dCache.org

dCache

- sync'n share
- Software defined storage (QoS)

Patrick Fuhrmann

On behave of the project team

What happens with ‘sharing’ ?

ownCloud

Nextcloud

Seafile

PowerFolder

Sharing requirements from DESY users

- Accessing data anytime from everywhere
- 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

And to cut a long storage short

OR

Simplified Technical Perspective

dCache.org

ownCloud Clients

Web Load Balancer

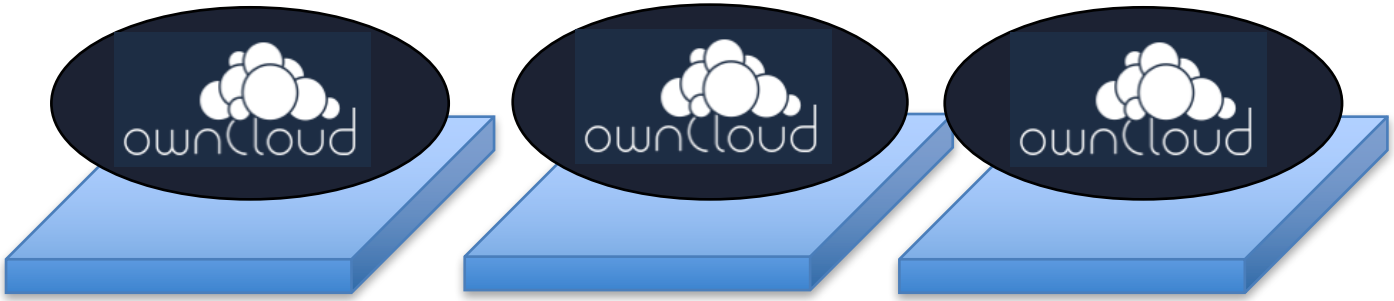

OwnCloud Server

pNFS Client

pNFS Server

Quick Update on the DESY installation (for details, see Birgit)

The Own Cloud Part

dCache.org

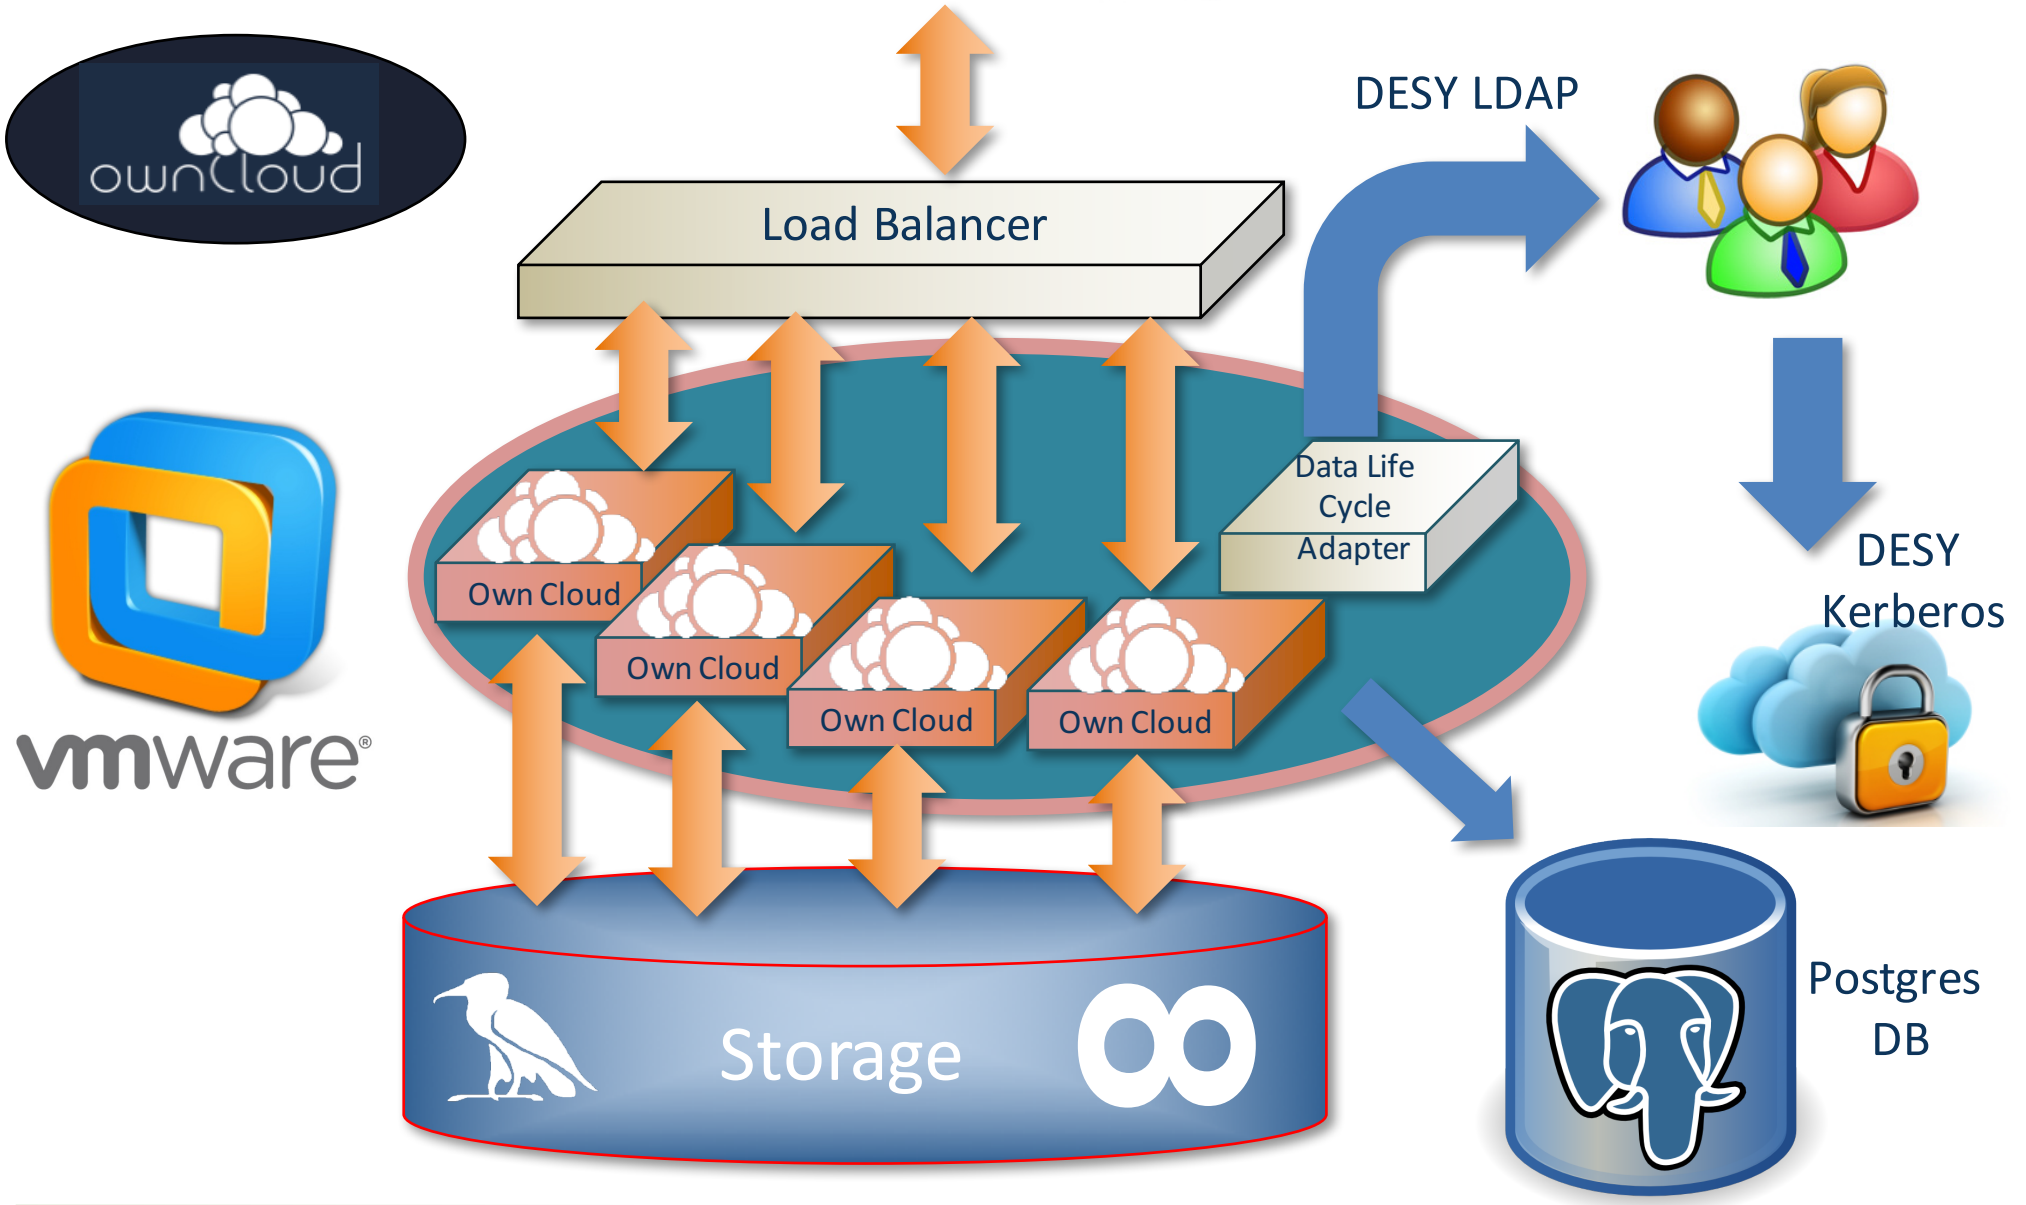

The dCache part

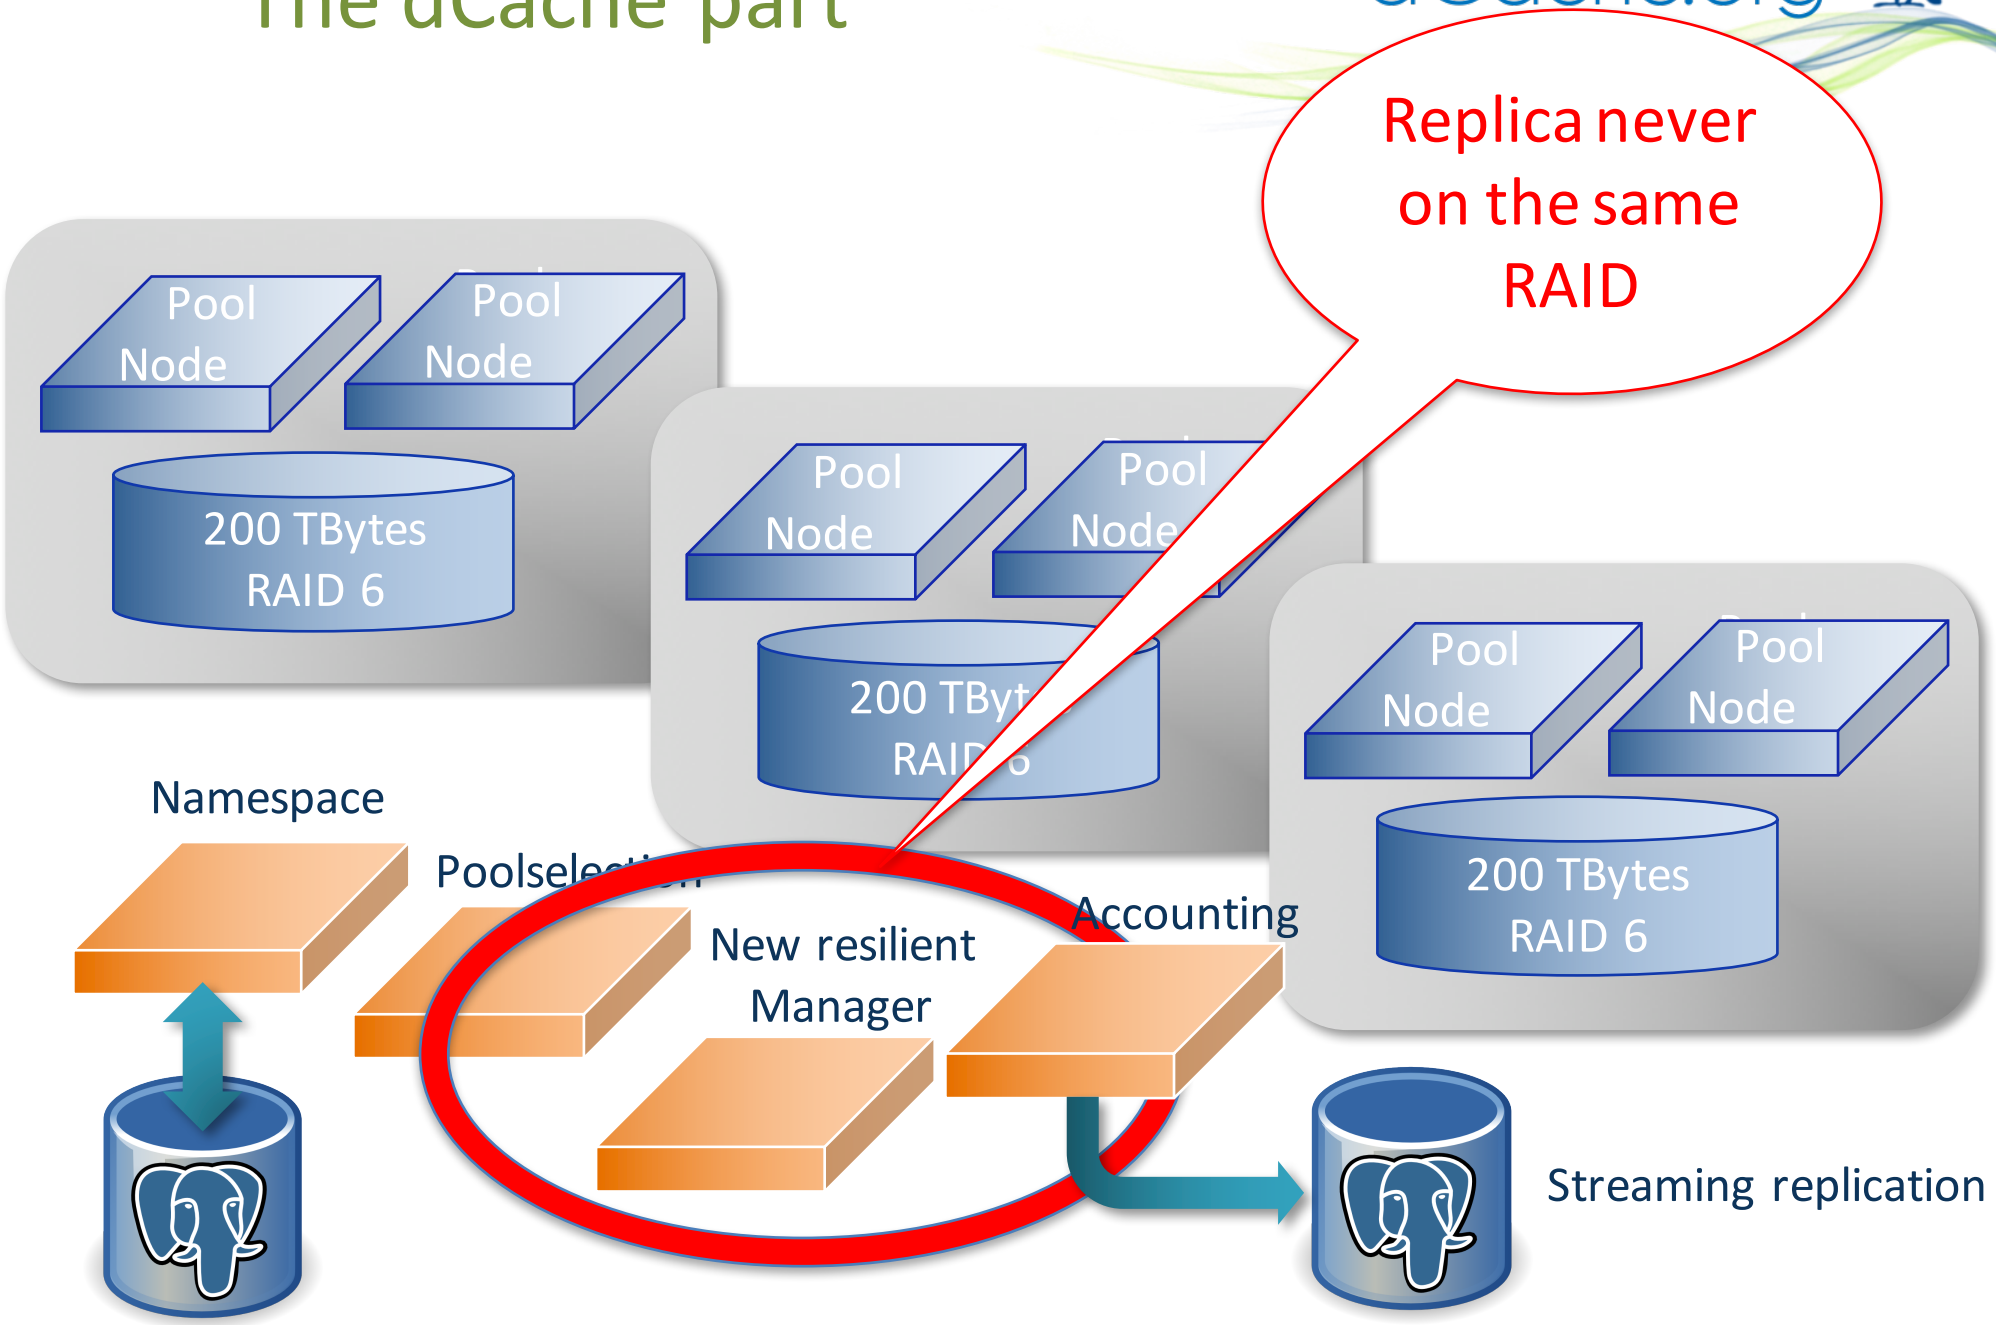

The horizontal scaling

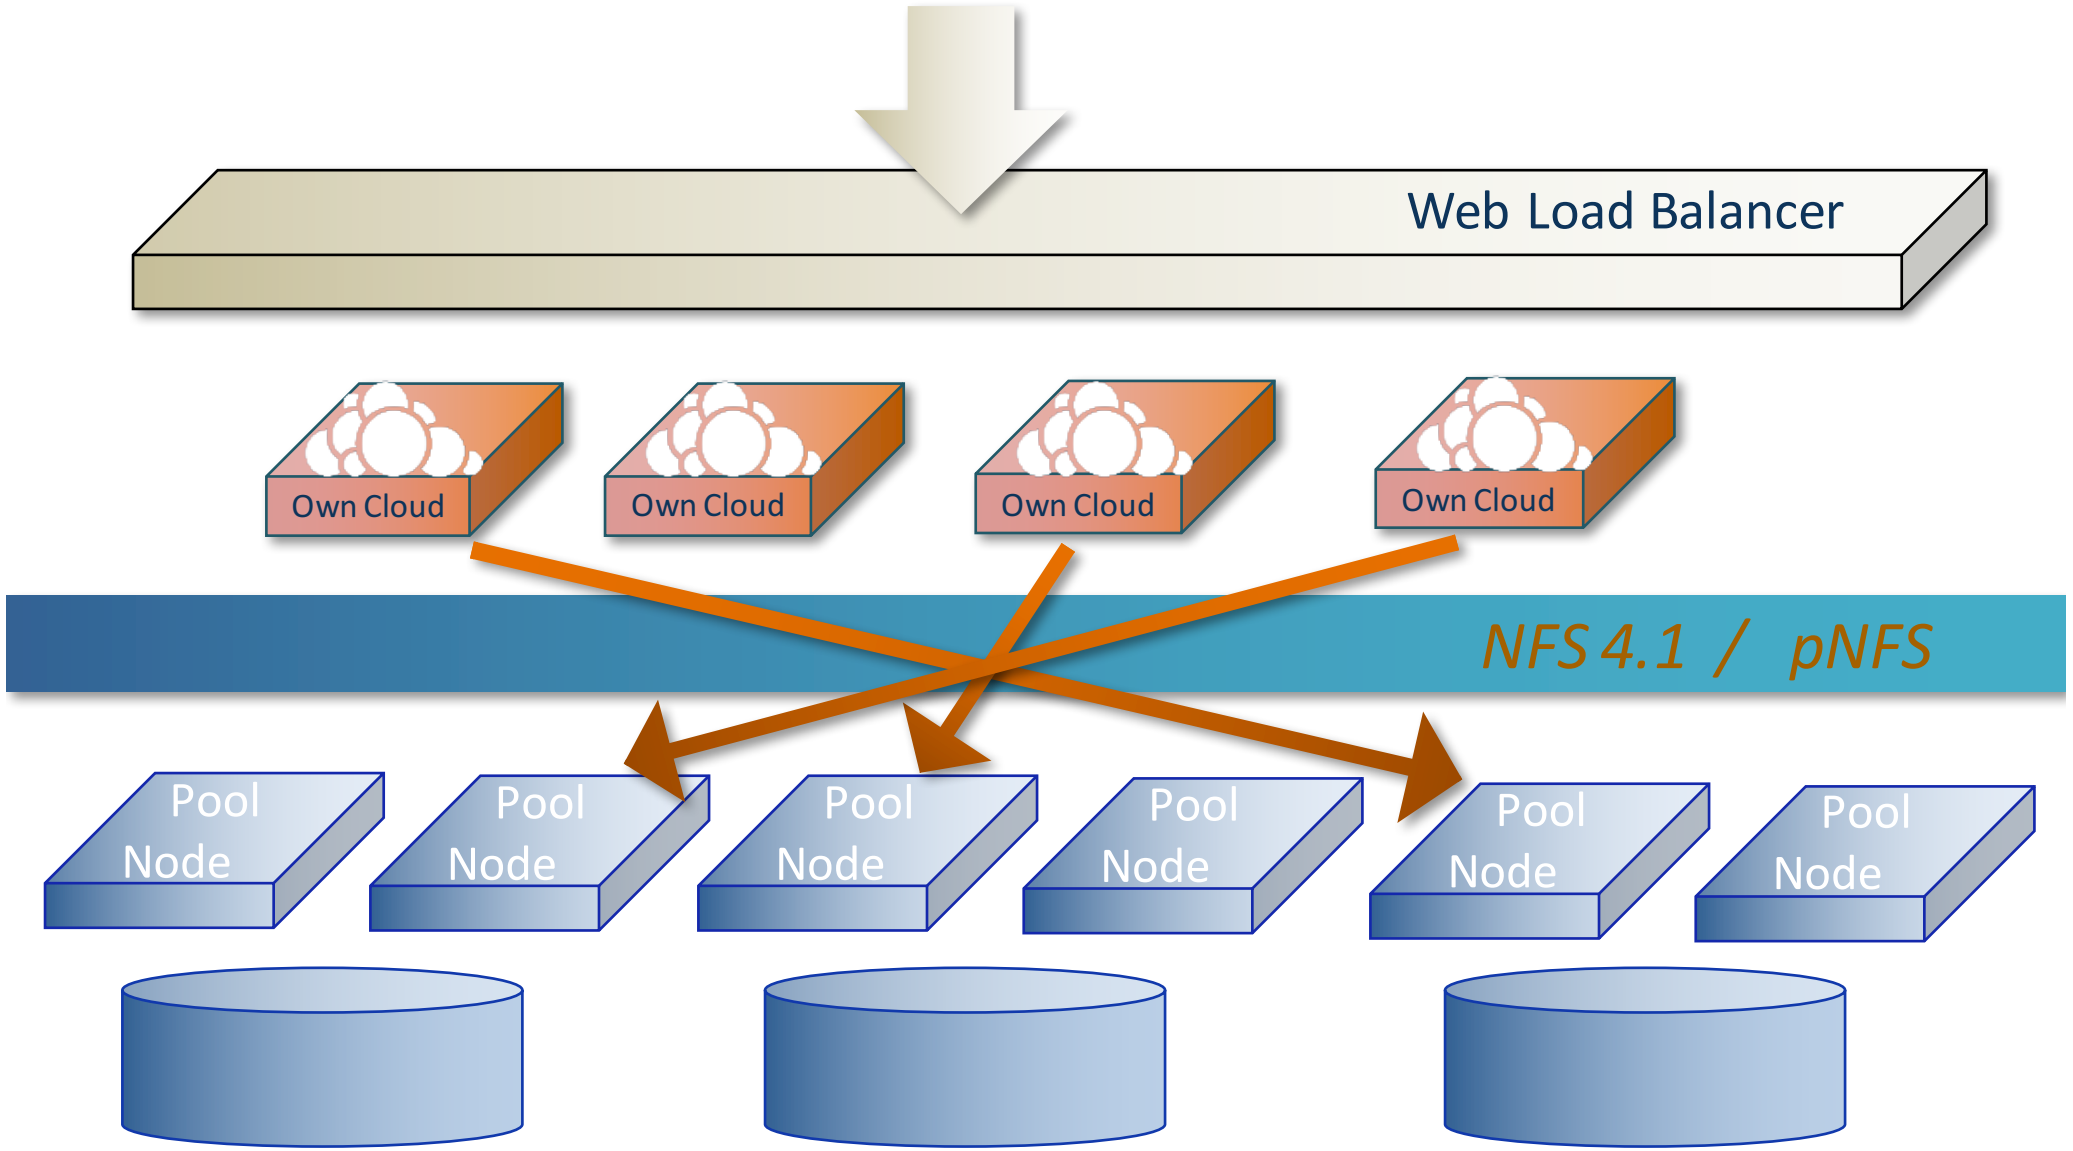

Current Usage Profile (@DESY)

dCache.org

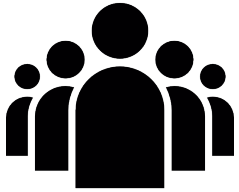

> 1000 Active Users (not just registered)

20 M Files

> 3 M Directories

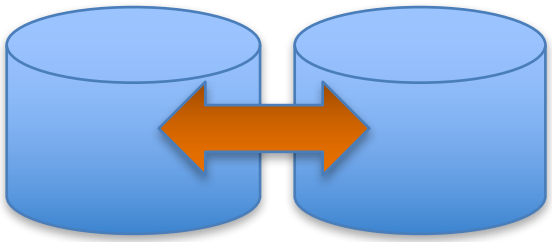

- 50 TB in total, mostly small files
- QoS : 2 replica, replica manager
- No tape copies yet

And the sharing part

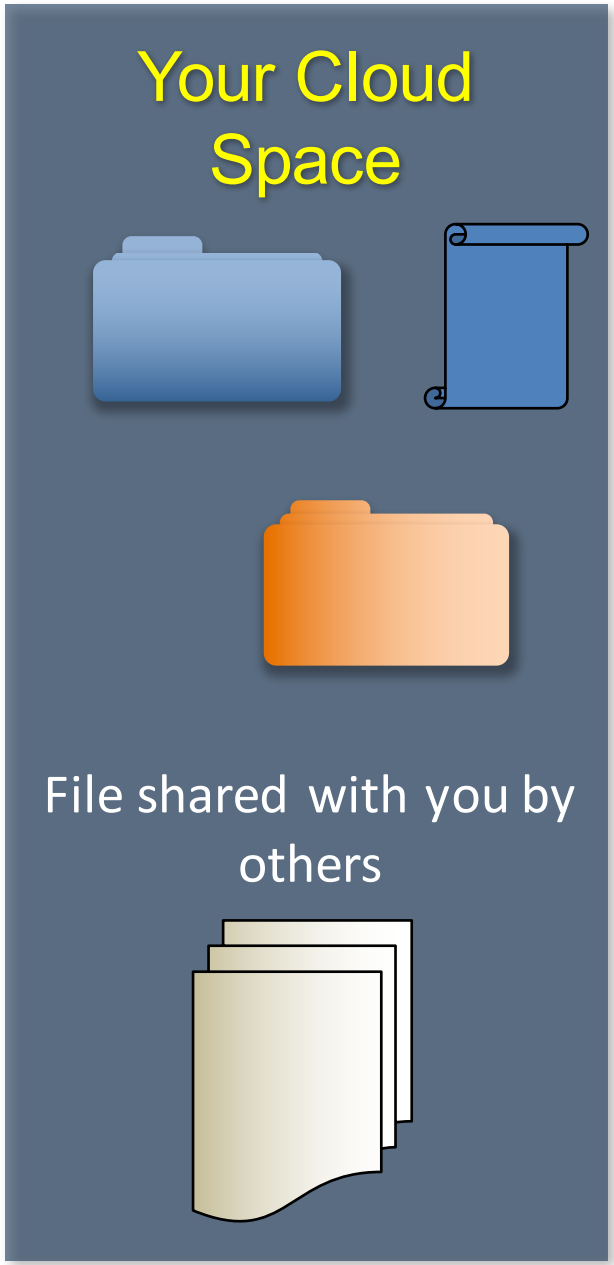

Share files/folders with individuals

Share files/folders with ‘groups’

Share with ‘public’ with and w/o password

(Shares can expire)

Share space(s) with others for upload

OwnCloud X with nice new ‘sharing features’

Others sharing data with you (in your home)

There were still some issues

So far so good… issues
In the order of importance

dCache.org

- ONE: NFS4.1 server and client incompatibilities needed to be solved.
- TWO : To support Quality of Service, the resilient manager needed to cooperate with ’custodial’ (e.g. Tape or Cloud).
- THREE: dCache Kernel would need to understand ‘shares’
- FOUR: Best would be a ‘read/modify/write’ dCache.

Issue ONE : The NFS stability

- We had quite some ‘I/O’ errors reported from dCache, which ownCloud really didn’t like.
  - These are essentially solved.
- Still an issue with proper interpretation of errors between
  - dCache – NFS – ownCloud (python)
  - Tigran is on it
- I’m convinced it’s essentially OK now.
  - Proof : All my work related data is in DESYCLOUD as the only copy.

Issue TWO : QoS

In a Jiffy

Issue THREE : ‘shares’ in dCache

dCache.org

Views

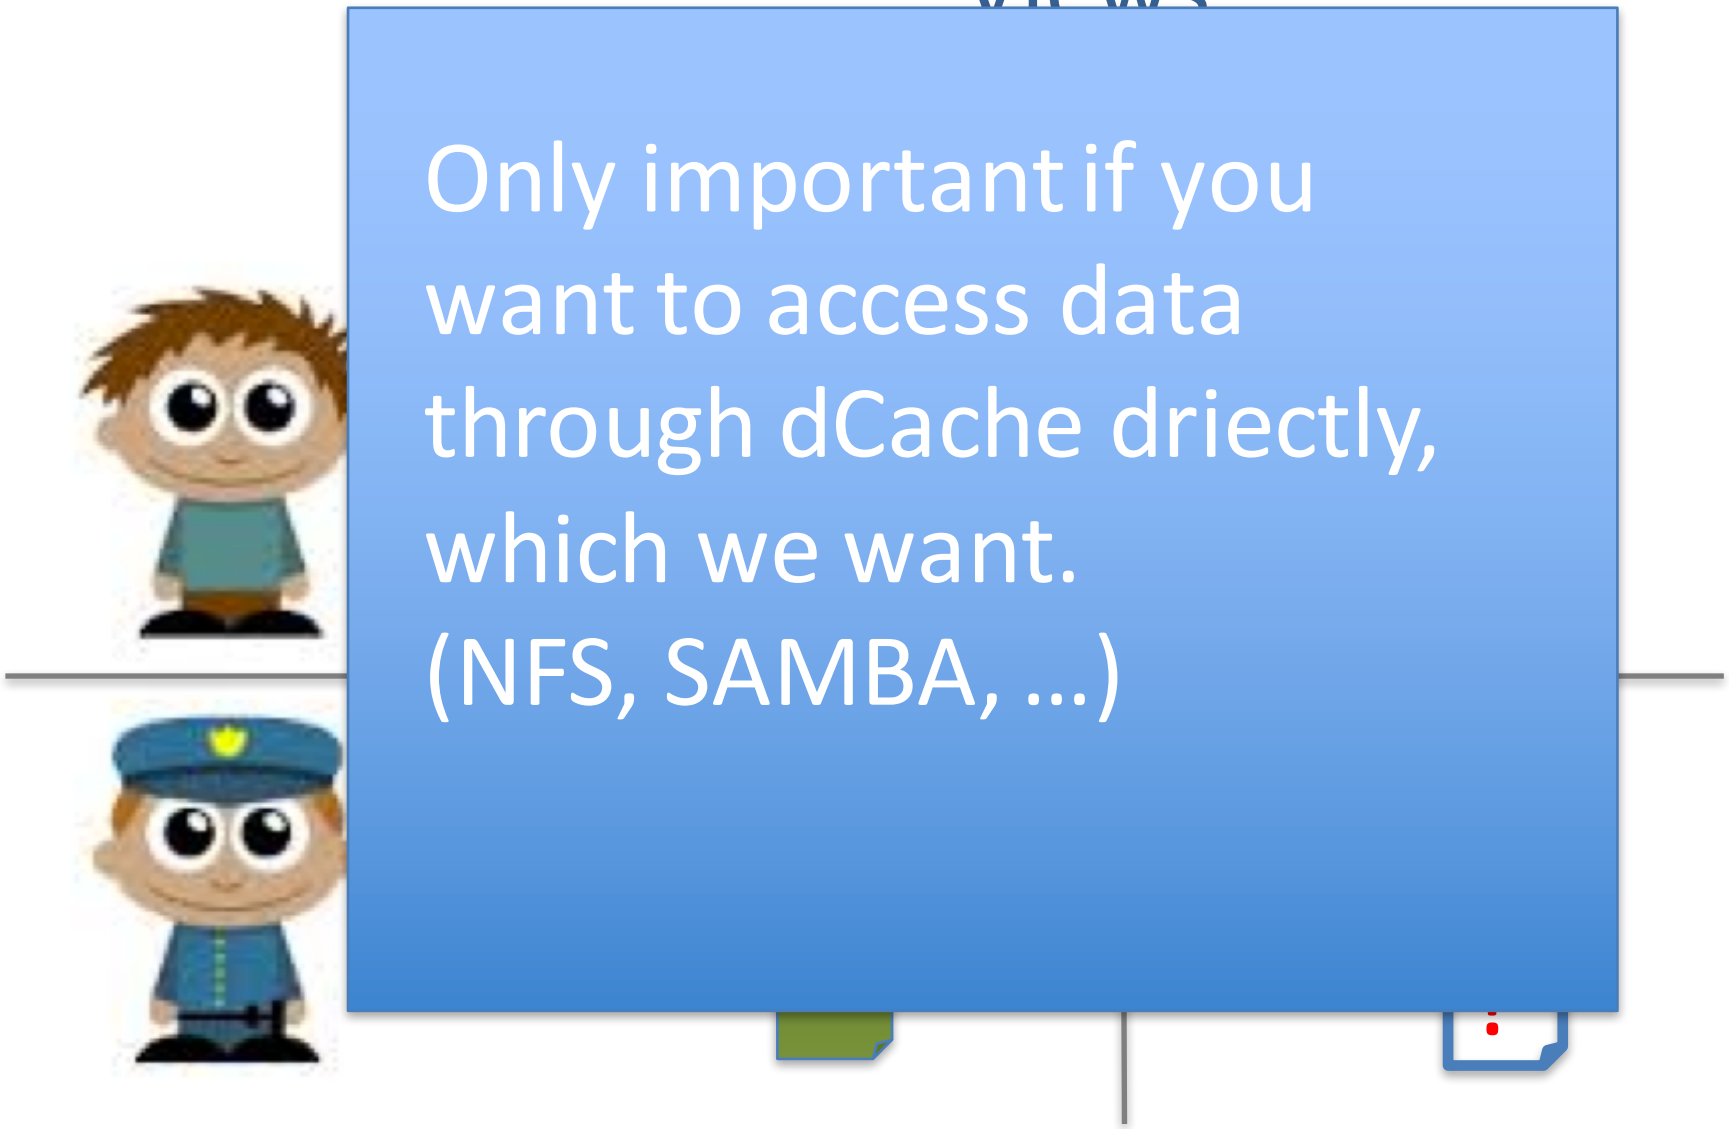

Shares in dCache

- Was significantly delayed.
- We made am initial design, however
- ownCloud – next Cloud issue.
  - After the split, the future was unclear.
  - Now it seems both are doing OK.
- Still need some discussion on shares, as with ownCloud X, the idea again changed slightly.
- And we need to find some time and efforts to work in this.

Issue FOUR : read/modify/write

More requirements

dCache.org

- Request for *unlimited, indestructible storage*.
- 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.
- User defined *Data Life Cycle*
  - Move data to tape after ‘n’ months.
  - Remove from random access media after ‘m’ months.
  - Make public after ‘x’ month.
  - Remove completely after ‘y’ months.
- Controlled by Web or API (*Software defined storage*)

Coming back to QoS

dCache.org

What we would need

dCache.org

- Offering predefined storage qualities
  - Scratch
    - Disappears any time after ‘n’ hours.
  - Regular
    - Multiple copies on disk (differed hardware)
  - Permanent
    - Same as Regular plus Media break
  - Archive
    - Write once read never
    - Media break

Now available

- This is now possible with the ‘new resilient manager’
  - Number of disk copies bound to ‘storage group’
  - Coexistence of ‘resiliency’ and ‘external storage’
    - E.g. two disk copies and one copy on tape.
- For details, see Dmitry’s Talk

Static Configuration (Example)

dCache configuration

		Resilient	AL/RP
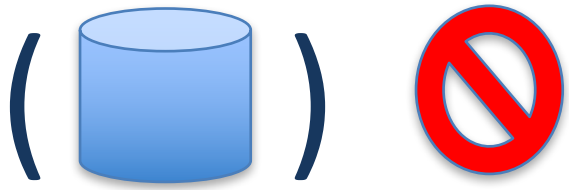	Scratch	NO	Replica/Nearline (pool: Volatile)
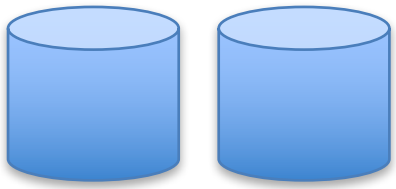	Regular	N=2	Replica/Online
	Permanent	N=2	Custodial/Online
	Archive	NO	Custodial/Nearline

QoS Transitions are in preparation (see Marina)
but not yet available for the Cloud.

Summary and outlook

- DesyCloud is one of the most outstanding DESY services.
- Likely to replace home directories and afs.
- Essential gaps:
  - QoS is essential to allow customers to select the implementation of their use cases.
  - Access besides OwnCloud, e.g. nfs and SAMBA
- Still bits and pieces missing but we'll move on in this direction.
- Future : integration with OpenStack @ DESY

The END

further reading

www.dCache.org