# DPM: Configuration and Testing

A. Sartirana

# Introduction…

- **DPM Workshop: 23-24 Nov. 2016** LPNHE, Paris
  - https://indico.cern.ch/event/559673;
  - report at LCG-FR Tech on 16/12/2016
    - https://indico.in2p3.fr/event/13392/contribution/4/material/slides/0.pdf;
- Thu **24th morning, focus** on
  - **configuring** DPM: puppet/manual config.
  - **testing** DPM: functionality and performance;
  - **clinic**: checking/fixing DB consistency;
- only very **quickly mentioned** in the 16/12 report as "Other Business".

# …Introduction.

- this is just a **(slightly) more detailed report** about these topics
  - « usual » questions:
    - how do I **configure** my DPM instance?
    - how do I check if it is **working/performing well**?
    - how do I **monitor** this (and detect problems);
    - how do I **fix inconsistences/dark data**/etc…;
  - **related talks** at the workshop + some **personal reflection** (and examples from LLR);
- to start with: **documentation pages**
  - https://twiki.cern.ch/twiki/bin/view/DPM/;
  - still under construction.

# Configuration…

- General **presentation on DPM setup by A.Manzi**
  - https://indico.cern.ch/event/559673/contributions/2277499/attachments/1372787/2085924/DPM_Setup.pdf

- **puppet: official way** for installing and configuring DPM
  - **easy to use and well documented**
    - https://twiki.cern.ch/twiki/bin/view/DPM/DpmSetupPuppetInstallation ;
  - **active support** from the developers;
  - puppet 4 and CentOS7 supported;
  - compliant with passing parameters via Hiera;
  - available to quattor sites via ncm-puppet.

# Configuration…

- Very handy using **lcgdm-dpm module:**
  - https://forge.puppet.com/lcgdm/dpm;
  - "meta-module": **manages deps** from other modules;
  - wraps up the configuration;
  - soon also in **rpm format**
    - solution tested at GRIF;
    - « yum install puppet-module-lcgdm-dpm ».

```
class{"dpm::headnode":
  configure_vos        => true,
  configure_gridmap    => true,
  disk_nodes           => "dpm-puppet02.cern.ch",
  localdomain          => "cern.ch",
  webdav_enabled       => true,
  ...
}
```

# Configuration…

```
[root@llrpp01 ~]# cat /etc/puppet/manifests/quattor_default.pp
hiera_include('classes')


[root@llrpp01 ~]# cat /etc/puppet/hieradata/quattor.yaml
---
classes: dpm::headnode
dpm::params::db_pass: xxxxxxxx
dpm::params::disk_nodes:
- llrpp02.in2p3.fr
- llrpp03.in2p3.fr
dpm::params::dpm_xrootd_fedredirs:
  cms:
…
dpm::params::headnode_fqdn: llrpp01.in2p3.fr
dpm::params::localdomain: in2p3.fr
dpm::params::memcached_enabled: 1
…
```

- Supports hiera:
  - hiera (+ rpms): yaim-like;
  - compliant with foreman management;
  - this is the way ncm-puppet workd;
  - no need to actually deal with puppet.

# Configuration…

- **manual configuration** documentation is still under construction:
  - https://twiki.cern.ch/twiki/bin/view/DPM/DpmSetupManualInstallation
- **"legacy" Quattor** configuration still maintained:
  - actually ~all quattor sites are still using it.

# …Configuration.

- Besides DPM services it is also important to **correctly configure** the underlying **OS**
  - **https://twiki.cern.ch/twiki/bin/view/DPM/DpmSetupTuningHints**
- **hints** for configuring a **performing system**
  - **HW** dimensioning;
  - **file descriptors** and system **limits**;
  - **mysql** tuning;
  - **memcache** setup;
  - dpns/dpm/srm **threads limits**;
  - xrootd **log levels**;
  - **apache** config.

# Testing DPM…

- Once configured my DPM (or upgraded) **how can I test** that it works correctly?
  - **dpm-tester.py** tool for functional (and perf) tests on a DPM instance;
  - presentation/docs
    - https://indico.cern.ch/event/559673/contributions/2277498/attachments/1376982/2091297/gbitzes-dpm-workshop-2016-11-24.pdf;
    - https://twiki.cern.ch/twiki/bin/view/DPM/DpmAdminDpmTester (I know... still empty ☺ );
  - 'yum install dmlite-dpm-tester'
    - requires gfal2 with all plugins (xrootd, etc…), which should be standard now;

# …Testing DPM…

- Sets of **functional tests** for protocols
  - gsiftp, davs, root, srm;
  - create, remove (also recursive) files/directories;
  - upload/download files;
  - checksum compute (davs and ftp);
- also (sort of) **performance tests**
  - upload N files in parallel;
  - create, remove file N time;
- test combining different protocols.

# …Testing DPM.

```
[phedex@llrphedex ~]$ dpm-tester.py --host llrpp01.in2p3.fr --path /dpm/in2p3.fr/home/cms
--hammer-parallel-uploads 5 --upload-delete-loop 2 --tests gsiftp root
PASS => gsiftp :: Verify base exists: /dpm/in2p3.fr/home/cms/                    0.41495 sec
PASS => gsiftp :: Verify testdir does not exist: /dpm/in2p3.fr/home/cms/dpm-tests 0.21427
sec
PASS => gsiftp :: Create testdir: /dpm/in2p3.fr/home/cms/dpm-tests               0.11425 sec
PASS => gsiftp :: Upload to testdir: services                                    0.71556 sec
PASS => gsiftp :: Download from testdir: services                                1.01609 sec
PASS => gsiftp :: Verify downloaded contents are identical                       0.00146 sec
PASS => gsiftp :: Verify size is 640978                                          0.06410 sec
PASS => gsiftp :: Verify md5 checksum: f86cd03ad51a92c286202876fd81630d          5.32504 sec
PASS => gsiftp :: Remove: services                                               0.61531 sec
PASS => gsiftp :: Upload to testdir: evil filename-!@#%^_-+=:][}{><'" #$&*)(     6.22798 sec
PASS => gsiftp :: Download from testdir: evil filename-!@#%^_-+=:][}{><'" #$&*)( 0.91602 s
ec
PASS => gsiftp :: Verify downloaded contents are identical                       0.00131 sec
PASS => gsiftp :: Verify size is 640978                                          0.06437 sec
PASS => gsiftp :: Verify md5 checksum: f86cd03ad51a92c286202876fd81630d          0.56503 sec
PASS => gsiftp :: Remove: evil filename-!@#%^_-+=:][}{><'" #$&*)(                4.72404 sec
.... => gsiftp :: Hammer test - upload 5 files in parallel: /dpm/in2p3.fr/home/cms/dpm-tes
PASS => gsiftp :: Hammer test - upload 5 files in parallel: /dpm/in2p3.fr/home/cms/dpm-tes
ts 5.08257 sec
PASS => gsiftp :: Upload and delete the same file 2 times                       12.58767 sec
PASS => gsiftp :: Recursively remove contents: /dpm/in2p3.fr/home/cms/dpm-tests2.96455 sec
PASS => gsiftp :: Remove directory: /dpm/in2p3.fr/home/cms/dpm-tests             0.06434 sec
PASS => root :: Verify base exists: /dpm/in2p3.fr/home/cms/                      0.11452 sec
PASS => root :: Verify testdir does not exist: /dpm/in2p3.fr/home/cms/dpm-tests0.03196 sec
PASS => root :: Create testdir: /dpm/in2p3.fr/home/cms/dpm-tests                 0.11433 sec
PASS => root :: Upload to testdir: services                                      0.66544 sec
PASS => root :: Download from testdir: services                                  0.06427 sec
PASS => root :: Verify downloaded contents are identical                         0.00154 sec
PASS => root :: Verify size is 640978                                            0.01603 sec
PASS => root :: Remove: services                                                 0.61552 sec
PASS => root :: Upload to testdir: evil filename-!@#%^_-+=:][}{><'" #$&*)(       2.42048 sec
PASS => root :: Download from testdir: evil filename-!@#%^_-+=:][}{><'" #$&*)(   0.96627 sec
PASS => root :: Verify downloaded contents are identical                         0.00148 sec
PASS => root :: Verify size is 640978                                            0.03235 sec
PASS => root :: Remove: evil filename-!@#%^_-+=:][}{><'" #$&*)(                  7.52924 sec
.... => root :: Hammer test - upload 5 files in parallel: /dpm/in2p3.fr/home/cms/dpm-tests
```

# Monitoring.

- **Detect problems** or **critical situations**?
  - **VOs** functional **tests** are a good reference;
  - **need** monitoring for **site-level issues**;
- example of probes (check_mk LLR)
  - **services**: gsiftp/srm/xrootd/dpm/dpns/mysqld…;
  - gsiftp **# procs**/xrootd **# threads**;
  - dpm filesystems/pools **occupation**;
  - **RAIDSETs health** and DS **fs occupation**;
  - **dpm-tester**, can also be made a probe;
- **tricky** part: thresholds and **pbs signatures**
  - IMHO this is quite **site-specific**.

# Monitoring.

- What about **performance monitoring**?
  - personal view: quite **tricky** at site level on a production system (useless?);
  - again **VOs monitoring** can be very helpful
    - **instant** perf monitoring: e.g. PhEDEx;
    - **a posteriori** perf analysis: e.g. jobs efficiency;
  - at **site-level**, see if infrastructure are exploited in a coherent way e.g. **network** usage **vs** diskservers **load**;
  - for other things – e.g. HN ops rate – no idea…

# DB Checking.

- At the workshop were also presented a **couple of tools for checking and fixing the DB** consistency: dpm-cleaner, dpm-dbck;

- **dpm-cleaner:**
  - by **Alex Mikula;**
  - https://github.com/samuraiii/dpm-cleaner
  - https://indico.cern.ch/event/559673/contributions/2348409/attachments/1376054/2089571/dpm-cleaner.pdf
  - works on a **running instance;**
  - focused on **re-aligning DB and filesystems.**

# DB Checking.

- **dpm-dbck:**
  - by **Andrey Kiryanov;**
  - yum install dpm-contrib-admintools
  - **needs a downtime**
  - checks db against **inconsistencies**
    - detached subtrees;
    - replicas on non-existent fs;
    - LFN's with no (primary) replicas;
    - stuck deleted LFN and replicas;
    - reference link counts for dirs;
    - spacetoken storage accounting;
    - stuck entries in requests logs;
  - also **synch DB and filesystems.**

# Summing up.

- Last workshop. Last day: **focus on configuration, testing, DB checking**
  - here are few **info/pointers/remarks;**
- **configuration:** official tool is **puppet**
  - lcgdm-dpm metamod. Hiera support. Soon RPMs;
  - **manual conf** and **legacy quattor still maintained;**
- **testing** and **monitoring**
  - **dpm-tester.py** useful and easy to use;
  - **functionality monitoring:** VO mon, Nagios, …;
- tools for DB check: **dpm-cleaner, dpm-dbck.**

# Questions?