via a second ownCloud App Password. This allows us to hide our backend storage away from direct access, giving better security, while providing a seamless user experience.

In addition to greater security, we have customised SWAN to be very generic, so that if required, we can deploy SWAN at a remote site, even on a different network allowing for future expansion options. CloudStor SWAN really only needs to be able to communicate with ownCloud itself thus making it generic enough that any ownCloud operator can deploy it.

In order to use SWAN, users upload data as per normal into CloudStor and then start up a notebook where they can interact with data either in their home directory or from a shared group drive. By keeping data in the cloud and providing tools for data processing and analysis, in the right conditions, users' workflows become more efficient. The benefit of Jupyter Notebooks in this way is it allows users to keep data analysis functions near the data which encourages reproducible analytics.

CloudStor Minio: Improving S3 performance in CloudStor

Michael D'Silva¹

¹ *AARNet*

Corresponding Author(s): michael.dsilva@aarnet.edu.au

We at AARNet as well as the research community in Australia need bulk data access to our synch servers because one-off ingest of seriously large datasets performs subpar across the webdav/synch pathway. This presentation will discuss AARNet's experiences, journey and many iterations to achieve high-speed data transfers via S3 protocol (the de facto standard) and the challenges and improvements made along the way.

Minio helps some users interact with CloudStor using the S3 protocol. In the beginning we mounted EOS, CloudStor's storage backend via FUSE and ran Minio over the top. We found that by using this approach transfers were very slow for large files and our FUSE mount kept crashing due to overloaded metadata queries. Fortunately Minio is Open Source and written in the Go programming language which means we can start hacking it!

We then modified Minio to locally stage file uploads, then use xrootd's xrdcopy command as a background task which (from the users point of view), increased uploads from ~4mb/s to ~800mb/s. Later on we had a user group uploading a dataset with many tiny files (>100,000) in the one bucket. This uploaded without issue, but doing an object listing on the bucket took over 2 hours. We then modified Minio again so that file listing was done via EOS's /proc/user interface rather than listing via the EOS FUSE mount reducing the time down to 40 seconds. This worked well but the code was no longer maintainable with modifications all over the code base.

From here we made the decision to start over, but rather than hacking the Minio code, we instead decided to write a separate EOS gateway module for Minio. The goal was to fork Minio and improve it to work with EOS in a way that is easy to maintain and update. The other goal of the EOS gateway was to remove the need to use EOS's FUSE connector as it is a source of slowness.

The EOS Gateway for Minio communicates to EOS via EOS's WebDAV, EOS's /proc/user web services and xrootd's, xrdcopy. We also added ownCloud hooks so that files coming in and out of S3 are file scanned so that users can share uploaded data with other users and groups.

AARNet's Minio modifications provides an S3 implementation that allows for tighter integration to Open Source and commercial products that are already integrated in the users' workflow. These includes collaboration and backup products such as FigShare, LabArchive, Alfresco, Commvault to name just a few. This enables users to upload, download, view and share as well as use S3 on the same data area.

Open Data Ecosystem and CS3

Open, FAIR, and GO FAIR

Erik Schultes¹

¹ *GO FAIR International Support and Coordination Office*

Corresponding Author(s): erik.schultes@go-fair.org

The 15 FAIR Principles have found unusually rapid uptake among a broad spectrum of stakeholders, from research scientists who make data, to publishers who distribute data, to science funders who track impact of data. Erik will describe the FAIR Principles, their relation to Open data, and review example implementations. Erik will also present a set of core FAIR Metrics that can help gauge the level of FAIRness of any digital resource. This discussion, and these examples will be presented in the context of the International GO FAIR Initiative. GO FAIR is a voluntary community of stakeholders devoted finding consensus on standards and solutions that comprise an emerging Internet of FAIR Data and Services.

Building a Scalable, Flexible, Interoperable Federated Cloud to Power Australian Research

Author(s): Joanne Morris¹

Co-author(s): Paul Coddington¹

¹ *Australian Research Data Commons*

Corresponding Author(s): jo.morris@ardc.edu.au

The Nectar Research Cloud provides a self-service OpenStack cloud for Australia's academic researchers. Since its inception in 2012 it has had rapid growth, with approximately 37,000 virtual CPUs now being used at any given time (more than 95% of the available resource), with over 7,000 virtual machines being run by approximately 2,000 researchers and used by thousands more. The number of registered cloud users is now over 14,000 and continues to grow at over 200 per month. In the last year, over 4,000 researchers ran a virtual machine in the Nectar cloud. Over 2,500 project allocations have been approved, with about 50 new applications per month. More than 700 allocations are to support multi-institutional collaborations, and over 300 support national research grants.

The Nectar Cloud is different to many clouds in being a federation across seven organisations, each of which runs cloud infrastructure in one or more data centres and contributes to the operation of a distributed help desk and user support. A Nectar core services team runs centralised cloud services. The Nectar Cloud federation is continuing to be extended with the inclusion other institutional cells and also the recent partnering with Auckland University. The offering of services that Nectar Research Cloud continues to grow with most recently releasing a Simple Container Service (Docker) and an Container Orchestration Service (Kubernetes). The Nectar Research Cloud is very focused on enhancing the user experience through integration with other services and providing simplified user interfaces.

Nectar recently merged with the Australian National Data Service (ANDS) and Research Data Services (RDS) to form the Australian Research Data Commons (ARDC). ARDC is funded by the Australian Government through the National Collaborative Research Infrastructure Strategy (NCRIS) and is currently developing its five year strategy and capital plans. As part of the strategic planning cycle ARDC is engaging with a range of resource providers, institutions and research organisations, e-infrastructure providers, peak bodies and government. ARDC's vision is "Transforming Australian Research together; Towards an Agile, Interoperable & Sustainable eResearch Infrastructure Ecosystem for Australia".

One of the motivations and challenges of the merger of Nectar, ANDS, and RDS into the ARDC is to provide more seamless integration of the national cloud and compute services with data and storage services. One example is recent work to make it easier for users of the Nectar cloud to access their data from Cloudstor, a national sync-and-share storage service.

This presentation will give an overview of the experiences, challenges (both past and present), and benefits of running a federated OpenStack cloud and an overview of the some of the integrations we have implemented to improve the user's experience. It will also describe future plans to extend the infrastructure and services, integrate with other services and data storage providers, as well as extending interoperability with the growing number of international science and research clouds through initiatives such as the Open Research Cloud, along with improving access for international researchers to assist in international collaborative research.

Open Data analysis with EOSC-hub services

Baptiste Grenier¹ ; Enol Fernández¹

¹ *EGI Foundation*

Corresponding Author(s): baptiste.grenier@egi.eu, enol.fernandez@egi.eu

The EOSC-hub project mobilises providers from the EGI Federation, EUDAT CDI, INDIGO-DataCloud and major research e-infrastructures offering services, software and data for advanced data-driven research and innovation. These resources are offered via the Hub – the integration and management system of the European Open Science Cloud, acting as a single entry point for all stakeholders.

Several of the use cases in EOSC-hub will enable scientific end-users to perform data analysis experiments on large volumes of data, by exploiting a PID-enabled, server-side, and parallel approach. Users expect also easy to use interfaces like Jupyter Notebooks for interacting with the system.

This talk presents an ongoing effort to turn these needs into a production-ready service that provides a FAIR-approach to the researchers analysis workflow by leveraging and integrating the following services from the EOSC-hub catalogue:

- Check-in for handling user authentication and authorisation to services with federated AAI
- DataHub for providing a data access and management
- B2HANDLE for PID minting for research objects available in DataHub
- Notebooks as an interactive interface to carry out the end-user tasks, including data analysis of a public dataset referenced by a PID and publication of the results back with a new PID.
- B2FIND for making data discoverable by harvesting data using OAI-PMH

Thanks to the use of federated AAI and interoperable protocols, this work will be further extended to integrate new services from EOSC-hub to support computing intensive workloads, other data-related services, and potentially any other service with community-scoped functionality as required by the different use cases.

Hybrid Cloud storage deployment models

João Fernandes¹ ; Bob Jones¹ ; Marion Devouassoux²

¹ *CERN*

² *Polytech Lille (FR)*

Corresponding Author(s): joao.fernandes@cern.ch

Currently, many research projects struggle to manage their data, as the archiving and preservation services are inadequate and fall below expectations while data stewardship costs are frequently underestimated during the planning phase. Using the EC Pre-Commercial Procurement instrument, the ARCHIVER project goal is to fulfil these data management promises in a multi-disciplinary environment, allowing each research group to retain ownership of their data whilst leveraging best practices, standards and economies of scale.

ARCHIVER will combine multiple ICT technologies, including extreme data-scaling, network connectivity, service interoperability and business models, in a hybrid cloud environment to deliver end-to-end archival and preservation services that cover the full research lifecycle. The use-cases driving the consortium's need for research and development of innovative data preservation services will extend the preservation ecosystems of the procurers to create more dynamic solutions using a hybrid model allowing to combine in some cases on-premise capacity with external services operated by commercial suppliers or, alternatively extending existing scientific preservation software to additional sciences domains and workflows integrated with capacity supplied by commercial cloud services in order to ensure service long-term sustainability. In both cases, the approach will be enhanced to comply with the OAIS reference model and the relevant series of standards (ISO 14721).

Commercial service providers are being considered by the European research community as part of a hybrid cloud model to support the needs of their scientific programmes. However, research institutions and end user researchers struggle to incorporate these commercial services into their activities in a structured manner. Service discovery and procurement are difficult and take up an inordinate amount of time. Service providers find it difficult to reach and meet the needs of the research community in areas such as legal, financial and technical compliance. The Open Clouds for Research Environments (OCRE) consortium aims to change this: to bridge this gap between the supply and demand sides and enable and facilitate researchers to use these cloud services in a safe and easy manner, in order to make it part of their day-to-day scientific activity.

This presentation will explore the use-cases driving ARCHIVER and introduces the benefits of the Open Clouds for Research Environments (OCRE) for providers and for researchers.

Connecting and Supporting the Research Workflow with OSF

Nici Pfeiffer

Corresponding Author(s): nici@cos.io

The non-profit Center for Open Science (COS) seeks to connect and streamline the research workflow through use of its web application, OSF (<https://osf.io>). Free and open source, OSF manages the entire research lifecycle: planning, execution, reporting, archiving, and discovery. The OSF represents a research coordination center: a hub that connects many services a researcher uses, allows for customized access controls to manage many collaborators, and maintains an accurate record of the project via a history log and built-in version control.

The main functions of the OSF are in managing and streamlining scholarly and scientific workflows and providing a platform for openly communicating scientific knowledge. The OSF streamlines workflows by connecting tools researchers already use — examples include ownCloud, Dropbox, Google Drive, figshare, and Dataverse — allowing resources housed in different services to be displayed in one central location. Users can also choose to store files in default OSF Storage, configuring the physical location of storage from options in North America, Europe, and Australia. Institutions can leverage OSF workflows with a view layer on top that aggregates and showcases public research outputs affiliated by researchers at that institution or university.

This presentation will discuss the features of the OSF, and ways it can improve a researcher's workflow. It will also include discussion of where the product is going: how feedback from the community can inform feature development and how partnerships with other services and repositories can improve functionality.

Sync/share Technology&Research

Designing and Implementing for Scale, Distribution & Async

Jörn Dreyer¹ ; Tom Needham²

¹*ownCloud GmbH*

²*ownCloud*

Corresponding Author(s): jfd@owncloud.com

ownCloud has traditionally used the database to store metadata including the file hierarchy, shares, comments and tags. Keeping this database metadata in sync with the actual filesystem becomes challenging when requests time out or fail for other reasons such as stale locks, partially traversed trees, or changes on storages that are not propagated, ultimately leading to filecache corruptions.

In order to power the next generation of storage services, we're rewriting the fundamental layers that ownCloud was built on using seven years of experience and combining it with the latest developments in technology.

A new event driven storage architecture makes background processing a first class citizen, allows storing all metadata in the storage, brings commands and queries to replace the synchronous nature of request processing, and an asynchronous protocol opens up new possibilities for distributed storages, federation and client communication.

On the way to OCM 1.0, a first stable version ready for mass adoption

Björn Schießle¹

¹ *Nextcloud GmbH*

Corresponding Author(s): bjoern@nextcloud.com

During the last year Nextcloud took huge efforts to implement the Open Cloud Mesh API. We see this as a important building block to bring Cloud Federation to a large group of users in a vendor neutral way. As the principal authors of the concept of Cloud Federation and driving force behind it, we looked not only on what we have but also on what we want to achieve in the future. This led to many suggestion of small changes and improvements which we think should form the base for a first stable version. Throughout the year we discussed our ideas already with the broader community, gathered feedback and other vendors started to implement our proposal. This talk will give a overview of the changes we made and explains the rationale behind it. It will show how the adoption of this API enabled Nextcloud to move Cloud Federation forward and gives an outlook over our plans for the future.

A journey to the mesh : microservices-based sync and share using gRPC and Protobuf

Hugo Gonzalez Labrador¹

¹*CERN*

Corresponding Author(s): hugo.gonzalez.labrador@cern.ch

CERNBox is a multipetabyte-scale sync and share platform at CERN, storing close to 7 PB of data for more than 16K users. Over time we have identified different improvement points to increase development agility and maintenance costs for running the service. Last year we have evolved the architecture of the service from a monolithic stack to a decentralised model based on micro-services. This mesh is the core of the new system and relies on latest technologies (gRPC and Protobuf) for efficient inter-component communication and controlled API evolution. In this talk we present the process we followed and our observations during this journey, which led to the creation of the public consumable CS3 APIs for a distributed sync and share platform.

Scalable Storage Backends for Cloud, HPC and Global Science

Anomaly detection of large scale distributed storage system based on machine learning

Lu Wang¹

¹*Computing Center, Institute of High Energy Physics, CAS*

Corresponding Author(s): wanglu@ihep.ac.cn

In a large scale storage system which consists of hundreds of servers, tens of thousands of clients, a variety of devices, anomaly detection is a nontrivial task. Traditional solutions which are still working in our cluster operations include setting static thresholds on KPIs, searching key words in system logs and so on. These methods highly depend on experience of system administrators, cannot adapt to new anomalies in the cluster. The machine learning communities has developed a wide range of algorithms which is able to do anomaly detection of high dimensional data by statistically learning over a large training data set. The monitoring system of IHEP cluster has accumulated billions of performance metrics entries in its database. It provides possibilities to train those machine learning models. This presentation will show how our preliminary works on anomaly detection of ganglia time serials monitoring data by machine learning algorithms including LSTM, HTM and Isolation Forest.

From performance monitoring to predictive maintenance.

Jean-Thomas Acquaviva¹

¹*DDN Storage*

Corresponding Author(s): jacquaviva@ddn.com

Storage architectures are becoming larger and deeper. This leads to a quadratic increasing of systems complexity. To address the issue more ambitious monitoring schemes are arriving to the market. A similar trend is observed for both HPC and general purpose data centers.

In this talk we will position the DDN monitoring efforts in this context, and open a more generic discussion on the advantages of applying analytics to collected measurements, the trade-off related to probes, and failure predictions.

Converge heterogeneous legacy filesystems at the Pawsey Supercomputing Centre

Luca Cervigni¹ ; Gray Mark¹

¹*Pawsey Supercomputing Centre*

Corresponding Author(s): luca.cervigni@pawsey.org.au, mark.gray@pawsey.org.au

The Pawsey Supercomputing Centre is an unincorporated joint venture to support Australian researchers, funded by the four major universities in Western Australia, local government, and the CSIRO. The centre works with many different types of filesystems such as Lustre, GPFS, NFS, Mediaflux and Ceph, that for legacy reasons do not talk very well with each other. Recent trends in research are leading us toward a more seamless integration between computing resources and storage. Researchers need solutions that work out of the box, so they can spend their time purely focusing on their research and not with IT problems. Pawsey Sync is the new pilot sync and share service at Pawsey that aims to solve these issues, to empower researchers with a seamless integration between storage systems and allow heterogeneous workflows between HPC, cloud, and data sharing. Pawsey Sync will also allow a smoother data migration to future and more converged filesystems that will come with the 2019-2022 capital refresh funding.

dCache - active storage for science and cloud

Birgit Lewendel¹ ; Dmitry Litvintsev² ; Marina Sahakyan ; Paul Millar³ ; Tigran Mkrtchyan³ ; Vincent Garonne⁴ ; Peter van der Reest³ ; Quirin Buchholz³

¹ *Deutsches Elektronen-Synchrotron (DE)*

² *FNAL*

³ *DESY*

⁴ *University of Oslo (NO)*

Corresponding Author(s): tigran.mkrtyan@desy.de

Over five years storage team at DESY provides a reliable data-cloud service. While service is still officially in a pilot phase, it's has the same support and priority level as any other production services provided by the IT group. Our choice of technology used to provide data-cloud service is nextCloud, as user facing front-end, and dCache, as back-end storage system.

The latest developments in dCache have added storage events - an ability to turn passive storage system into an active part of data coordination and processing workflow. In this presentation we will share our setup, lessons learned and integration of dCache with other services to provide fault tolerant and reliable service to end-users.

Cynny Space Cloud Object Storage integrated with Network Intelligence platform Vegoia

Stefano Baldi¹ ; Fabio Del Vigna²

¹ *Cynny Space*

² *Larthia*

Corresponding Author(s): s.baldi@cynnyspace.com, fabio.delvigna@larthia.com

The aim of this presentation is to show how Cynny Space cloud object storage solution integrates with Vegoia, a new Network Intelligence platform.

The presentation will outline the following topics:

1) Cynny Space cloud object storage

The object storage solution is the first storage built on fully-equipped ARM® based micro-servers. Thanks to 1:1 micro-server to storage unit ratio and the SwARM File System optimized for ARM® architecture, the solution delivers optimal levels of scalability, resilience, reliability, and efficiency. Every storage node is fully independent and interchangeable, providing a storage solution with no single point failure and very high durability levels (99,999999999%). To maximize the parallelization of the operations, every node is responsible for some chunks (and not the whole file). The use of Dynamic Hash Tables and a distributed Swarm Intelligence on a peer-to-peer network minimizes the cooperation overhead and allows a high level of scalability, fault-tolerance, and clusterization. The solution is API S3 compatible natively and accessible via NFS and SMB gateways.

2) Network Intelligence applications

Network Intelligence is a fundamental support for professionals operating in network management and defense because network awareness opens up to a safe and informed network control. Vegoia platform is a complete solution that can be applied in a wide range of applications where an accurate knowledge of traffic and apps is necessary: cybersecurity, network operation, maintenance and troubleshooting, business intelligence and marketing. It has been designed for a good accuracy in the protocol detection and for scalability to cope with high-speed links (10 to 40 Gbps) and can be configured both as a virtual probe (for statistics and punctual analysis, troubleshooting) or for an active role with user accounting and policing.

3) What is Vegoia: a new Network Intelligence platform

Vegoia is a Network Intelligence platform based on a Deep Packet Inspection (DPI) library and a multi-core modular engine used to analyze and classify in detail the network traffic. The engine performs the DPI operating through dissectors making the protocol parsing on standard protocols, while Machine Learning algorithms are applied where there are proprietary protocols and encrypted connections. The platform is not only able to detect the protocols but even the applications, producing a multi-layer, real-time classification with special reference to the smartphones apps (social, web, mail, ...). Results are stored in form of tickets that describe each flow and connection. They can be manipulated by means of Big Data techniques, so providing statistics and analytics, alarms, ...

4) The Vegoia and Cynny Space cloud storage integration use case

From the partnership between Cynny Space and the Vegoia designers , an integrated solution has been developed to join together the two technologies. In fact, since the DPI solution has the need to store a huge amount of data traveling at high speed the cloud storage model is particularly suitable for the purpose. The architecture of the integrated system and the available results will be presented during the CS3 conference.

A demo will be available for visitors at Cynny Space's desk during the conference.

Creating a load balancer for Ceph backed S3 service using open-source tools

Anders Bruvik¹

¹ *Infrastructure Engineer for Safespring*

Corresponding Author(s): anders.bruvik@safespring.com

Safespring are currently building a new version of our ceph service, and in this session, I will give an overview of the architecture of the load balancer for our S3 service, which we are building using open source tools. In our design, we use internet as a design pattern for the datacenter cluster network ensure scalability and predictable performance. That means we use BGP everywhere, even for the last hop to each server node. In the talk, I will outline how we use BGP per-flow ECMP routing to achieve redundancy and high throughput for the frontend services, BFD for failure detection, we use traefik as load balancer, and Prometheus and Grafana for monitoring.

This is a work in progress, but I will give an overview of our design decisions, and experiences so far.

iRODS: Managing Data from the Edge to HPC

Jason Coposky

Corresponding Author(s): jasonc@renci.org

Data management has historically started at the point of ingest where users manually placed data into a system. While this process was sufficient for a while, the volume and velocity of automatically created data coming from sequencers, satellites, and microscopes have overwhelmed existing systems. In order to meet these new requirements, the point of ingest must be moved closer to the point of data creation.

iRODS now supports packaged capabilities which implement the necessary automation to scale ingest horizontally. This shifts the application of data management policy directly to the edge.

Once your data resides in the iRODS namespace, additional capabilities such as storage tiering, data integrity, and auditing may be applied. This includes tiering data to scratch storage for analysis, archiving, and collaboration. The combination of these capabilities plus additional policy allow for the implementation of the data to compute pattern which can tailored to meet your specific use cases.

Onedata - Managing Data and Metadata in Hybrid Clouds

Michal Orzechowski¹ ; Bartosz Kryza² ; Lukasz Dutka²

¹ AGH University of Science and Technology, Academic Computer Centre Cyfronet AGH, Krakow, Poland

² ACC Cyfronet-AGH

Corresponding Author(s): michal.orzechowski@cyfronet.pl

Onedata [1] is a transparent, high-performance data management system, which provides transparent access to globally distributed storage resources and supports a wide range of use cases from personal data management to data-intensive scientific computations. Due to its fully distributed architecture, Onedata 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. Onedata comprises the following components: Onezone, distributed metadata management and authorisation component that provides an entry point for users; Oneprovider, which is the main data management component providing the transparent virtual filesystem over distributed heterogeneous storage resources; and Oneclient, which provides a virtual POSIX file system mountpoint on user worker nodes.

Onedata introduces the concept of Space, a virtual volume, owned by one or more users, where they can organize their data under a global namespace. The Spaces are accessible to users via an intuitive web interface, Fuse-based client providing POSIX file system, as well as REST and CDMI standard APIs. Each Space can be supported by a dedicated amount of storage supplied by one or multiple storage providers. Storage providers deploy Oneprovider instance near the storage resources, register it in the selected Onezone instance to become part of a federation and expose those resources to users. By supporting multiple types of storage backends, such as such as POSIX, S3, Ceph, WebDAV, dCache and OpenStack Swift, Onedata can serve as a unified virtual file system for hybrid-cloud environments. Using a comprehensive Onedata REST API it is possible to automatically extract metadata from exposed files and ingest them into Onedata. The data and metadata managed by Onedata are synchronised with any changes made to data directly on the underlying storage. In order to enable an easy way to expose existing data collection, dedicated deployment procedure called Onedatify is available, which provides a command-line wizard that guides a user through \op deployment procedure and exposes existing data from the specified legacy storage.

Currently, Onedata is used in Helix Nebula Science Cloud [2], eXtreme DataCloud [3], PLGrid [4], and European Open Science Cloud Pilot [5], where it provides data transparency layer for computation deployed on hybrid-clouds. Furthermore, in European Open Science Cloud Hub [6] it also serves as the basis of EGI Open Data Platform, supporting various open science use cases such as open data curation (metadata editing), publishing (DOI and PID registration) and discovery (OAI-PMH protocol).

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

[2] Helix Nebula Science Cloud (Europe's Leading Public-Private Partnership for Cloud). <http://www.helix-nebula.eu>.

[3] eXtreme DataCloud (Developing scalable technologies for federating storage resources). <http://www.extreme-datacloud.eu>.

[4] PL-Grid (Polish Infrastructure for Supporting Computational Science in the European Research Space). <http://projekt.plgrid.pl/en>.

[5] European Open Science Cloud Pilot (The first phase in the development of the European Open ScienceCloud). <https://eoscpilot.eu>.

[6] 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>.