



# Workshop on Cloud Services for File Synchronisation and Sharing

## Report sul Workshop

**Stefano Stalio**  
**Giacinto Donvito**

**CERN 17-18**  
**November 2014**

<https://indico.cern.ch/event/336753/>  
Abstract submission: 30 September  
Registration: 31 October

ORGANISING COMMITTEE

Miguel Branco  
Massimo Lamanna  
Jakub T. Moscicki



# 4 sessioni in 2 giorni

- Introduction and Keynotes
- Technology and research
- Site reports
- Software solutions

Sul sito dell'evento sono disponibili le riprese di molti degli interventi

<http://indico.cern.ch/event/336753/>

# Ma tutto è chiaro fin dall'inizio.....

*“....CERN adopted ownCloud....”*



# A Little History



- **Unison** is the only synchronizer based on a formal specification (AFAIK)
- Main designers
  - Trevor Jim, Jerome Vouillon, and BCP
- First version distributed in 1998
  - Earlier prototypes going back to 1995
  - Still widely used
- Open-source, multi-platform (Windows, OSX, Linux)
- *Very* few reported bugs with potential to lose data

# But the world has changed...

**Synchronization tools:** (bidirectional, sync operations explicit)

- Unison, ...



**Synchronization services:** (multi-directional, sync in background)


- Central server ("cloud-based") • • • concentrate on these
  - Dropbox, Google Drive, OneDrive, Owncloud, SpiderOak, Sugarsync, Box.net, Seafile, Pulse, Wuala, Teamdrive, Cloudme, Cx, Amazon cloud service, ...
  - Also distributed filesystems (Coda, GFS, ...)
- Peer to Peer
  - Bittorrentsync



## ▶ The Problem



Big Data Storage Technologies

- data growth faster than budget/technology
- 'traditional' systems not long-term scalable
- cost & complexity increase when scaling



## ▶ 'The Solution'

From Scale-Up to Scale-Out Storage



- structured data with central view & strong consistency
- predictable access
- SPOF
- slow growth rate

- best for unstructured data
- unpredictable growth rates
- MPOF resistant

many storage systems like HDFS use hybrid approaches with scale-up for meta-data and scale-out for data path



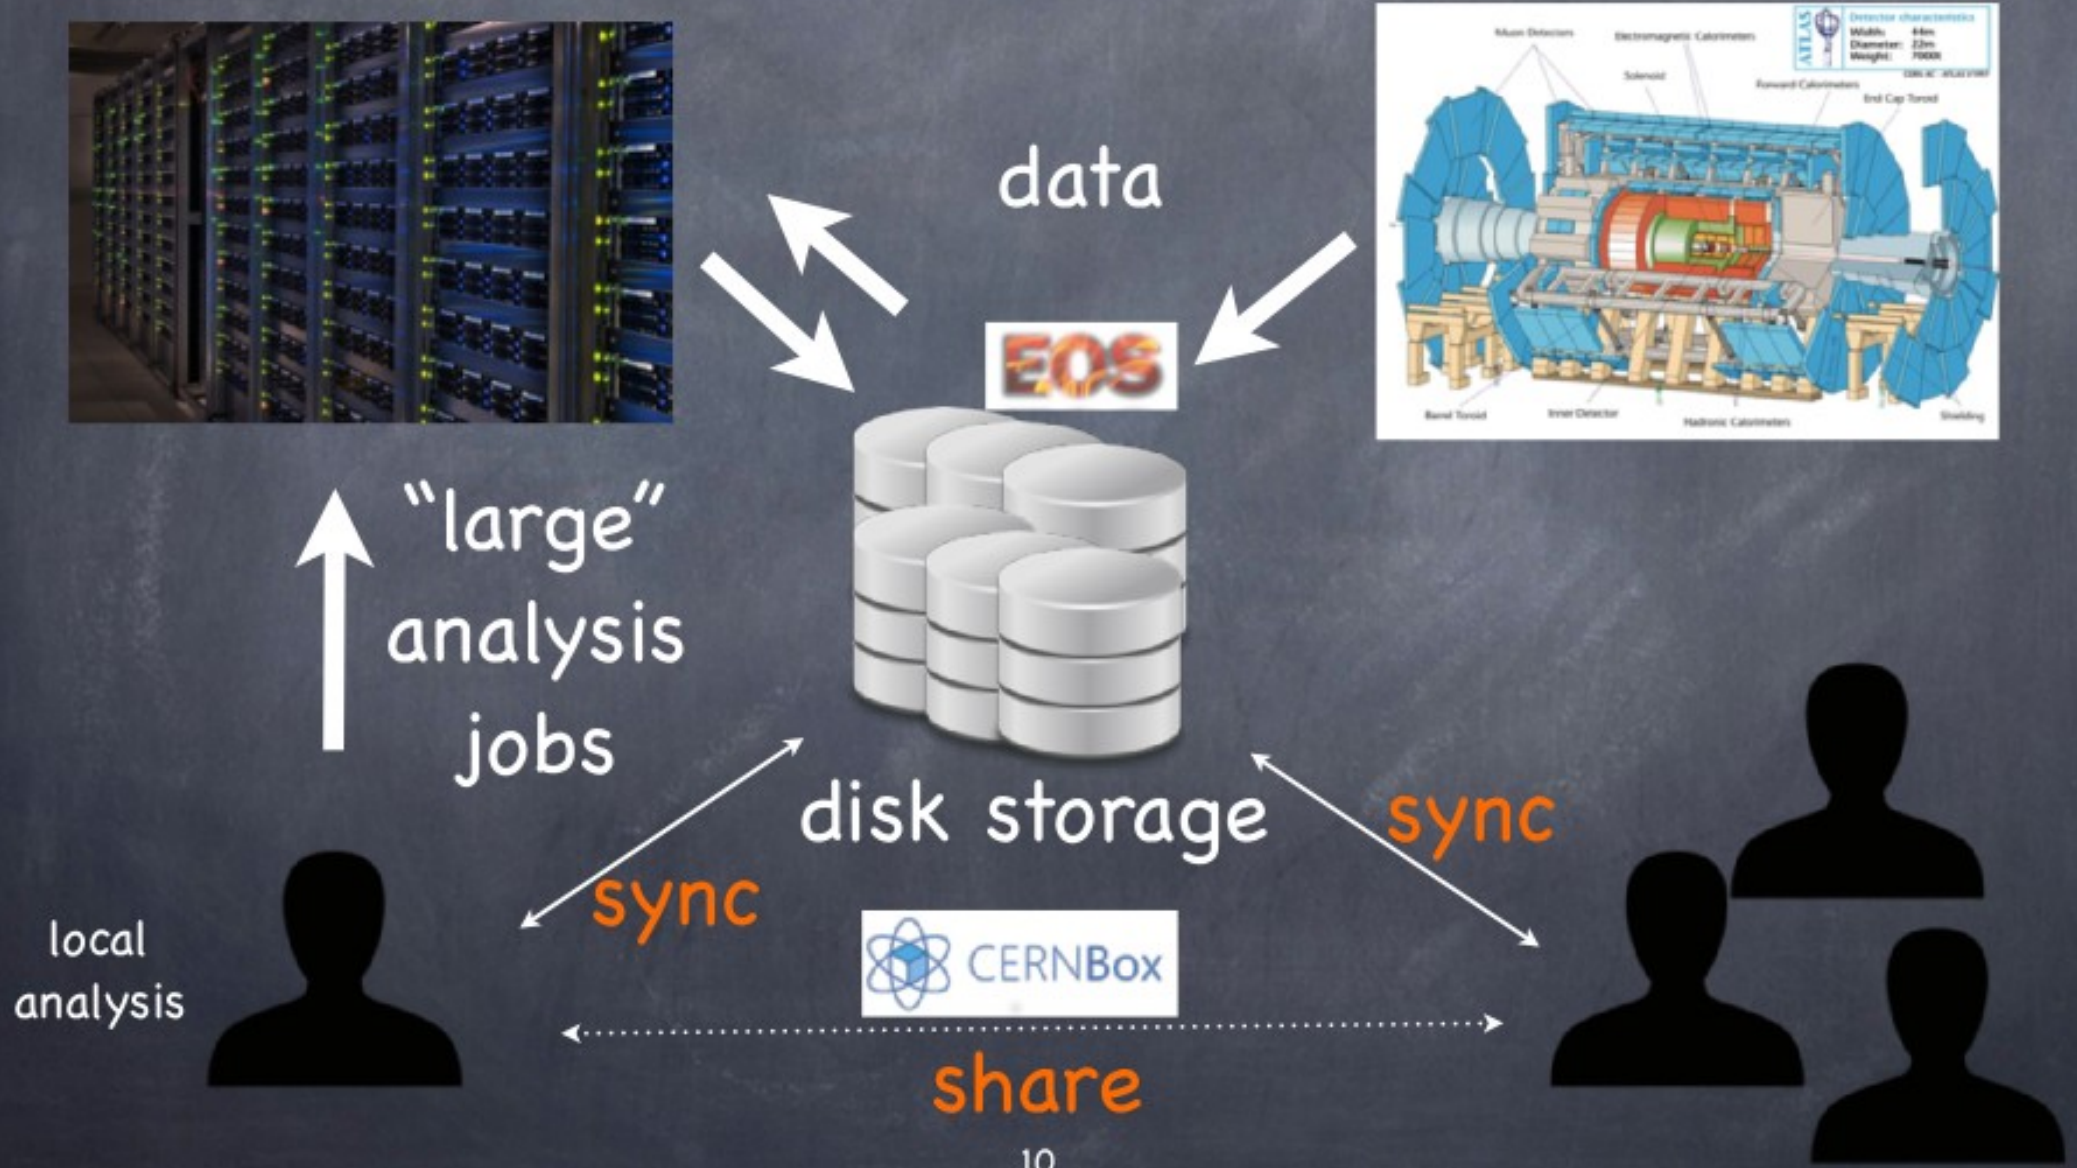


# “Dropbox” for science



- Provide “Sync and Share” Service for the core business of doing science
- Challenge: “highly nontrivial” use-case
- Reward: unique capability for our researchers perfectly integrated in their daily environment




# Use Case: Data Analysis

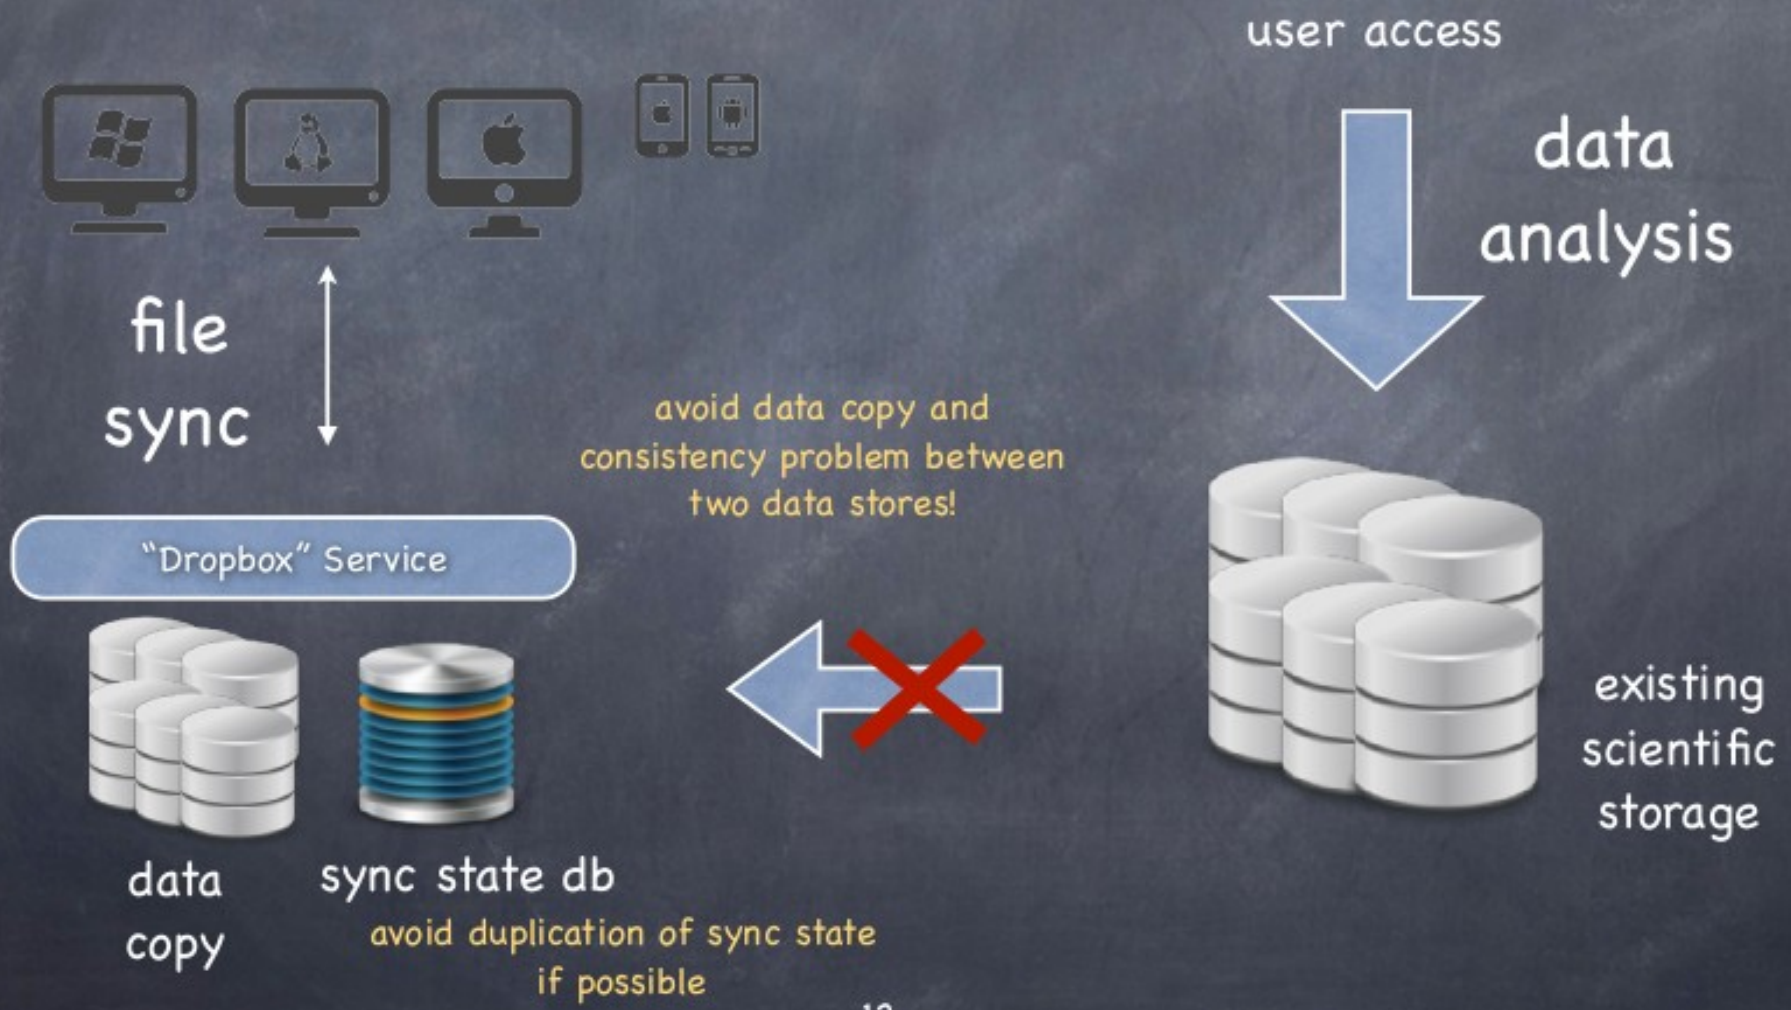




# Use Case: Global Science

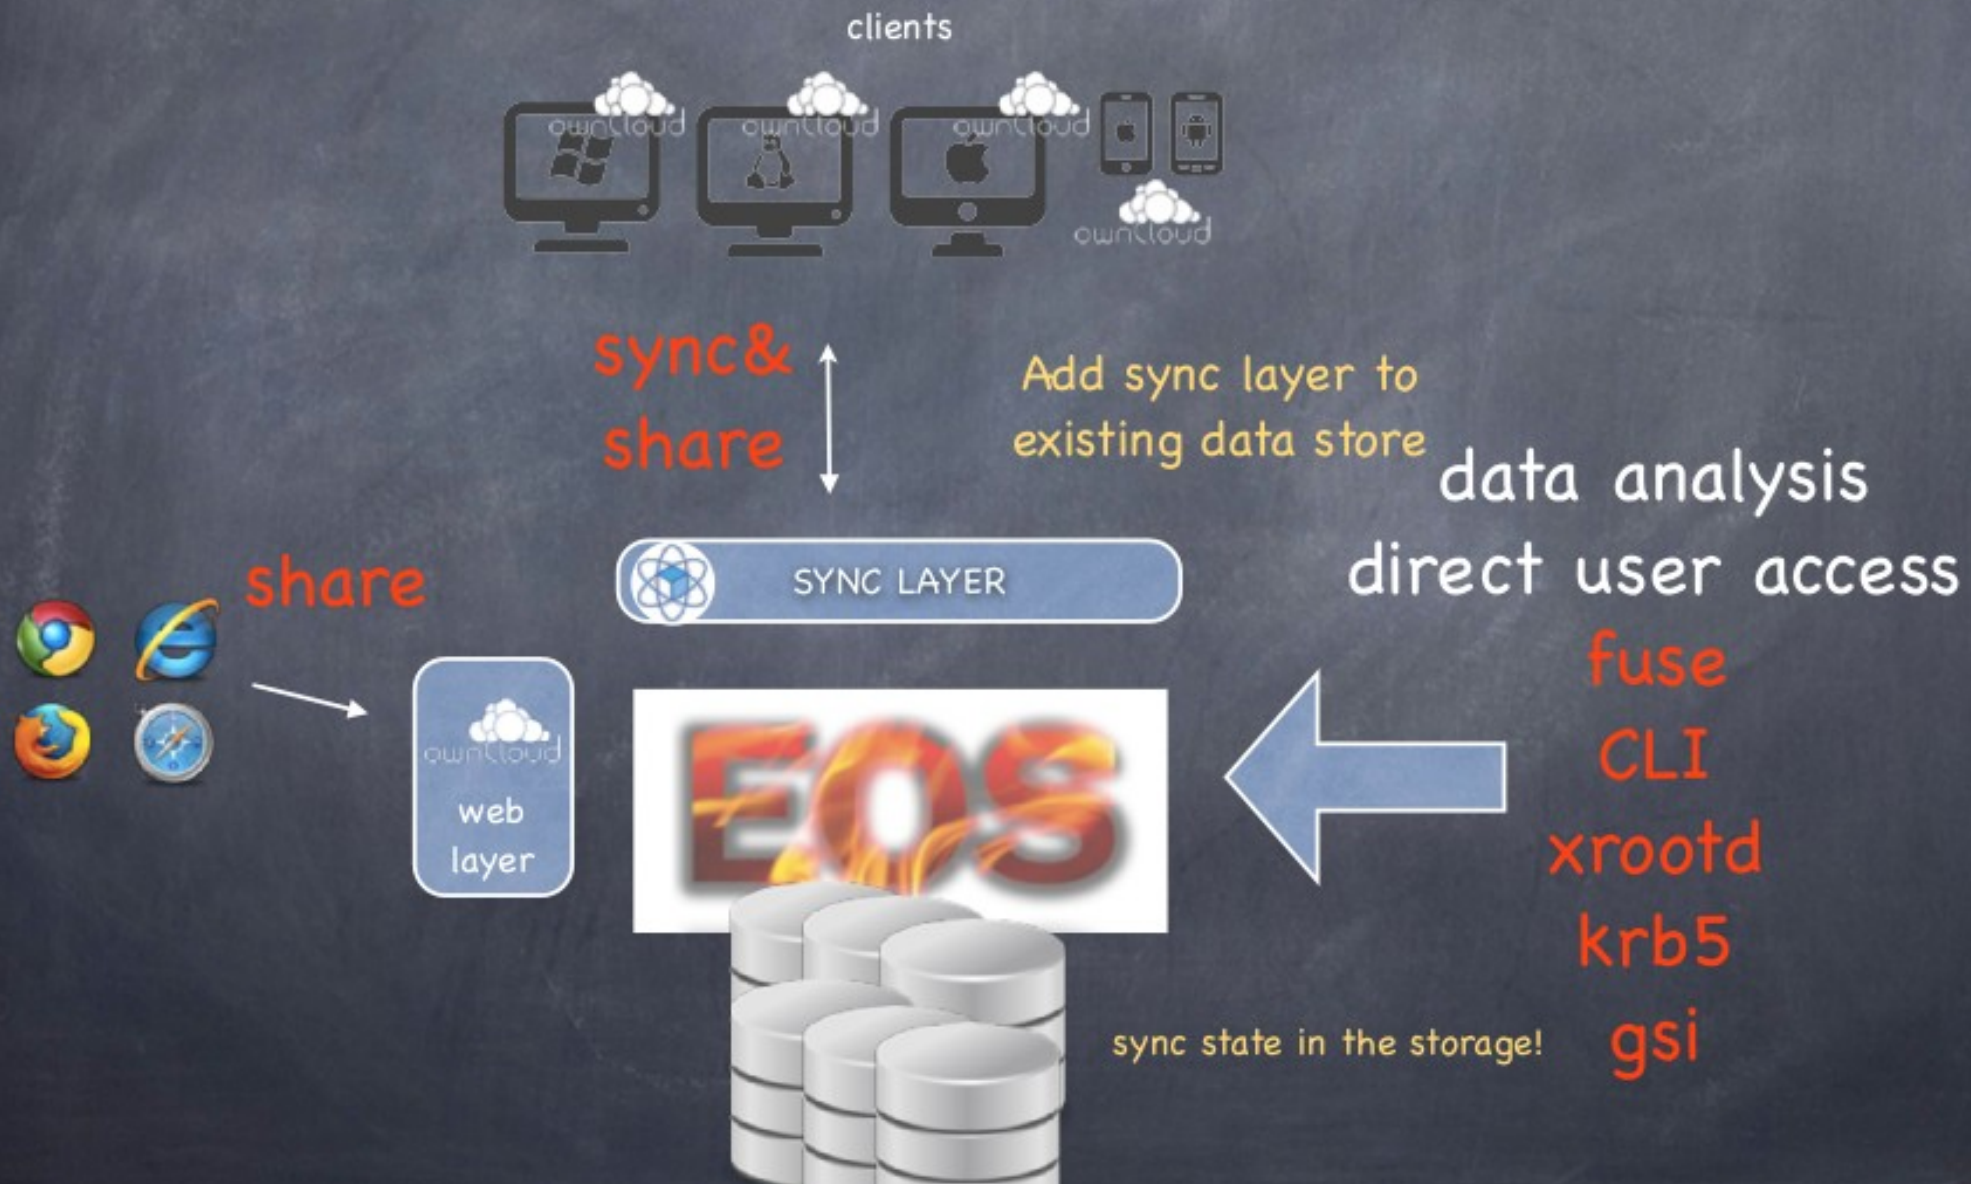


# No duplication of storage!

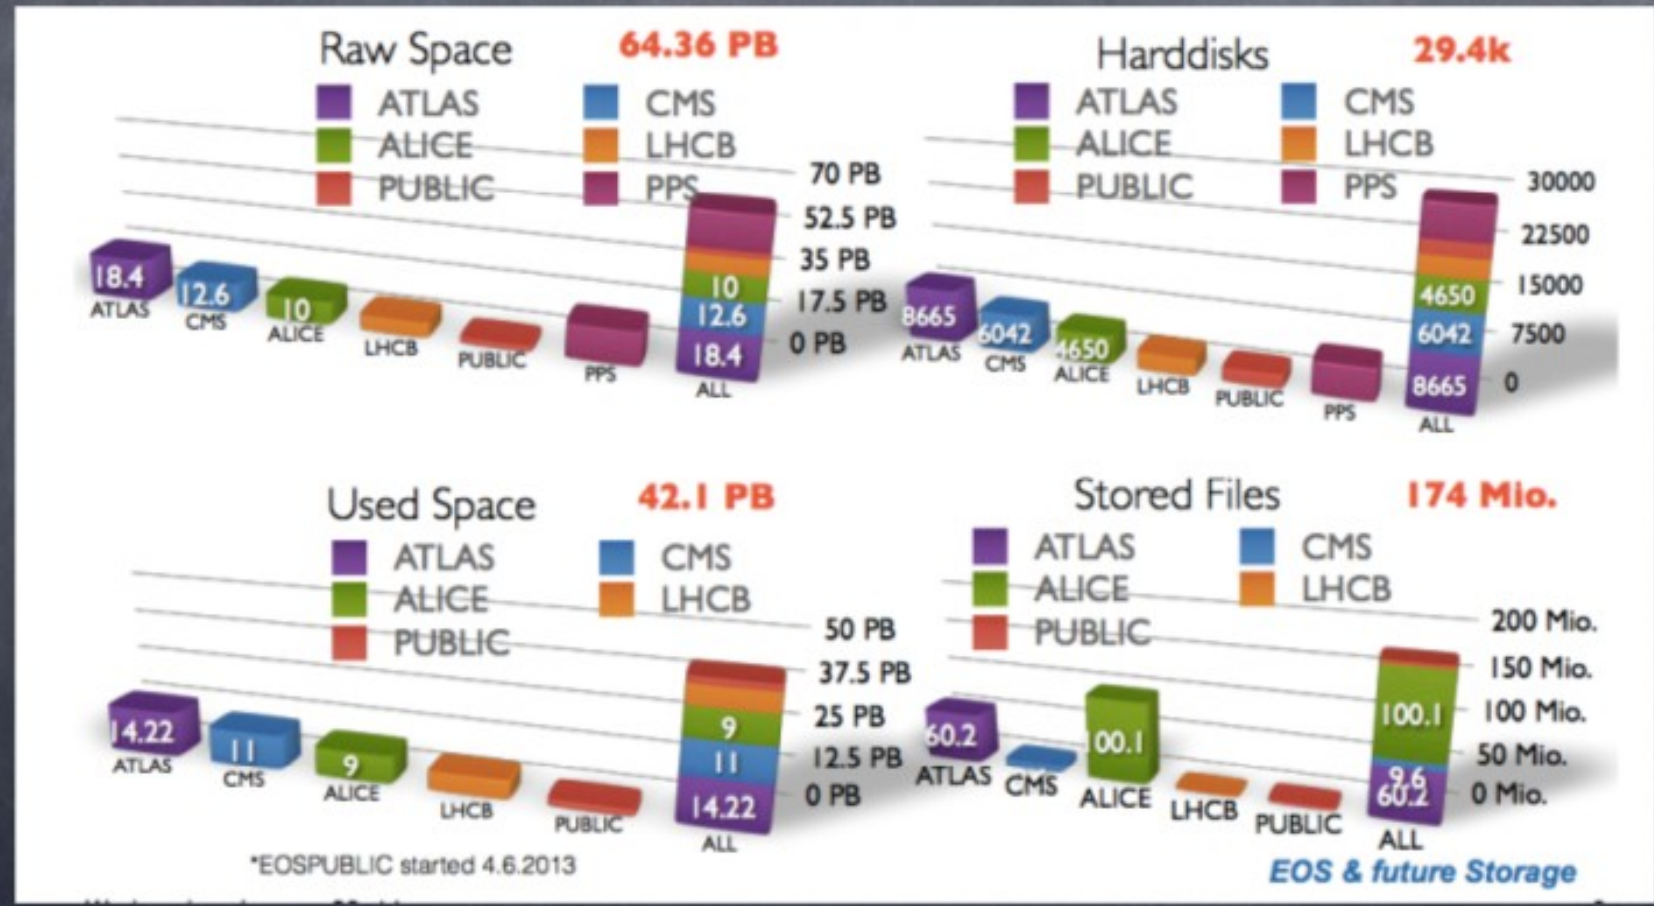




# CERNBox architecture



# Immediate access to all our data!



# Software Solutions

- Presentation of Pydio
- PowerFolder – Peer-to-Peer powered Sync and Share
- Seafile open source cloud storage, technology and design
- ownCloud
- IBM Software Defined Storage and ownCloud Enterprise Edition - a perfect match for hyperscale Enterprise File Sync and Share



# Mature solution, backed by a growing community...

in a Nutshell, Pydio...

... has had 5,772 commits made by 60 contributors representing 322,372 lines of code

... has a well established, mature codebase, maintained by a large development team with stable Y-O-Y commits


... took an estimated 84 years of effort (COCOMO model), starting with its first commit in June, 2007 ending with its most recent commit 2 days ago

Source: [OPENHUB.NET](https://www.openhub.net)

# pydio

PUT YOUR DATA IN ORBIT

- > AGPL v3
- > 80k+ active installs
- > 20k forum members
- > 10k Newsletter subscribers
- > 60 contributors
- > Source code on Github



# PydioSync - reloaded

**Sync client completamente riscritto**

Welcome to the Pydio Sync

The easiest way to keep your data in control

pydio

PUT YOUR DATA IN ORBIT


Per-workspace & per folder choice

Bidirectionnal, up or down only.

Ability to run headless

Limited resources even for huge number of files

Still in beta but promising



# Seafile

- **Seafile** ci è sembrato particolarmente interessante
- Ma gli sviluppatori sono 2!!!

# Seafile

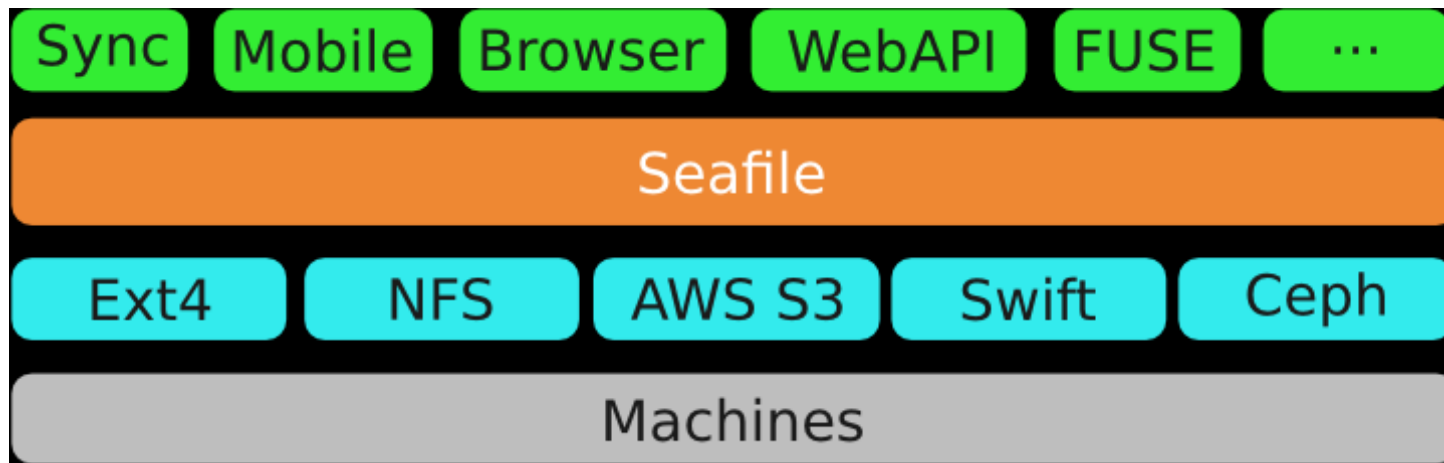
## Design Advantage

- Object storage is more scalable than file system
  - Heavy DB + Filesystem v.s. Light DB + Object Storage
- No database bottleneck
  - Metadata is in object storage
  - Filesystem level versioning v.s. File-level versioning
- File system designed for syncing
  - Storage/Network deduplication
  - No upload/download limit, fast upload
- Backend daemons implemented in C

# Seafile



## Server Architecture

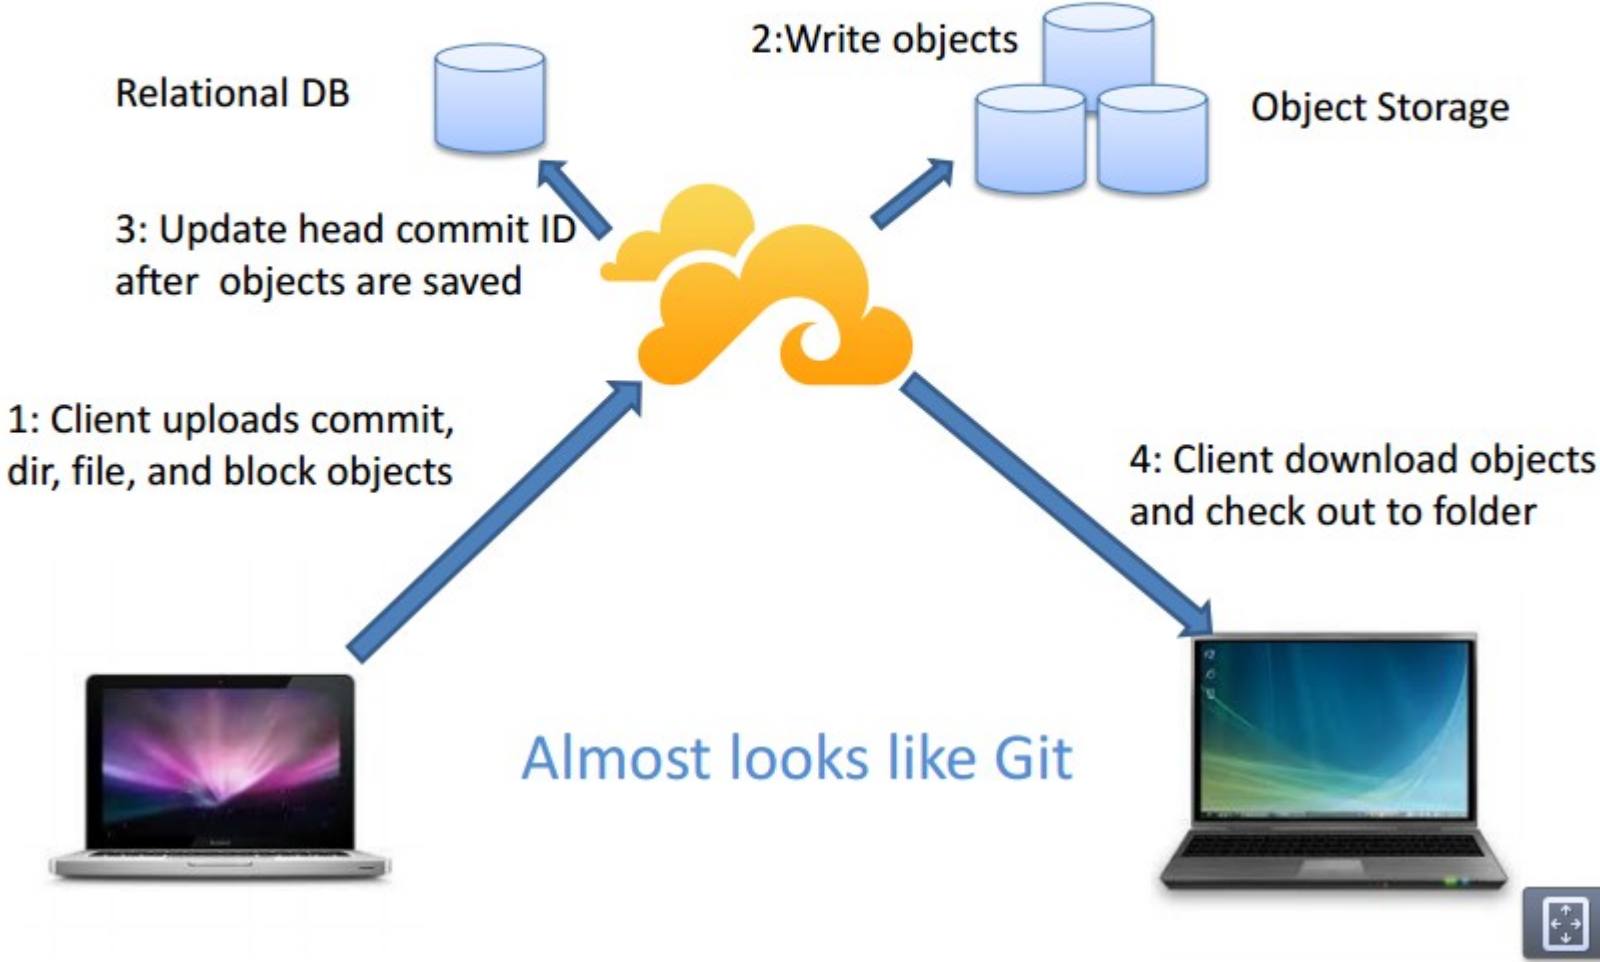


Seafile is a “file system” built on top of object storage

Non-POSIX, User space, Light weight

# Seafile

## How Syncing Works





# Technology and Research /1

- OpenStack Swift as Multi-Region Eventual Consistency Storage for ownCloud Primary Storage
- Programmatic access to file syncing services
- Data Management Services for VPH Applications
- Storage solutions for a production-level cloud infrastructure
- DataNet: A flexible metadata overlay over file resources

# Technology and Research /2

- Adaptive Query Processing on RAW Data
- The File Sync Algorithm of the ownCloud Desktop Clients
- Combining sync&share functionality with filesystem-like access
- Dynamic Federations: scalable, high performance Grid/Cloud storage federations
- The dCache scientific storage cloud
- WebFTS: File Transfer Web Interface for FTS3
- Current practical experience with the distributed cloud data services

# Technology and Research /1

- OpenStack  
Consistent  
Storage
- Programm
- Data Manag  
Applications
- **Storage solutions for a production-level cloud infrastructure**
- DataNet: A flexible metadata overlay over file resources

**Unico contributo INFN  
ed unico contributo  
italiano della  
conferenza.**

# AARNet – OwnCloud + Swift

With the addition of features to the ownCloud software allowing primary data storage on OpenStack Swift, AARNet has chosen to deploy Swift in a nation spanning multi-region ring to take advantage of Swift's eventual consistency capabilities and the local region quorum functionality for fast writes.

The scaling capability of Swift resolves the twin problems of geographic redundancy, and user proximal access while scaling into the petabyte range.

...

Additionally, the policy capabilities within Swift combined with the ability to grant a user multiple storage targets within ownCloud, allow us to honour data sovereignty rules with respect to the physical location of the data on a per top level folder basis.

Finally, using the combined read and write affinity features of the Swift proxy, AARNet is presently experimenting with deployment of flash cache backed site local application nodes, providing the perception to the user of near instant data ingestion, while the node trickle uploads data to the redundant bulk storage ring using private layer 3 networks.

By switching to true object store systems, AARNet is able to achieve two of its original design goals for cloud storage services, chiefly being user proximal data storage, and continent-spanning geographic redundancy, from hundreds of terabytes into the petabyte scale.

# AARNet – OwnCloud + Swift

With the addition of features to the ownCloud software allowing primary data storage on OpenStack Swift, AARNet has chosen to deploy Swift in a nation spanning multi-region ring to take advantage of Swift's eventual consistency capabilities and the local region quorum functionality for fast writes.

The scaling capability of Swift resolves the twin problems of geographic redundancy, and user proximal access. ... Additionally, the policy capabilities within Swift combined with the ability to grant a user multiple storage targets within ownCloud, allows us to honour data sovereignty rules with respect to the physical location of the data on a per top-level folder basis. Finally, using the combined read and write affinity features of the Swift proxy, AARNet is presently experimenting with deployment of multiple geographically stacked storage application nodes, providing the perception to the user of near instant data ingestion, while the node trickle uploads data to the redundant bulk storage ring using private layer 3 networks.

By switching to true object store systems, AARNet is able to achieve two of its original design goals for cloud storage services, chiefly being user proximal data storage, and continent-spanning geographic redundancy, from hundreds of terabytes into the petabyte scale.

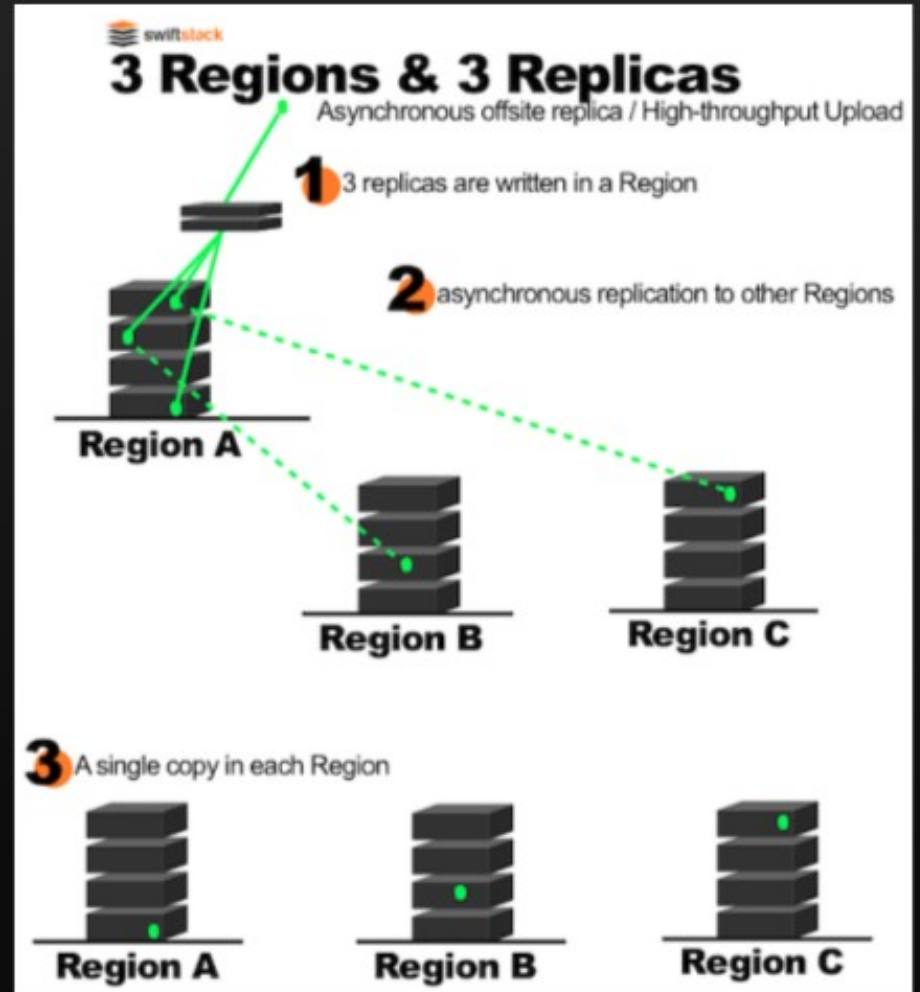
## On Swift: Swift and writing replicas

System “peels” off locally written replicas to remote regions

Step 1 very easily could be a single host with 3 PCI-e flash devices

Policies then have the data moved to cheaper bulk storage behind the scenes

Image source: <https://swiftstack.com/blog/2012/09/16/globally-distributed-openstack-swift-cluster/>






## ownCloud project at CNRS

CNRS will launch next November an ownCloud based service with the intend to serve CNRS research units. The first step is to deploy this service as a beta solution for 2 months and 2000 end users, and then to generalize this offer to the whole CNRS users (potentially 100000 users).

Our platform is based on **ownCloud 7** community edition, with **VMWare for virtualization**, a **Galera/MariaDB** cluster database and **Scality** for the distributed storage backend.

How it all started about a year ago

We want our own



## How it all started about a year ago

But,

- Privacy of data is important.
- No NSA allowed!
- Dutch institutes agreed on security & privacy rules/guidelines.
- Large number of wishes and requirements.

# The cloud storage service bwSync&Share at KIT

The **Karlsruhe Institute of Technology** introduced the bwSync&Share collaboration service in January 2014. The service is an **on-premise alternative to existing public cloud storage solutions for students and scientists** in the German state of Baden-Württemberg, which allows the synchronization and sharing of documents between multiple devices and users. The service is based on the commercial software PowerFolder and is deployed on a **virtual environment to support high reliability and scalability for potential 450,000 users**.


The integration of the state-wide federated identity management system (bwIDM) and a centralized helpdesk portal allows the service to be used by all academic institutions in the state of Baden-Württemberg. Since starting, approximately 15 organizations and 8,000 users joined the service. The talk gives an overview of related challenges, technical and organizational requirements, current architecture and future development plans.



## Service usage (Accounts)

- 9123 accounts (12.11.14)
- Potential:
  - 450 000 eligible users
    - 9 universities
    - >50 colleges
  - 55 000 expected users

bwSync&Share. Number of accounts over the year 2014.



# "Dropbox-like" service for the University of Vienna

The increasing popularity of dropbox and at the same time increasing awareness for data security did create the demand for an onsite "Dropbox-like" "sync and share" service at the University of Vienna. It has been decided that **ownCloud would be a good start**, since other academic institutions have been working on an ownCloud based solution as well.

Based on ownCloud enterprise Version 6 the service is currently in test operation with campus wide availability for staff only planned for 12/2014. Major concerns were the scalability of the storage backend. So instead of using an enterprise storage solution **we use Scality's RING as backend**. The RING is an object storage based solution using local storage nodes. Since the ownCloud architecture does so far not allow a REST based storage backend we use Scality's FUSE connector to simulate a virtually limitless filesystem (POSIX). Based on the experiences reported by other academic facilities and our own, our main concerns have been database performance-scalability, storage backend architecture and general software design. Some of which might already have been addressed by ownCloud community version 7. It's also noteworthy, that **ownCloud's support team responds properly to submitted bug reports**. The admittedly limited user feedback has been quite positive so far.

A other issue which have to be solved is the legal issue: What happened with the data, which are shared, after a staff member leaves the University? We want to establish special terms of use for this service, which everybody who want to use the ownCloud service have to accept.



# "Dropbox-like" service for the

our main concerns have been database performance and scalability, storage backend architecture and general software design.

Some of which might already have been addressed by ownCloud community version 7. It's also noteworthy that ownCloud's support team responds properly to submitted bug reports.

The  
crea  
has  
wor  
Bas  
ava  
back

RING is an object storage based solution using local storage. Since the ownCloud architecture does so far not allow a REST based storage backend we use Scality's FUSE connector to simulate a virtually limitless filesystem (POSIX). Based on the experiences reported by other academic facilities and our own, our main concerns have been database performance-scalability, storage backend architecture and general software design. Some of which might already have been addressed by ownCloud community version 7. It's also noteworthy, that **ownCloud's support team responds properly to submitted bug reports.** The admittedly limited user feedback has been quite positive so far.

A other issue which have to be solved is the legal issue: What happened with the data, which are shared, after a staff member leaves the University? We want to establish special terms of use for this service, which everybody who want to use the ownCloud service have to accept.

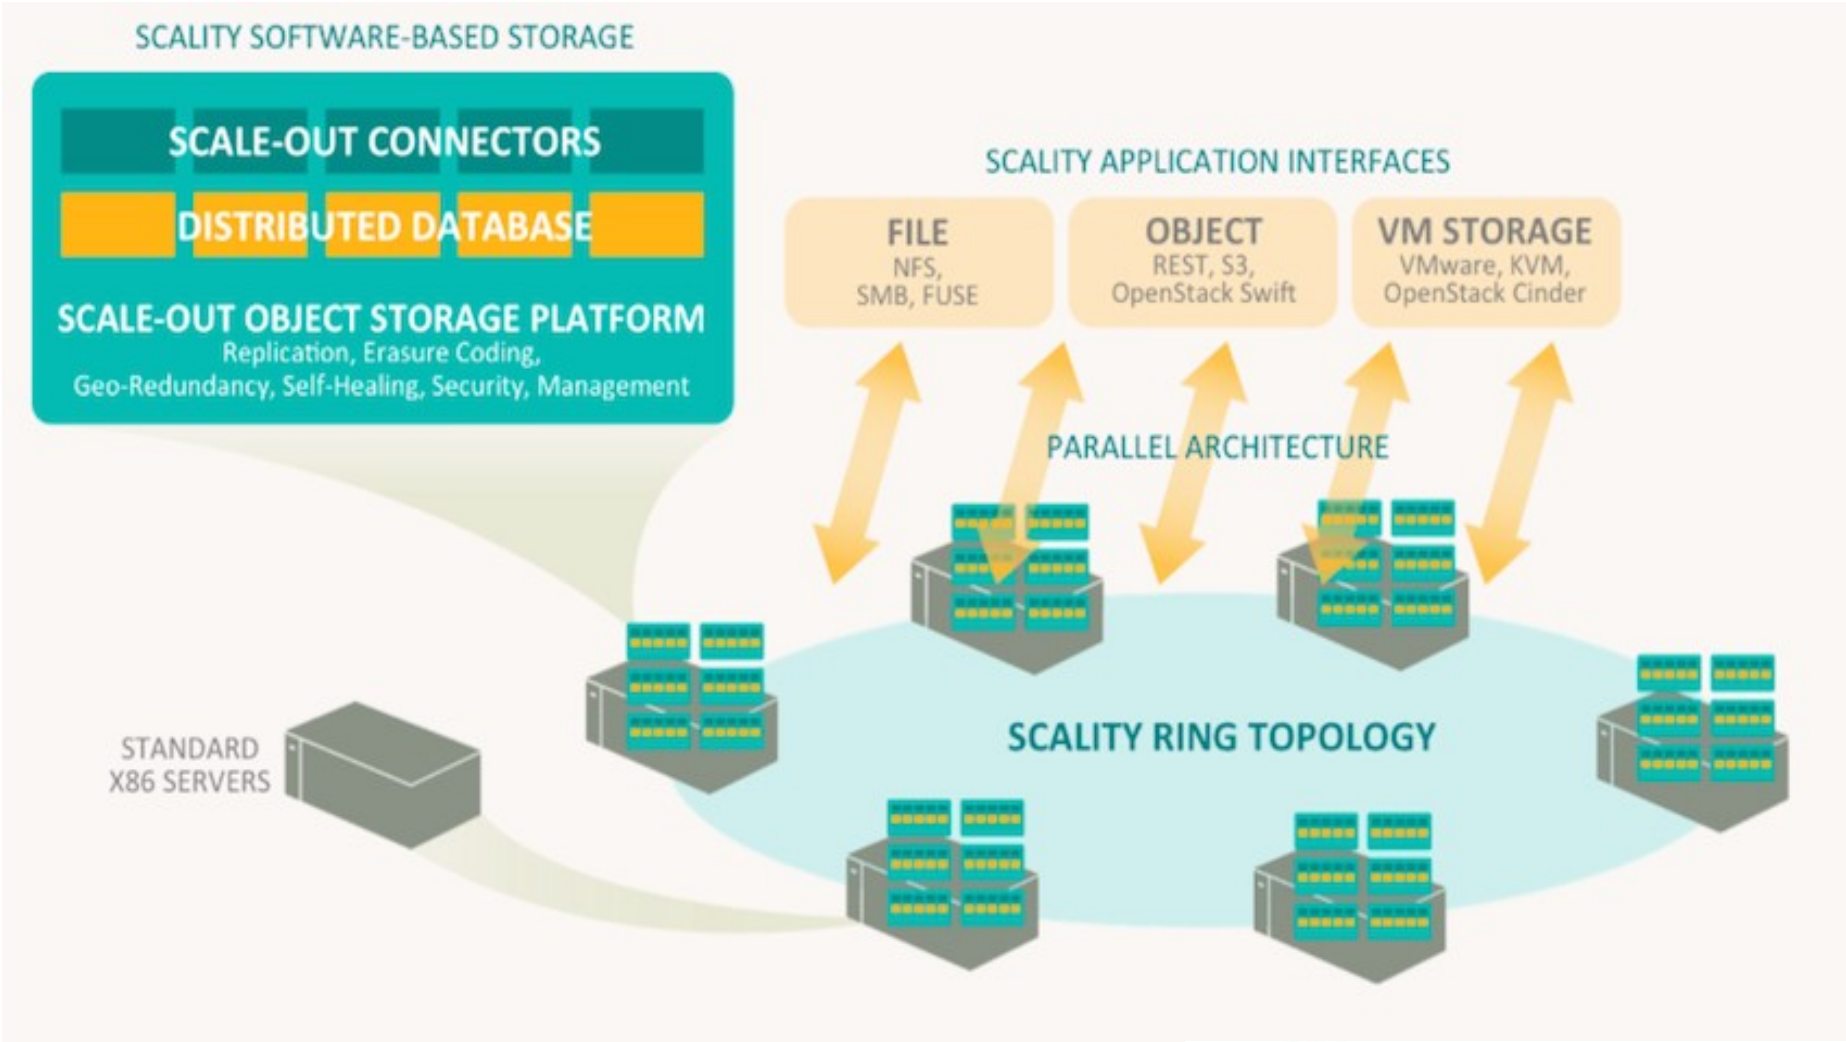
## Ma cos'è questo Scality ring d'oltralpe???

<http://www.scality.com/ring/>

The RING is a software-based storage that is built to scale to petabytes with performance, scaling and protection mechanisms appropriate for such scale. It enables your business to grow without limitations and extra overhead, works across 80% of your applications, and protects your data over 200% more efficiently at 50–70% lower cost.



# Ma cos'è questo Scality ring d'oltralpe???



# Altri siti che usano OwnCloud

- **Polybox at ETH Zurich**
- **CERNBOX**  
Cloud sharing and data synchronization over modern backend storages, like EOS ...
- **The Sync&Share project in North Rhine-Westphalia**

# The Sync&Share project in North Rhine-Westphalia – utenza



## > Scenario and Software

- Up to 500,000 users in North Rhine Westphalia (if all 30 universities and applied universities participate)
- We **will** start with 22 institutions and 350,000 potential users
- Will start in February 2015
- 30 GB/user with the possibility to increase the quota for the employees
- Project boxes for groups (with up to approx. 1 TB)
- ownCloud 7 Enterprise
- RedHat 6, later RedHat 7 (when GPFS support is available)
- Galera DB cluster with MariaDB



# The Sync&Share project in North Rhine-Westphalia – forze in gioco



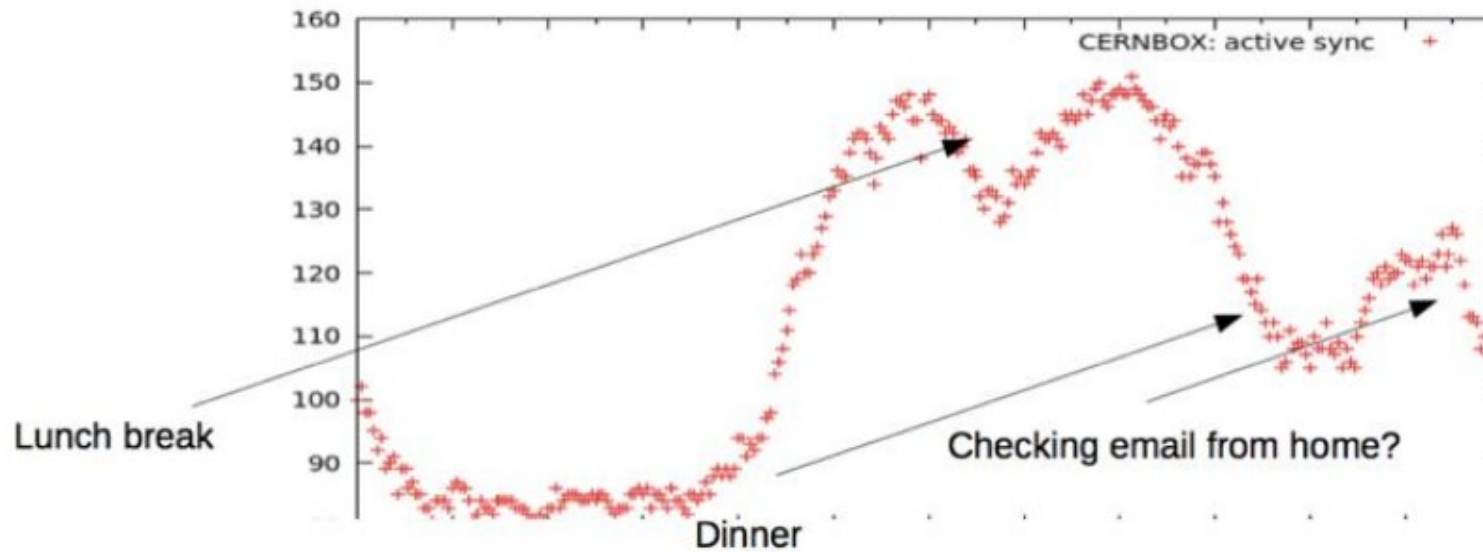

## > Hardware

- Storage: 5 PB at three sites (3 PB in Muenster, 1 PB in Essen and 1 PB in Bonn)
- Filesystem: GPFS 3.5 with declustered RAID and triple parity
- Application servers: 16 IBM NeXtScale nx360 M4 (16 cores) for ownCloud per site
- Database: 4 IBM System x3650 M4 with 6 800 GB SSD (RAID 10) per site
- Loadbalancers: 2 Linux machines per site (LVS with keepalive)
- Management: 1 server per site





# Site reports - CERNBOX







### User feedback

“I start using the cernbox since I'm a heavy user of Dropbox and I recently reached the limit of free disk space (5Gb). For work it will be great to have at least 50Gb of personal space “

“I would like to have is a free client for Android, which should be much more stable. “

“I find the service perfect to be able to get always the files/sources/documents I need independently of the place and connection. ”





### User feedback

- “On my Macbook Air I noticed that the battery was draining much faster than usual. I checked on the activity monitor and CERNBOX was consuming 80-95% of the total energy. “
- “What I would like to do in the future is to combine my private data like my photos for example on my home owncloud server, and my work data on the Cern owncloud server.”
- “I'm very glad that CERN has launched the service using the OwnCloud platform.”
- “I hope that you will be supporting this service officially soon!! ”

### Risk factors

- Stability of the sync clients across diff versions
- Resiliency of the system with “exotic” failure modes
  - Risk to lose files!!!
  - Risk to corrupt data!!!
- Product evolution
  - Stability of the sync protocol (controlled evolution)
  - Bulletproof core functionality
  - Market evolution



# Conclusioni

- I servizi *sync and share* Dropbox-like sono visti non come accessori ma come **parte integrante** dell'infrastruttura di calcolo di Università ed enti di ricerca;
- **OwnCloud** è il software più gettonato, probabilmente perché più maturo e perché l'azienda è più forte di altre e in contatto diretto con Università ed enti di ricerca, ma **ci sono anche altri approcci molto promettenti**;
- I sistemi di **storage ad oggetti** sono visti come backend naturale per i servizi *sync and share*;

# Conclusioni

C'è un fortissimo interesse **dall'alto** a mantenere l'informazione “in house” ed al sicuro dalle multinazionali potenzialmente controllate da NSA.

*the German Federal Ministry of the Interior published guidelines that require cloud providers entering into contracts with German Federal Government agencies to enter into “no spy guarantees” ensuring that no sensitive data will be shared with unauthorized 3rd parties, regardless of where the data resides.*