

DPM Admin Toolkit

A DPM Community Resource

So, what is it?

- Set of useful python scripts that talk to DPM for you.
- Add several generic filesystem-like functions (du, find)
- Also adds several utility functions that don't come with DPM yet.
- And now provides some data consistency checking tools as well.

Tool types.

- Namespace conversion tools
 - dpns <-> filesystem (SURL <-> SFN)
- File listing tools
 - list files owned by user, group, spacetoken, hotness
- Space usage tools
- Integrity checking tools (NEW)

Namespace Conversion tools

- dpm-disk-to-dpns
 - ls of a whole filesystem with equiv. SURLs
- dpm-dpns-to-disk
 - returns the filesystems mapped to a SURL
- dpm-sql-pfn-to-dpns
 - reverse lookup of the above

File listing tools

- `dpm-list-disk`
- `dpm-sql-dpns-by-replication`
- `dpm-sql-files-by-vo-user`
- `dpm-sql-spacetoken-list-files`
- `dpm-sql-list-hotfiles`
- `dpns-find`

Space usage tools

- `dpns-du`
 - very simple “space allocated in DPNS” tool
- `dpm-sql-usage-by-vo-user`
 - nicer formatted breakdown of space usage
- `dpm-sql-spacetoken-usage`
 - nicely formatted spacetoken usage

Miscellany

- “Auto” file replication
 - `dpm-sql-spacetime-replicate-hotfiles`
- Manual dereplication
 - `dpm-delreplica`
- Usage accounting (via cron and `dpm-usage-by-vo-user`).
 - Deprecated in favour of DPM Nagios plugins (which promised this function)

Consistency Checking

- DPNS to/from filesystem integrity checking
 - With optional checksumming (adler32)
- dpm-sql-diskfs-to-dpns-chk
 - checks if files on disk match dpns
- dpm-sql-dpns-to-diskfs-chk
 - checks if files in dpns match disk

Consistency Checking (2)

- Interest in LFC/SE consistency checking.
- Whole service consistency “background scan”
 - for VOs too small to do their own checks
 - Need embedded checksums for all files.
- There are questions of what level this should be integrated at (filesystem level, namespace level?)

In Progress

- Consolidation phase for older tools.
 - All of slide 5 could be one tool:
 - dpm-sql-ls
 - Tool that “lists space” with selection options
 - dpm-sql-du
- Beta versions of tools are in the svn.

Consolidated tools v GFAL v HTTP/Dav Tools

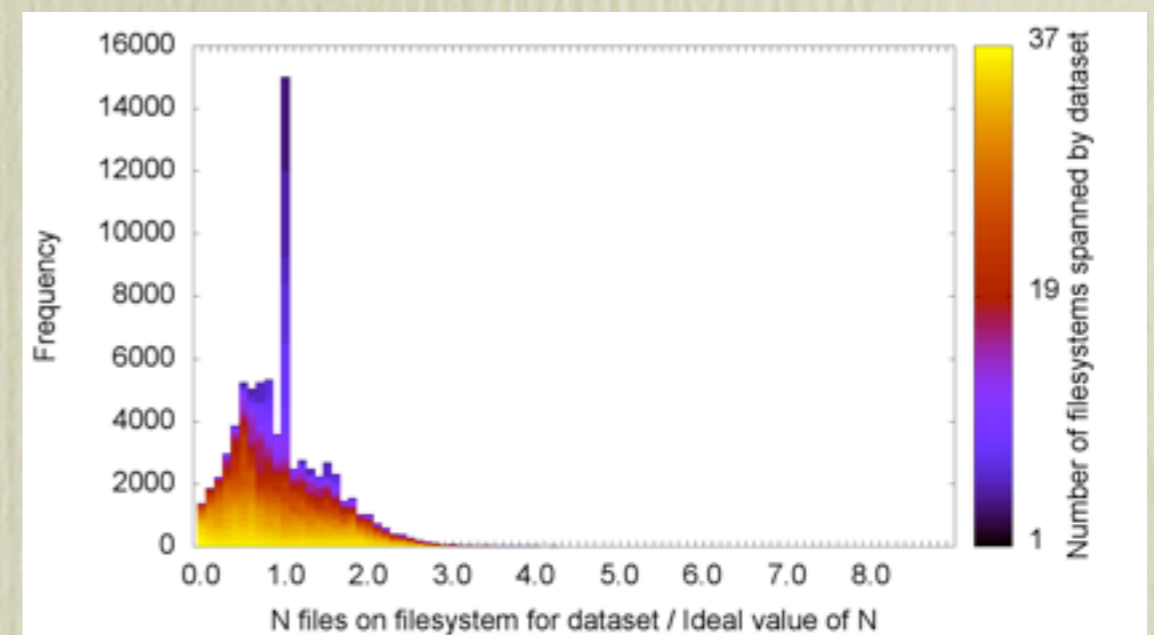
- A question for the audience:
 - Is there a need for the tools in the previous slide, given the availability of POSIXish mounting options for DPM now?

Filesystem Rebalancing

- DPM currently can't cope with overloaded filesystems. (Situation known for past year.)
- Can ameliorate by rebalancing file distribution.
 - Analysis of filesystems shows some datasets v poorly distributed (10:1 ratios in worst case)
- Tool in progress, see CHEP Poster / Paper.

Filesystem Balancing (II)

- Only enabled for DPM-native filesystems
 - (HDFS, S3 etc can manage their own file distribution)
- Waiting on some api changes in WebDAV support for direct replica movement.



DMLite API

- DMLite exposes a new Python API.
- Significantly more power than the old, poorly documented, SWIG API.
- Currently in process of understanding the new API, and considering migrating non-deprecated tools to it.

Aside: LFC/Metadata Tools?

- Not strictly a DPM tools issue
- Imminent retirement of lcg-utils in favour of gfal-tools
 - Reduced LFC / metadata catalogue support
 - Non-LHC VOs case?

Aside: RAIN/RAIS implementations

- Also somewhat out of scope of DPM admin tools
- Space-efficient resiliency implies use of Reed-Solomon codes (or other erasure codes).
- This can be implemented at the LFC level easily with a simple wrapper.
- (Test version exists, based on Tahoe-LAFS)

RAIN/S in DPNS?

- As DPNS and LFC are very similar, would also be possible to apply Reed-Solomon resiliency to a DPM.
- Would need extensive work on transfer tools (need to understand file block maps).

Acknowledgements.

- The DPM admin toolkit has benefitted from contributions from many people:
 - Grieg Cowan (originator)
 - Wahid Bhimji (hot files contributor)
 - Alessandra Forti (improvements to user/group selection code)
 - Ricardo Rocha + DPM developers in general

Contact

- The DPM toolkit is hosted by the DPM svn repository, and is available from the community repo.
- Contact:
 - sskipsey [at] gmail.com
 - GridPP-Storage and DPM Devel lists
- Documentation:
 - <http://www.gridpp.ac.uk/wiki/DPM-admin-tools>