



dpm-dbck **Swiss Army knife for DPM database**

Andrey Kiryanov
DPM Workshop, CERN

8/12/2015



Rationale

ONCE UPON A TIME I HAD A NEGATIVE VALUE PUBLISHED IN MY BDII...

Firstly blaming dpm-listspaces I realized that the problem was actually in spacetoken usage counters, so I wrote a script that recalculates them.

At the same time I decided to look out for other possible inconsistencies that accumulated in my DPM database over eight years (DPM on my site was installed in 2006 and I have never reinstalled it from scratch ever since)

AND I FOUND PLENTY

Checks

- Detached sub-trees in the namespace that have no parent
- Replica entries with nonexistent filesystems
- LFNs with no replicas
- Files that are on disk but not in the database and vice versa
 - Needs direct access to disk nodes, preferably via *ssh*
- Replica counts and LFNs with missing primary replicas
 - Normally there should be only one replica per LFN. While being technically legitimate secondary replicas are usually leftovers from failed *dpm-drain*.
- Stuck "deleted" LFNs and replicas
- Reference link counts for directories
 - Ever seen *dpns-rm* refusing to remove empty directory?
- Spacetoken storage accounting values
- Stuck entries in DPM request logs

Synopsis

dpm-dbck [-h] [-f] [-y|-n] [-s <n>] [-c <config dir>]

- h Show this help.
- y Answer 'yes' to all questions (auto fix).
- n Answer 'no' to all questions (dry run).
- s Number of parallel filesystem traverse streams.
- c Specify directory with DPMCONFIG and NSCONFIG files.
- f Use sftp instead of rfio for disk node access.

It's strongly advised to use **-f** switch in combination with passwordless *ssh* access to disk nodes. It's faster and more robust than *rfio*.

Future improvements

- An option to opt-out from full filesystem scan
 - The most time-consuming operation
- See what can be done to avoid shutdown of DPM daemons
 - Transaction isolation of individual checks
- You name it

That's all
Thank you!