The DESY Storage Cloud

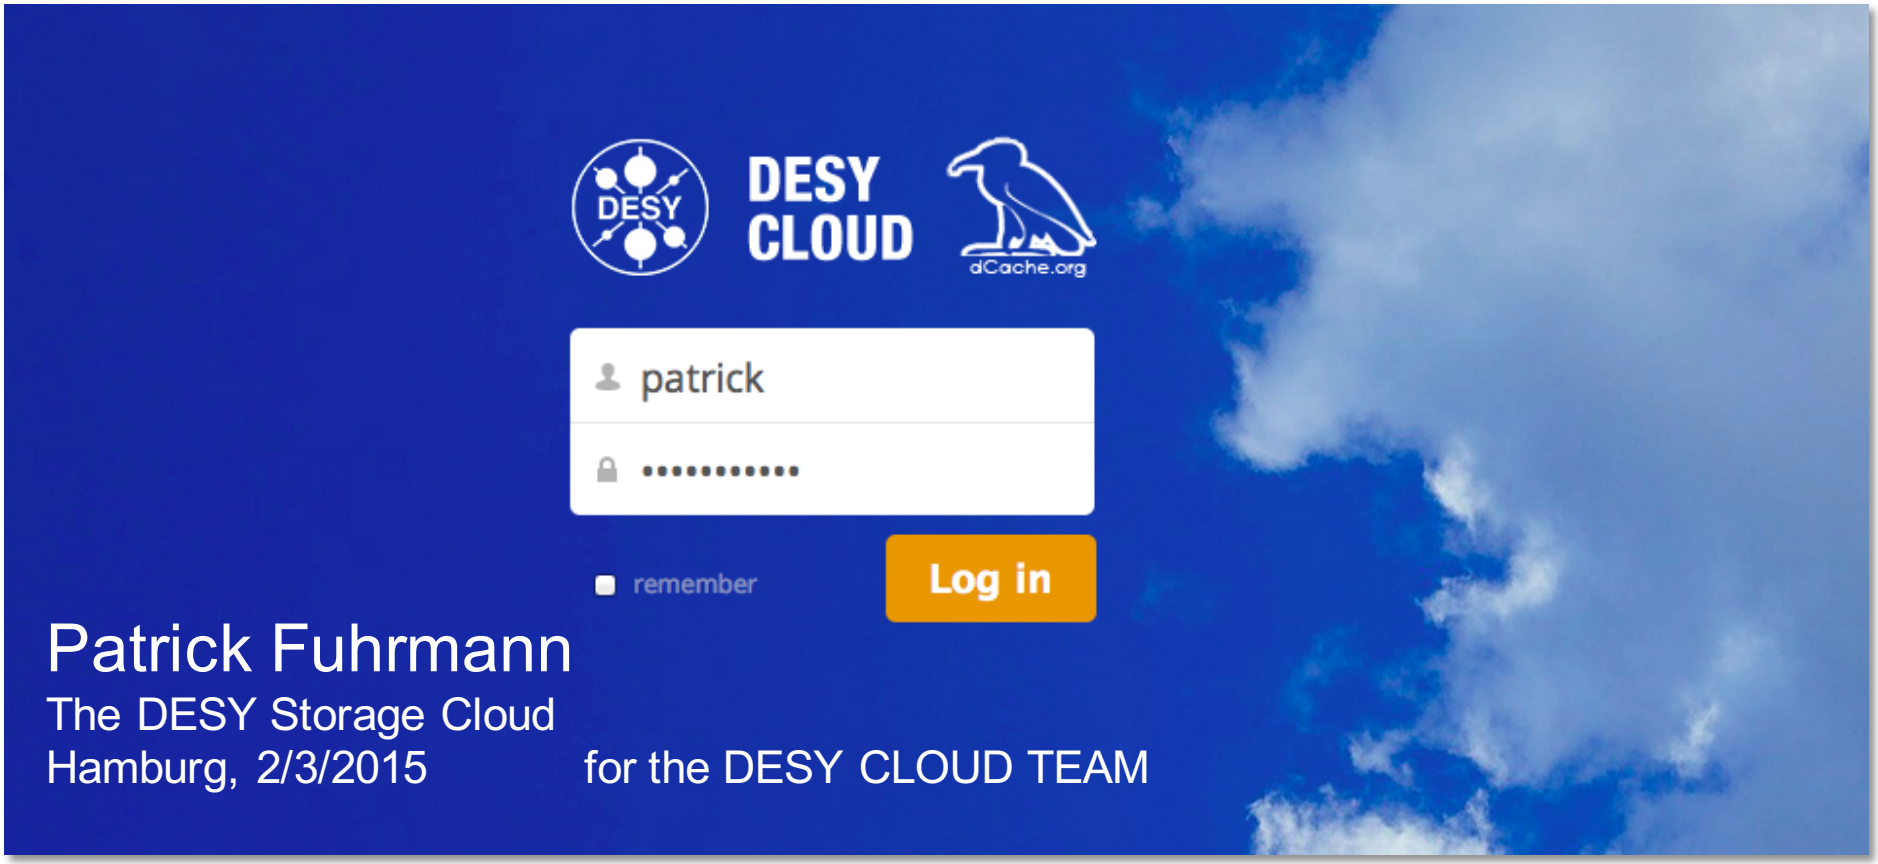

Patrick Fuhrmann

The DESY Storage Cloud

Hamburg, 2/3/2015 for the DESY CLOUD TEAM

HELMHOLTZ ASSOCIATION

DESY

Content

- Motivation
- Preparation
- Collaborations and publications
- What do you get right now ?
- Who is involved, how is it integrated ?
- How is the system implemented ?
- How do you get access ?
- The bigger picture.

What is this about … ?

It's about on how modern scientists (people) want to manage their data.

Requirements (Web2.0,Sync'n Share)

> New requirements from DESY users :

- New model in accessing data
  - Anytime from everywhere
  - From mobile devices
  - Bidirectional sync'ing between your cloud space and your local devices
- New model in sharing data
  - Fine grained sharing with individuals and groups.
  - Sharing via intuitive Web 2.0 mechanisms (Apps or Browser)
  - Sharing with 'public' with or w/o password protection
  - Sharing of free space (upload)
  - Expiration of shares

Requirements (QoS, Big Data)

- Request for unlimited, non destroyable storage space including data life cycle management.
- Request for different quality of services (SLA), coming with different price tags and controlled by customer.
  - Data Loss Protection (non-user introduced), e.g.:
    - One copy.
    - Two copies on independent systems.
    - Two copies in different buildings.
    - Two copies at different sites (e.g. Hamburg and Zeuthen)
    - Some of above plus ‘n’ tape copies.
  - Access latency and max data rate, e.g.:
    - Regular sync and web access.
    - Worker-node access: High throughput
    - Low latency (e.g. on SSD) for HPC.

Requirements (final)

- But all this should be just one huge repository. Not split into different systems.
- Access should be possible with different credentials
  - Username / password
  - X509 Certificates
  - SAML (Single Sign On)
  - Kerberos

DESY CLOUD dCache.org

Big Data, Qos Part

IT is doing this for decades. ~ 10 PetaBytes

Web 2.0

? Here we need some help

DESY

Decision

- For the web 2.0 interface we needed some experts.
- We went for the most popular solution: Own Cloud
  - Reduce likelihood for ‘product disappearing’
  - Possibly building a user-community (like today)
    - TU-Berlin, FZ-Jülich, TU-Dresden ****
    - CERN, United Nations
  - CERN is evaluating a similar approach and we are in contact anyway (WLCG)

Collaborations and publications

- Collaboration with HTW Berlin (LSDMA)
- Pre-evaluation of cloud solutions by “InFa” -> Q3/2013
  - Erarbeiten und Umsetzen eines firmeninternen Online-Speicherdienstes in einer Teststellung. (Quirin Buchholz)
- Presenting the concept at HEPIX.
- Information exchange with CERN. (CHEP’13) Oct 13
- Berlin Cloud Event, (mostly OwnCloud and PowerFolder) in Mai 14 (we published first paper)
- Participating the CERN Cloud Event (Nov ‘14) including a presentation of our proposed solution.
- Various papers submitted and accepted at ISGC in Taipei in March and CHEP’15 in Japan.

Last Preliminary Remark

- The DESY Storage Cloud is NOT in the CLOUD.
- Your data is stored on DESY premises and is regulated by DESY agreements. (German and EU law).
- We only provide interfaces to/into your storage space, which are similar to those provided by Cloud Providers.

So, what do you get right now ?

Sync'ing and Web access

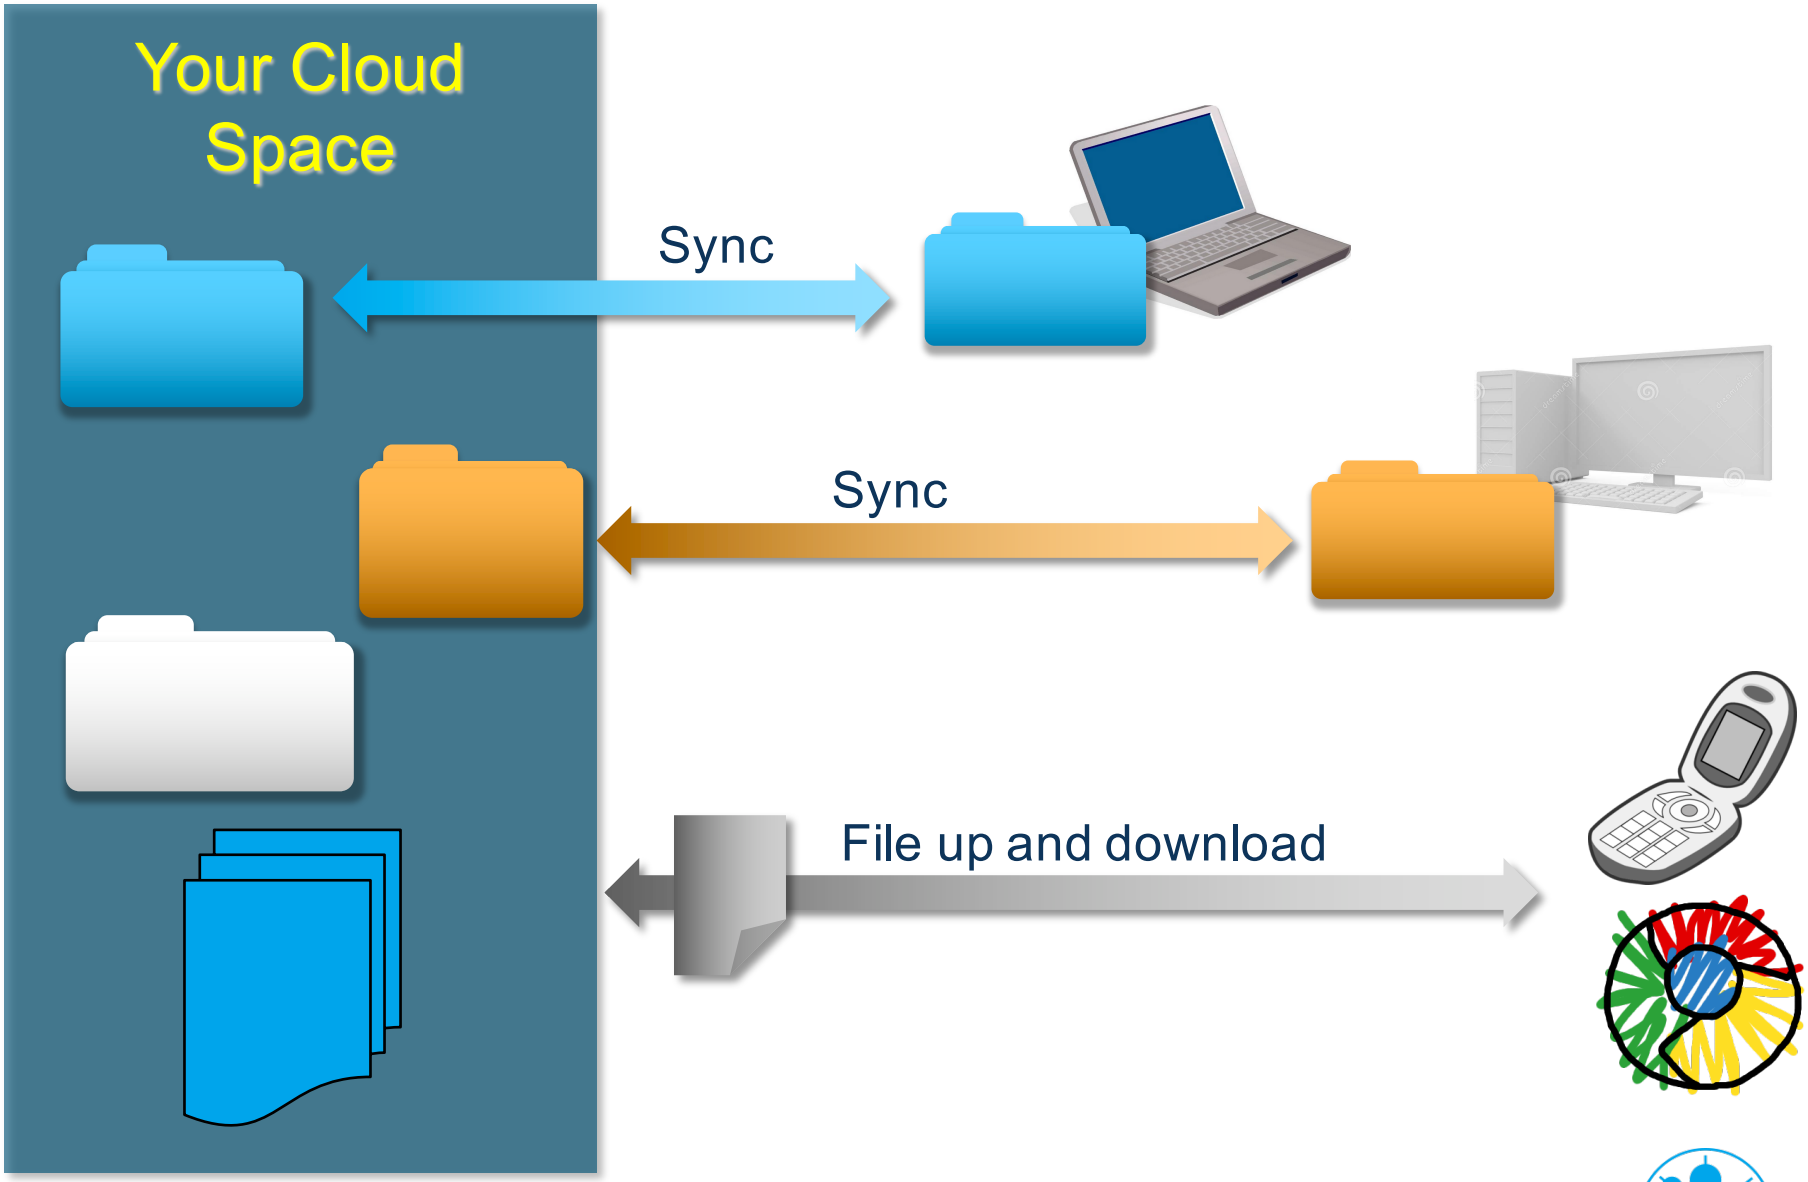

Sharing

Share files/folders with individuals

Share files/folders with ‘desy groups’

Share with ‘public’ with and w/o password

(Shares can expire)

Share space(s) with others for upload

Others sharing data with you (in your home)

Initial Quality of service

- Each file in your space has two copies on a completely distinct hardware setup.
  - Different disks
  - Different controller
  - Different host system.
- Checksum is calculated on arrival of backend storage.
- Continuous ‘scrubbing’ of disks to detect wrong checksums.

Who do we do that ?

Integration into the DESY infrastructure

DESY CLOUD

dCache.org

Monitoring

Accounting

User Management
Registry
LDAP

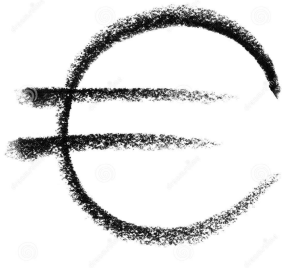

Virtualization

vmware®

Authentication
Kerberos

ownCloud

Local and Wide
Area Network
Load Balancing
Firewalls

Unlimited
Persistent
Storage

DESY

IT Groups involved

- Almost all IT groups are involved (See Design Picture)
  - Technical Coordination OwnCloud : Quirin
  - Technical Coordination DOT: Lusine and Ralf
  - Technical Coordination dCache.org : Tigran
  - And many more In Network, VM Management, Web Office …
  - Please don't contact those people directly, use weboffice@desy.de to report issues.

The Own Cloud Part

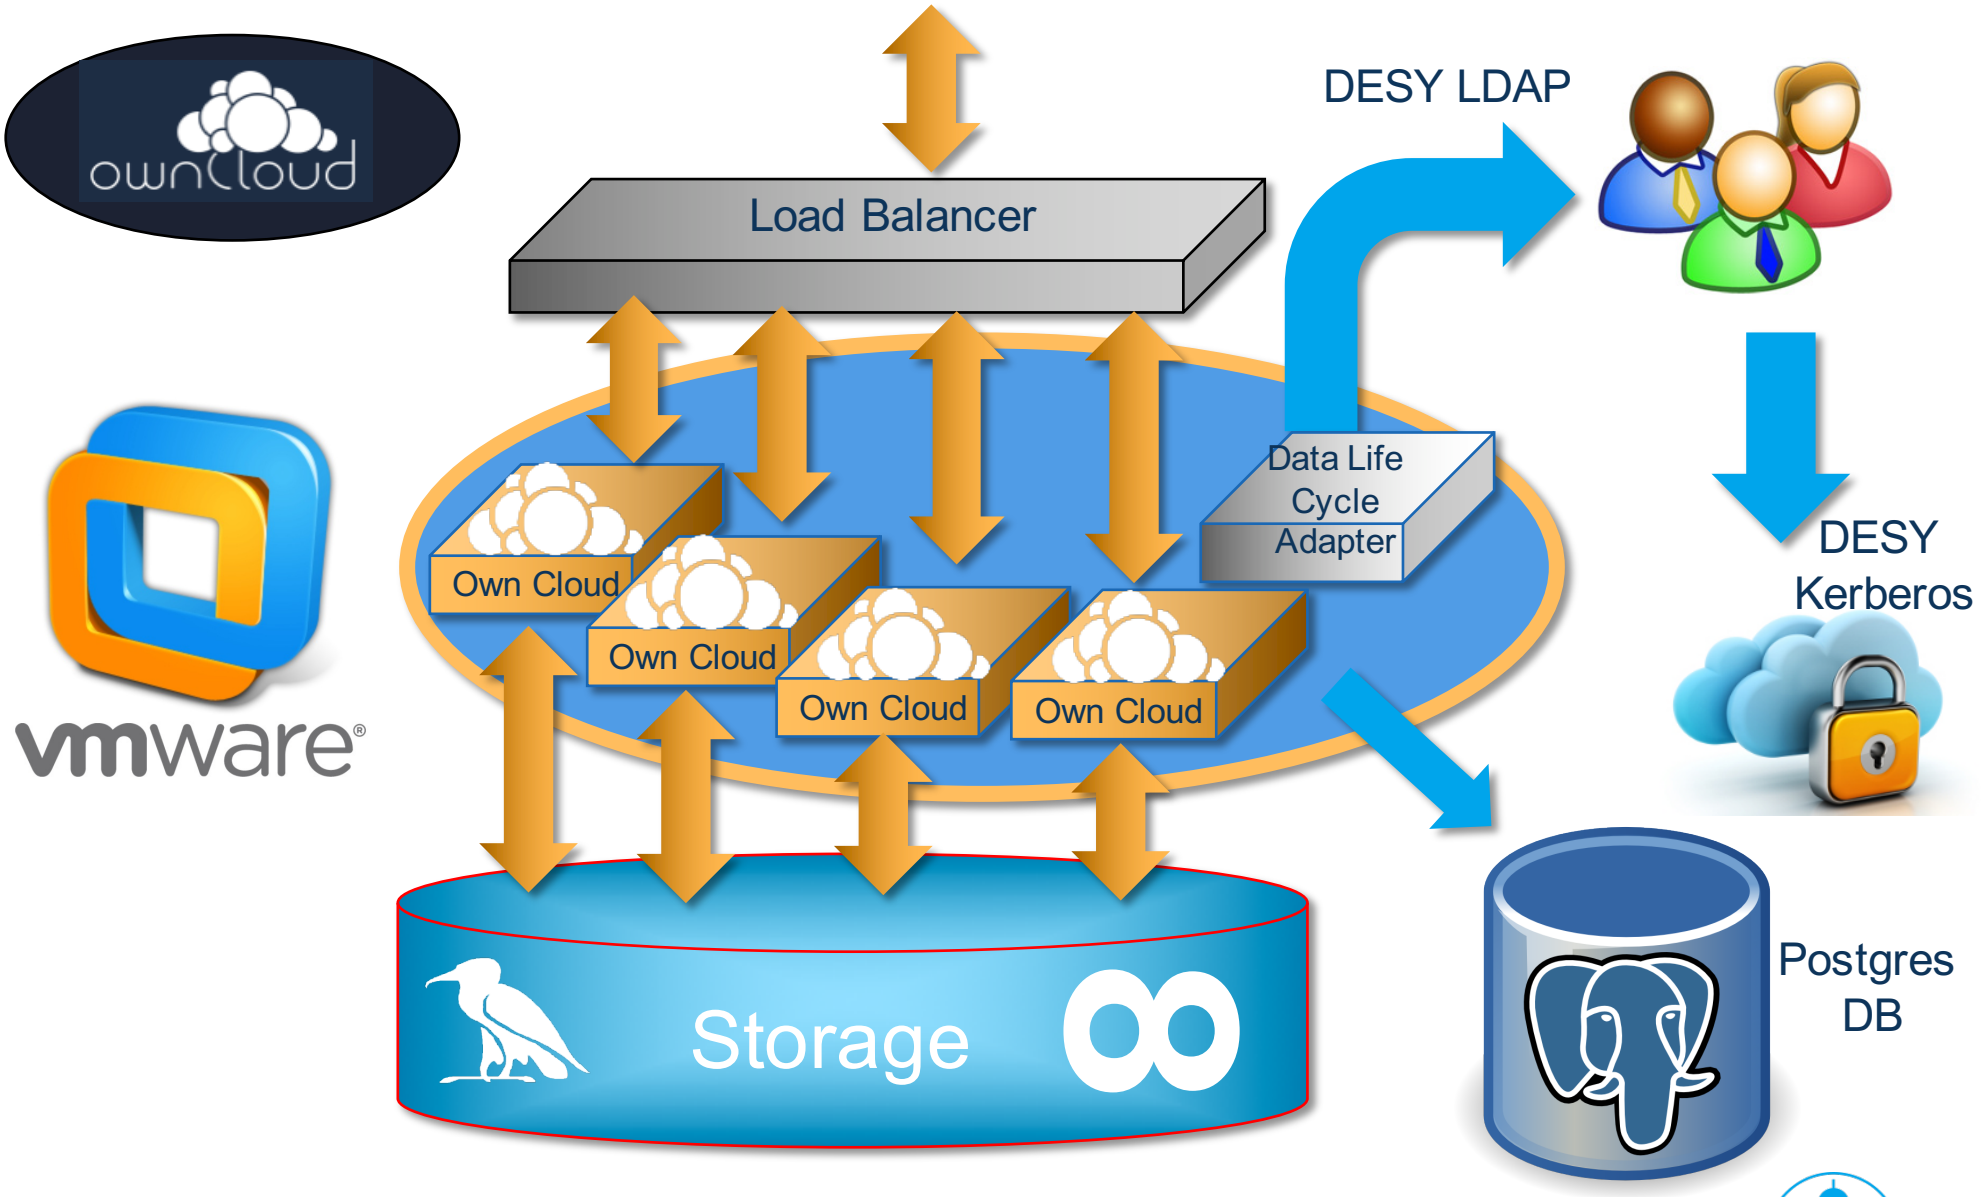

The dCache backend

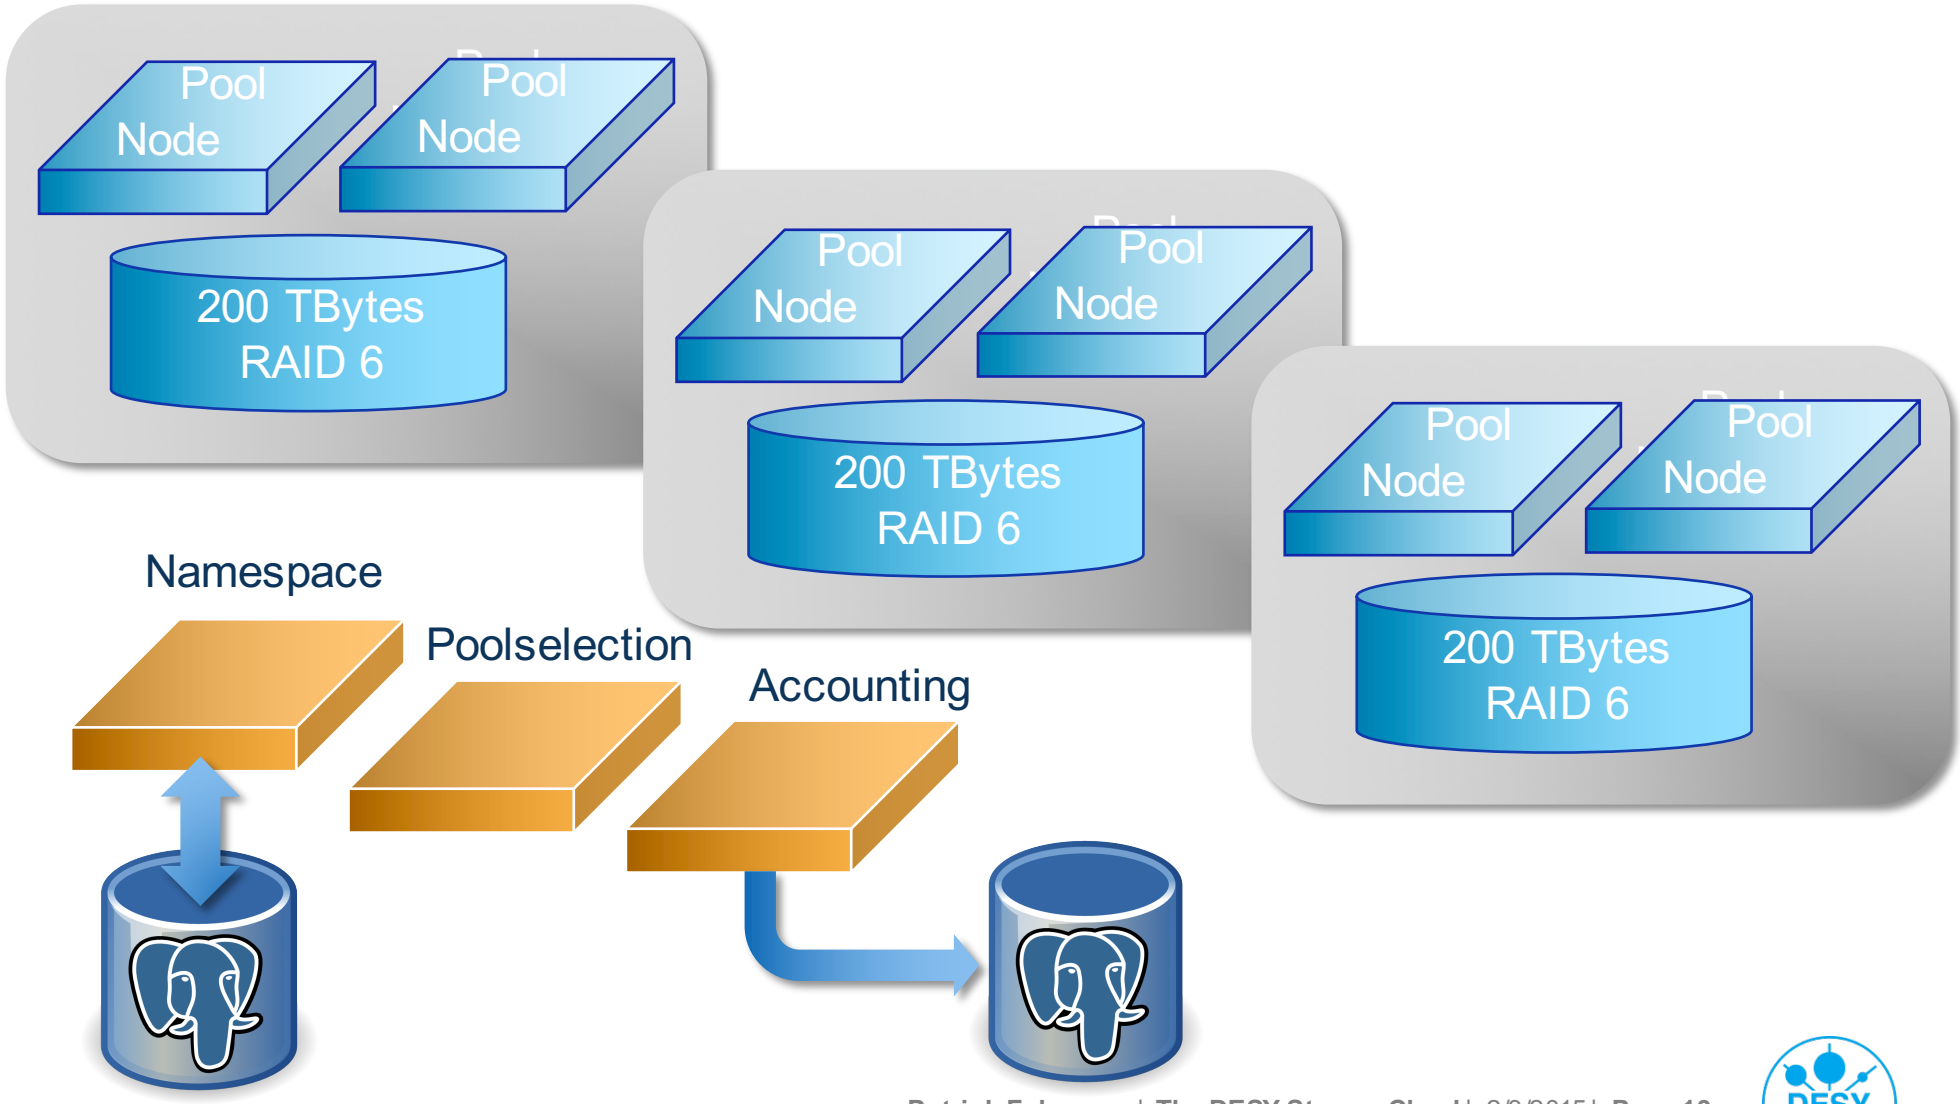

Horizontal Scaling

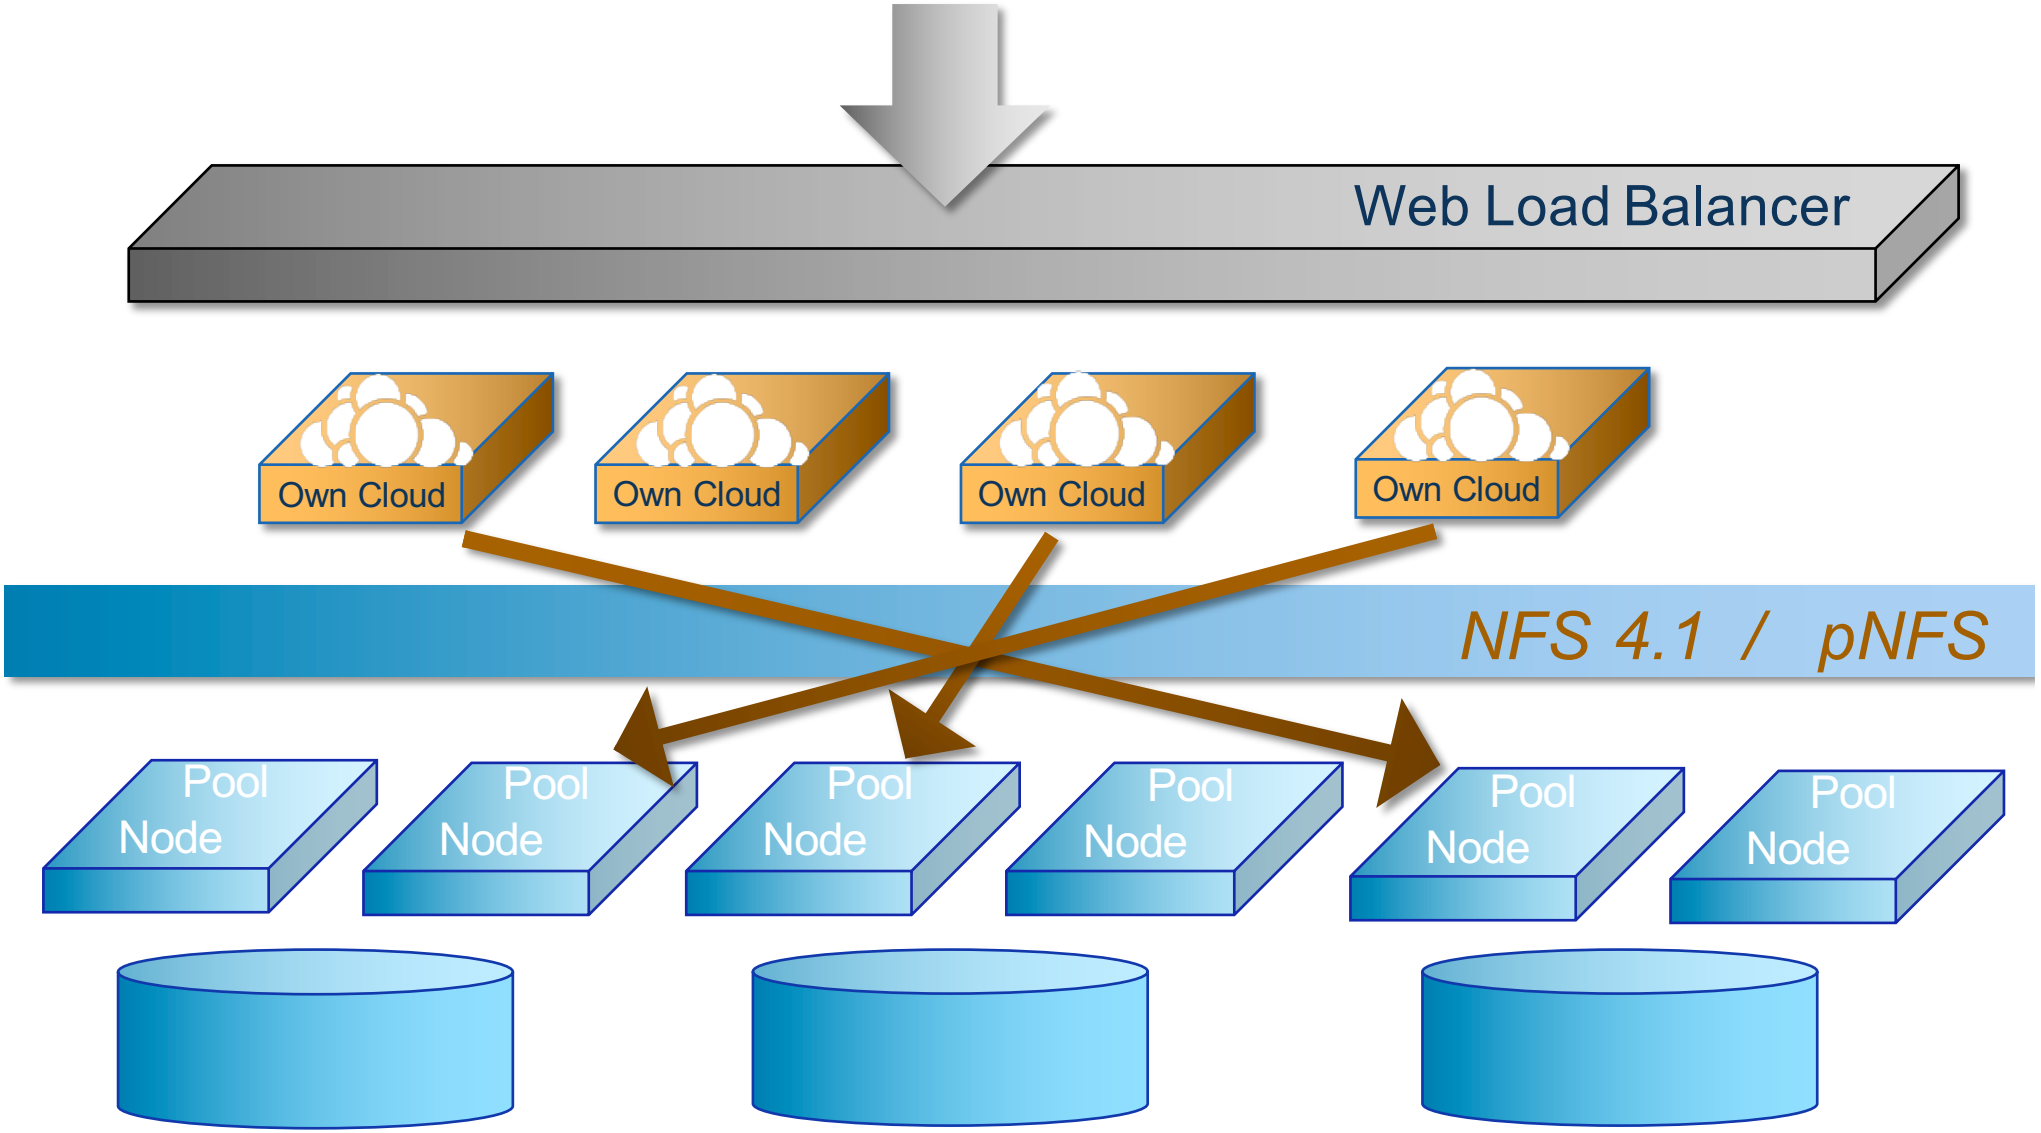

How do you get access ?

What you need to do, to get access

Get the OwnCloud resource from the DESY registry.
(Here vdR wanted to say something.)

Read the intro @
It.desy.de -> Dienste -> StorageDienste -> DesyCloud

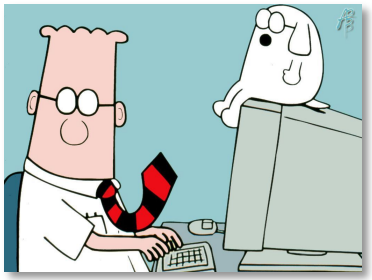

1. Log into desycloud.desy.de with your DESY account credentials
2. Install the OwnCloud client for your preferred OS
3. Configure your sync's and shares

Report issues to weboffice@desy.de

Things you might want to consider

Syncing might need some thinking

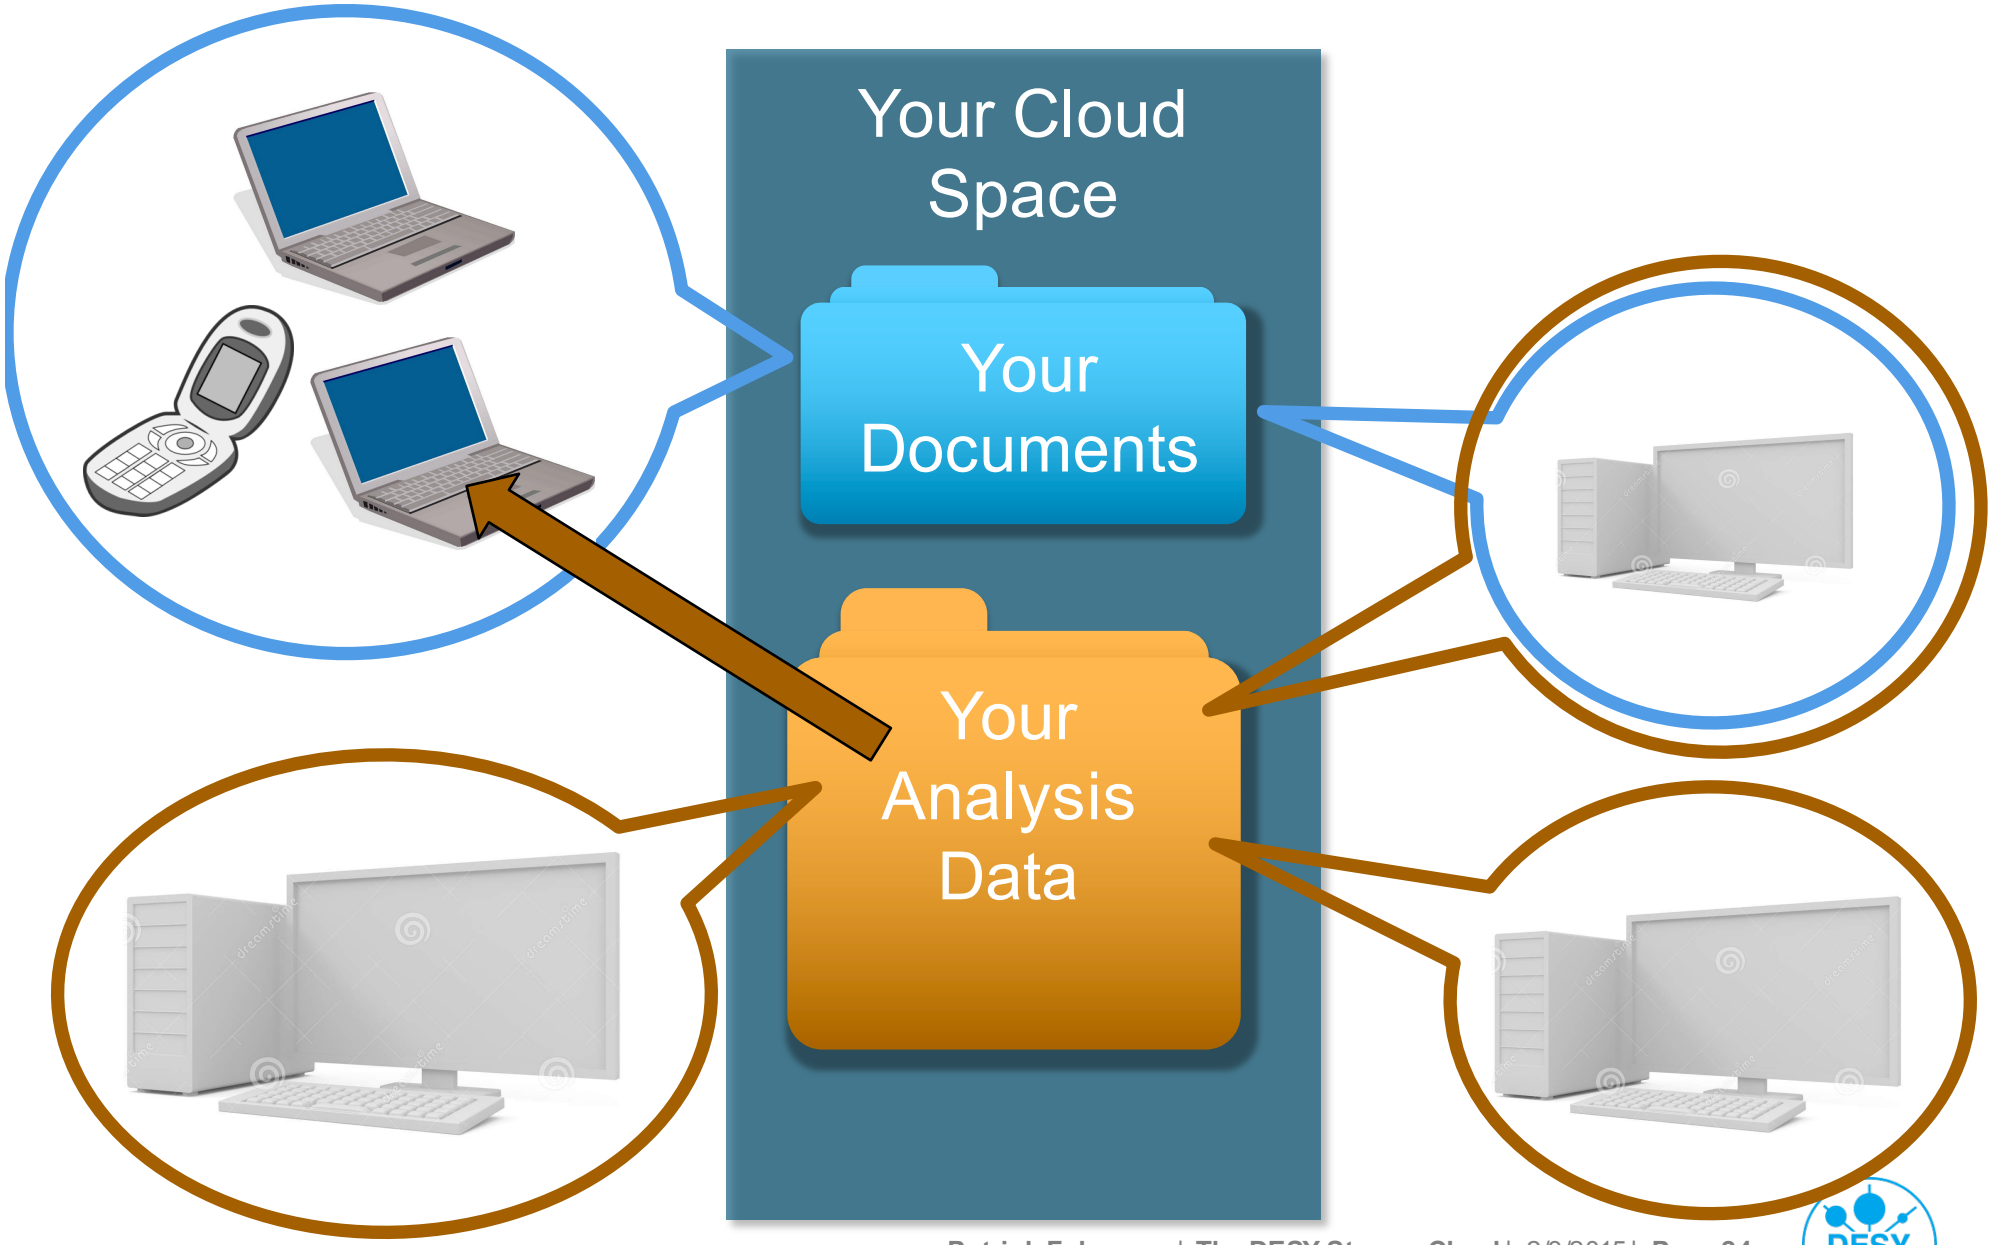

Watch OUT !

- !!! Do not sync your entire Own Cloud Space with your laptop -> chose a subdirectory instead.
  - Own Cloud Space is potentially unlimited … your laptop isn't
  - Everything someone else shares with you would be sync'ed to your laptop. Shares always end up in our home first.

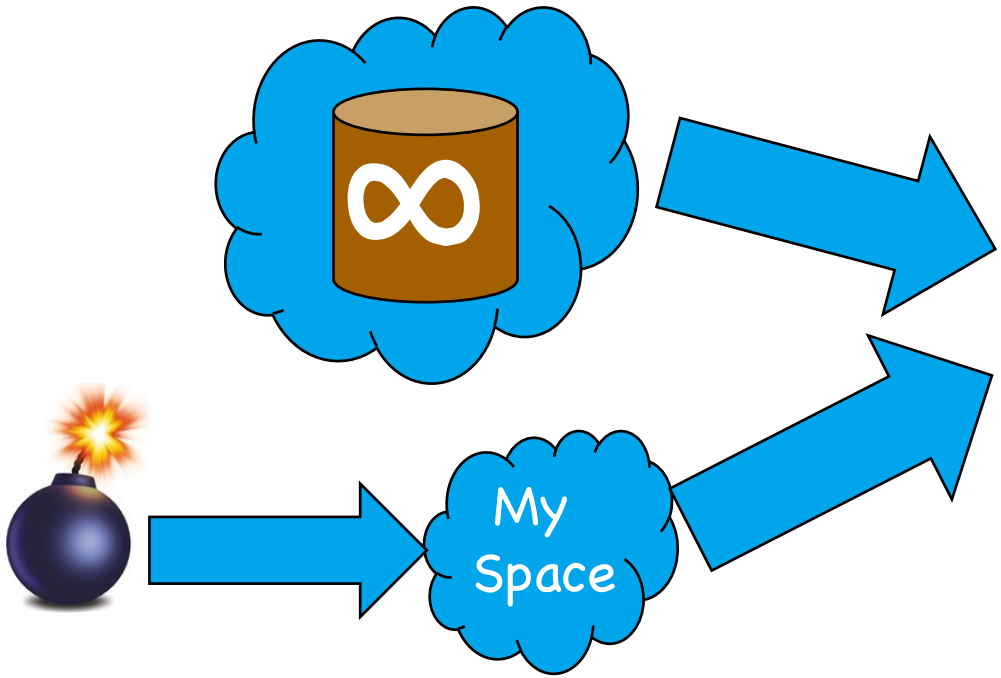

Watch out !!!

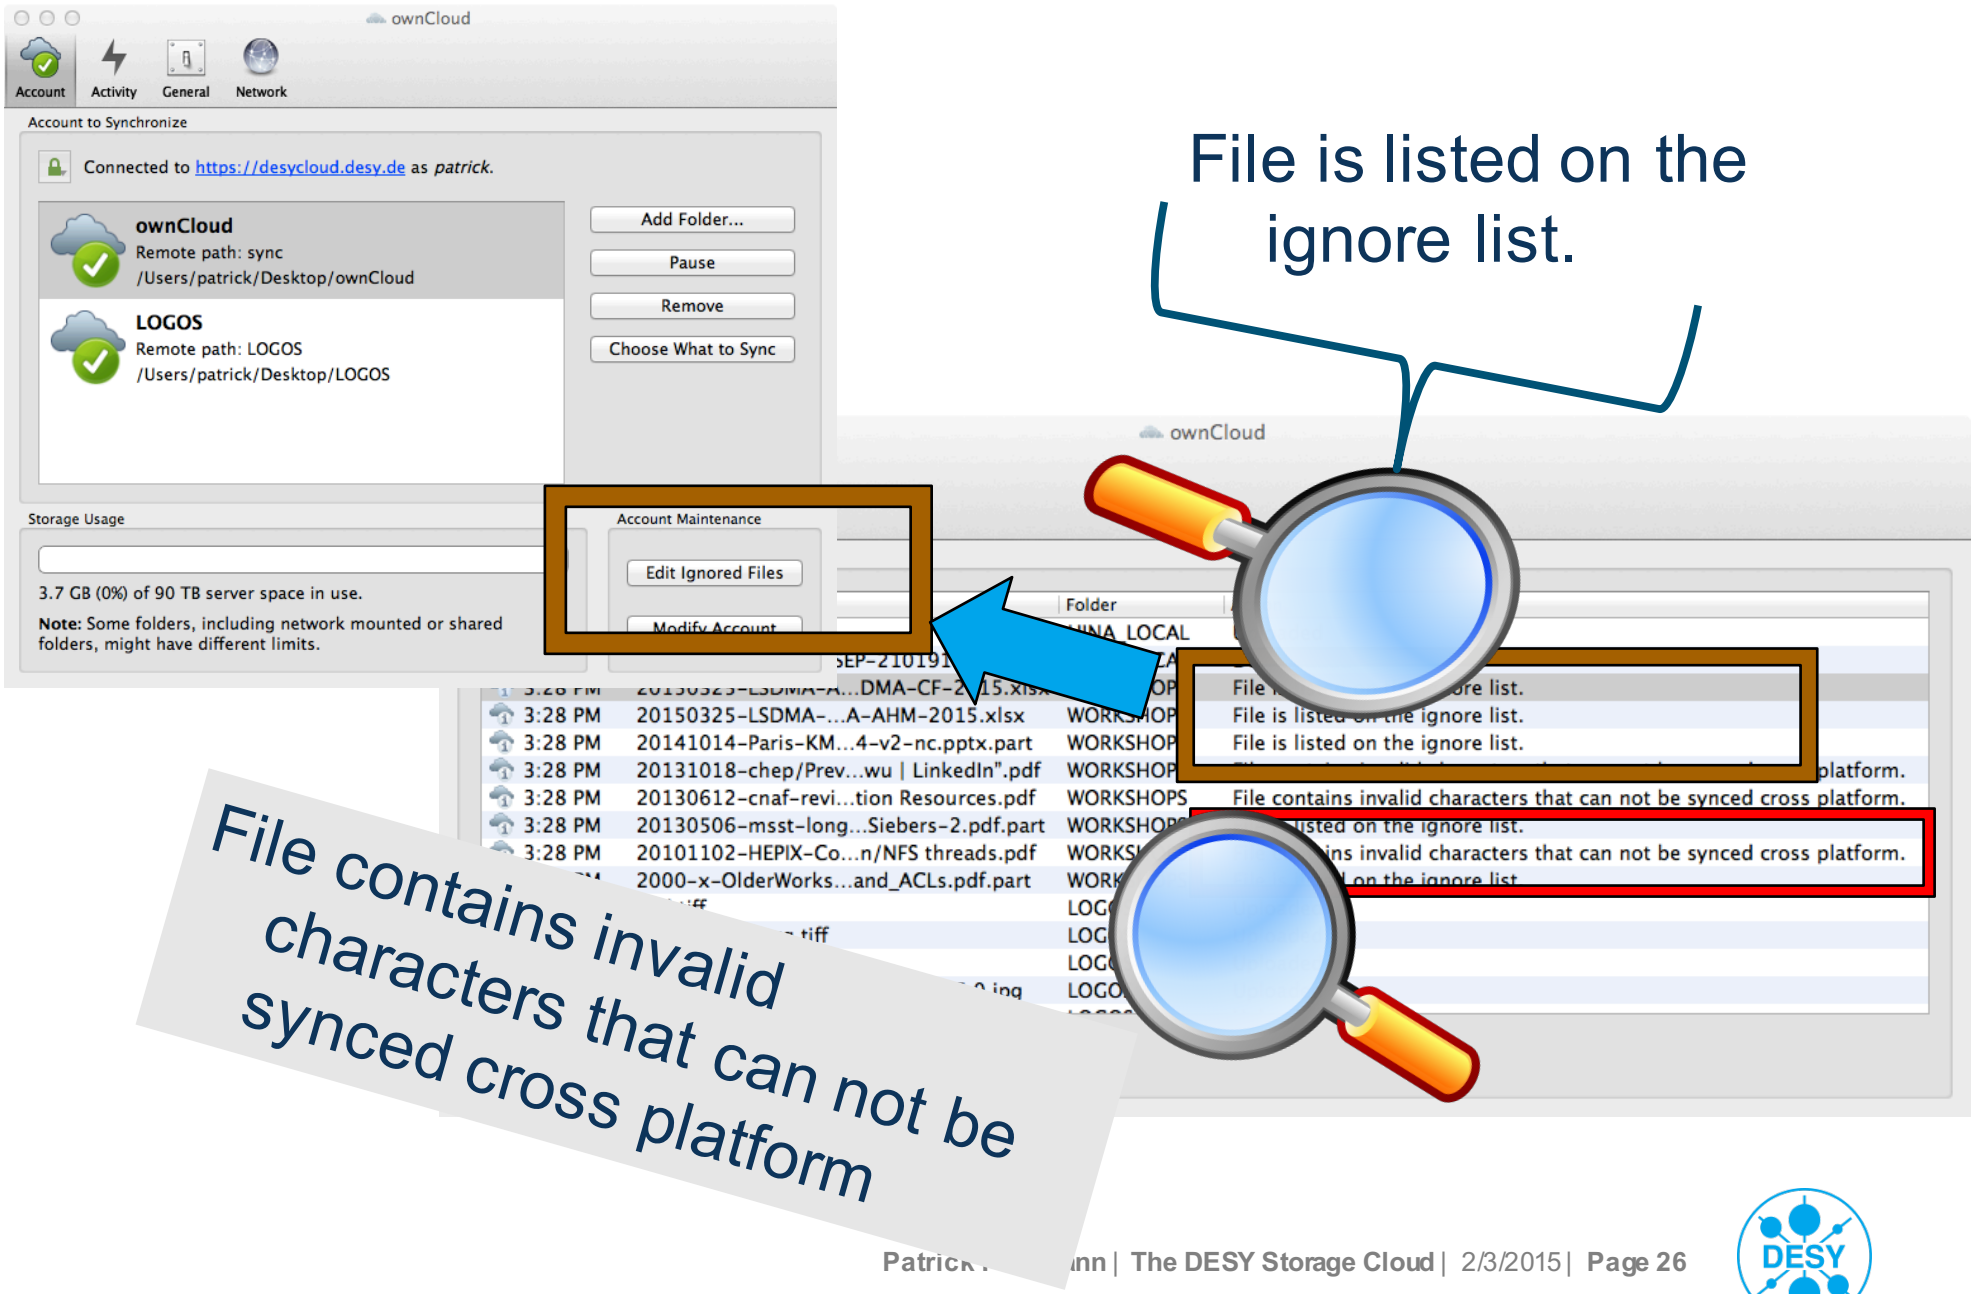
ownCloud
Account
Activity
General
Network
Account to Synchronize
Connected to https://desycloud.desy.de as patrick.
ownCloud
Remote path: sync
/Users/patrick/Desktop/ownCloud
LOGOS
Remote path: LOGOS
/Users/patrick/Desktop/LOGOS
Add Folder...
Pause
Remove
Choose What to Sync
Storage Usage
3.7 GB (0%) of 90 TB server space in use.
Note: Some folders, including network mounted or shared folders, might have different limits.
Account Maintenance
Edit Ignored Files
Modify Account
File is listed on the ignore list.
ownCloud
Folder
3:28 PM 20150325-LSDMA-...A-AHM-2015.xlsx
3:28 PM 20141014-Paris-KM...4-v2-nc.pptx.part
3:28 PM 20131018-chep/Prev...wu | LinkedIn".pdf
3:28 PM 20130612-cnaf-revi...tion Resources.pdf
3:28 PM 20130506-msst-long...Siebers-2.pdf.part
3:28 PM 20101102-HEPIX-Co...n/NFS threads.pdf
2000-x-OlderWorks...and_ACLs.pdf.part
File is listed on the ignore list.
File contains invalid characters that can not be synced cross platform.
File contains invalid characters that can not be synced cross platform

Folder Sharing Hierarchy

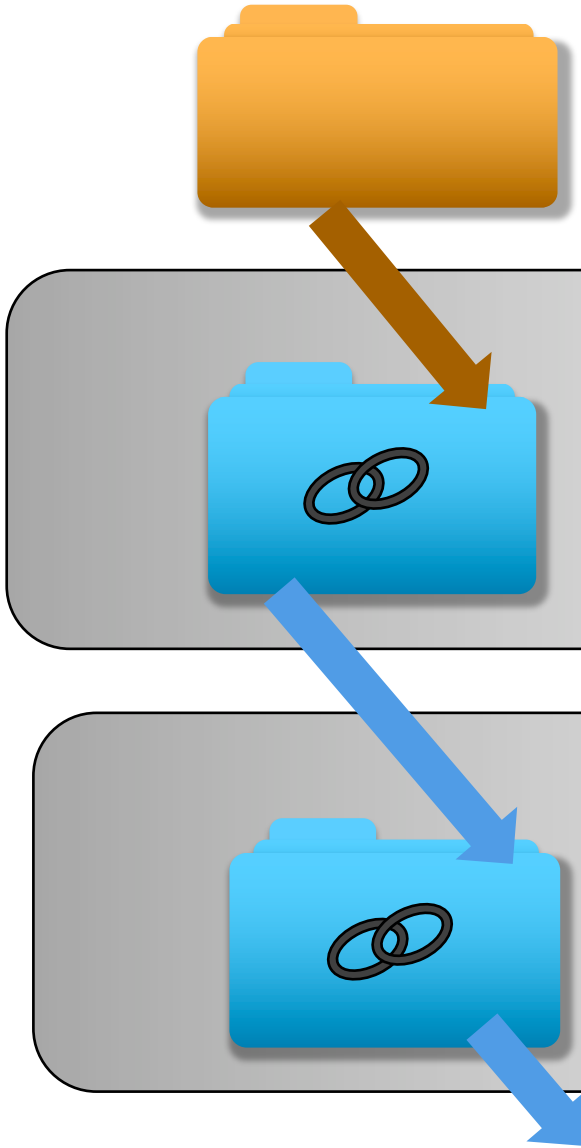

Your very private folder

Actively Created public link

https://desycloud.desy.de/public.php?service=files&t=3bd49ff864c465d362d697a54584a6e3

Implicitly Created public link

https://desycloud.desy.de/public.php?service=files&t=3bd49ff864c465d362d697a54584a6e3

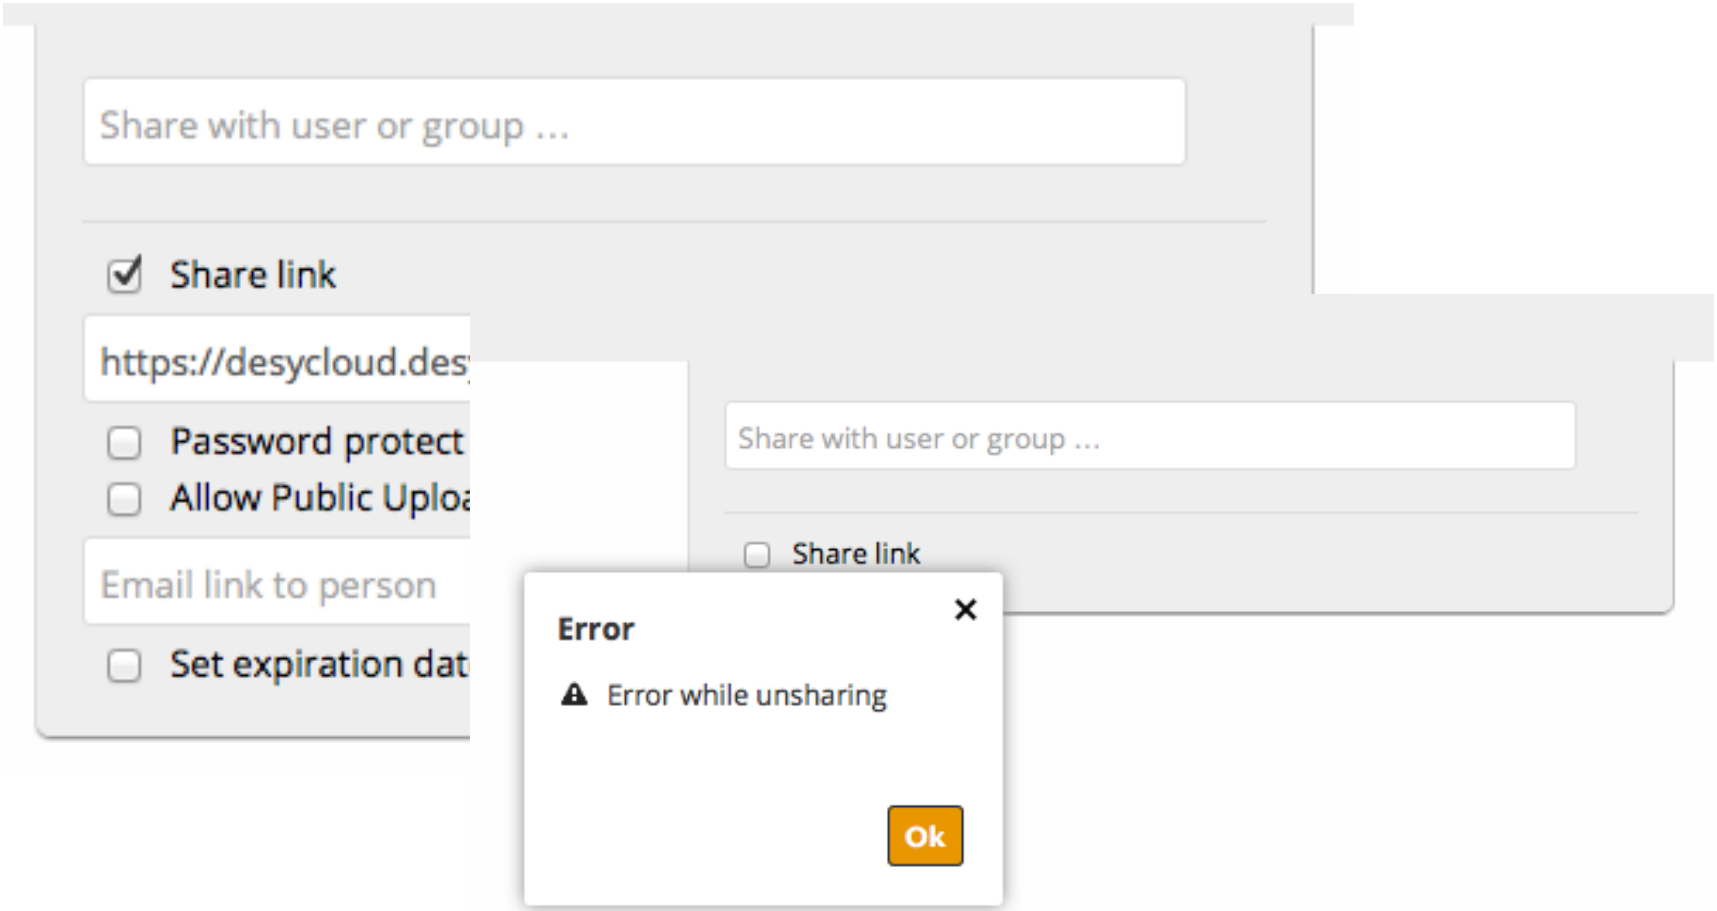
Share with user or group …
Share link
https://desycloud.des
Password protect
Allow Public Uploa
Email link to person
Set expiration dat
Share with user or group …
Share link
Error
Error while unsharing
Ok

More Wath Out

We'll try to keep the documentation up to day on issues we find.

Towards the scientific storage cloud

We are not running OwnCloud storage

BUT

We are using OwnCloud to provide a particular VIEW into your DESY storage space.

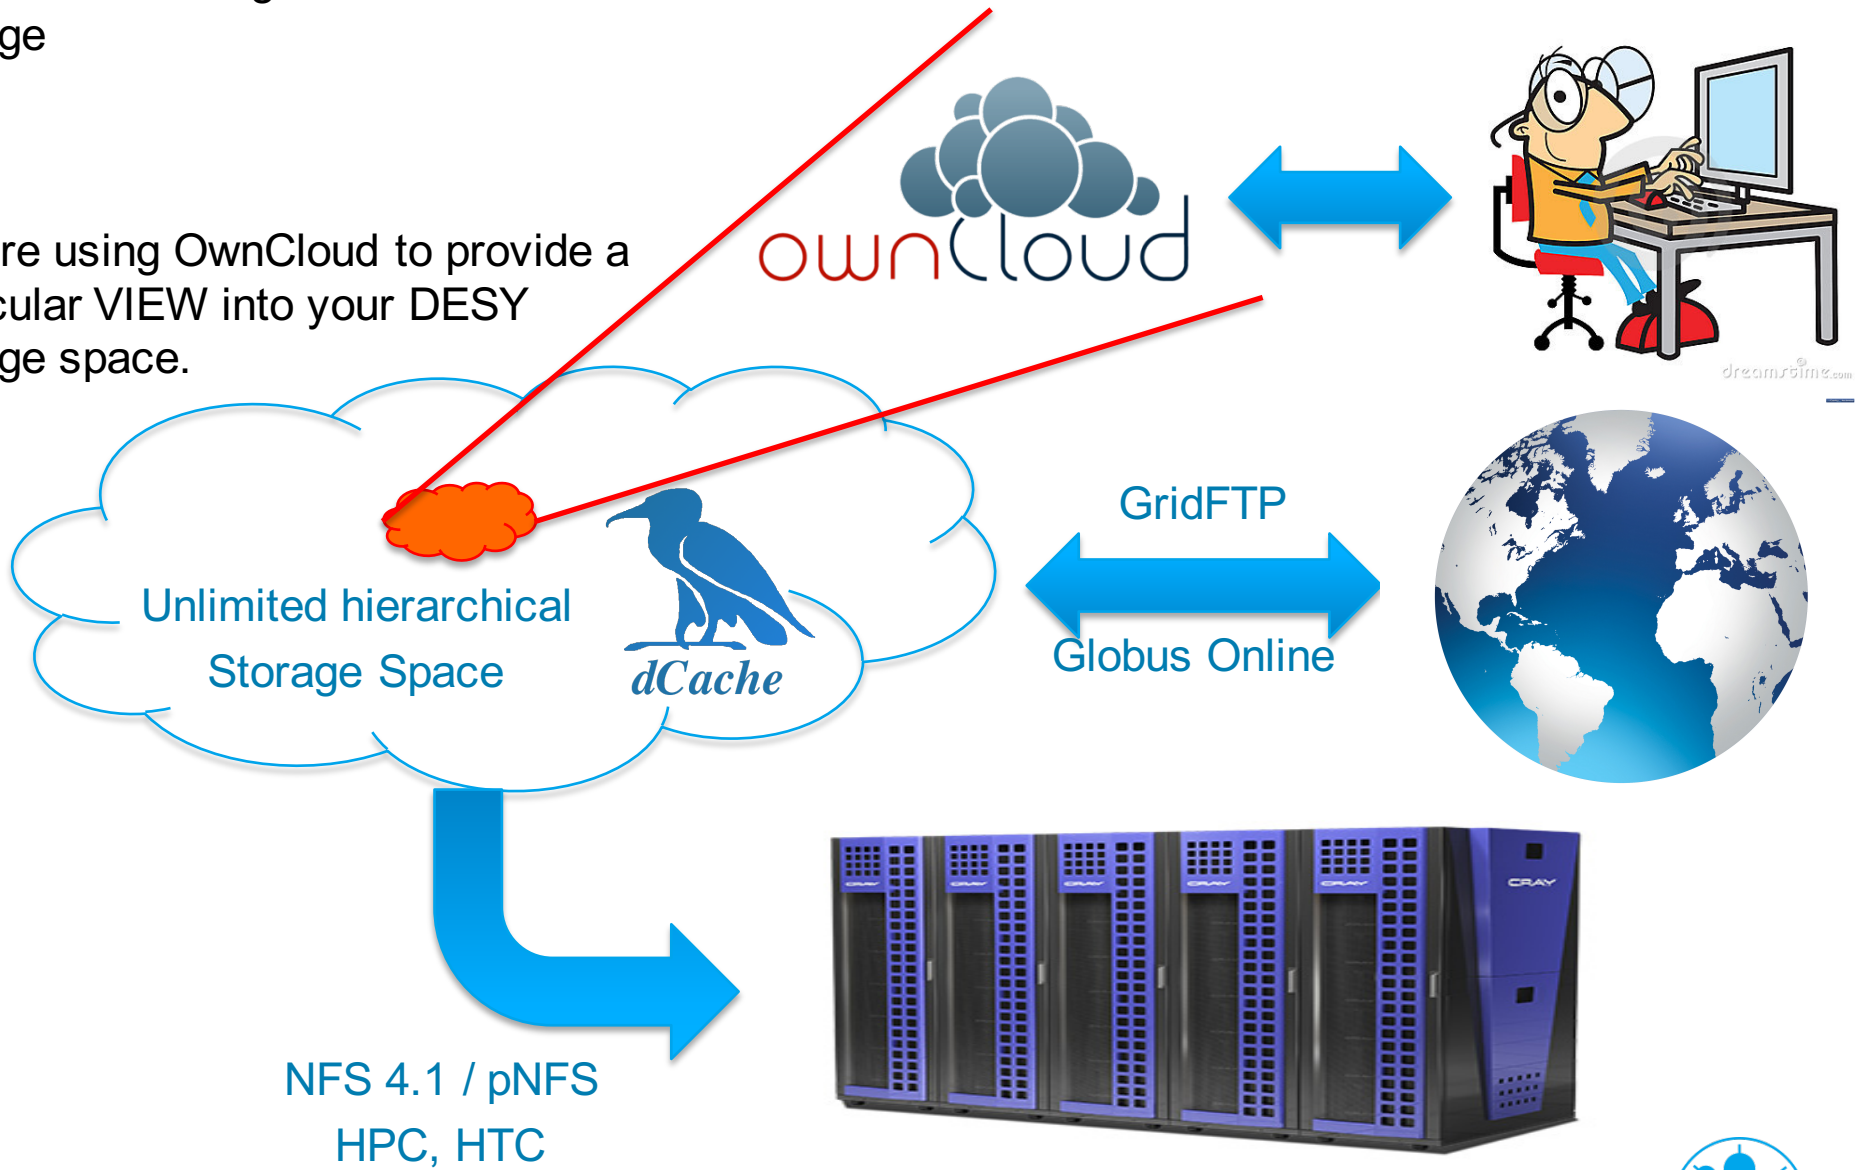

The big picture

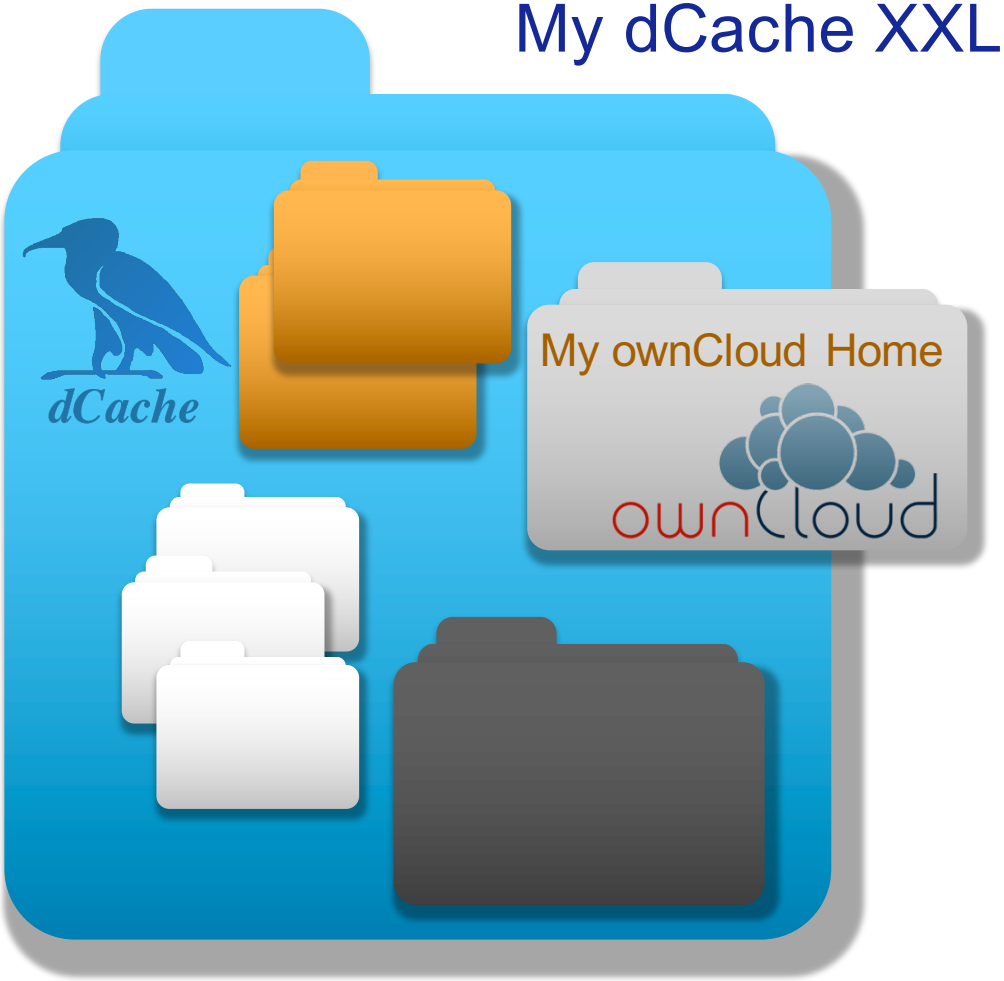

Sync
Share
Web 2.0

NFS 4.1/pNFS
GridFTP
WebDAV
SRM
(some private
Grid Protocols)
dCap
xRootD

High performance access

NFS mounted,

High performance,
Low latency

On workgroup server

Or worker nodes,
Compute servers

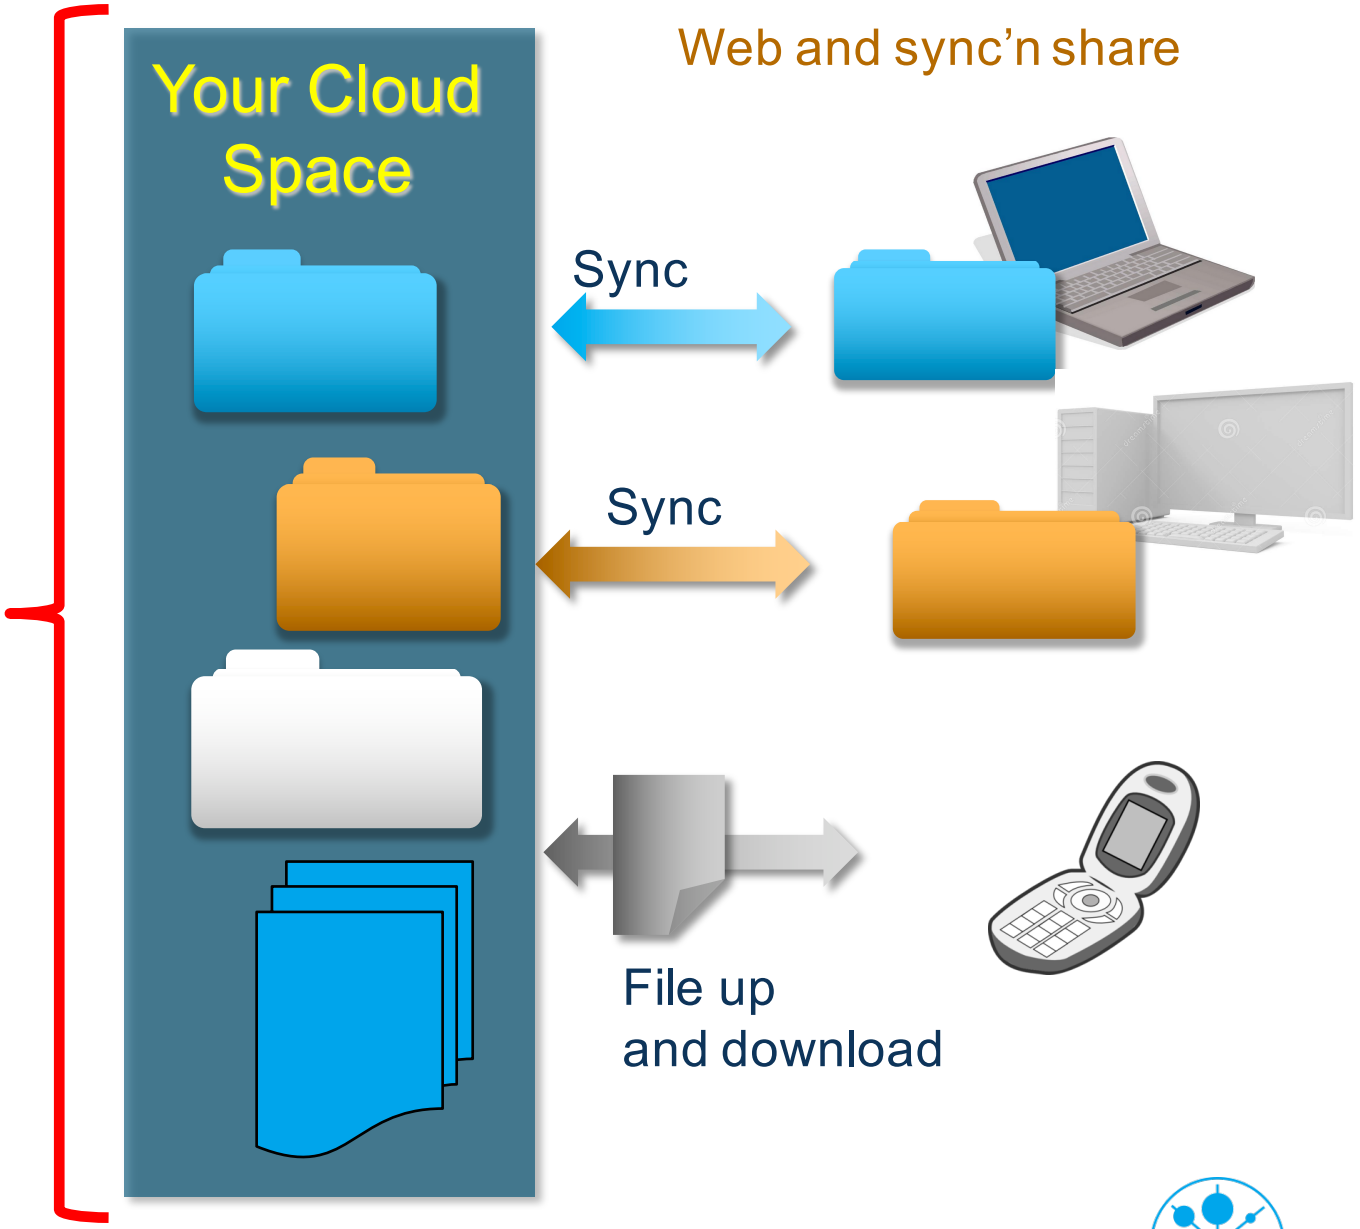

Web and sync'n share

Wide area file transfer

Wide are transfers

With Globus Online
Or WLCG Transfer tools
(using GridFTP)

Your Cloud Space

Web and sync'n share

Sync

Sync

File up
and download

Quality of service management

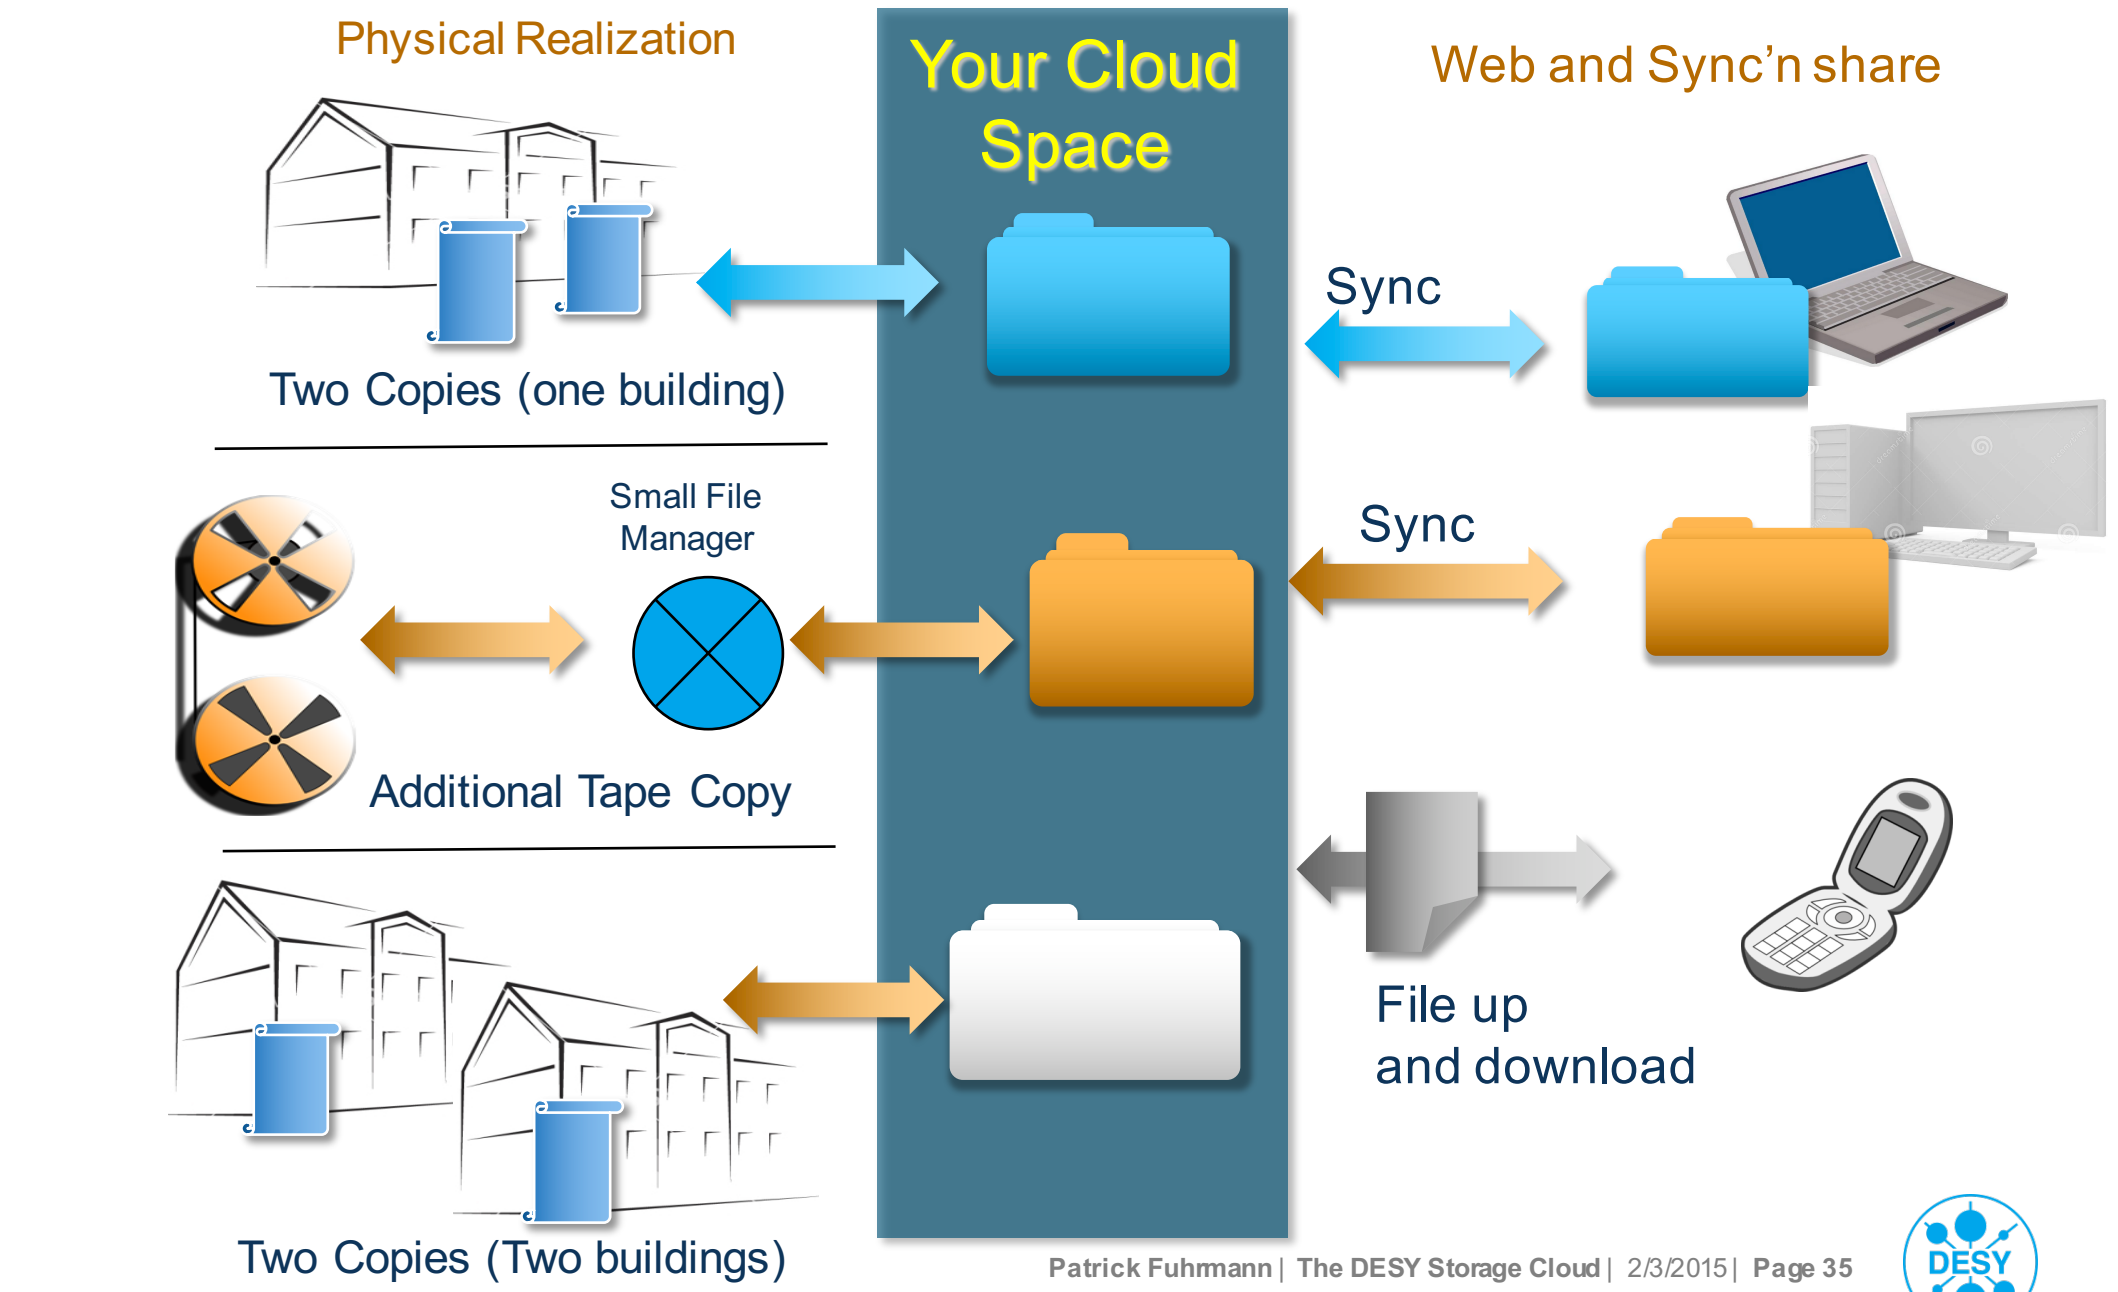

Even bigger system

Hamburg *Zeuthen*

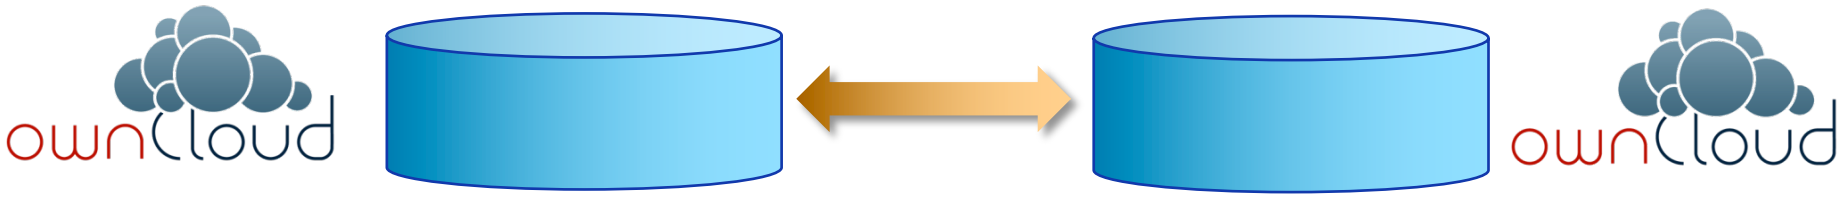

- One Namespace -> Transparent Sharing
- Zeuthen is ‘Backup’ for Hamburg, Hamburg is ‘Backup’ for Zeuthen.
- Zeuthen is ‘Cache’ for Hamburg storage if accessed in Zeuthen and vise versa.
- This is based on dCache features. Those setups already exist in North Europe and University of Michigan.

- All those nice features are already available in the used backend. (dCache)
- They only need to be enabled.
- This will happen step by step.
- More a political issue, concerning resources.

Just that I don’t forget …

Just not to forget

- *The system is still a pilot*
  - *Which should only influence the overall availability.*
  - *E.g. : Downtimes on short notice.*
  - *We already put significant efforts into keeping your data safe.*
- Payment mechanisms for mobile App's not clear.
  - Less than 1.00 Euro.
- Sharing with groups
  - Only people with the registry "OwnCloud resource" will get access.
  - Otherwise you would have to share with public link.
- In monitoring but not nightalarm yet.

Just not to forget

- *WebDAV access.*
  - *Due to limitations in the definition of WebDAV and the implementation of some OS clients, we are still trying to find a client which works sufficiently satisfying for everyone.*
  - *BTW : With WebDAV, currently 'unlimited storage space' means 90 TBytes.*
- *Issues with the NFS connection between OwnCloud and dCache.*
  - *We observe hick-ups. (Remote I/O Errors)*
    - *Currently taken care of by script*
    - *You shouldn't notice*
    - *Will be fixed.*
    - *Side effect: Web interface might ask you to log-in again.*

Scientific Storage Cloud

DESY CLOUD
dCache.org

Cloud User

WebDAV
HTTP(S)

Globus Online

Cloud

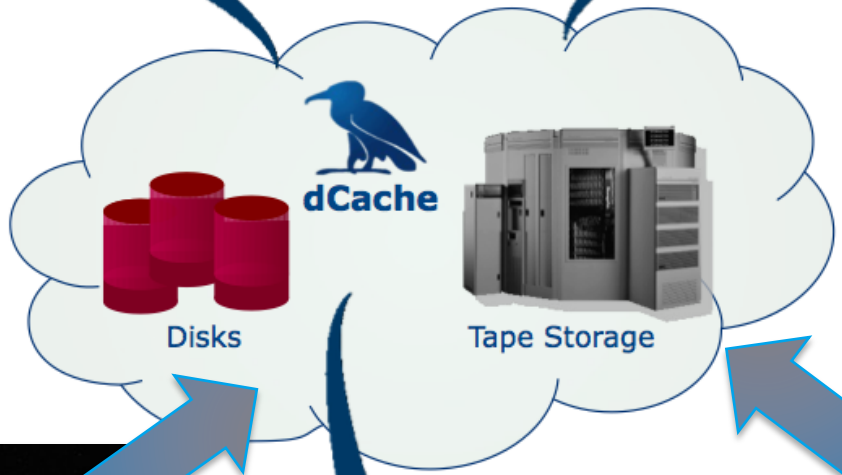

European XFEL

X-FEL
(Free Electron Lasers)
Fast Ingest

LOFAR antenna
Huge amounts of data

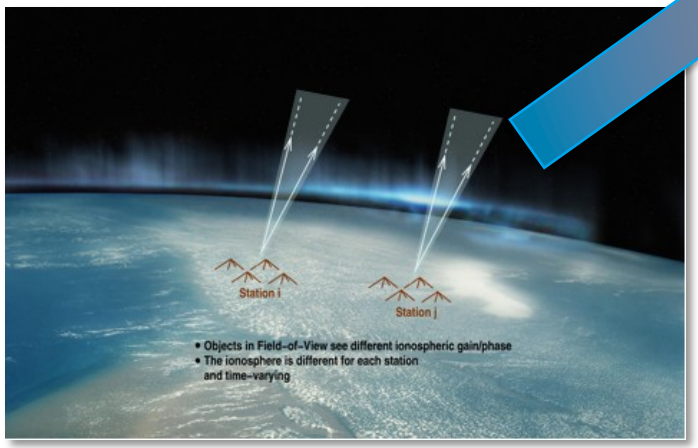

Mounted POSIX FS
(NFSv4.1, pNFS)

Computer farm

DESY