

Supporting the scientific data lifecycle

Patrick Fuhrmann

On behalf of the project team



Content

- How are software features selected.
- How are software features funded.
- Hardening new features.
- Exploring new communities.
- Responding on new technologies HW and SW
- Something about INDIGO-DataCloud
- Essentially a random walk focusing on things I thought might be interesting.

Some words on
why and when
dCache does what it does.

How are software features selected ? dCache.org



- Scientific communities believe, that Open Source Software is growing on trees.
- Consequently they are not willing to contribute to the development and software management at all.
- They assume that complains are very valuable contributions.
- Next consequence is that Open Source teams mainly implement software features, which are required by the labs, where the core team members are hosted.
- In order to explore new communities and satisfy their software requirements, Open Source Projects need external money.

How are new features funded ?



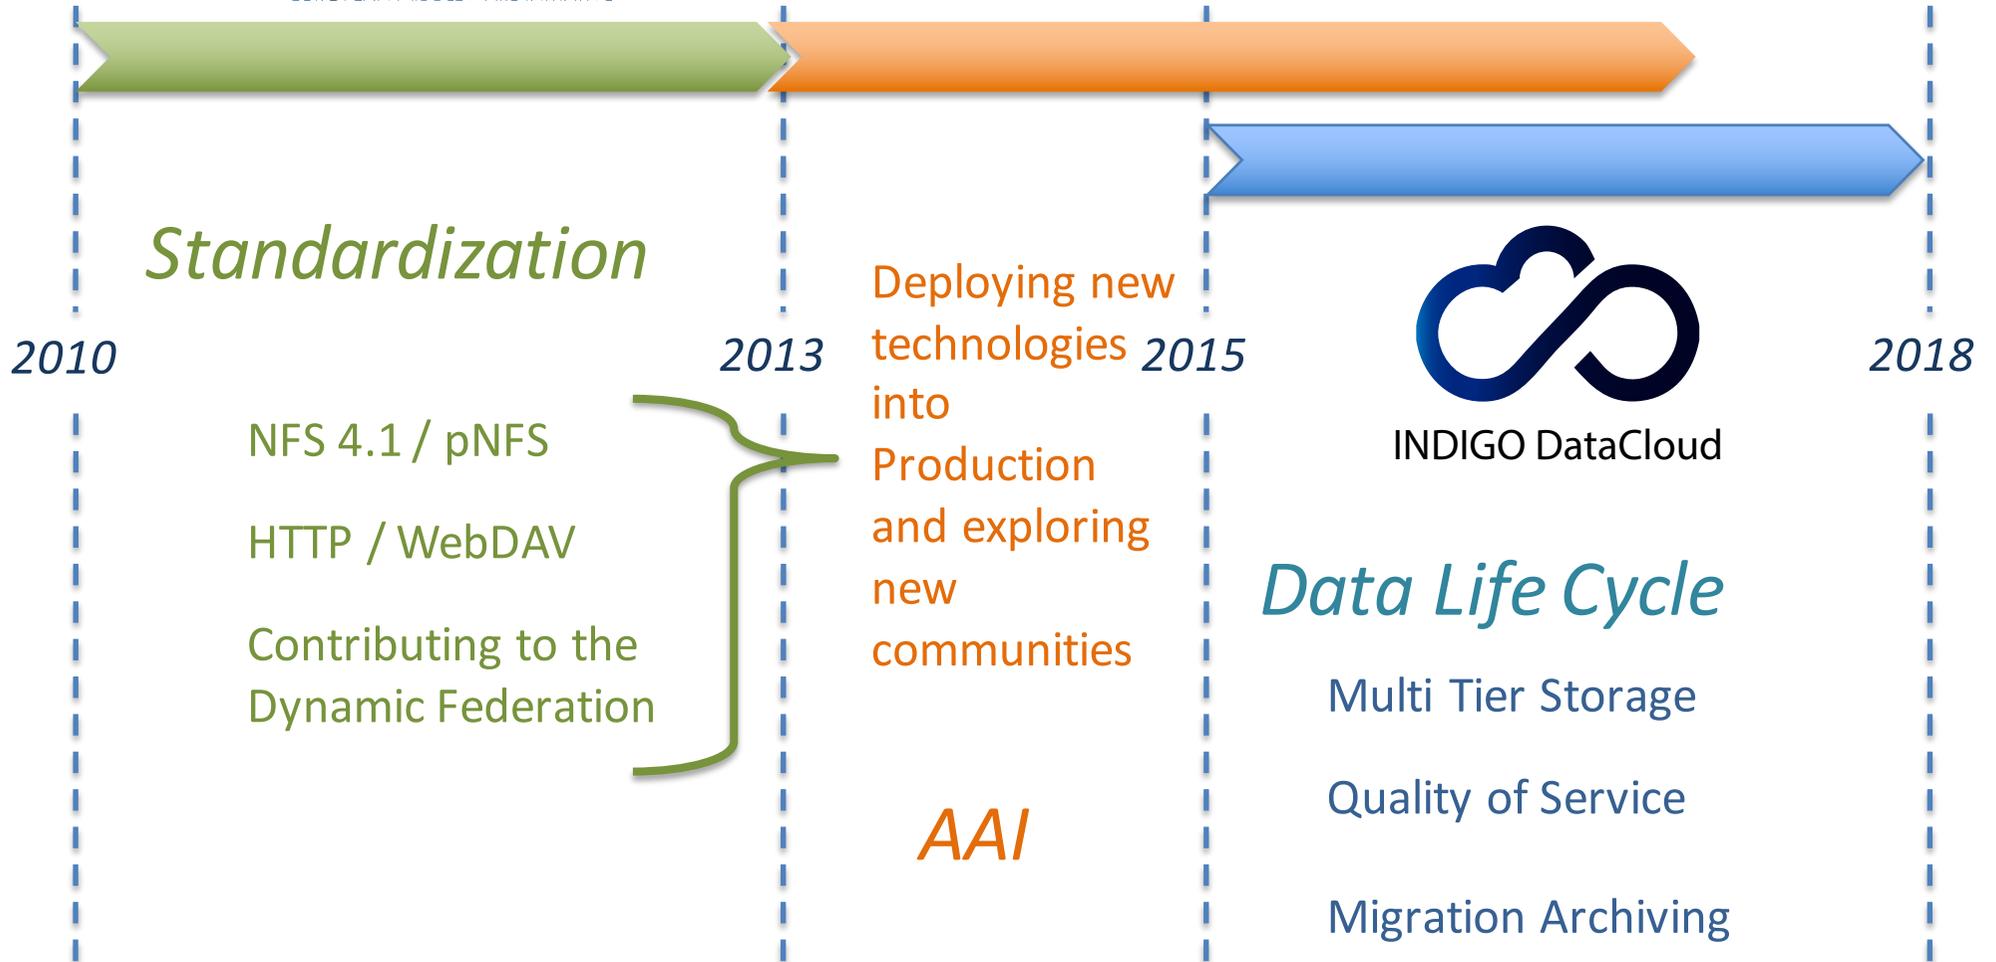
- This is where “National” and “European” projects come into play.
- For dCache, this :
 - was EMI
 - is the German national LSDMA project
 - and will be INDIGO-DataCloud
- The drawback: They tell you what they want to see in your code.

Funded features are not necessary those you need ?



- However, dCache has some invariant objectives:
 - The master plan (last slide of this presentation)
 - Be up to date on new technologies, either software or hardware.
 - Attract new communities as their specific requirements, if they can be fulfilled, make dCache even better.
- It can be a bit tricky to tune the funding projects exactly into the direction of our objectives.
- So, let's see how dCache managed/es that

Funding influences dCache development topics



From 2013 to now, we slowed down development between two very demanding development projects, EMI and INDIGO-DataCloud, to :

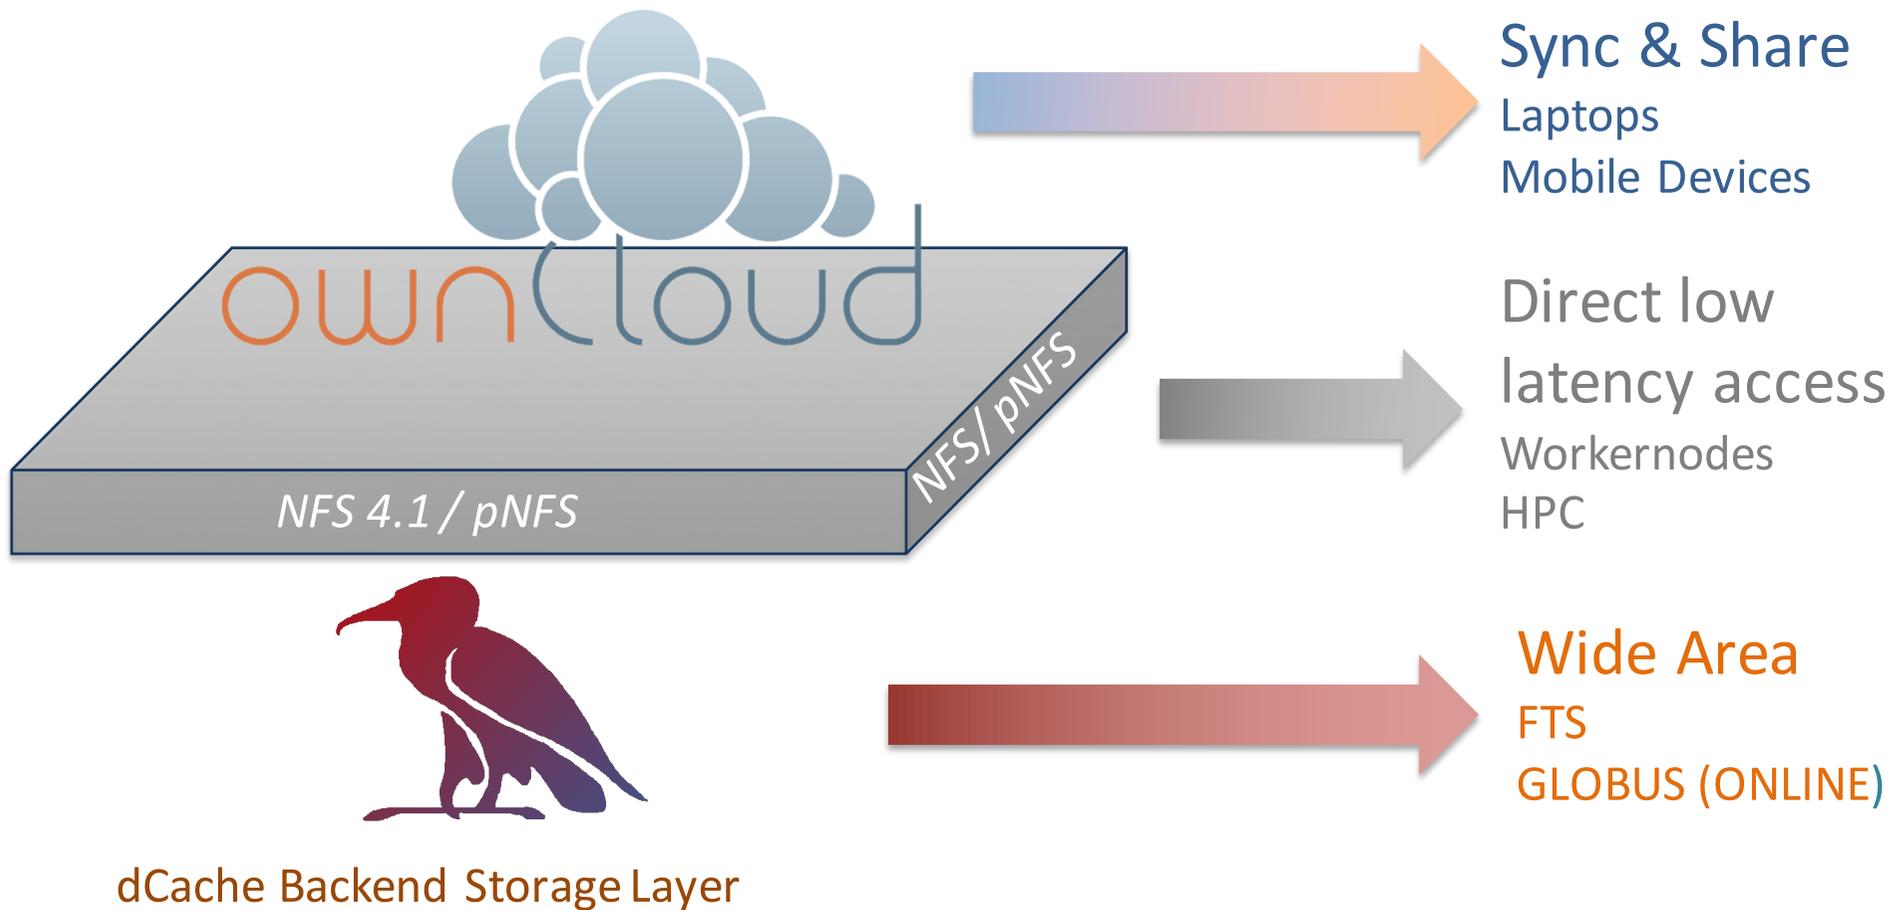
- Deploy newly implemented technologies into production.
- Explore new communities and learn about their needs.

- CMS Grid Infrastructure @ DESY
- The Desy-Cloud
- FERMIlab (various Intensity Frontier)
- And the issues

New Production Systems based on dCache NFS.

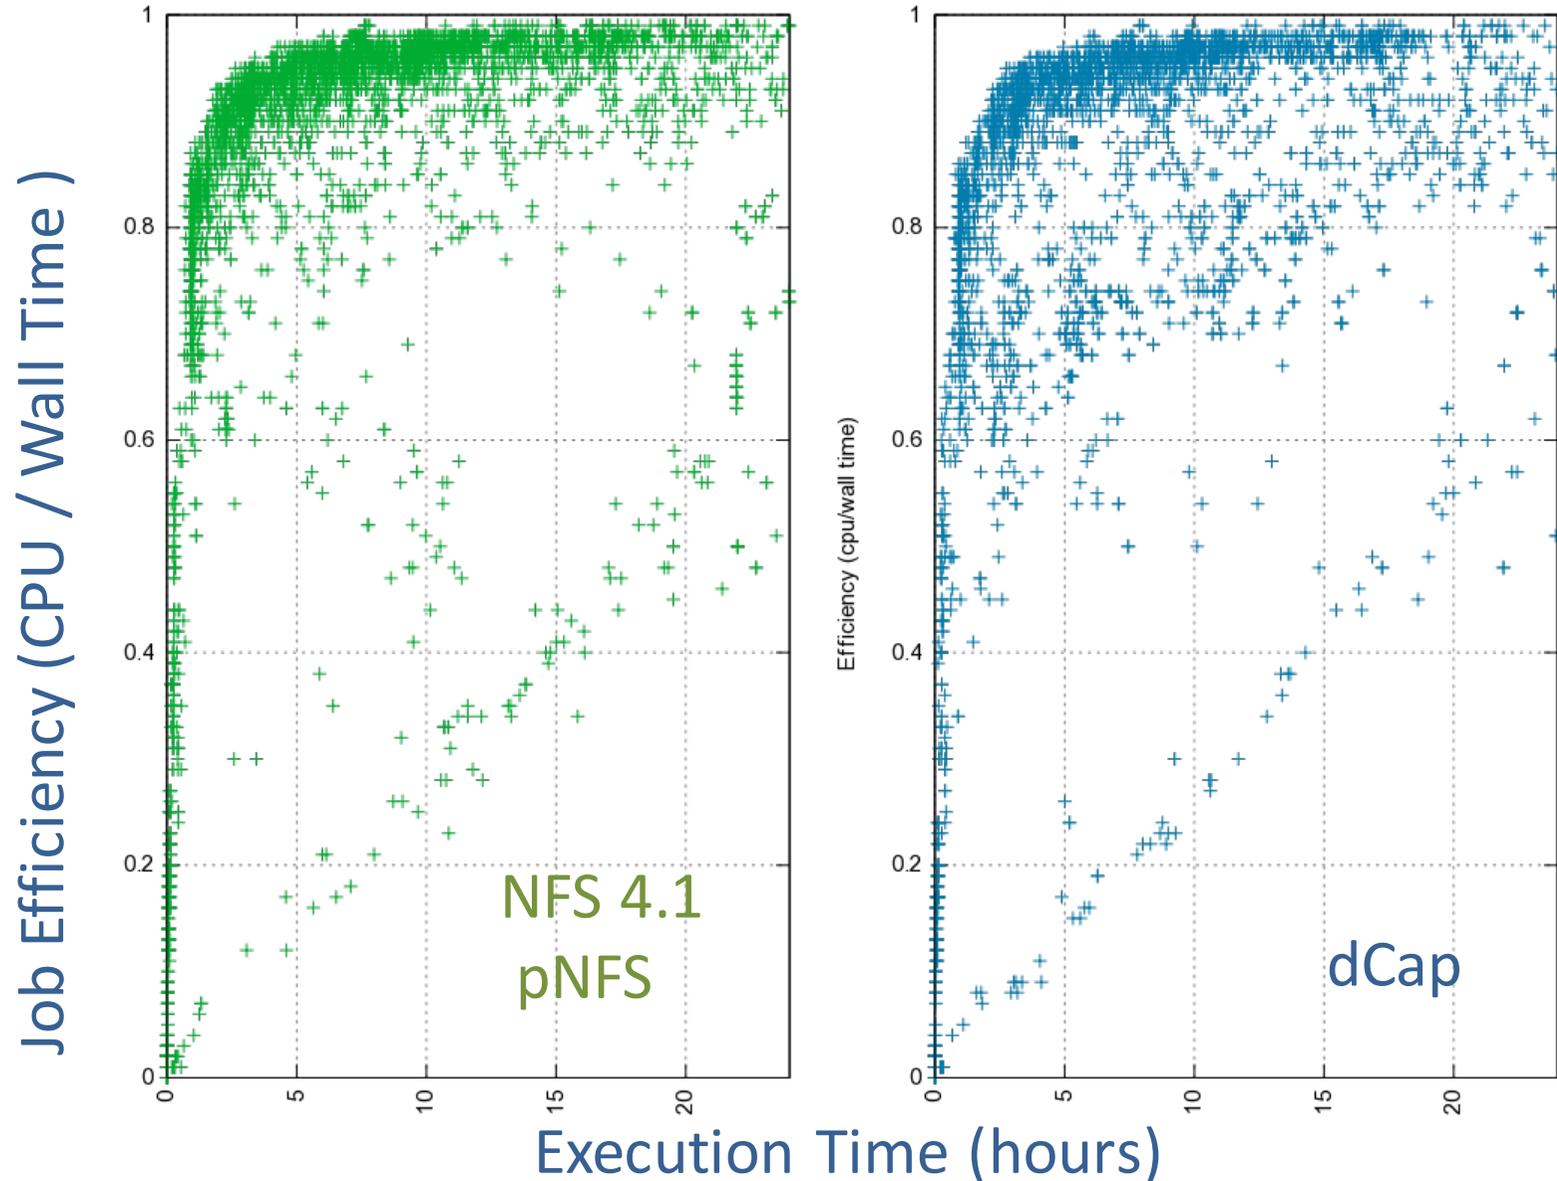


See Paul's presentation on Thursday



- Slowly migrating CMS Grid worker nodes to NFS4.1 data access.
- Good experience as long as the network is stable.

Job Efficiency (NFS – dCap)



As with all news spec's, there
are issues

- Network problems cause the system to be behave unpredictable.
- Data Server behind firewalls
- Weak clients on VM's
- Specification Violation
 - infinite state recovery with Linux kernel

Exploring new communities.

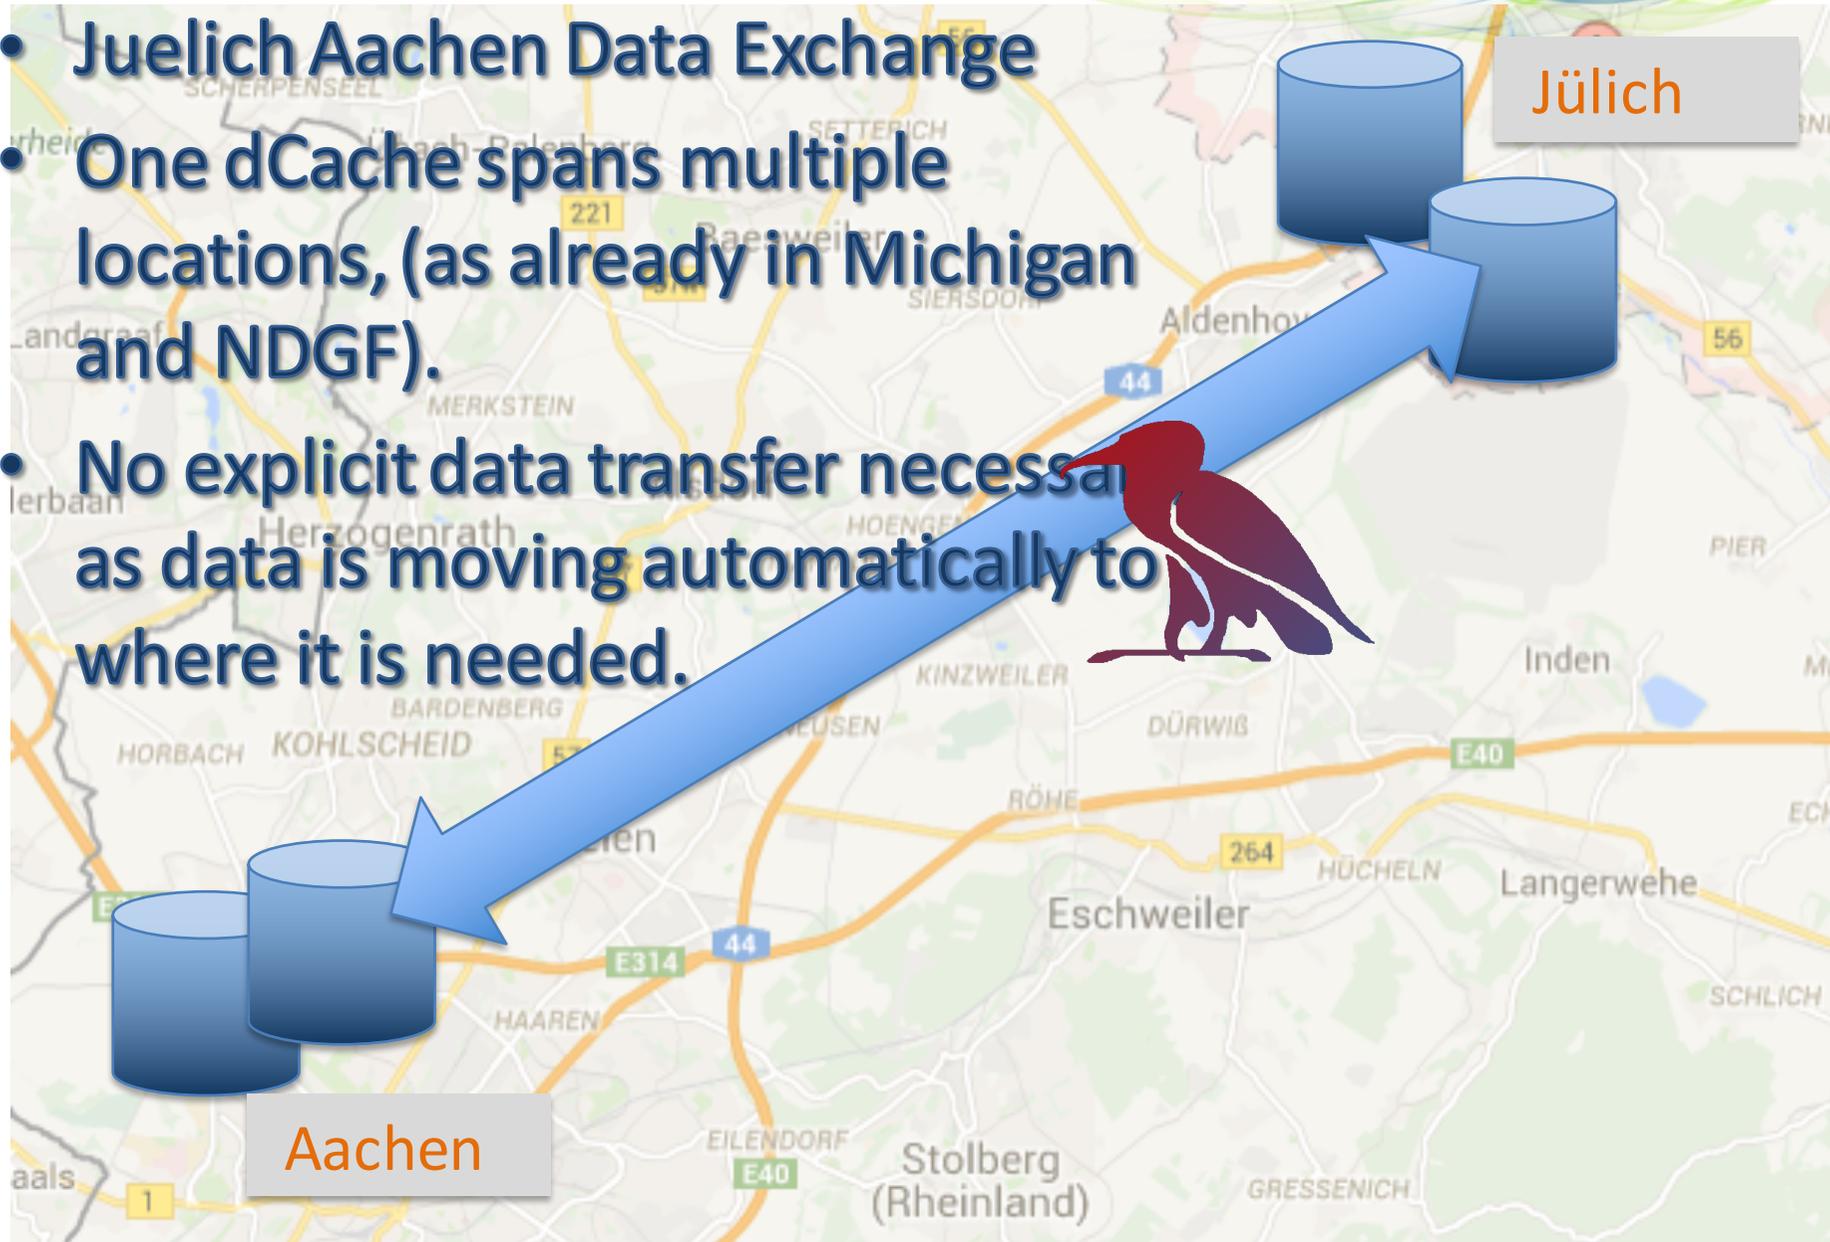
Exploring new communities.



- **Jülich – Aachen Research Association, JADE**
 - "Supercomputing and modeling for the Human Brain (SMHB)", associated to the European Human Brain Project (Plenary by KH Meier)
- **MoSGrid**
 - Scientific Gateway for molecular simulation.
- **VAVID**
 - Data Gateway for analyzing wind energy infrastructures

JADE

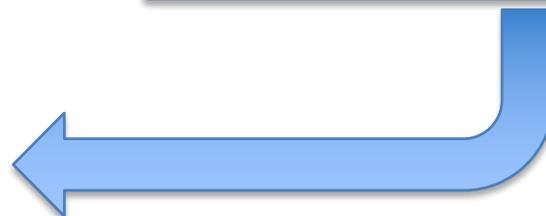
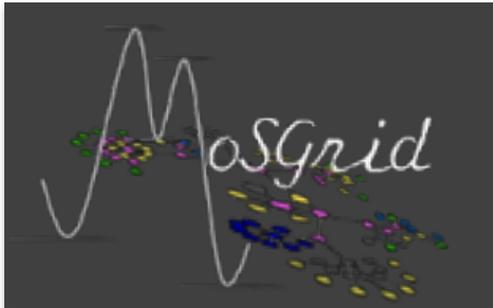
- Juelich Aachen Data Exchange
- One dCache spans multiple locations, (as already in Michigan and NDGF).
- No explicit data transfer necessary as data is moving automatically to where it is needed.



Projects in HPC



HPC jobs on supercomputer



HPC jobs get access to dCache storage.

With the start of INDIGO-DataCloud, its money
and a larger team (8+3) we can continue to
explore new horizons.
(Back to development mode)

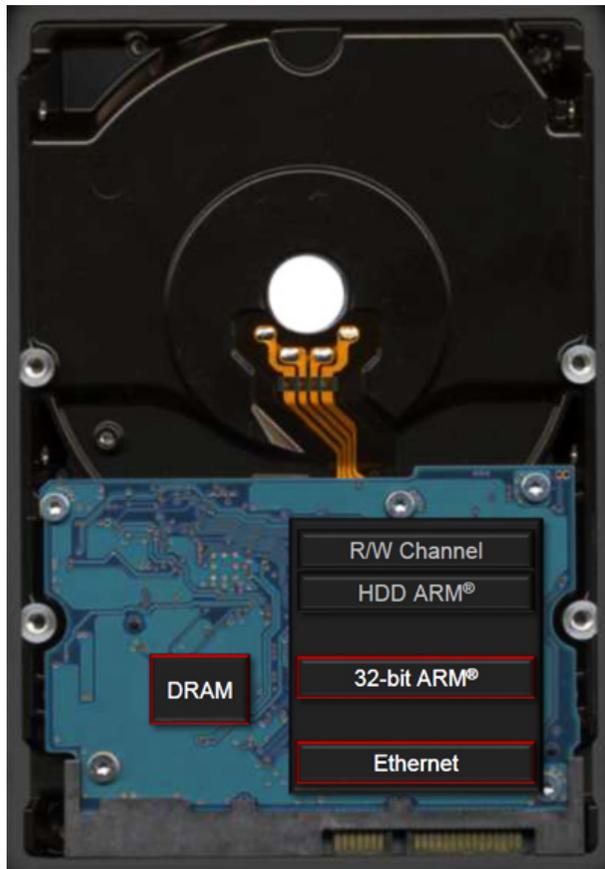
Responding to new technologies

- New Disk Technologies
 - Open Ethernet Disks (HGST)
- New Object-Store Back-ends
 - CEPH
- New European Projects (INDIGO DC)
 - Focusing on Data Quality of Service and
 - Data Lifecycle Management

HGST Open Ethernet Disks



- Small ARM CPU with Ethernet piggybacked on regular Disk.
- Spec:
 - Any Linux (Debian on demo)
 - CPU 32-bit ARM, 512 Level 2
 - 2 GB DRAM DDR-3 Memory
 - 1792 MB available
 - Block storage driver as SCSI *sda*
 - Ethernet network driver as *eth0*



- Additional CPU is not used by disk itself and can run arbitrary customer OS.
- Disk is seen as regular block device.
- Not yet on the market.
- dCache got 5 disks and we are evaluating to run pool nodes on the disk itself.
- See talk on Thursday.

Response to



- CEPH complements dCache perfectly.
 - Simplifies operating dCache disks.
 - dCache accesses data as object-store anyway already.
- dCache is evaluating a ‘two step approach’.
 - Each pools sees it own object space in CEPH
 - All pools have access to the entire space, which is a slight change of dCache pool semantics.
- Would merge CEPH and dCache advantages
 - Multi Tier (Tape, Disk, SSD)
 - Multi protocol support for a common namespace.
 - All protocols see the same namespace
 - All the dCache AAI features
 - Support for X509, Kerberos, username/password

INDIGO-DataCloud Cheat-Sheet

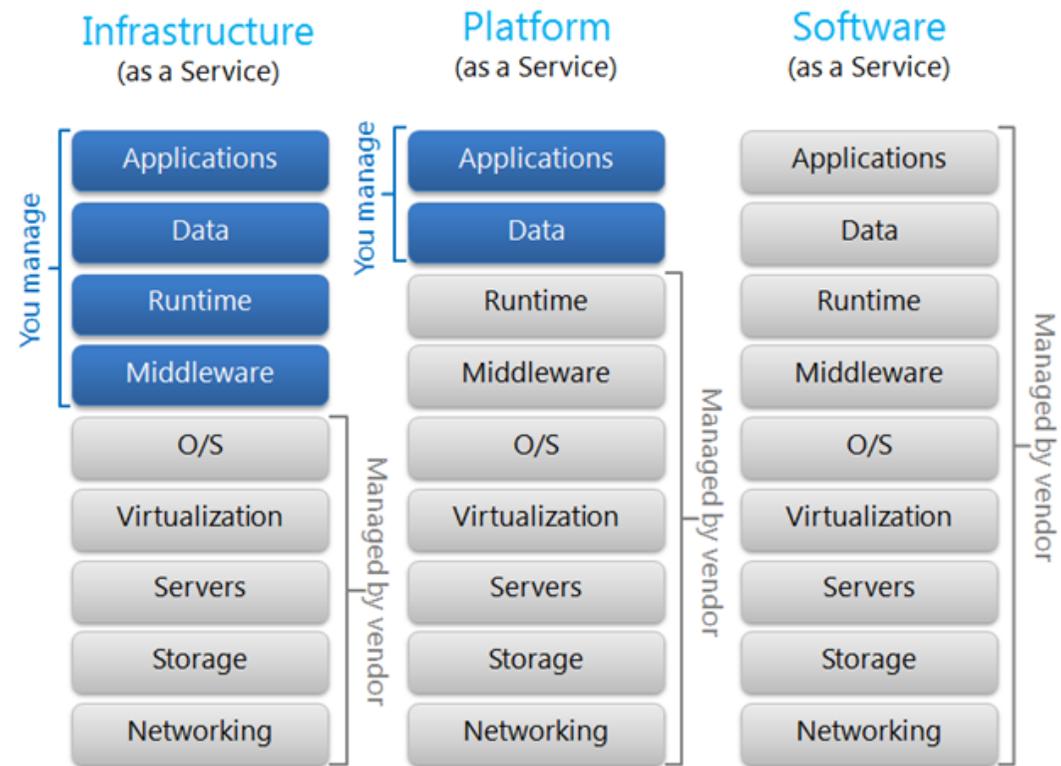
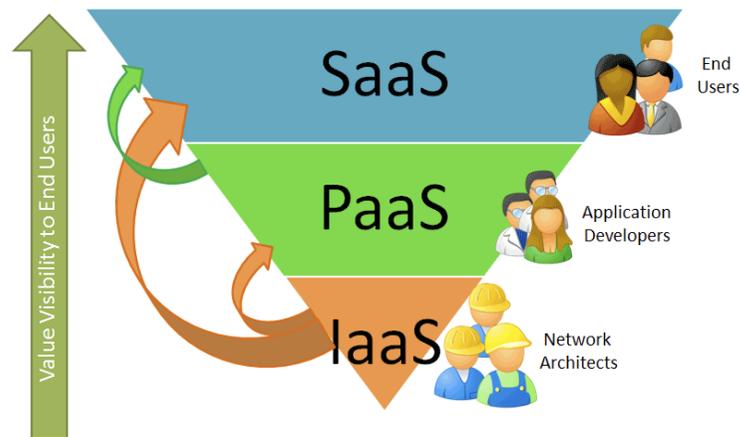


- Horizon 2020 project starting April or May
- Budget 11.1 Million Euros (800.000 for dCache)
- 26 Partners See Ludek's presentation on Wednesday
- Duration 30 months
- The project aims for an Open Source Data and Computing platform targeted at scientific communities, deployable on multiple hardware, and provisioned over private and public e-infrastructures.

INDIGO in a nutshell

1. Self-service, on-demand
2. Access through the network
3. Resource pooling
4. Elasticity (with *infinite resources*)
5. Pay as you go

In the end,
Applications Rule.



Stolen from Davide Salomoni (Project Director)

dCache involvement in INDIGO



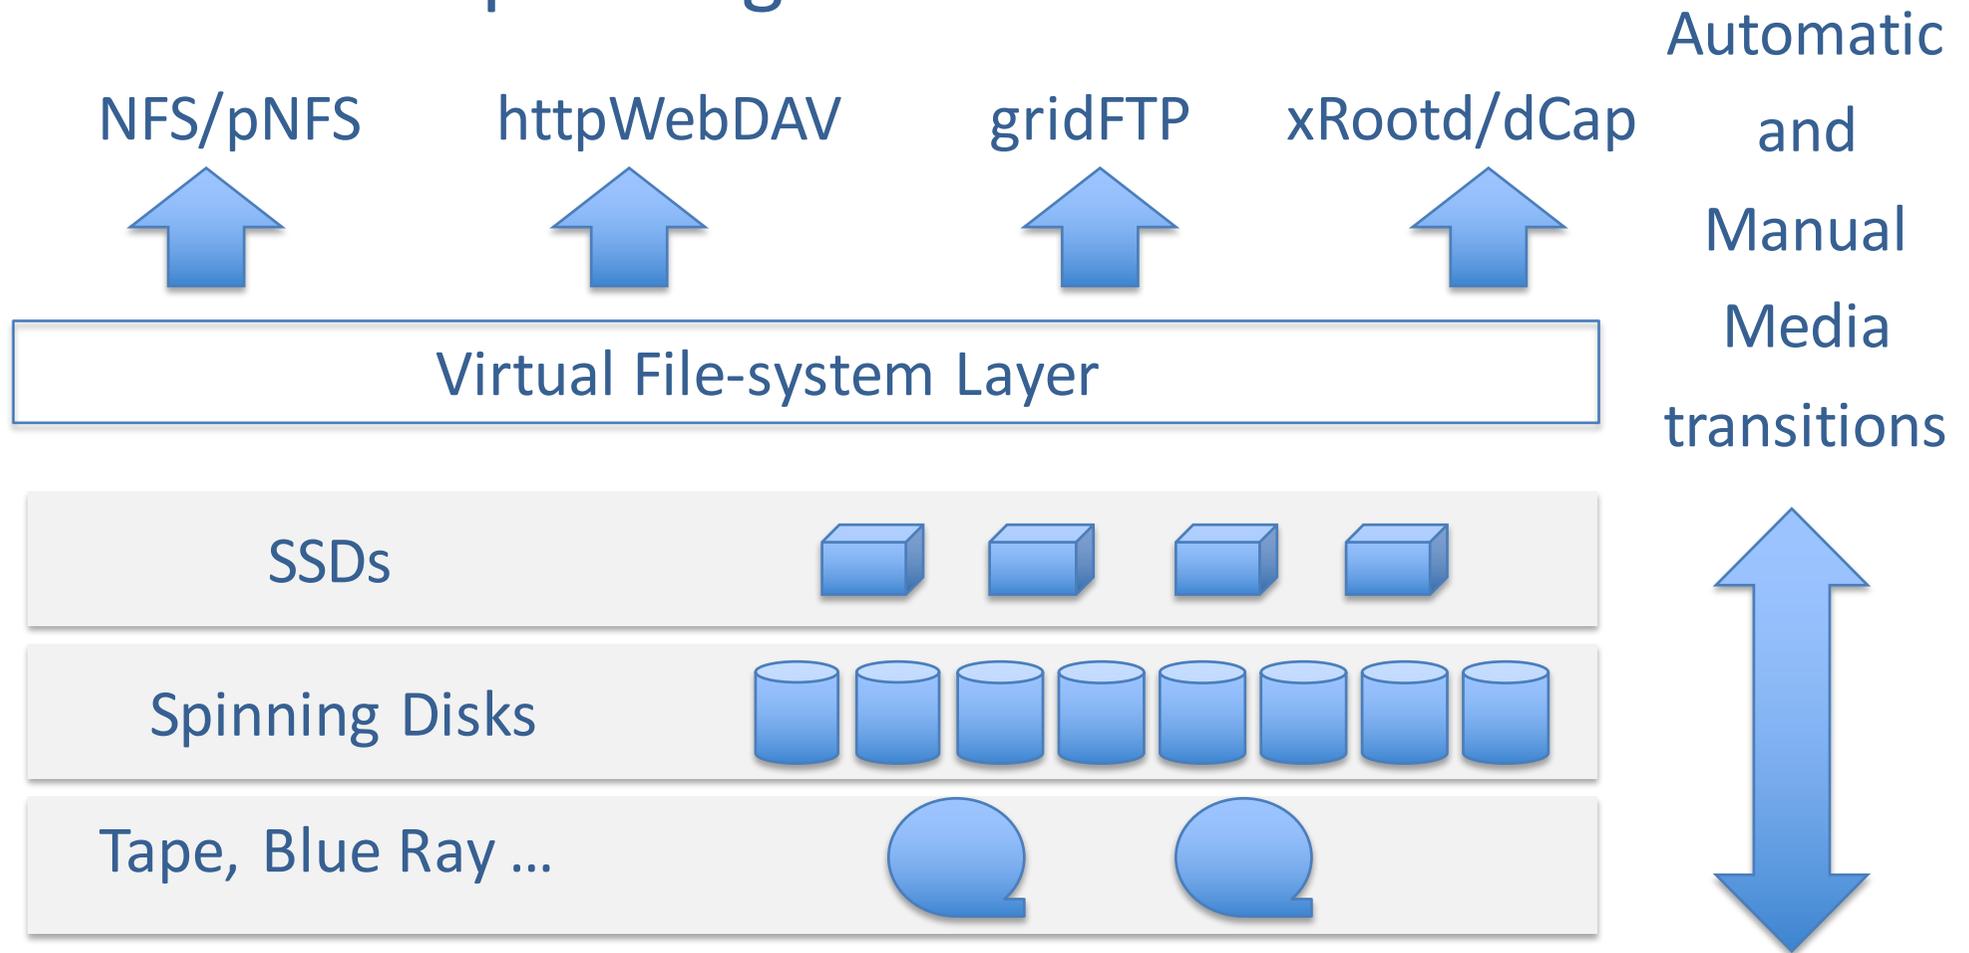
- dCache is mostly involved in WP4, which is about Virtual Infrastructures. (IaaS)
- For storage systems, like dCache, this essentially means SDS (Software Defined Storage), which according to Wikipedia is:
 - **Software-defined storage (SDS)** is an evolving concept for computer data storage software to manage **policy-based provisioning and management of data storage independent of hardware.**

- User/PaaS defined “Quality of Service” management
 - User/PaaS defined “Access Latency”
 - SSD or Tape depending from application requirements.
 - User/PaaS Defined “Data Protection”
 - On one disk, two disks or tree tapes depending on how precious your data is.
 - User/PaaS Defined “Data Migration Policies”
 - Like Amazon Glacier vers. S3
- Automatic Storage-Tier migration
 - Based on access profile
- All this wouldn’t be needed if SSD’s would be cheap and 100 % reliable.

dCache is well prepared

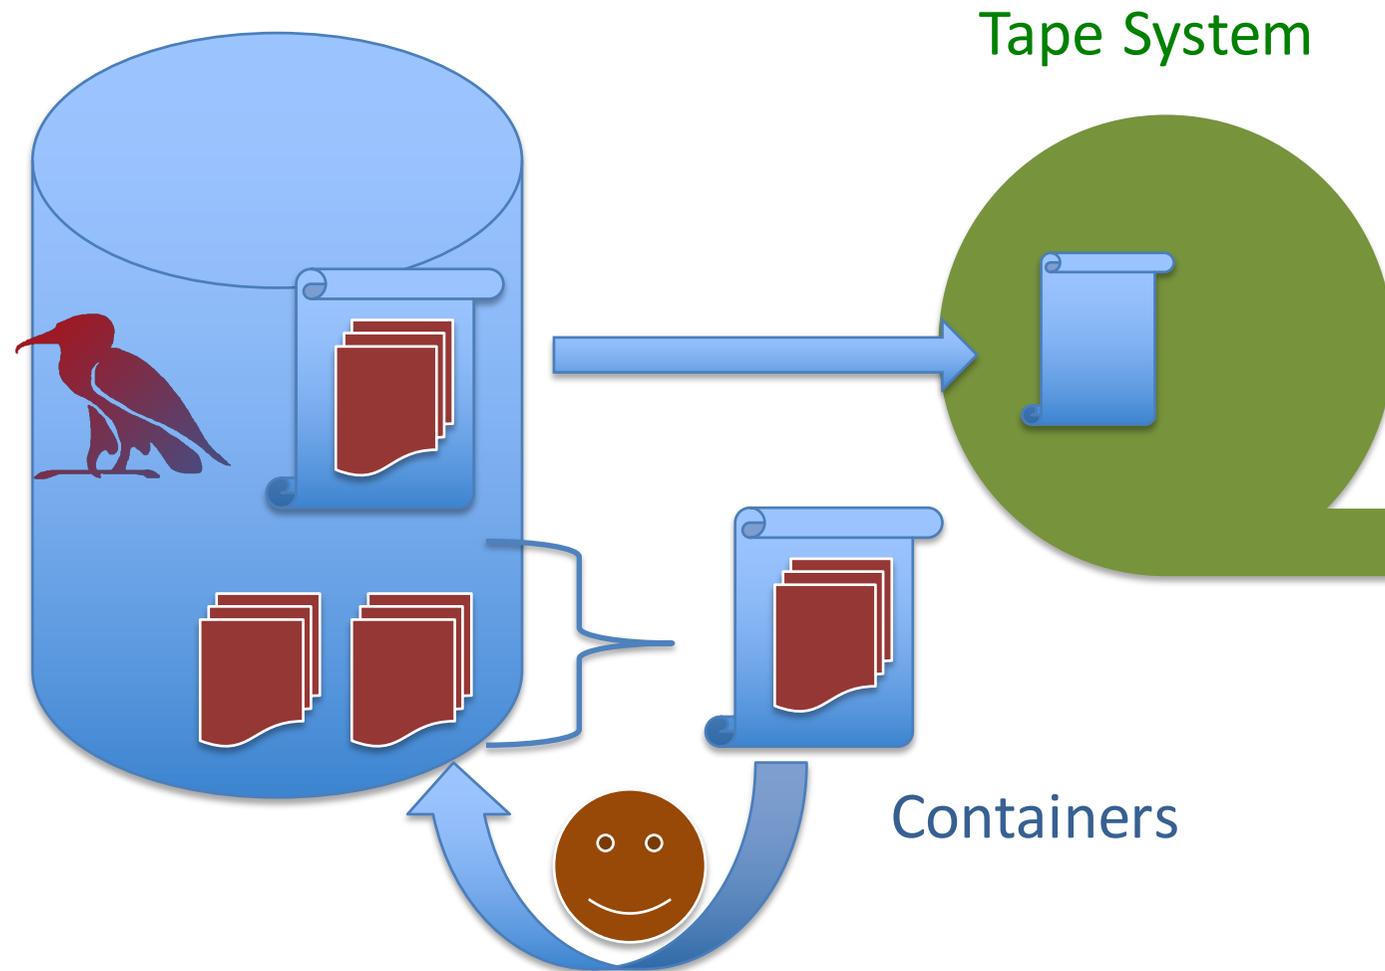


Historically dCache supports multi-tier storage and the corresponding transition.



Recently added

We optimized the 'small file' problem with disk <-> tape transitions.



What's missing

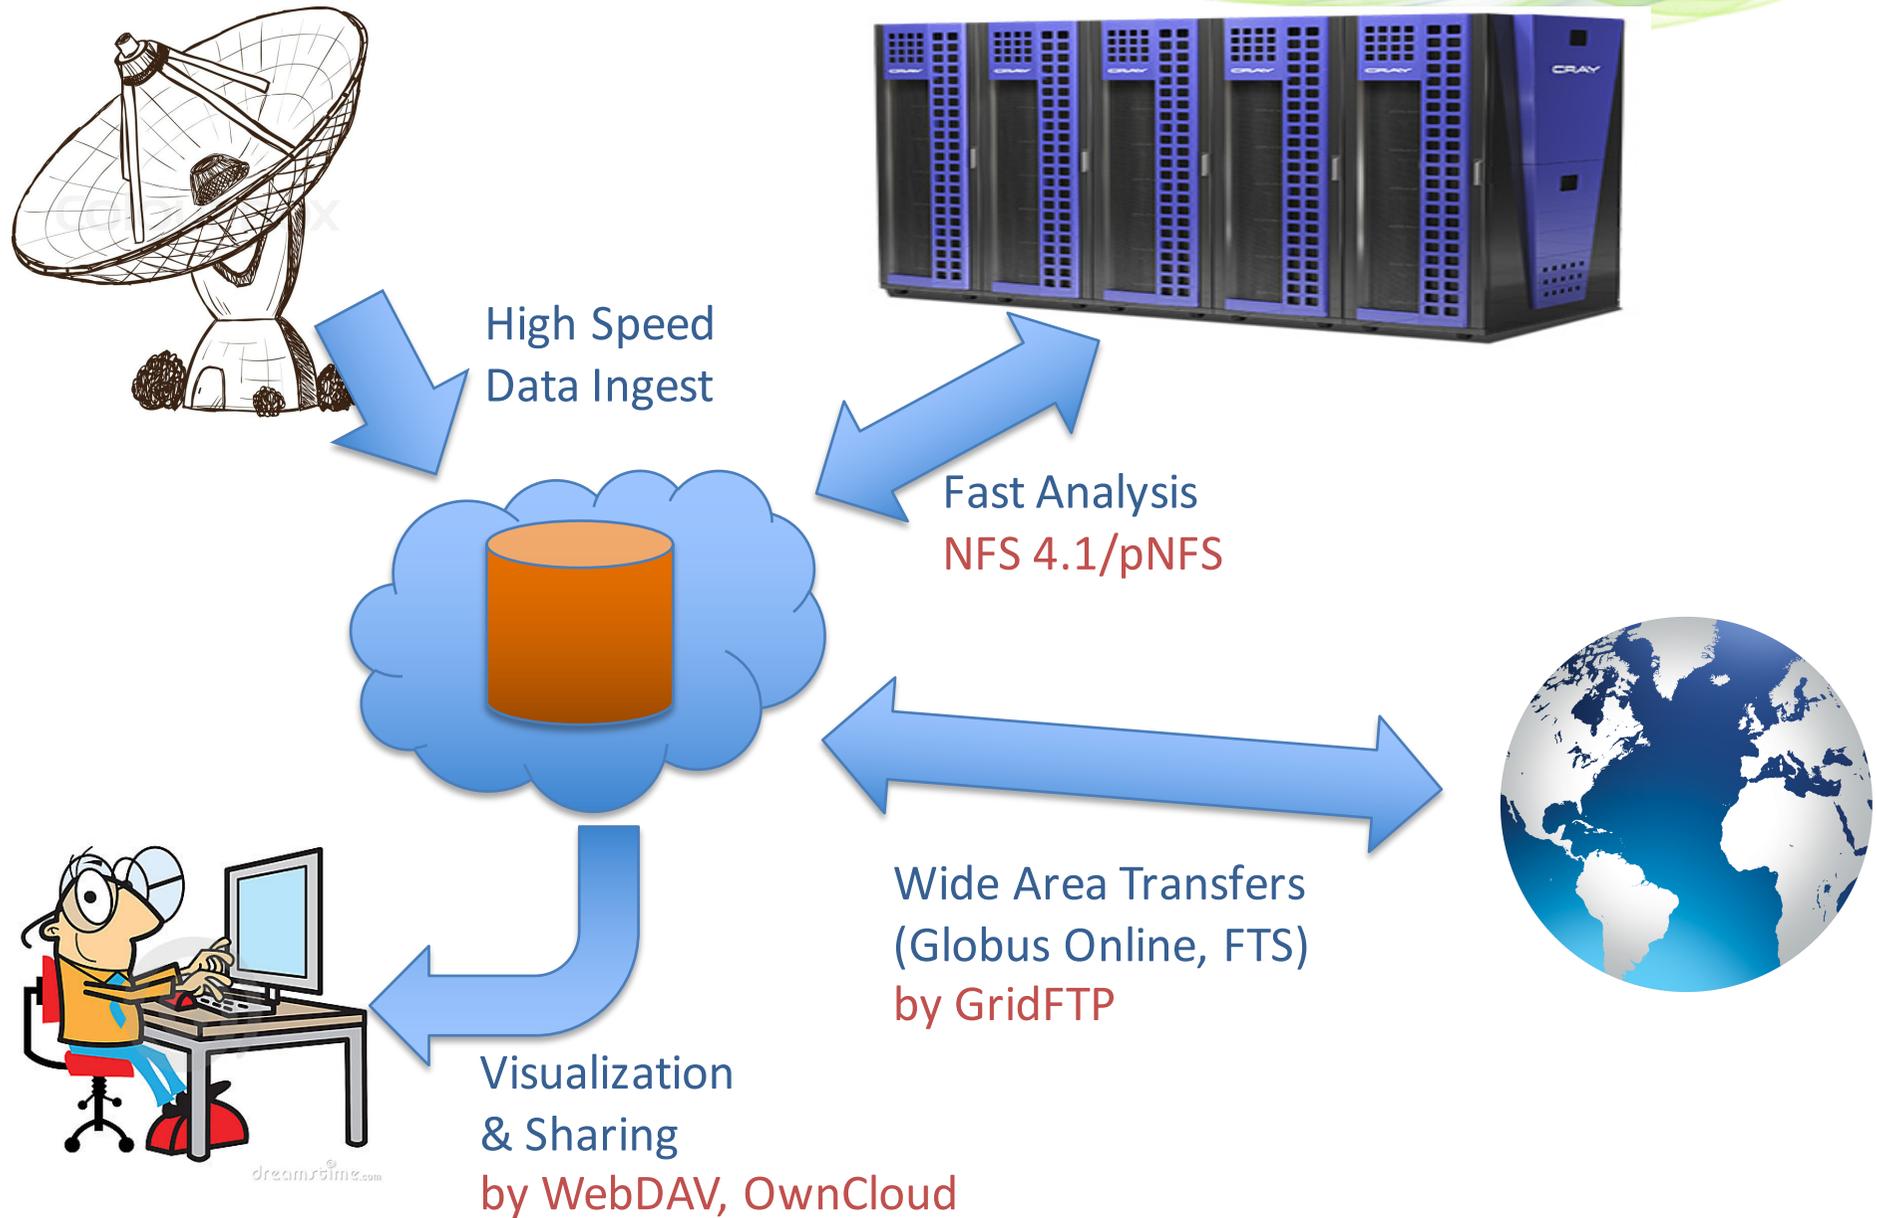
- Mainly a common agreement (standard) on how to trigger transitions. (Protocol, API ??)
- We have some experience with SRM, however it seems not to be suitable for this purpose.
- Another candidate is CMDI (SNIA), which is an industry standard.
- Migration Policies are already discussed, documented and implemented within RDA (Practical Policy Working Group).
- Details will only be available after the INDIGO kick off meeting end of April '15.

Summary

Magically, up to now, at the right moment, there was always an EU or National Project, funding dCache exactly for those features or activities, dCache was planning to do anyway and with that they helped us following our master plan :

The support of the
Complete Scientific Big Data Life Cycle
Management.

Scientific Data Lifecycle



Don't forget

Upcoming dCache Workshop

18 – 20 May 2015

Amsterdam, Science Park

Visit www.dCache.org for details

The END

further reading
www.dCache.org

