

dCache : Big Data and HEP

Paul Millar

(on behalf of the dCache team)
CHEP 2013, Amsterdam



Too much happening to report everything

Choose three areas I like:

- Managing storage
(controlling what's available)
- Data Analysis
(getting information from data)
- Establishing your identity
(the starting point for a secure system)

Managing storage



Managing storage with CDMI

- CDMI is the **ISO standard** for cloud storage
- Non-WebDAV features of SRM **supported by CDMI**
 - Discovering access protocols (FTP, NFS, ...),
 - File integrity,
 - Third-party copies
- **No direct support for tape**, but companies are investigating this:

NetApp plan to expose tape access via CDMI.


New possibilities with CDMI

- Metadata
 - Store and retrieve user-supplied metadata against files
 - Do loose skimming analysis to avoid opening uninteresting files
- Object-store
 - Don't have to auto-generate meaningless filenames
 - Faster opens
- Notification
 - Discover which files are being created,
 - Discovering when files have been recalled from tape is easier.

CDMI: status and plans

- Working on providing **initial implementation** with WebDAV-like functionality:
 - upload & download files, rename, move, deleted files and directories.
- Next, add support for storing user-supplied metadata:
 - Storing and retrieving,
 - Support for querying matching files

Data Analysis



One slide summary: NFS and pNFS

- NFS v4.1 introduced an optional feature: pNFS.
- pNFS means that HEP-proprietary LAN protocols (dcap, xrootd, rfio) become redundant:
 - Don't have to maintain a client
 - Build-in support for client-side caching
 - Lots of exciting innovations from others
- For WLCG, only just become feasible
requires WNs running Scientific Linux v6.

NFS experience: Photon Science

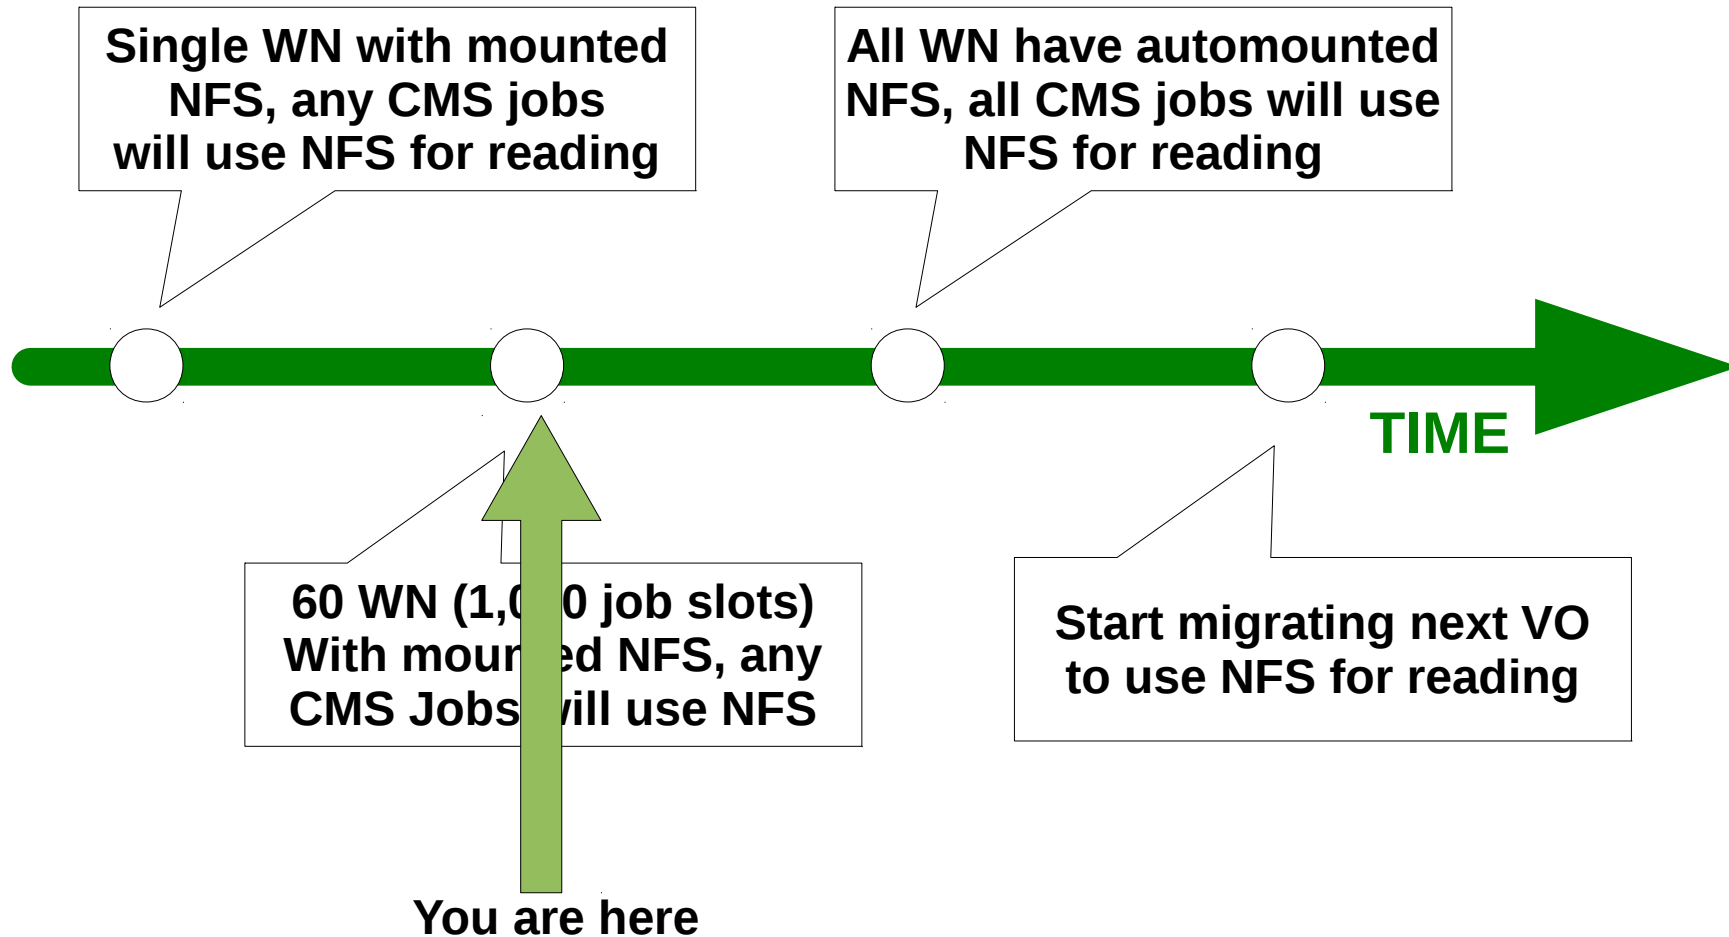


Belle experiment support at DESY

- Belle & Belle II are experiments at KEKB facility
- Belle analysis applications **require mounted filesystem**
- DESY provides considerable resources for Belle & Belle II
 - Tier-2-like site for Belle, providing analysis facilities
 - Tier-1-like site for Belle II, producing Monte Carlo results
- Interactive analysis of Belle data?
 - simply NFS mounted dCache**
- Users are so happy that they are pushing for NFS mount on the batch nodes
 - Plan to do this ASAP



Rolling out NFS for WLCG at DESY



Preliminary results

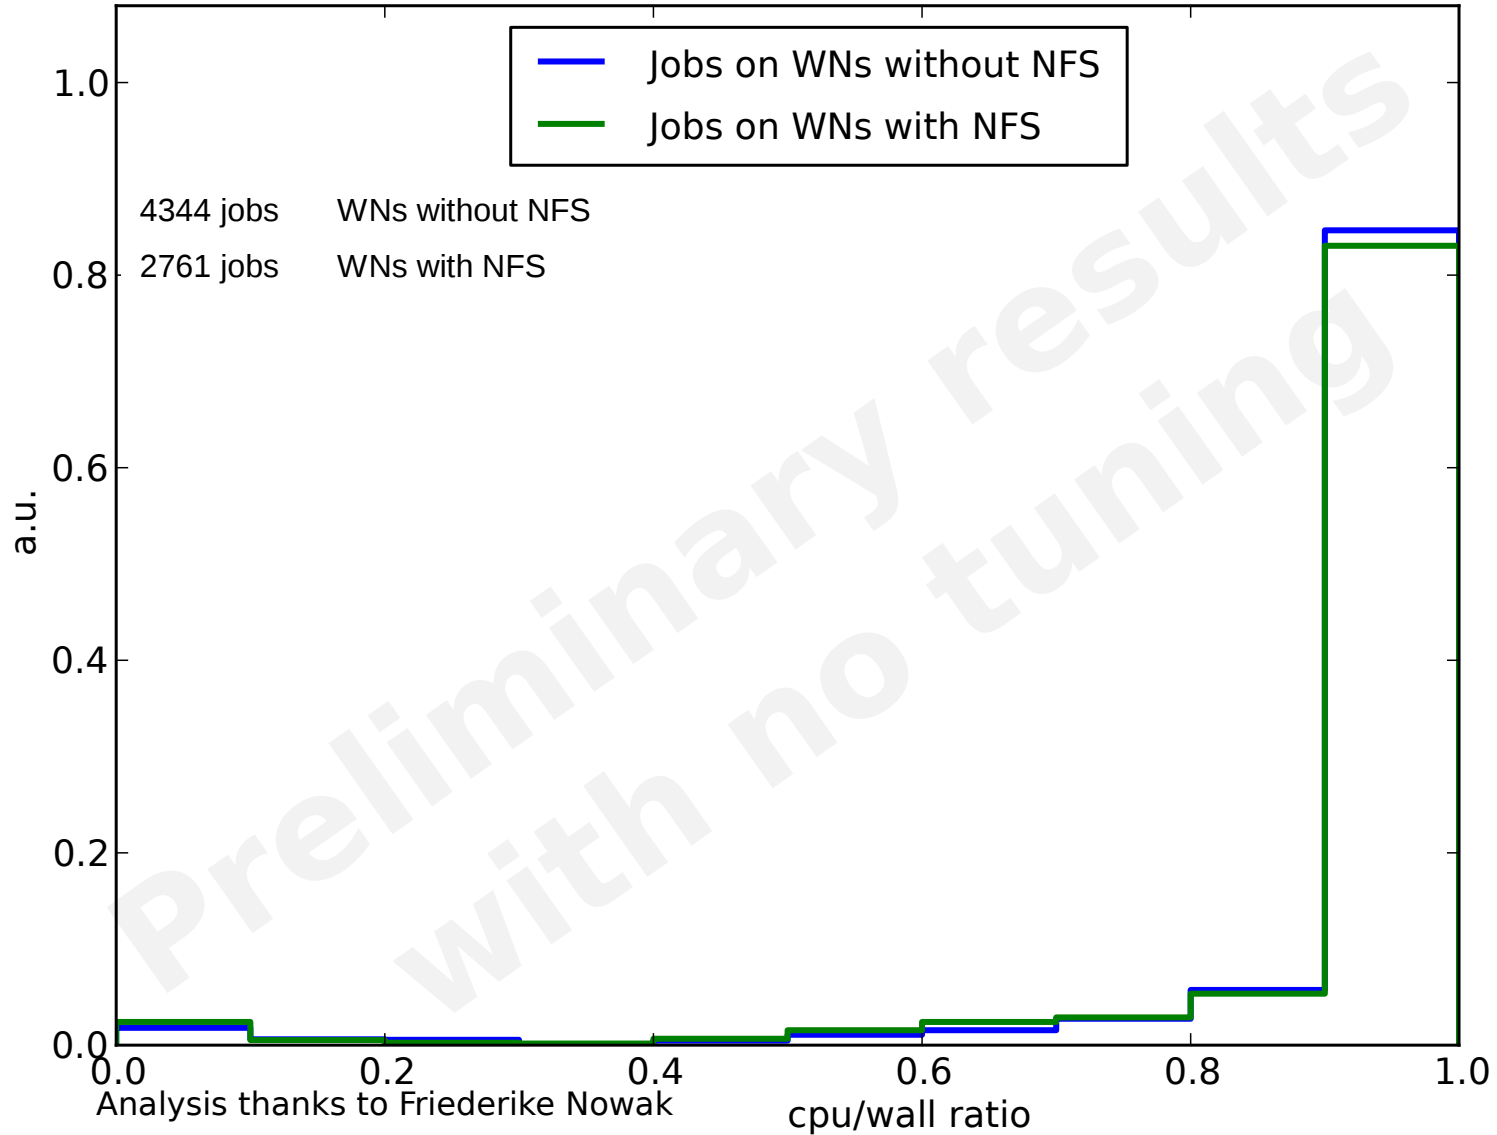
- Very rough analysis based on job efficiency (ratio of CPU- & wall-clock- times)
 - Comparing two sets of WNs with similar hardware:
(one with NFS mounted, other without)
 - Bin the measured efficiency and normalise the histogram so the area is the same.

NB #1: not enough time to capture good statistics.

NB #2: No tuning of NFS client or ROOT

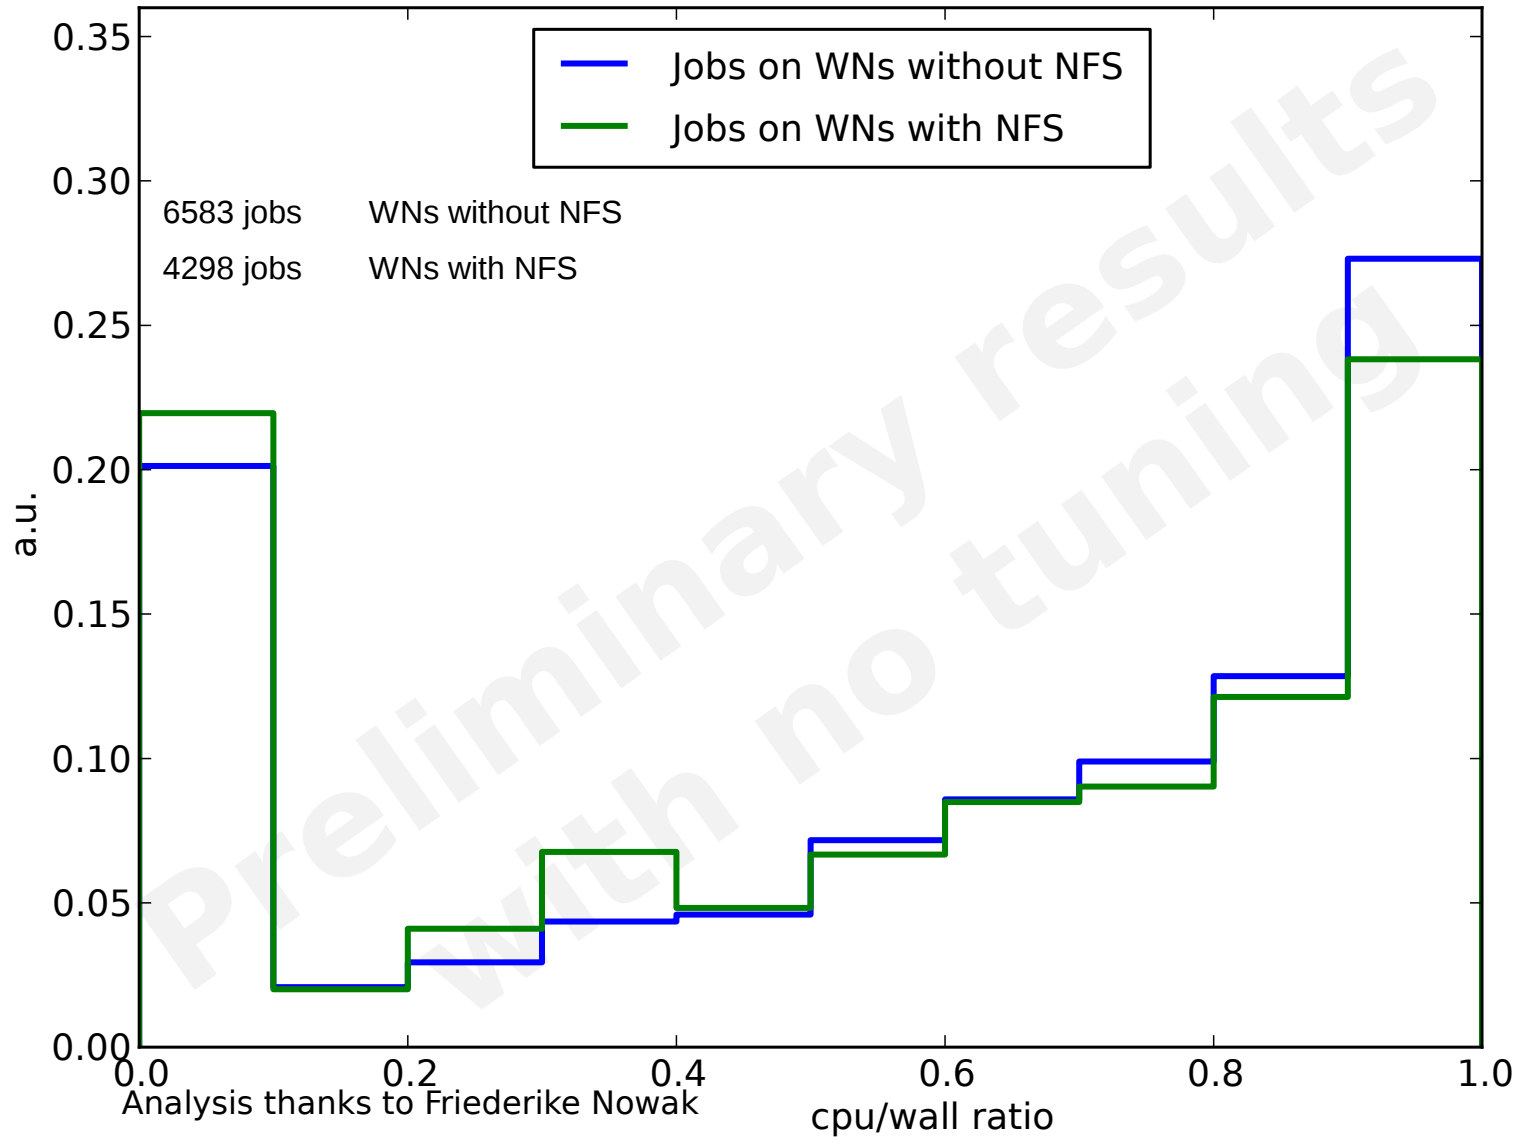
Preliminary results

ATLAS Production

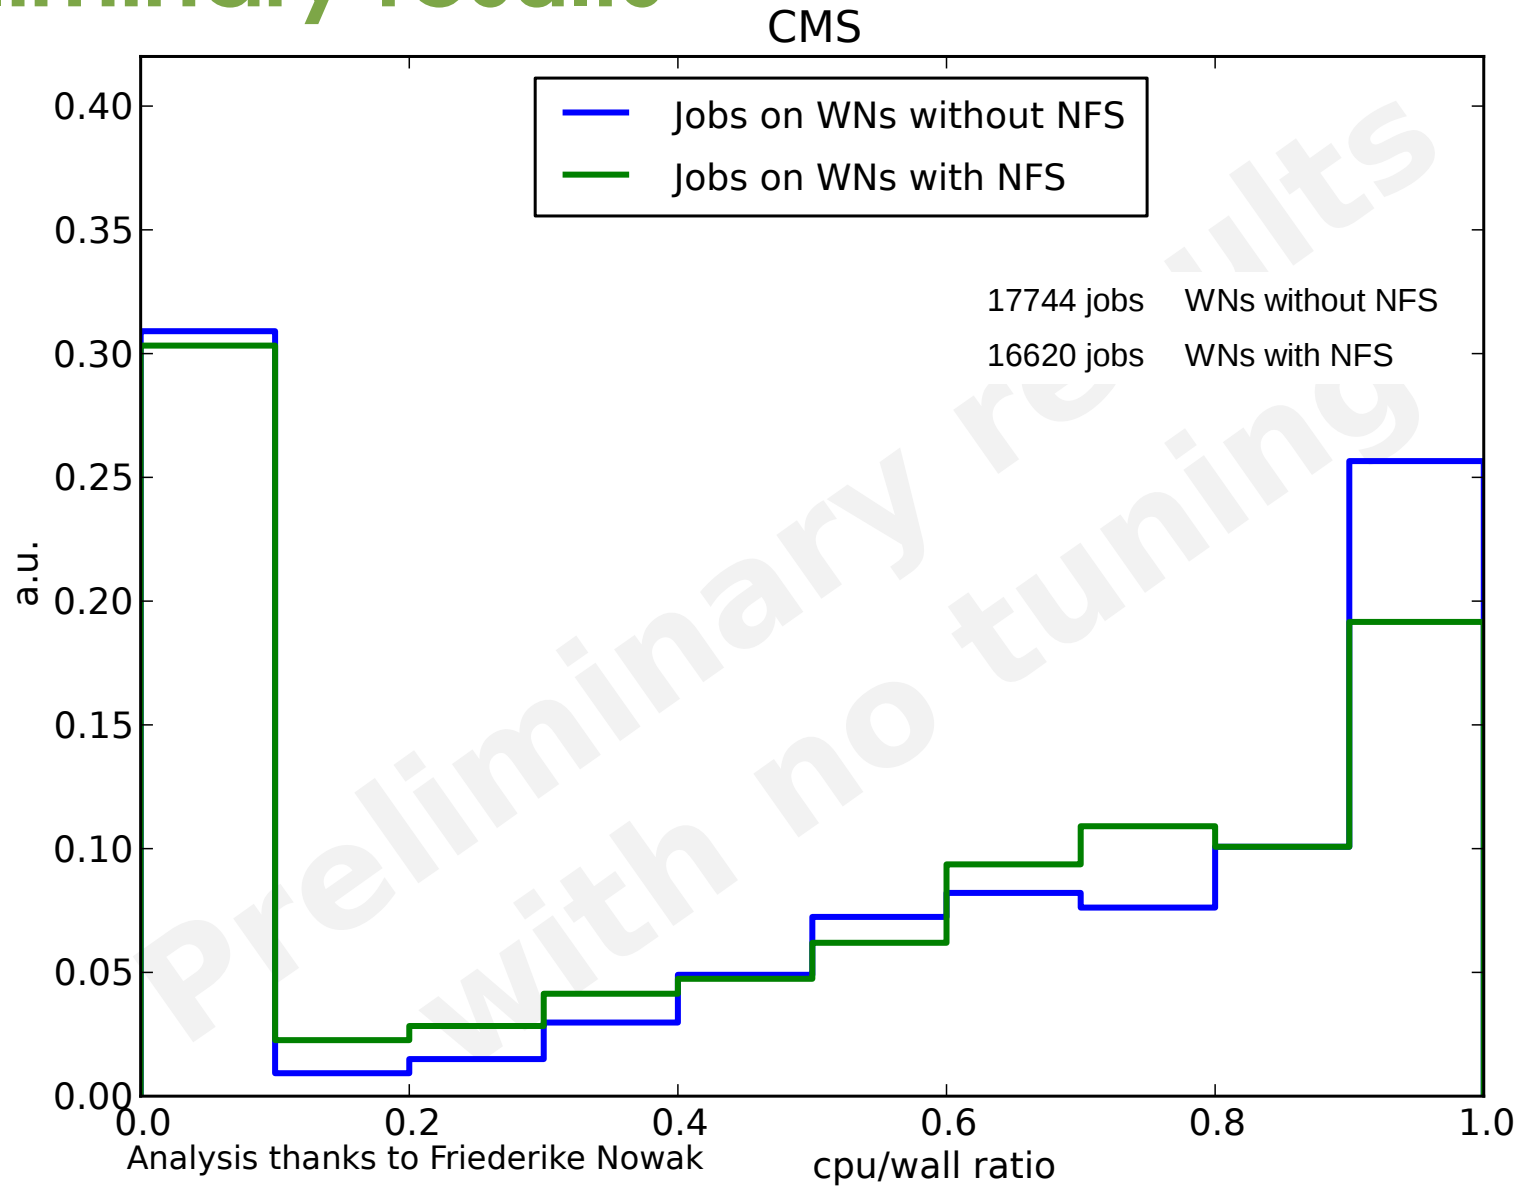


Preliminary results

ATLAS Pilots



Preliminary results



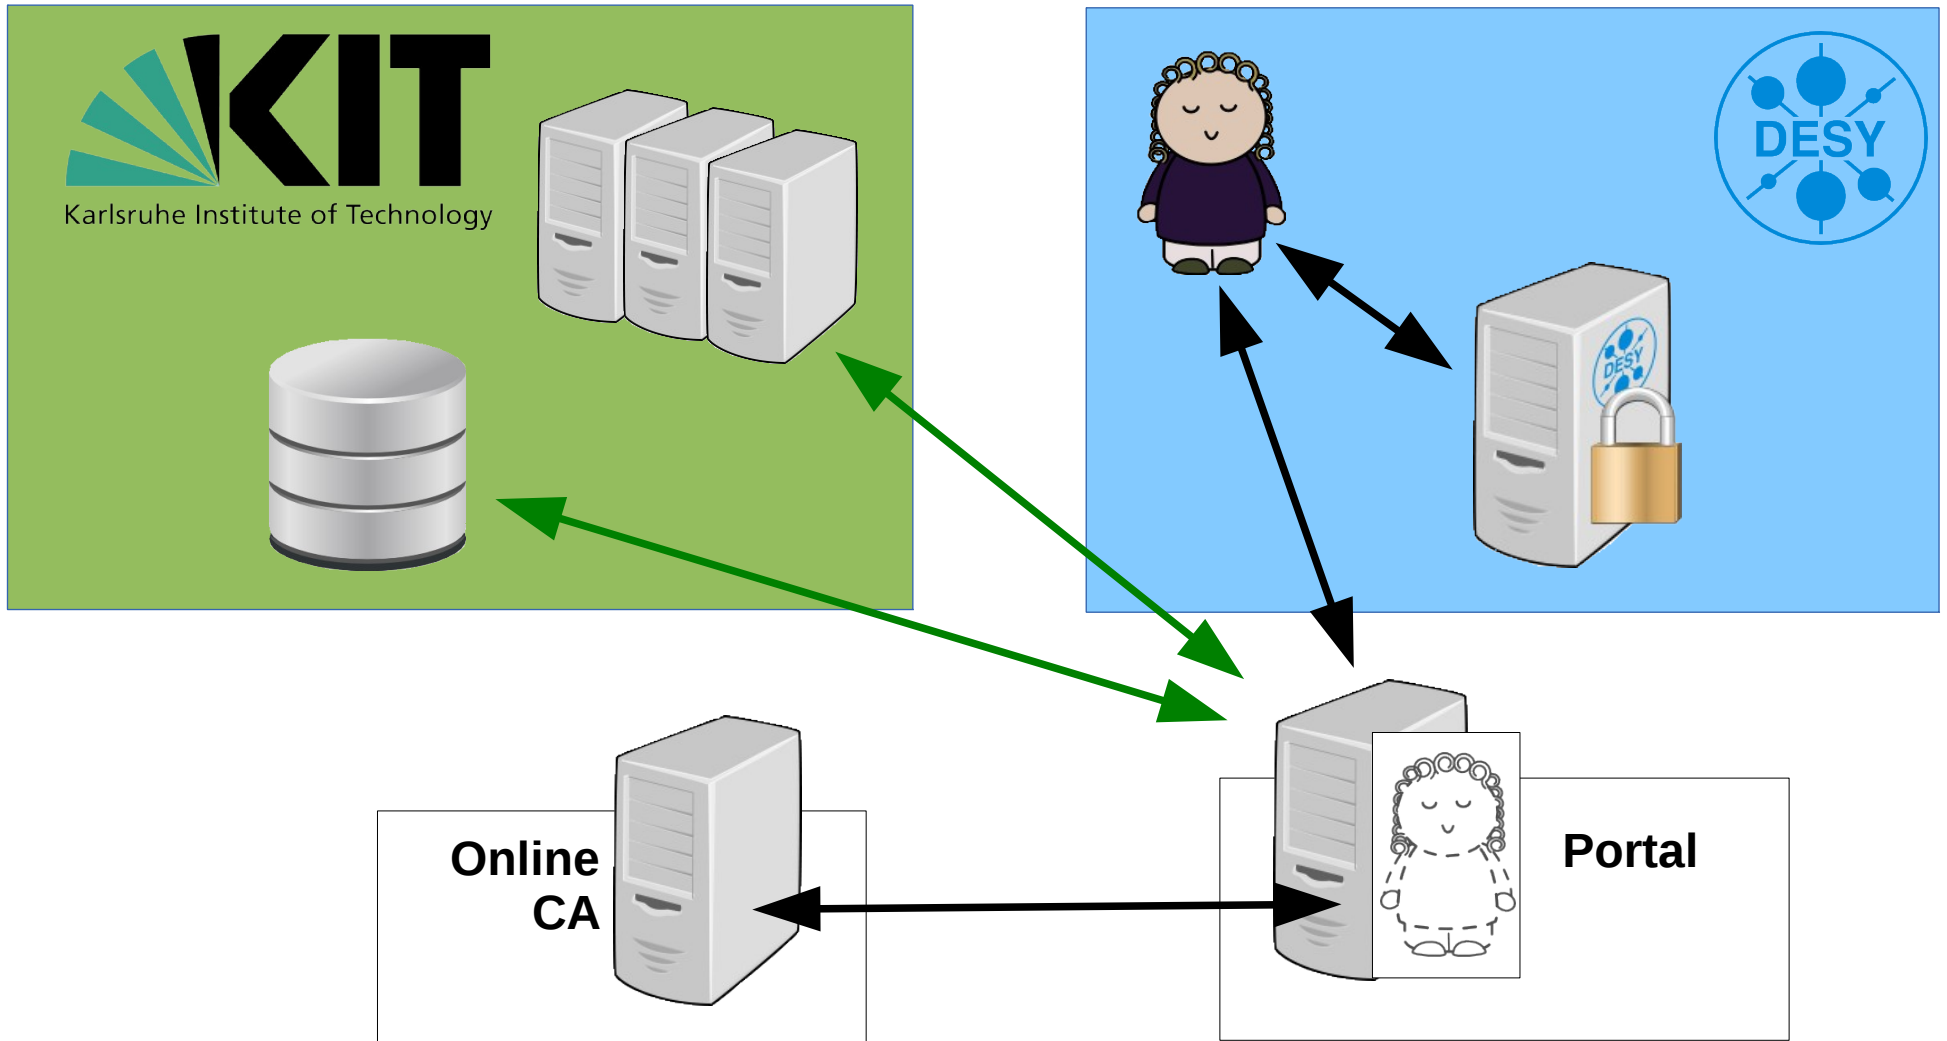
When NFS client attack!

- We discover a new failure mode.
discovering new problems is **expected**.
- Problem was identified and fixed **within 2 days**,
- Door already updated to be more robust, we will upgrade dCache CMS pools next scheduled down-time

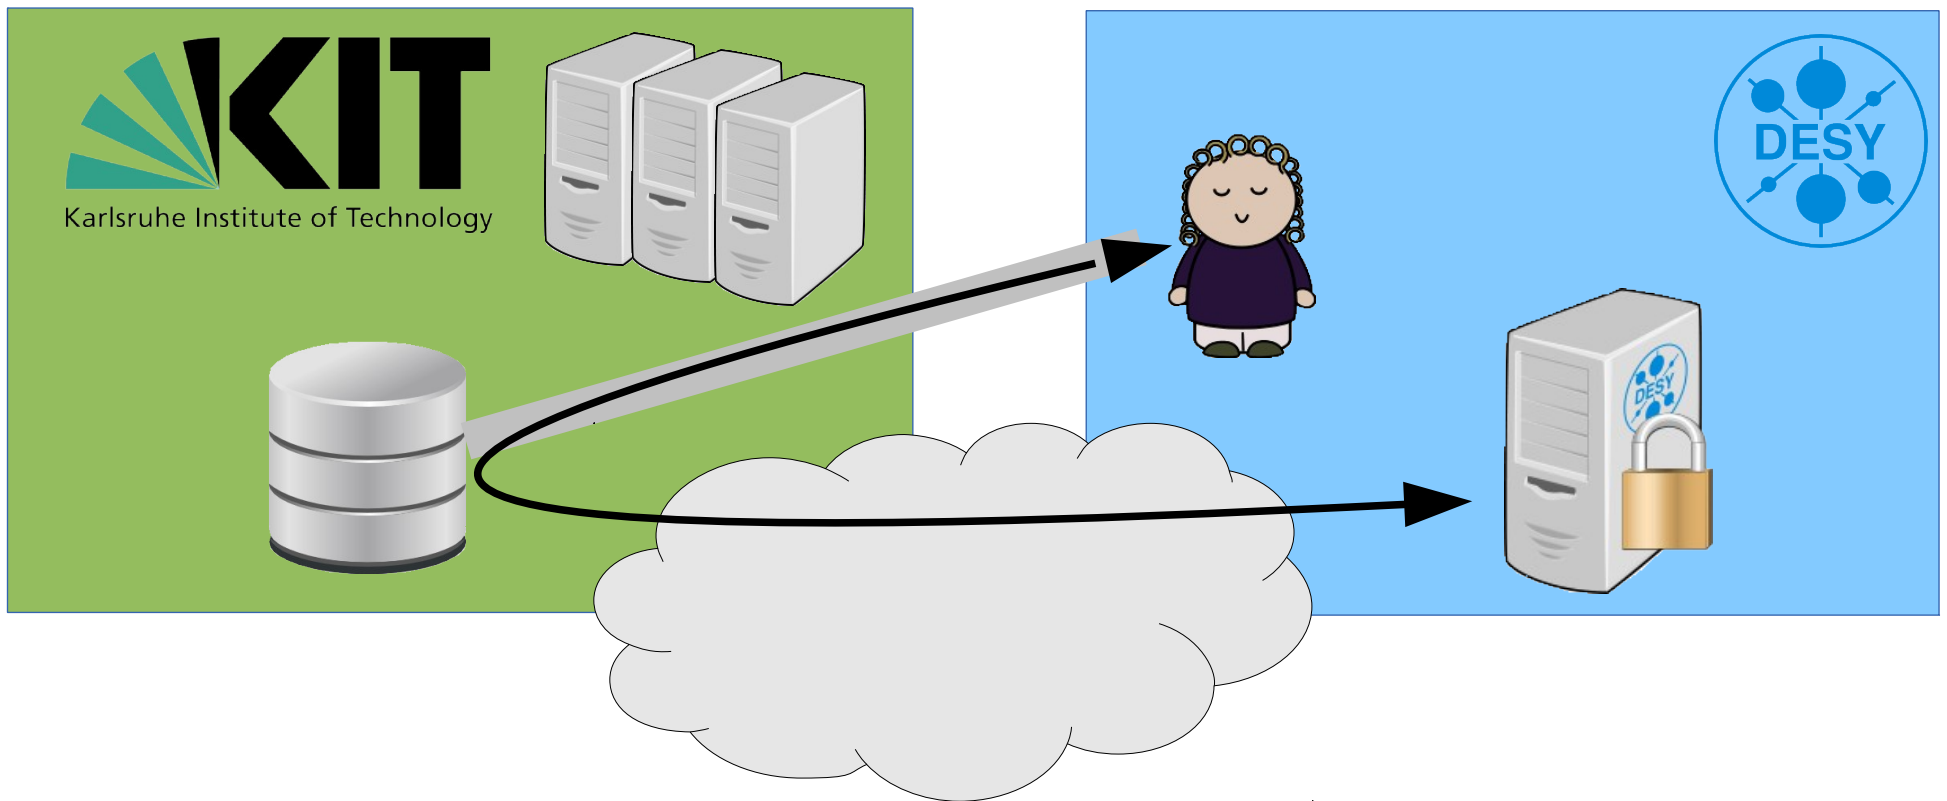
Establishing your identity



Working without X.509 certificates



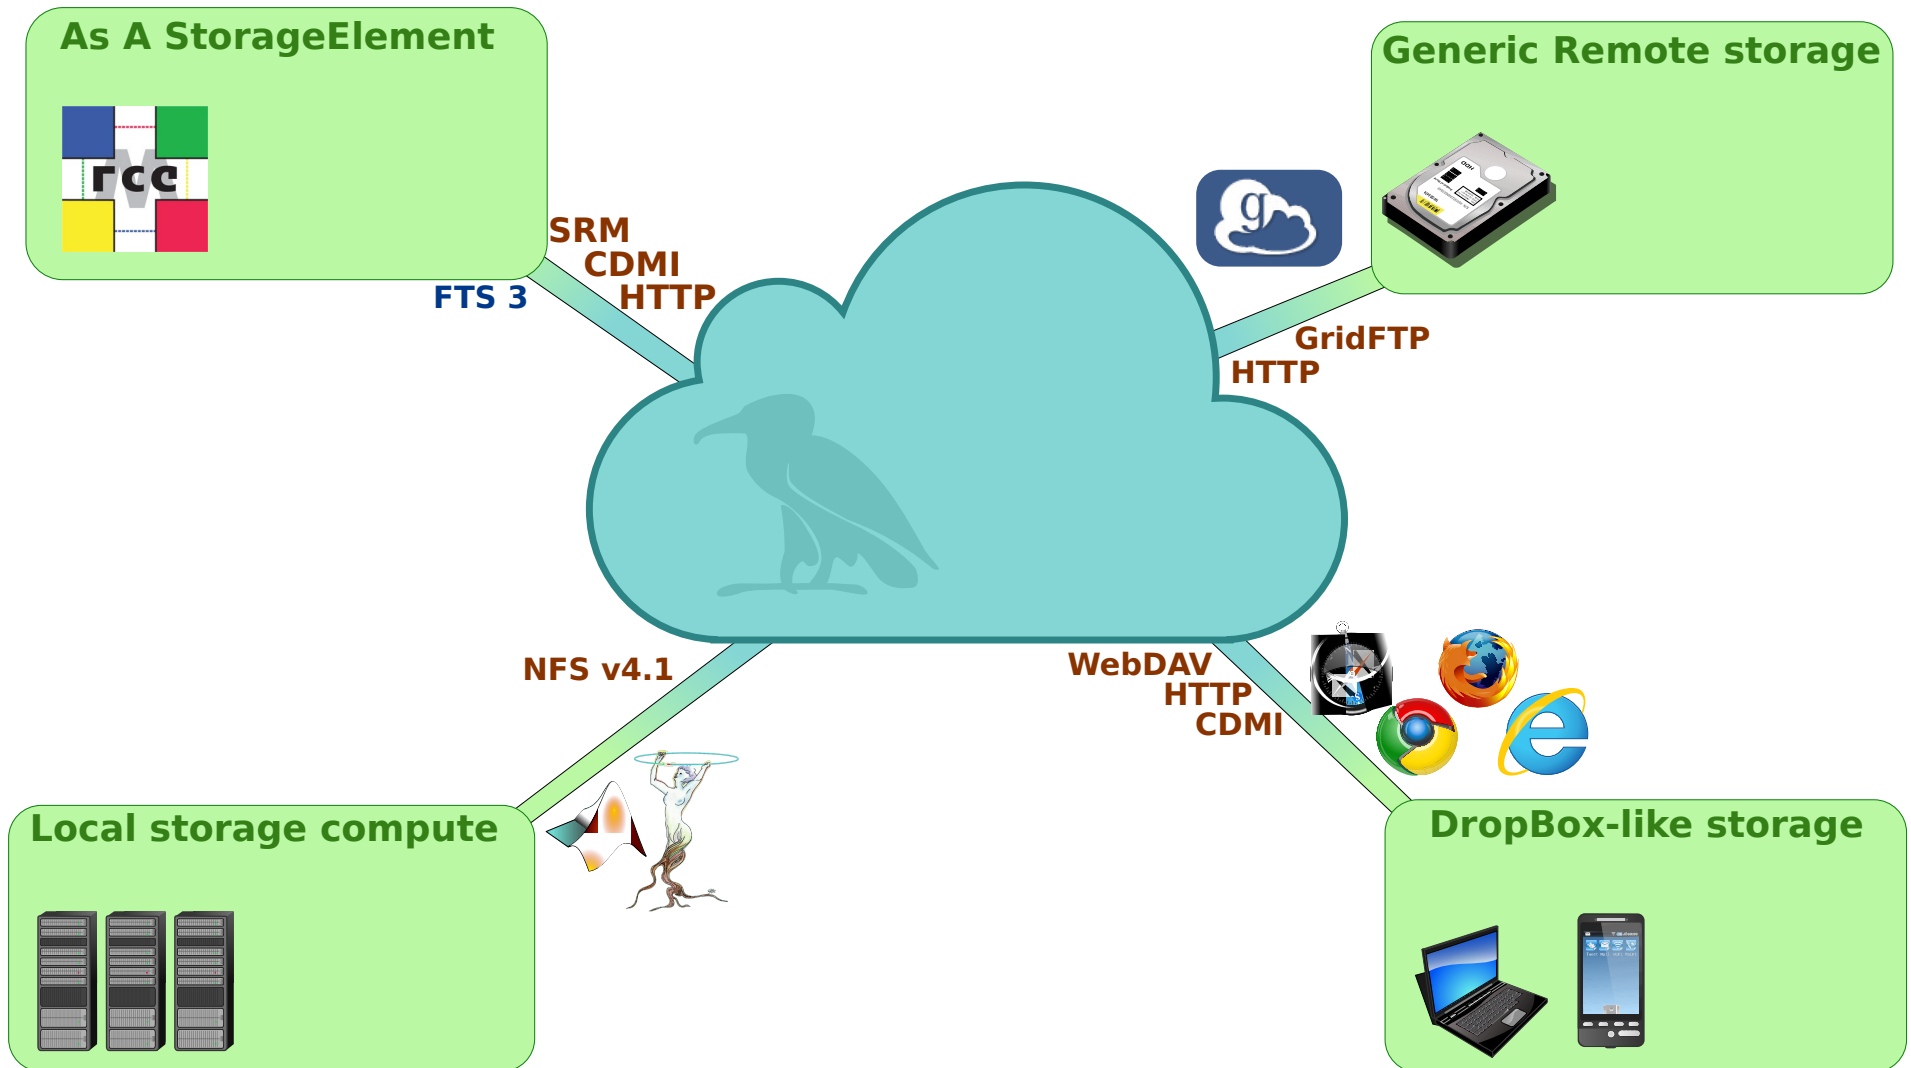
What about direct access to the data?



Bonus topic



dCache Scientific Cloud



Next steps within dCache collaboration

- Continue rolling out dCache NFS support
 - CMS, other WLCG VOs, Belle, ...
- Continue investigation into CDMI:
 - Finish current work, add metadata support, roll out for instances to satisfy demand
- Continue investigating alternatives to X.509
 - Work on supporting data migration, management and access without grid certificates within LSDMA
- Continue to innovate to match demands of our ever increasing user community

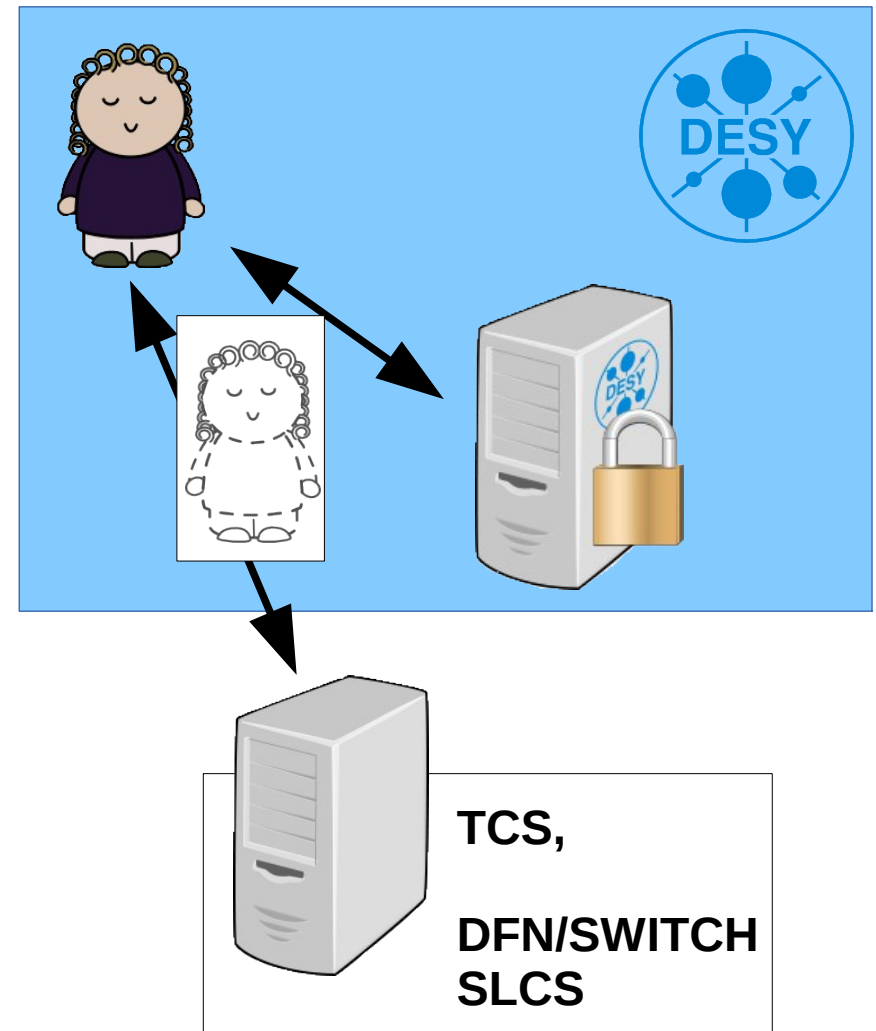
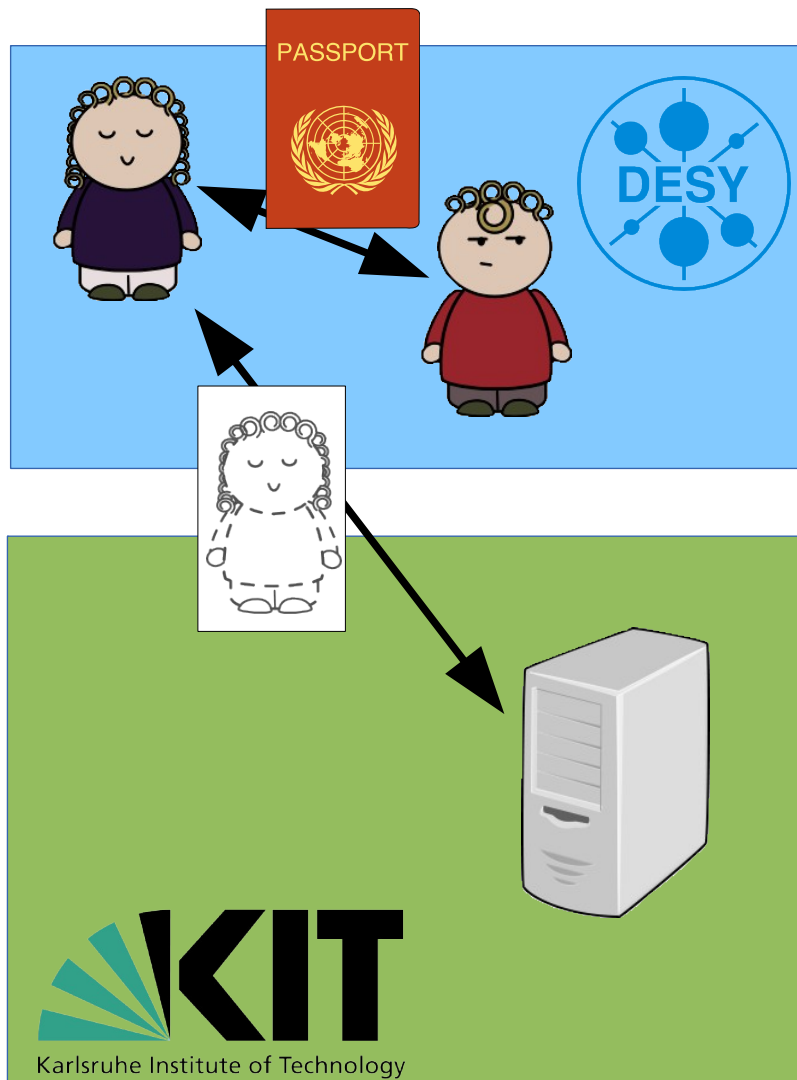
Thanks for listening

Backup slides

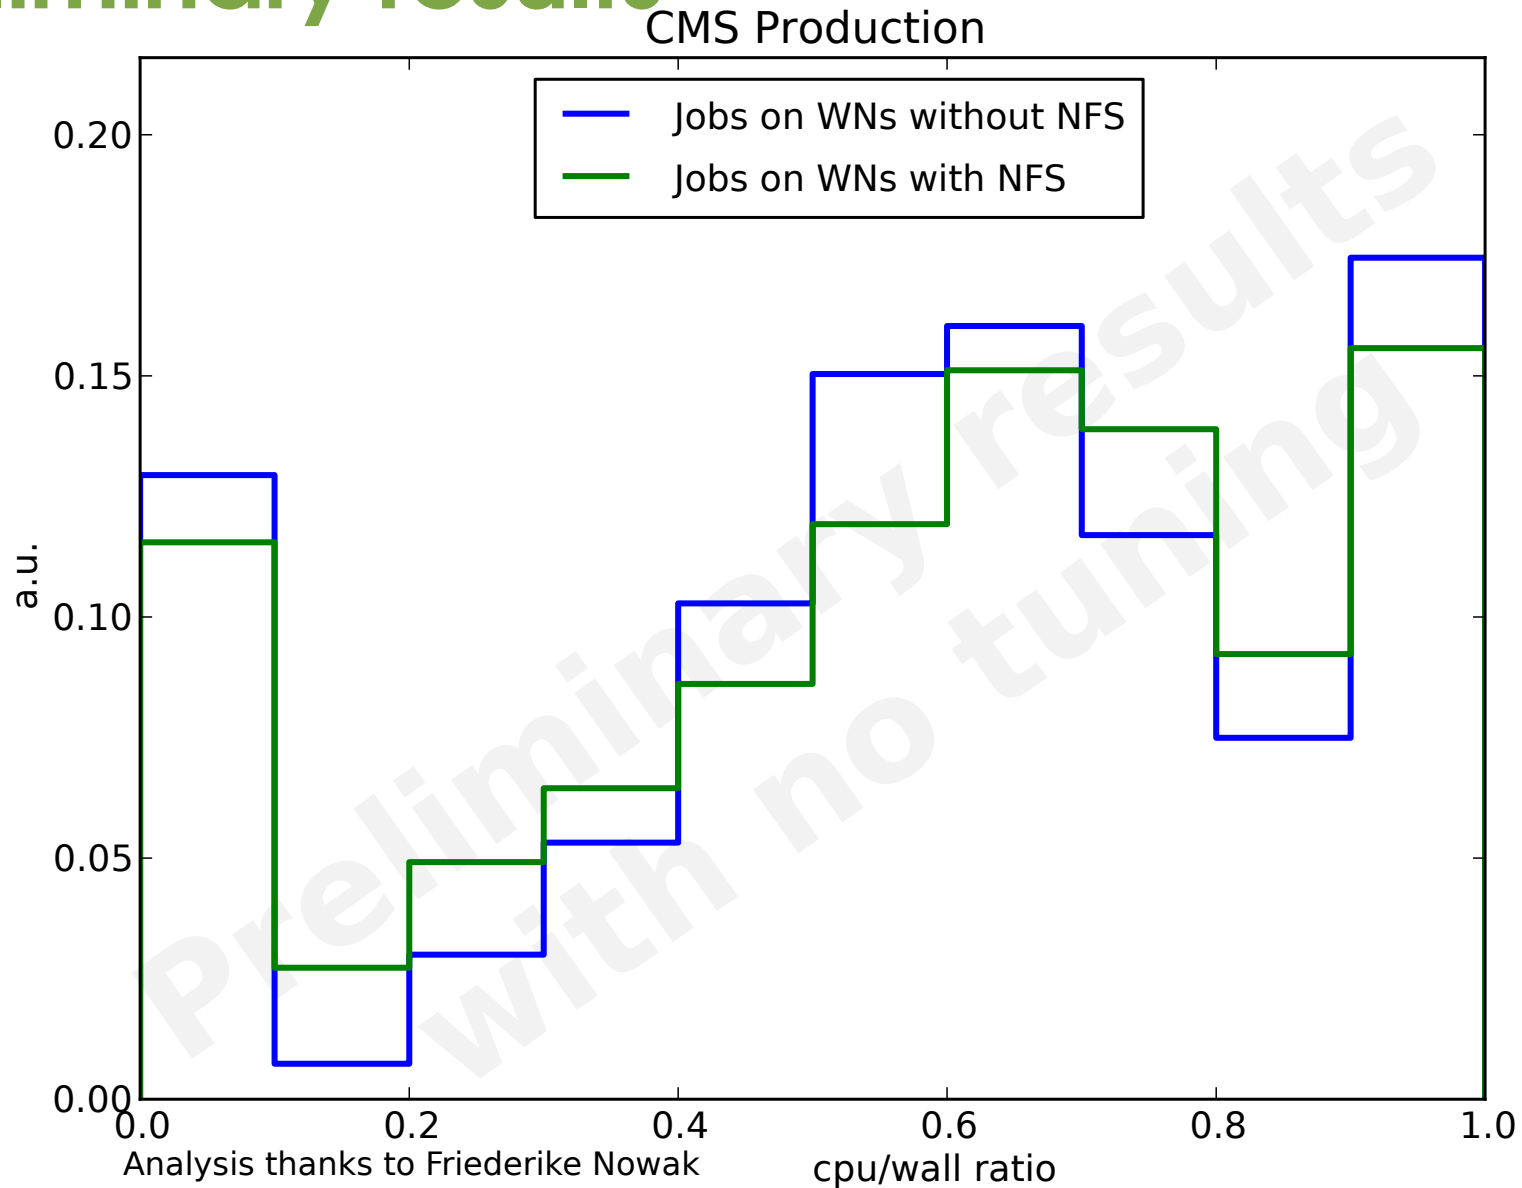
Managing storage

- Currently done with SRM, with some experiments also using WebDAV
- SRM features **in use** but not directly supported by WebDAV:
 - Dealing with tape
 - Space-accounting
 - Discovering access protocols (FTP, NFS, ...),
 - File integrity,
 - Third-party copies (*)
- Many of these can be added by **extending** WebDAV
This is not bad, but loses some advantages of using a standard.

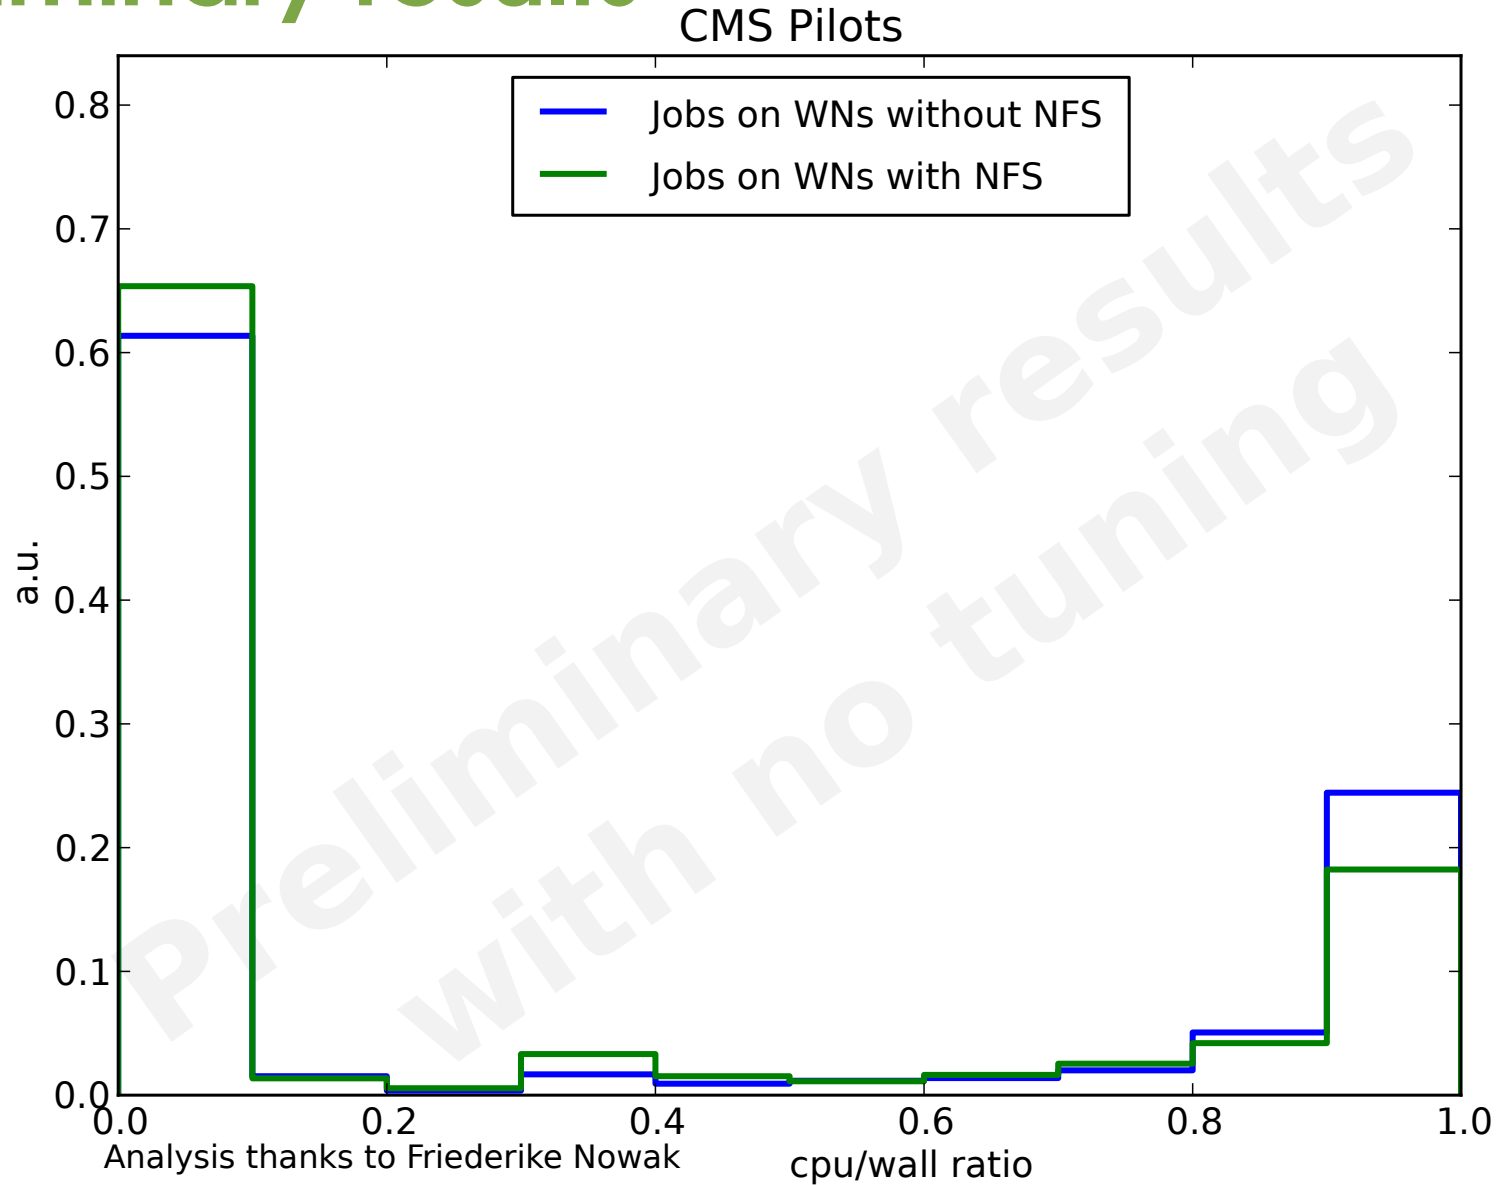
Authentication: getting your X.509 cert



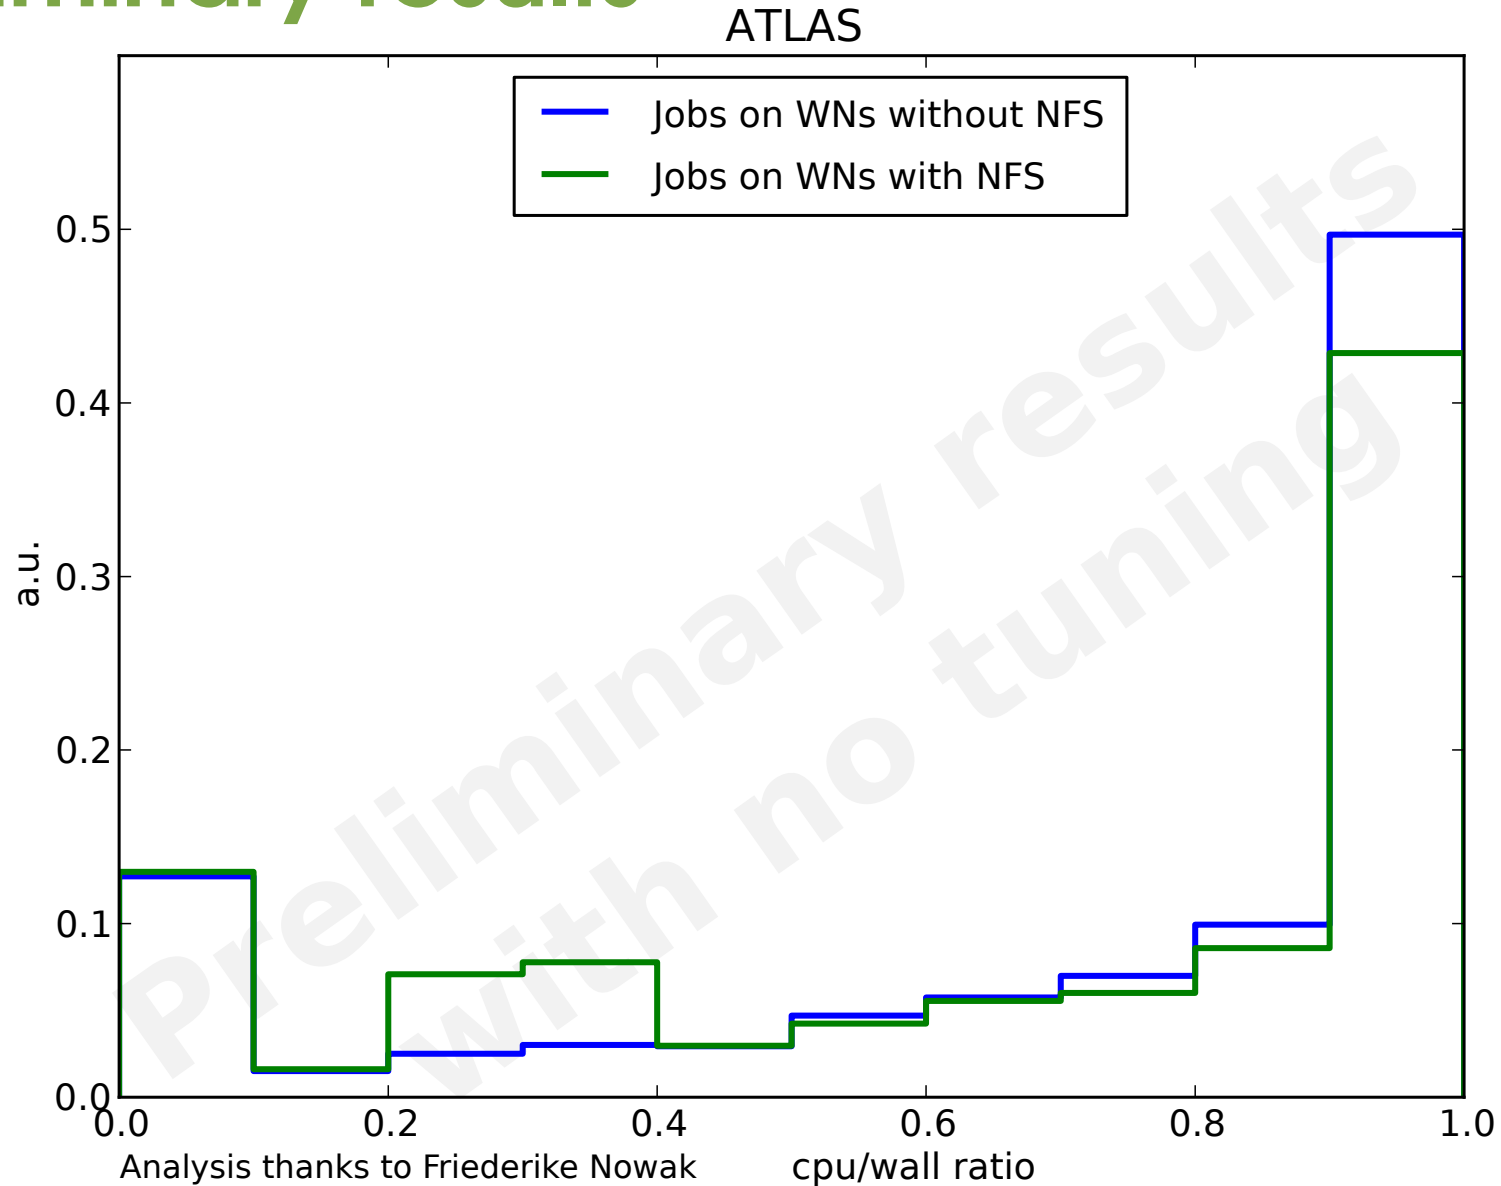
Preliminary results



Preliminary results

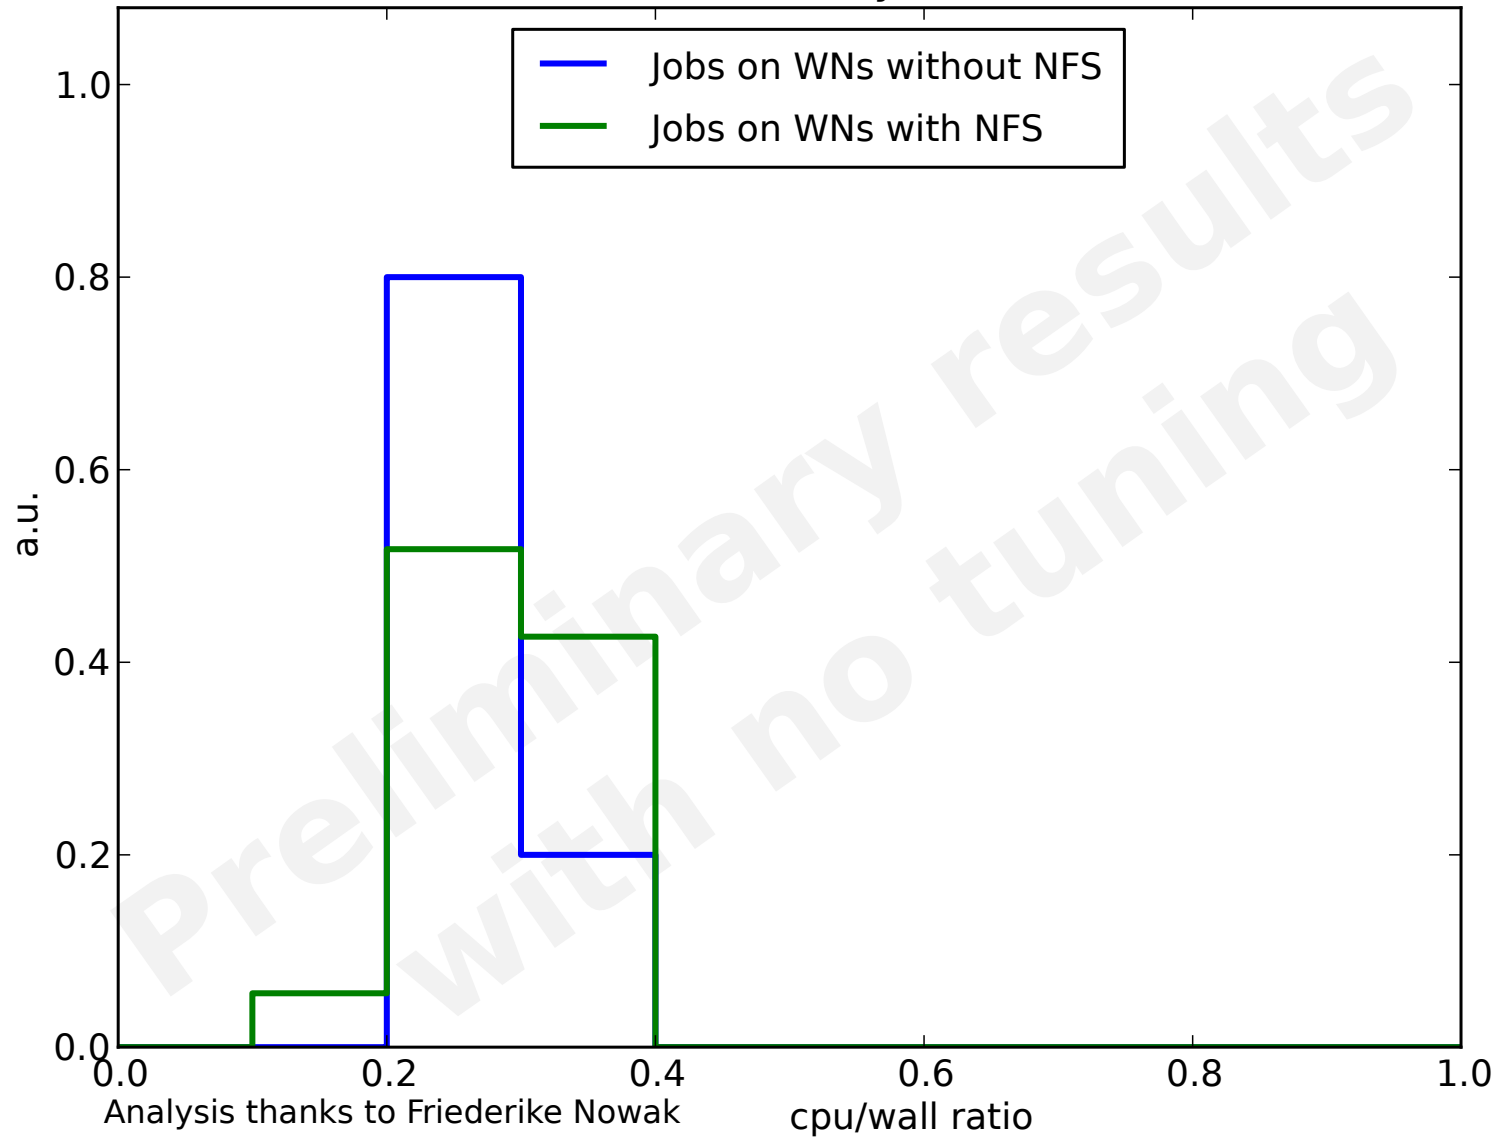


Preliminary results



Preliminary results

ATLAS Analysis



Analysis thanks to Friederike Nowak