

# Prodsys-Rucio integration

TIM2014 Chicago

# Topics to discuss

-----  
10:50 **Prodsys-Rucio integration 1h40'**  
Speaker: Dr. Peter Love (Lancaster University (GB))

-----  
Placement and Replication: RPG 30'

-----  
Data Deletion and obsolescence 30'

-----  
Popularity zoo: Rucio, Panda, xrootd, local access. File vs dataset access 40'

Lunch - outside

These topics have low profile in the zoo of things to do over TIM period of interest. If issues, then collect them here, now.

EVERYTHING IS  
AWESOME!



# Placement and replication

- Transition period for DQ2 GP\_PHYS datasets
- Stephane/David's Migration.py script (RPG)
- Latest was waiting upon AML tags, have these now.  
In meantime replicas not needed now.

# Placement and replication

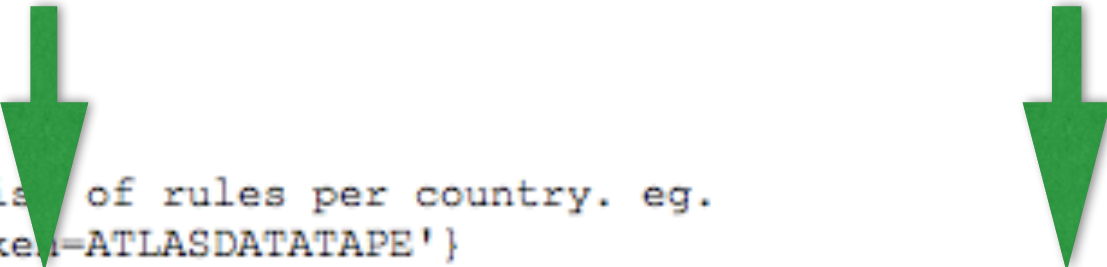
- Rucio subscriptions explored, currently using rucio-client from git. API docs being improved.
- Example below, decisions to be made how to use metadata, rather than deal with DID patterns
- Push this more with latest rucio\_test tasks

```
ddmadmin: Stress_test_t2_scratch UPDATED
  filter: {"project": ["step14"], "scope": ["tests"]}
  rules: [{"weight": "stresstestweight", "rse_expression": "tier=2&spacetoken=ATLASSCRATCHDISK&physgroup=None\rucioetes
tsite=1", "copies": 2, "source_replica_expression": "tier=1|tier=2\\site=CERN-PROD", "activity": "User Subscriptions",
"lifetime": 172800}]
  comments: Stress test from T1 to T2 (part2)
```

# Placement and replication

- Santa Claus
- Rucio subscriptions with weighted rules
- Static subscriptions filtering on metadata

```
RAW project=data12_8TeV
Implement MoU weights per T1 site, requires a list of rules per country. eg.
{'copies': 1, 'rse_expression': 'tier=0&spacetime=ATLASDATATAPE'}
{'copies': 1, 'rse_expression': 'tier=1&country=UK&spacetime=ATLASDATATAPE', 'weight': 0.10}
{'copies': 1, 'rse_expression': 'tier=1&country=FR&spacetime=ATLASDATATAPE', 'weight': 0.15}
{'copies': 1, 'rse_expression': 'tier=1&country=US&spacetime=ATLASDATATAPE', 'weight': 0.25}
...
```



# Placement and replication

- Discussion points
  - Transition period using Migrate.py, drop this given current situation
  - DDM ops tested subscription creation, exposed rucio-cli problem, plenty of JIRA tickets with quick response from DDM dev
  - More stress-testing needed, most likely as we go
  - Which subscriptions need to be in place for start of migration next week?
  - Filter on metadata, avoid patterns where possible

# Deletion

## Deletion in Rucio

- In Rucio, the daemon responsible for deletion is called the reaper.
  - It simply takes the files replicas with a "tombstone" and mark them as being deleted.
  - It sends a deletion request to the Storage Element. For this multiple protocols can be used (SRM, WebDAV...)
  - When the replicas has been physically deleted, it is removed
- With this workflow, the probability of having files still registered in the catalog but not on the Storage (or the other way around) is close to 0.
- The use of other protocols like WebDAV allows to have multiple Storage Cleaner for each endpoints.
- Can also be used for SRMless sites (the so-called gridFTP sites)
- Rucio now send deletion events to the DDM Dashboard.

Cedric - August





## SRMless deletion

- For datasets and files registered in Rucio, deletion with WebDAV was tried.
- Sites already tested :
  - T1s : BNL-OSG2, CERN-PROD (EOS endpoints), FZK-LCG2, IN2P3-CC, NDGF-T1, NIKHEF-ELPROD, PIC, TAIWAN-LCG2.
  - T2s : AGLT2, MWT2-UC.
- Right now deletion only limited to datasets from Rucio stress-tests and deletion of Dark Data.
- As soon as the production system starts writing some new datasets in Rucio (coming soon), the potential number of datasets deletable via WebDAV will grow.





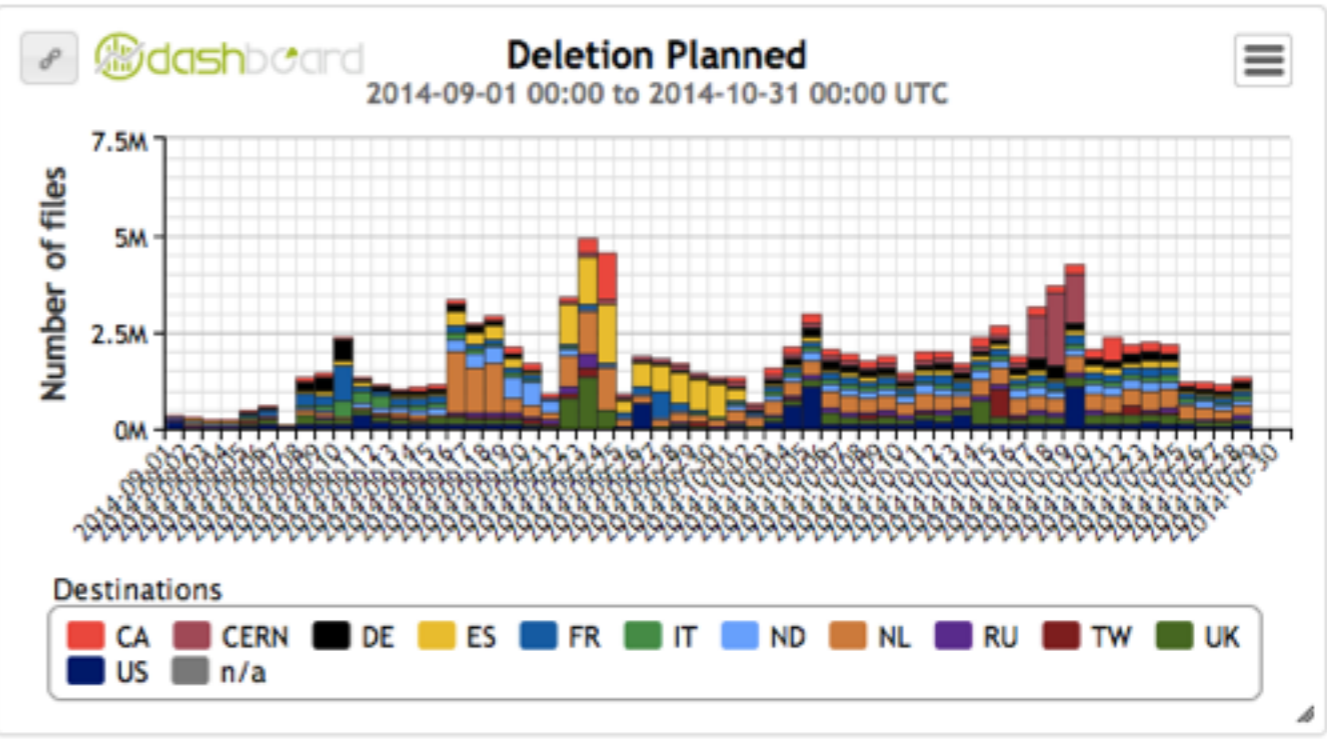
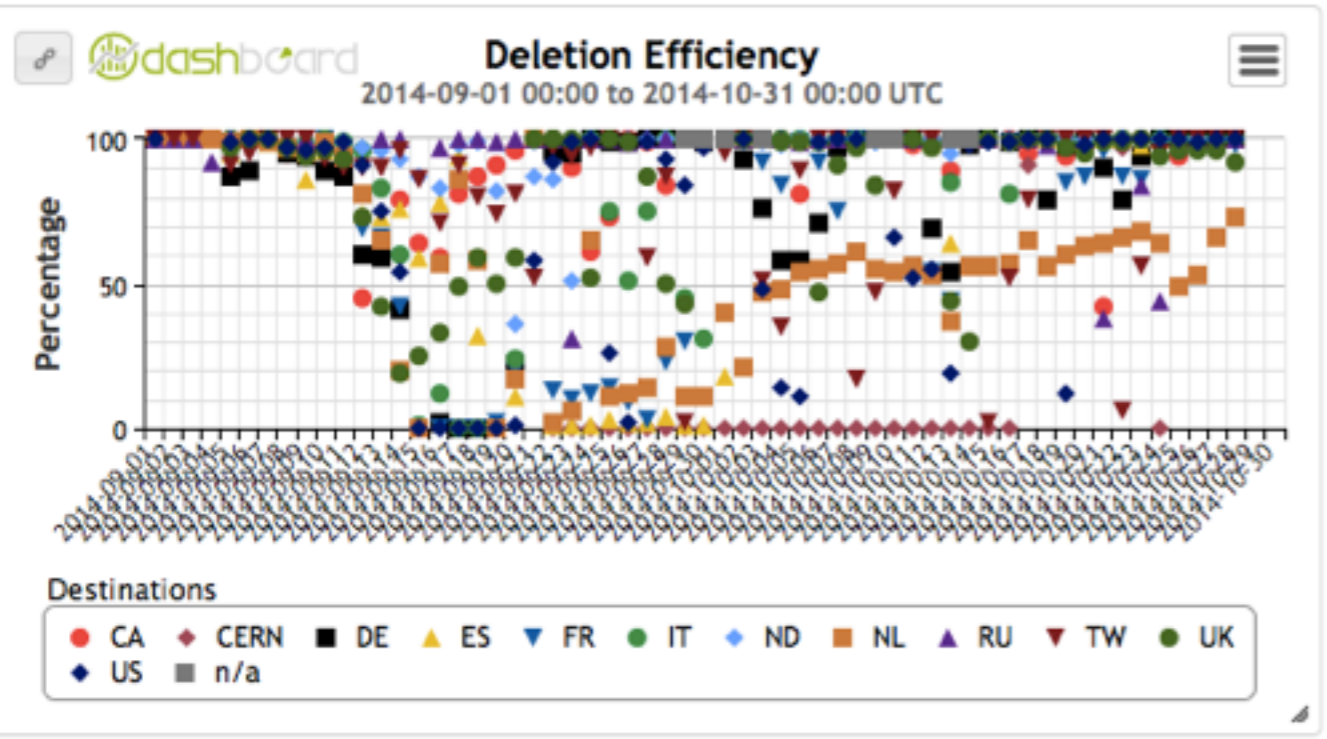
# Deletion - last 48 hrs

Matrix Transfer Plots Staging Plots **Deletion Plots** Centric Plots Details



# Deletion - last 2 months

Matrix Transfer Plots Staging Plots **Deletion Plots** Centric Plots Details



# Obsolescence

- Reprocessing campaigns: update subscriptions with current N, N-1, N-2 (using campaign metadata?)

```
AOD_data12_8TeV*physics_STREAM*merge.AOD*r4065*
STREAM=Egamma, Muons, JetTauEtmiss
---
The proposed 'current' policy:
{'copies': 1, 'rse_expression': 'tier=0&spacetime=ATLASDATATAPE'}
{'copies': 1, 'rse_expression': 'tier=1&spacetime=ATLASDATADISK'}
{'copies': 2, 'rse_expression': 'tier=2&spacetime=ATLASDATADISK'}

The proposed 'N-1' policy:
{'copies': 1, 'rse_expression': 'tier=1&spacetime=ATLASDATADISK'}
{'copies': 2, 'rse_expression': 'tier=2&spacetime=ATLASDATADISK'}

The proposed 'N-2' policy:
delete (considering lifetime)
```

- Group production: workflows in place, need changes? JIRA->prodsys2->reaper
- Abort: identify requests and abort at this level. Prodsys2 handles.
- Users: as discussed yesterday rucio-ui / api

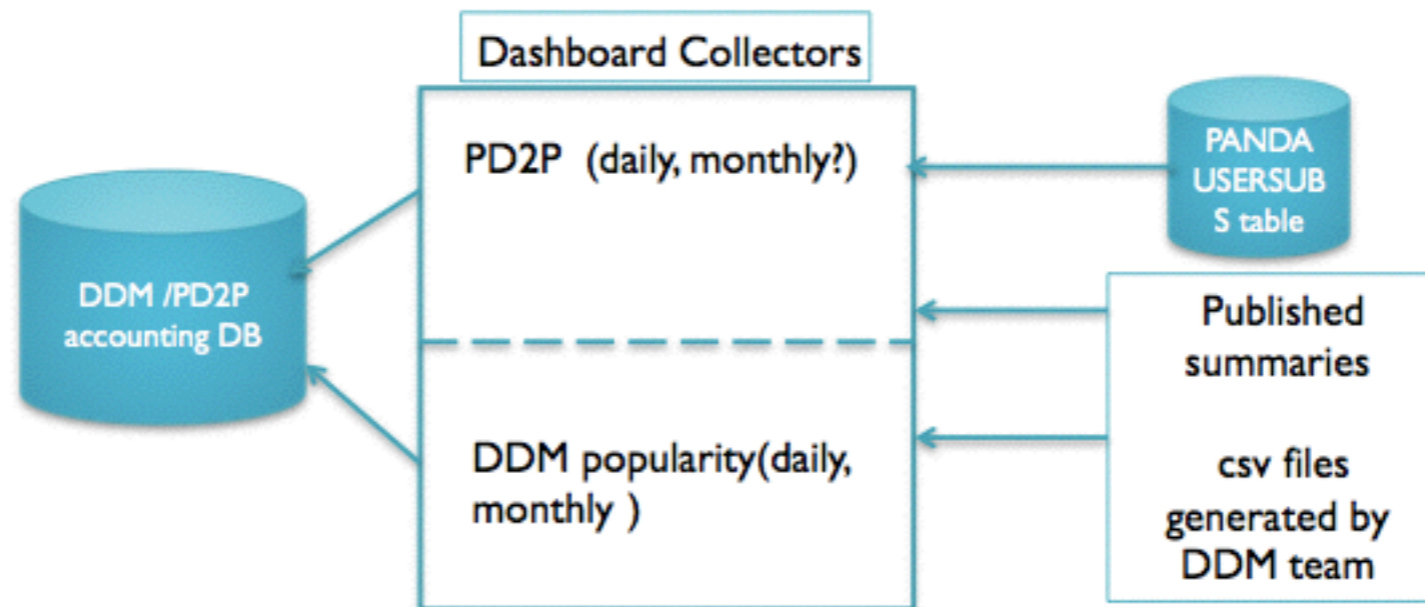
# Deletion/obsolescence

- Discussion points
  - Deletion service looking good. Jog on.
  - Update reprocessing subscriptions - routine part of DDM-ops
  - Obsolescence workflow review (for DPD), via prodsys2.
  - Rucio handling of obsolete DID, reaper

# Popularity

- Collected from: panda, dq2, cli (only grid access)
- Who uses it?
  - ad-hoc reports
  - essential for Disk Residency Priority (DRP) metric
  - useful for sites when recovering disks
- PD2P doesn't, has it's own tables from panda and only user analysis

- PD2P/DDM popularity collectors



Dashboard collectors, web UI had been moved to a puppet managed machine.

# Popularity

- Discussion points (longer term, beyond TIM scope)
  - PD2P production
  - analytics