



## The dCache technology

Patrick Fuhrmann

On behalf of the project team



Courtesy of Ron Trompert, SARA



## Content

- More a random walk
- Organizational
  - Funding
  - People
  - Distribution Channels
  - Release Policies
- Special Topics
  - Software Defined Storage
  - Sync-n-share
  - The scientific storage cloud
  - Federated Identity Management

# What is dCache



- **Open Source** software for aggregating heterogeneous storage.
- **Immutable file-system** with its own namespace **independent of data location**.
- Supports “tiered” storage with automatic or manual transitions between media (**SSD, spinning disk, Tape**)
- Sophisticated **data-placement**
- Built-in support for **multiple protocols** (NFS, FTP, HTTP/WebDAV)
  - Consistent and coherent view of the files
- **Pluggable** authentication / identity system
  - Supports X.509 client certificates, username/password and Kerberos
  - Integrates with site IdM: NIS, LDAP, Active Directory, Kerberos, ...

- dCache.org members:
  - FERMIlab
  - DESY
  - NDGF
  - HTW Berlin (through LSDMA, mostly students)
- **Projects:**
  - LSDMA (Germany) completely replacing EMI
  - Involved in three H2020 proposals

# People

- **FERMIlab:**
  - Dmitry and Al
- **NDGF:**
  - Gerd
- **DESY**
  - Karsten, Christian, Paul, Tigran and Patrick

## The Consortium


- **THW Berlin (Students)**
  - Leonie, Jana and Tom
- **German Support**
  - Aachen (Oleg), KIT (Xavier), Munich (Christoph),  
Göttingen (Gen, Jordi ?)
- **Considering to join: Jülich (HPC)**
  - Bastian, André

- EMI distribution still supported for some months
- Through EGI (UMD).
  - Provides interoperability and staged role-out.
- Downloads directly from the dCache.org web pages.
  - includes feature releases
- Mid term goal is to be integrated into an official distribution (e.g. EPEL)

# Deployments

- **WLCG**: 50 sites (world-wide) together provide 100 PB, satisfying about 50% of LHC current requirement.
- **DESY HEP**: HERA, ATLAS, CMS, LHCb, Belle
- **DESY etc**: **CFEL**, **XFEL**, **IceCube**,
- **FERMILab**: CMS, **Intensity Frontier**, general store
- **BNL**: **ALTAAS** and **RHIC**
- **SNIC**: **SweStore**
- **NDGF**: geographically largest single instance, spread over 4 countries
- **SARA** and **Juelich**: **Lofar** (telescope)
- And many more ...

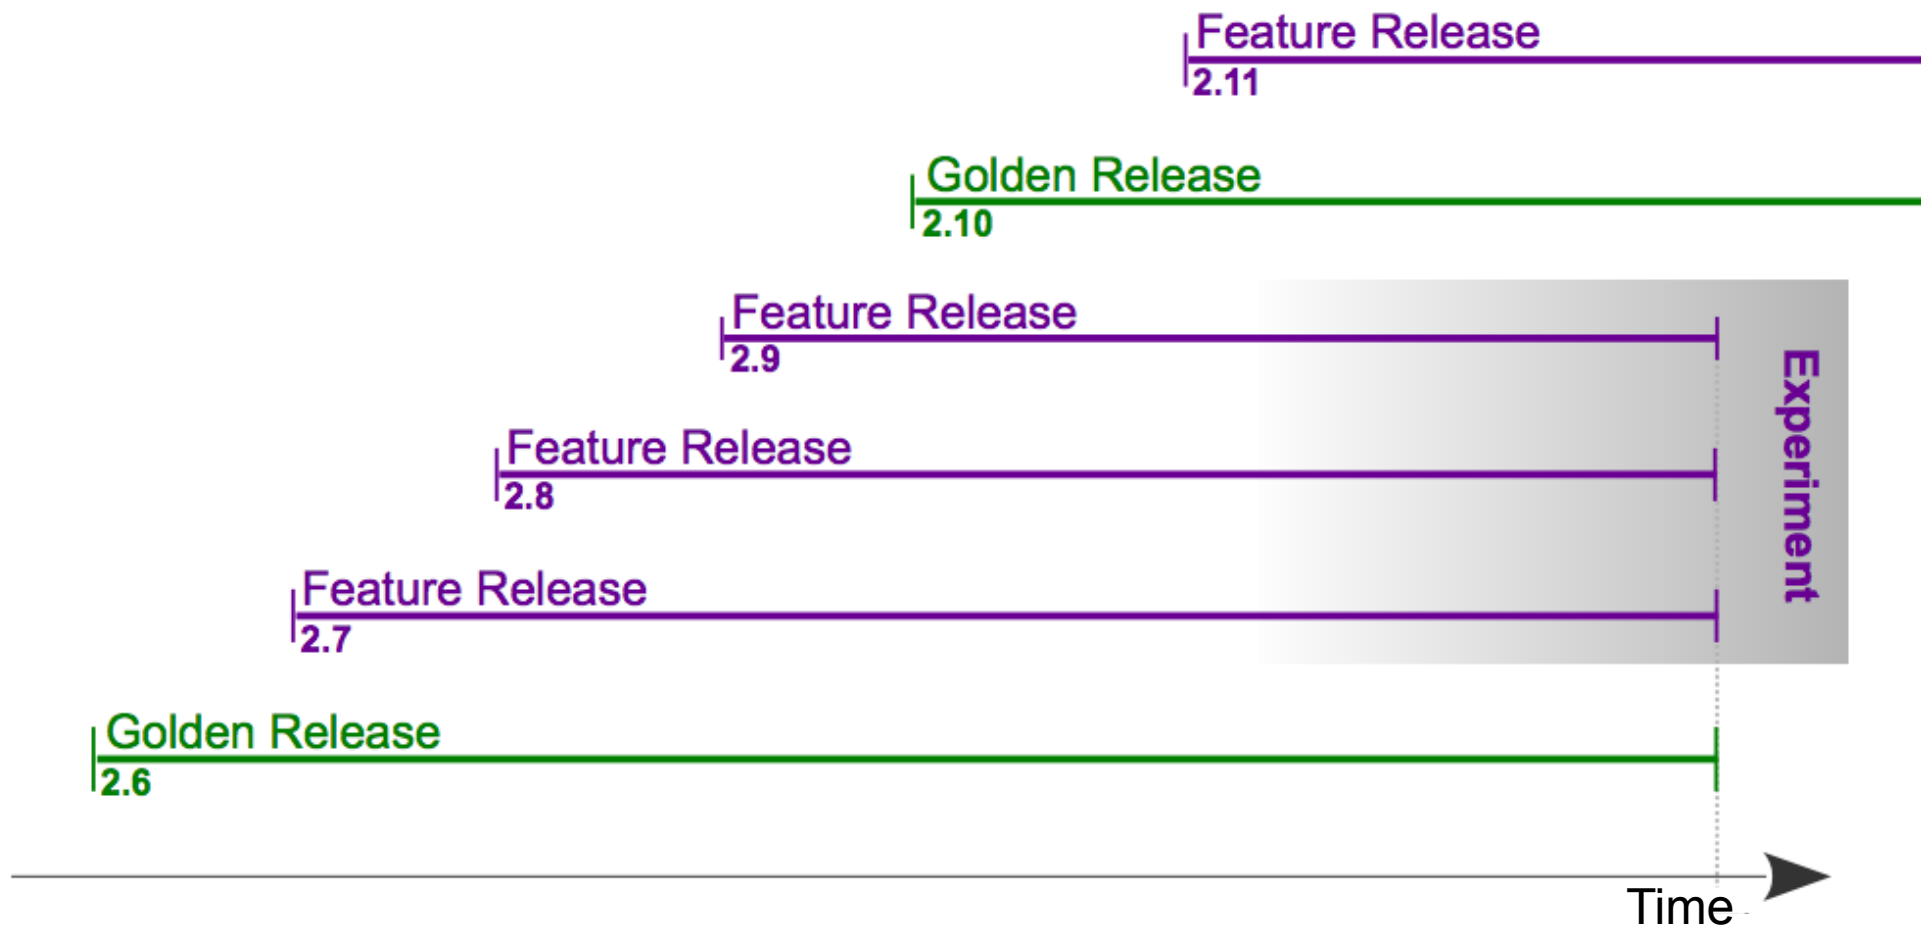
# Golden Release Policy



Stolen from Paul Millar Presentation @ 8<sup>th</sup> dCache Workshop 2014



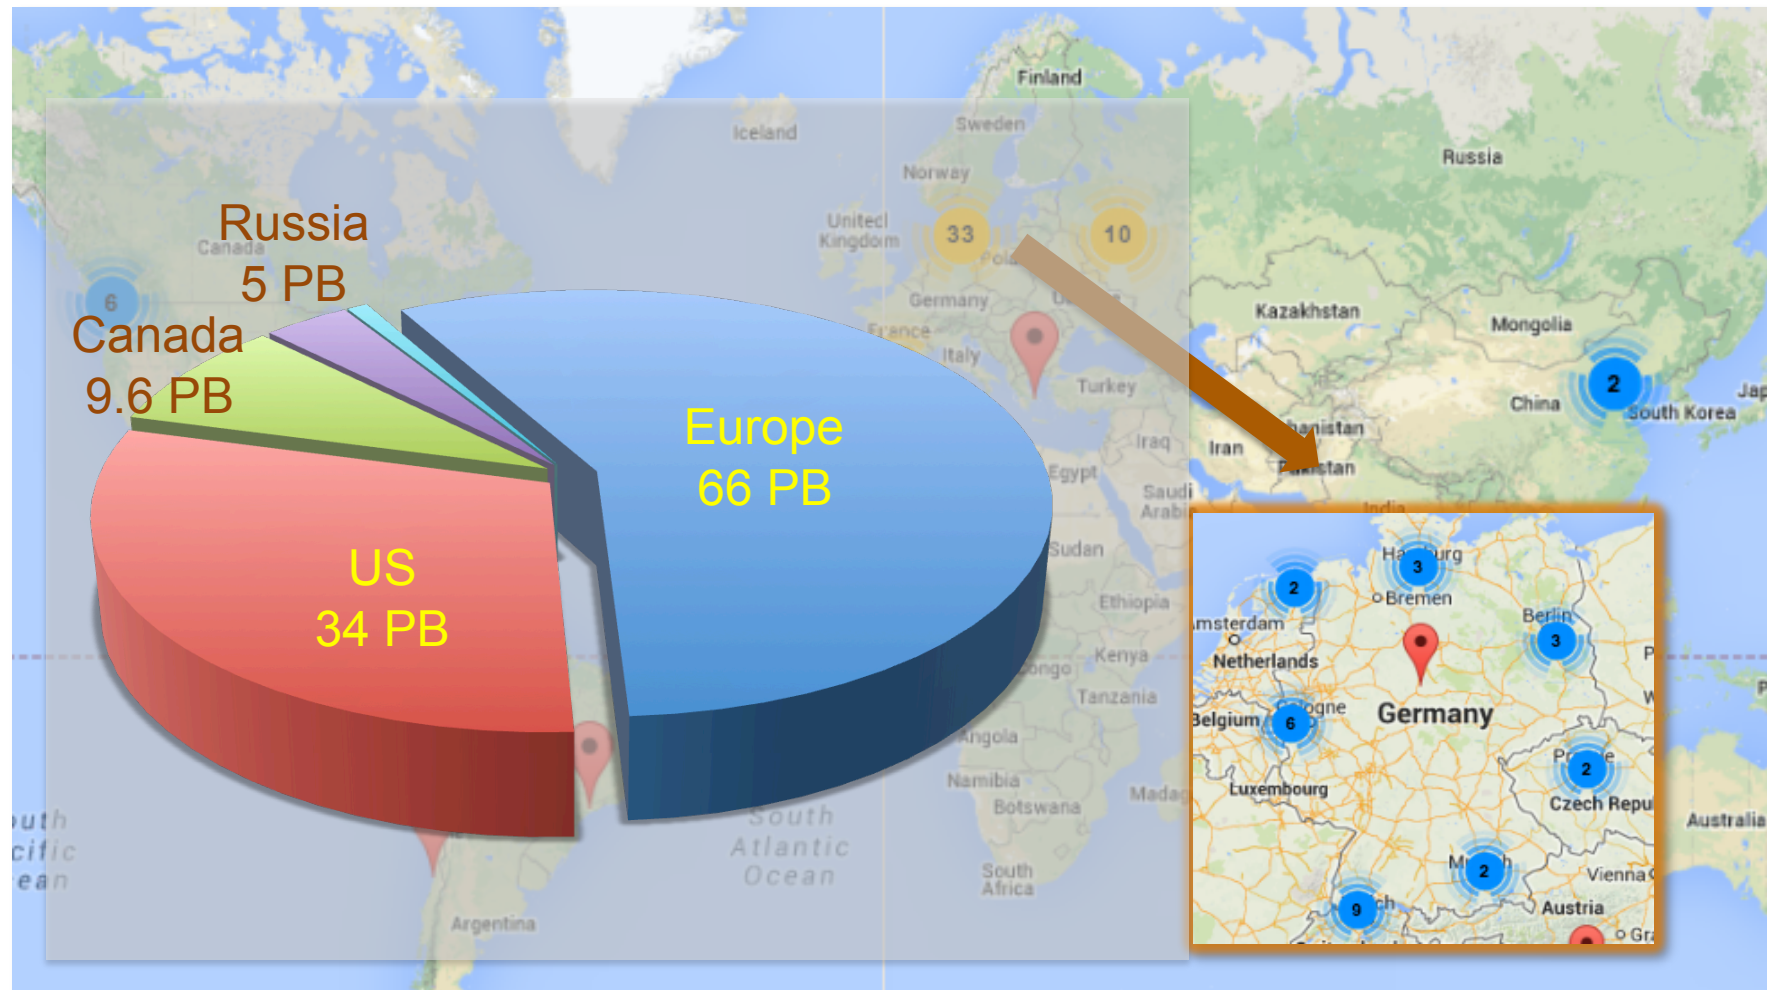
# New Feature Release Policy



Stolen from Paul Millar Presentation @ 8<sup>th</sup> dCache Workshop 2014

# Installations

There are about 80 dCache instances around the world, managing more than 120 Petabytes for the LHC, **which is about 50 % of their Higgs'es**



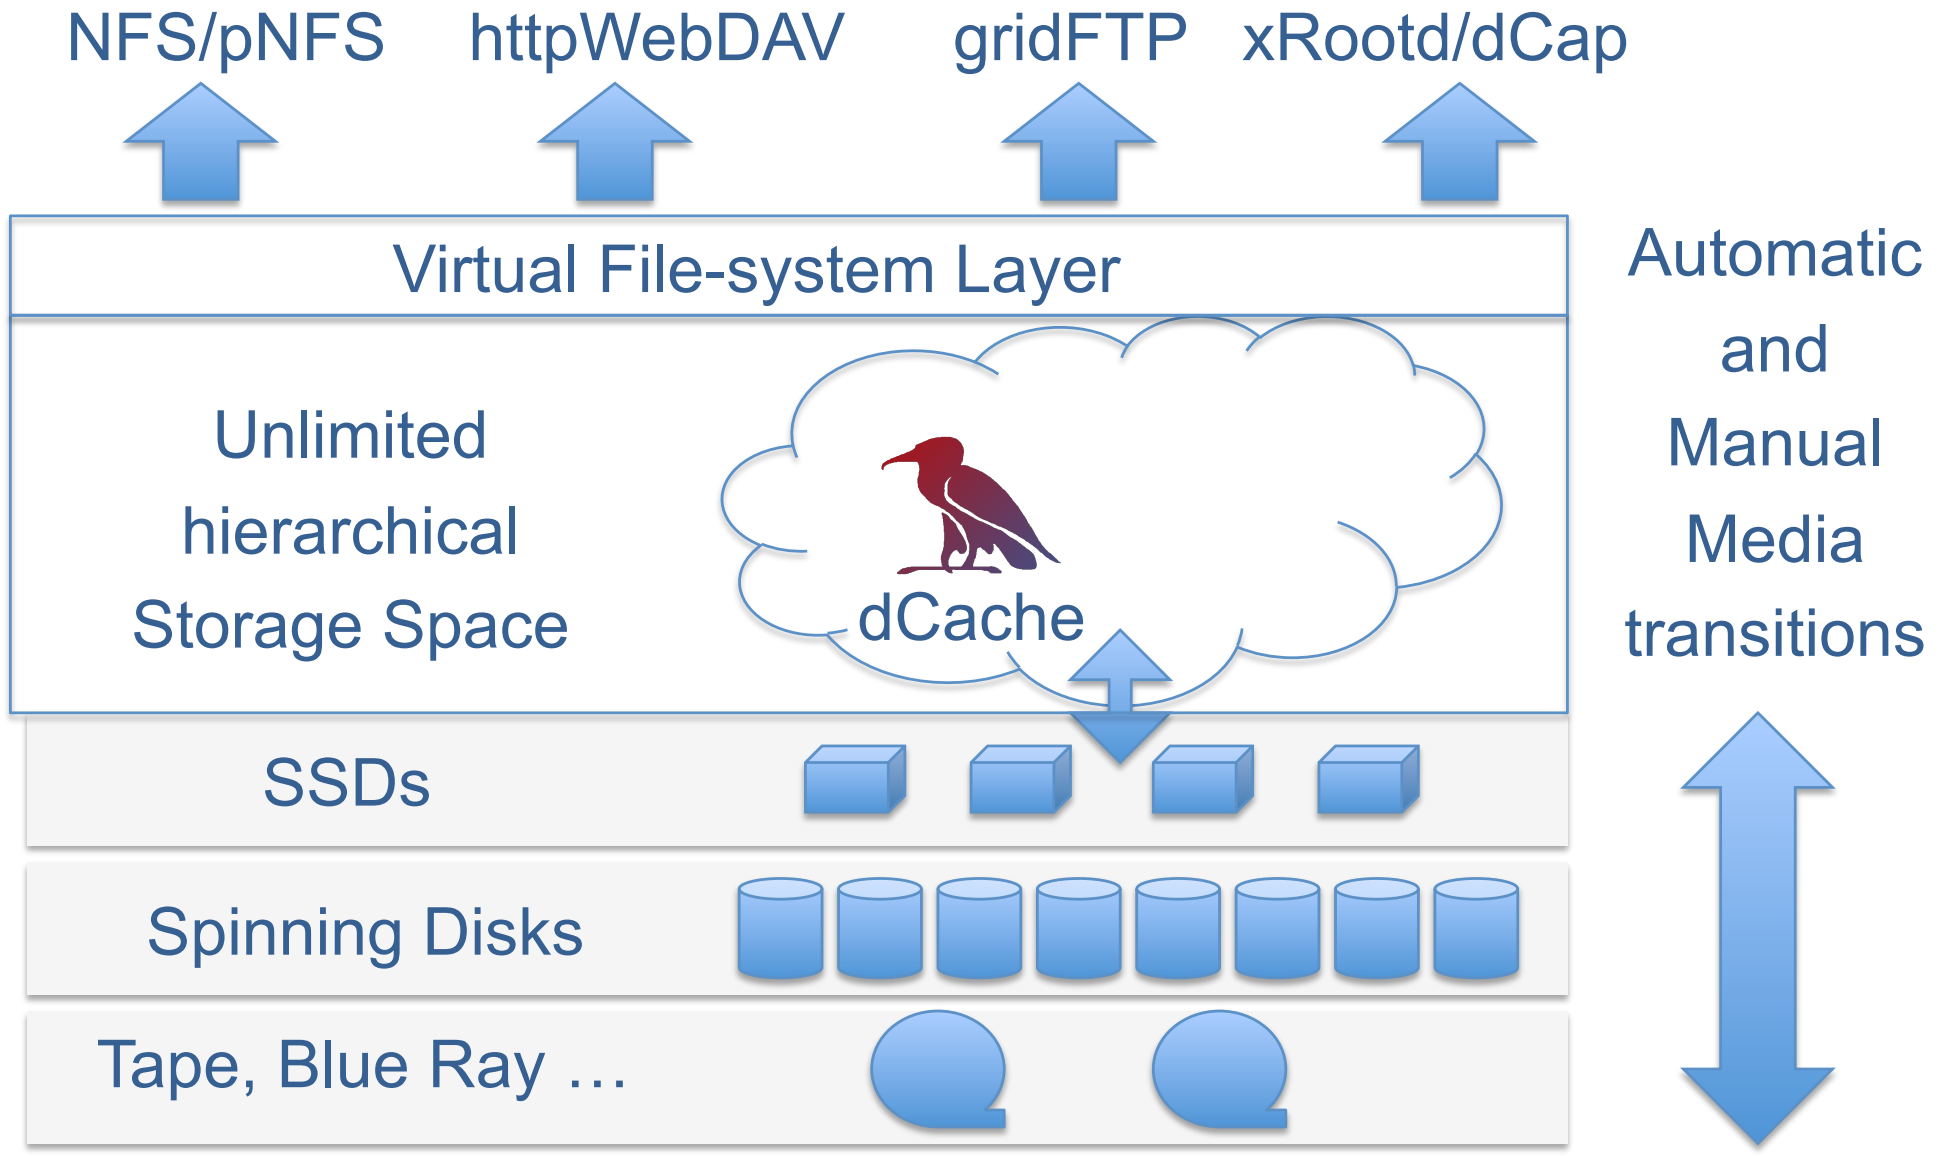
## Installations (cont)

- Smallest:
  - On my raspberry Pi
- Biggest:
  - Fermilab
    - 40 Petabytes on Tape
    - 14 Petabytes on Disk
- Largest, geo area, NDGF:
  - One dCache in Sweden, Norway, Finland and Denmark : 1.000 sq km

## Having a closer look



# dCache spec for Dummies



## Software defined storage

- OR
  - Tiered Storage
  - Media Aware Storage
  - Quality of Service

# Software Defined Storage

- Storage Media (SSD/Tape/Disk) can be defined by
  - I/O operation : read or write
  - Protocol: NFS, WebDAV, GridFTP
  - File System location (path)
  - Manually
- dCache can be configured to change media if appropriately.
  - Fast Random Read: SSD
  - High Throughput Sequential: Spinning Disk
  - Long time no read: Tape



- Media Awareness (Example)
  - Small files are collected in container before written to tape (if configured)
  - Containers are expanded when small files are requested back from Tape (transparent for user)



For the DESY photon science community  
and the FERMIlab “Intensity Frontier” we  
need to support the full scientific data life  
cycle.

# Support for Scientific Data Lifecycle (one size fits all)



High Speed  
Data Ingest

Fast Analysis  
NFS 4.1/pNFS

Wide Area Transfers  
(Globus Online, FTS)  
by GridFTP

Visualization  
& Sharing  
by WebDAV



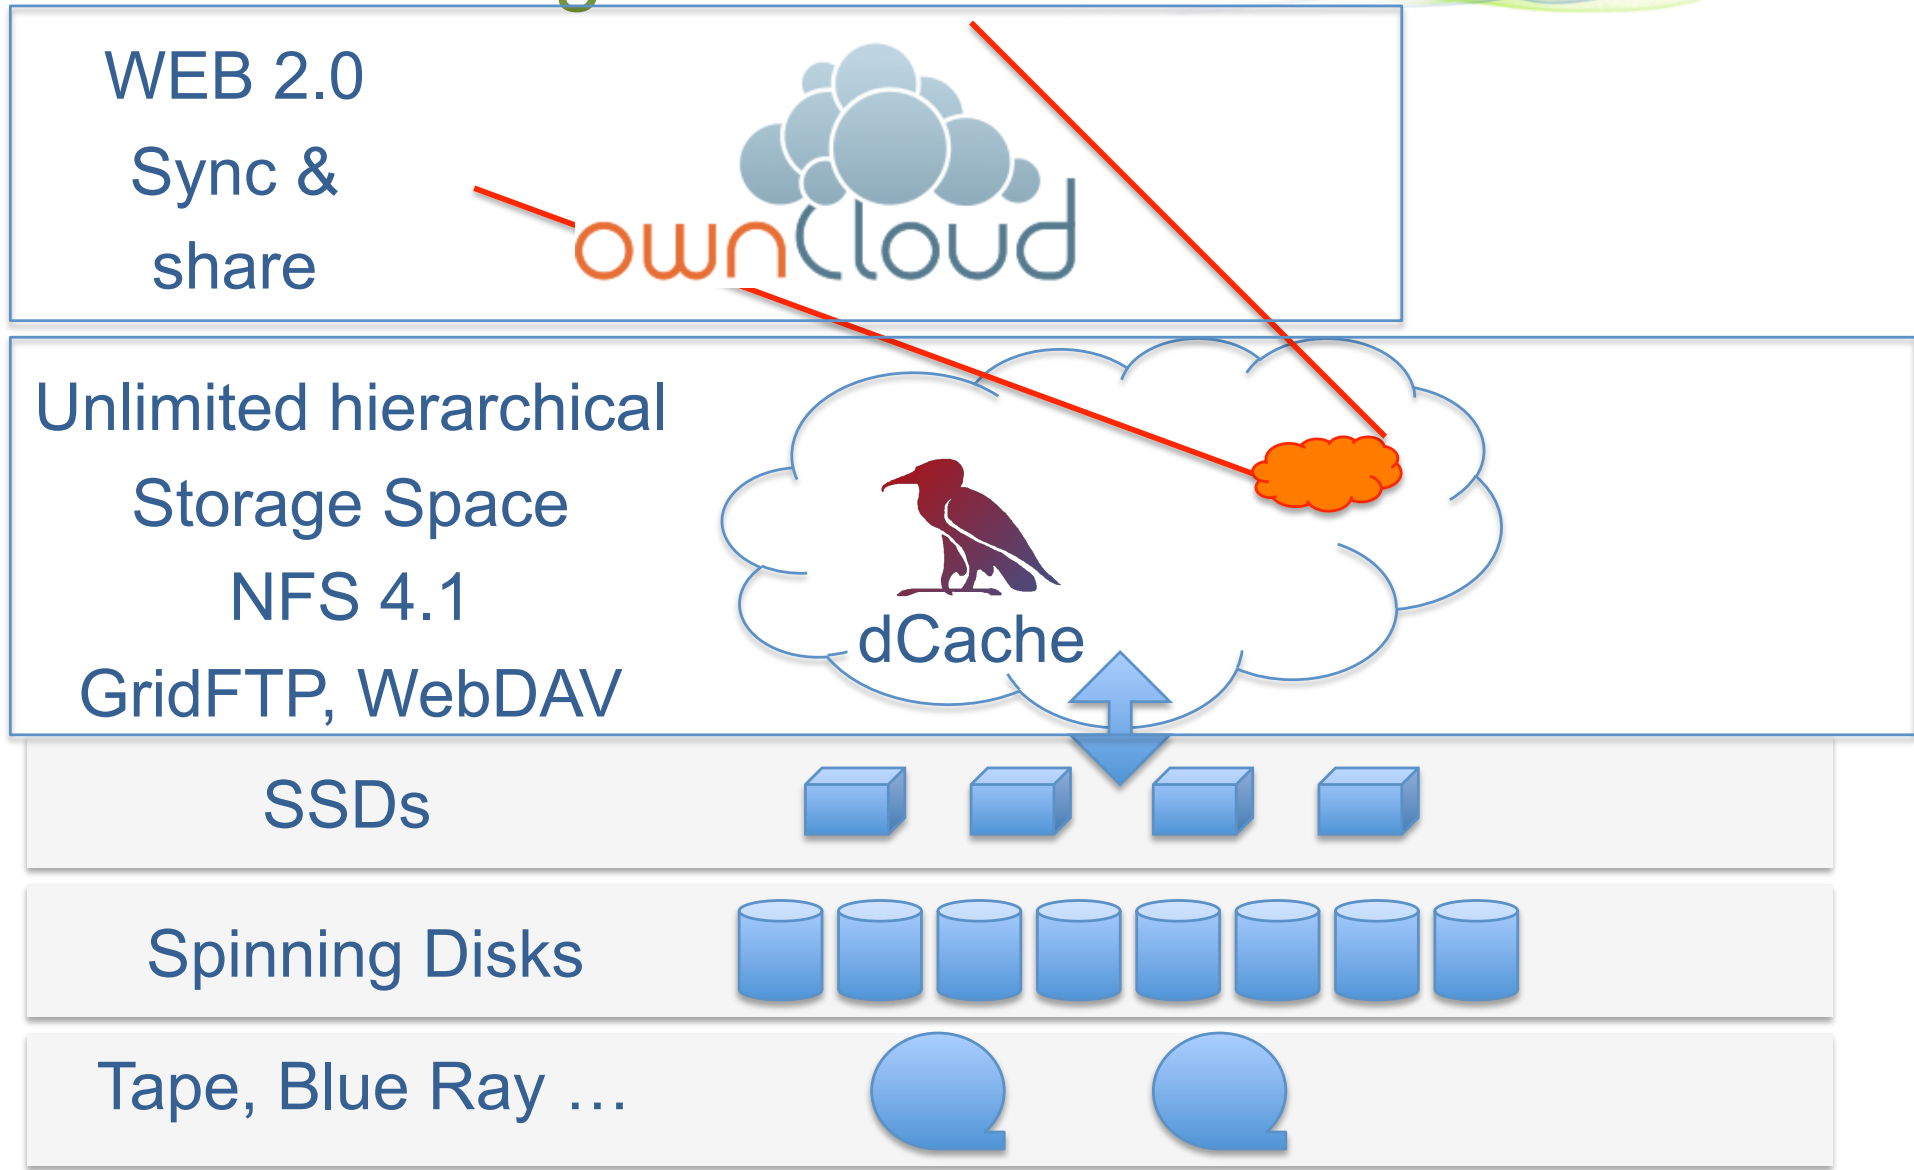
# Ongoing activities

- Sync-n-share interface
- Cloud Data Management Interface (CDMI)
- Federated Identity Management

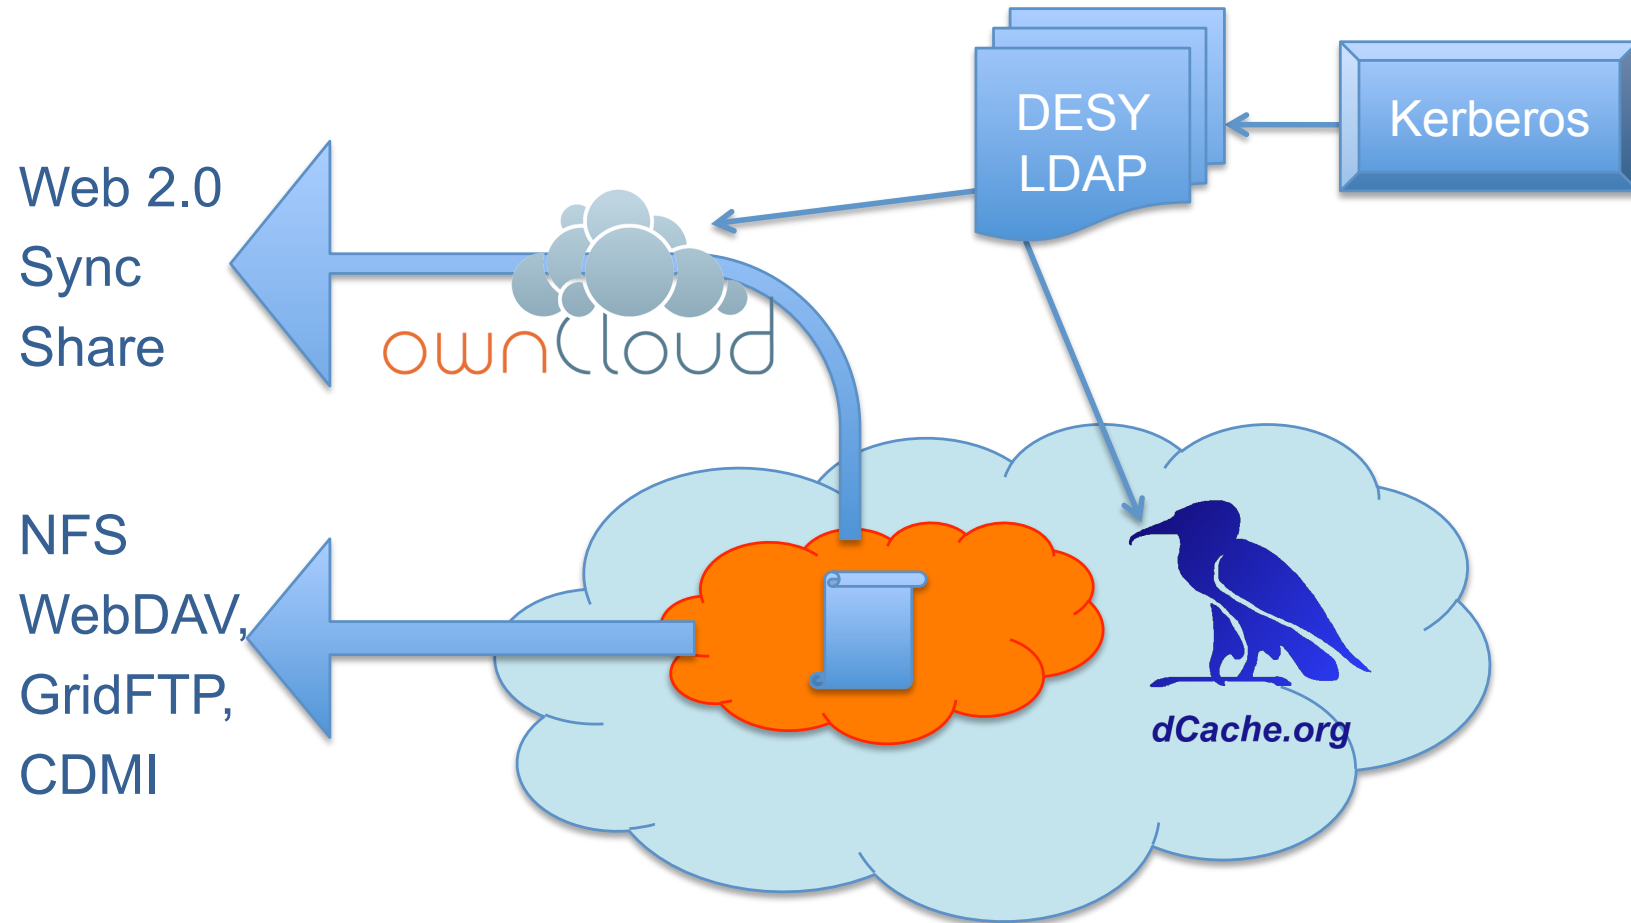
## Sync'n Share with Own Cloud

- We are using Own Cloud as one of our Doors, similar to NFS, ...
- This is for now only a Service for DESY users and not yet in the dCache distribution.
- We need to get experience with the Hybrid System

# dCache – ownCloud Integration



# Ownership/mapping issue



# Cloud Data Management Interface

- Network protocol for Cloud Storage
- Initially by SNIA, now an ISO standard
- Allows storage management
  - Quality of service
  - Access Latency and Retention Policy (SRM replacement)
- Proper Management of Meta Data
  - Arbitrary meta data bound to data, including search
- Support POSIX namespace and Objects
- Prototype version in dCache. (not production ready)
- Might become important : EGI Fed Cloud Protocol

- X509 credentials are not very popular in non HEP sciences.
- **SAML** seems prevalent system
  - OpenID Connect is also gaining traction
- With LSDMA: initial work on **credential translation** (SAML -> X.509)
- Later: add **native SAML** support
  - Initially with WebSSO, later maybe Moonshot or similar.



# Summary

- Funding and number of people is ok.
- Interesting new Developments
  - Big Data, fast analysis
  - Cloud semantics
  - Federated Identities
- Interesting new communities
  - Intensity frontier (FERMIlab)
  - Photon Science (DESY)
  - Human Brian people in Jülich (Germany)
- Putting efforts to get H2020 funding
- In general: Very agile community

# The END

further reading  
[www.dCache.org](http://www.dCache.org)

