



# DPM Configuration

**Andrea Manzi**  
CERN

DPM Workshop

08/12/2015

# Outline

- Puppet configuration
  - Status
  - Fixed problems
  - What's new
- Ansible
- Containers ( Docker)
- Hands on

# DPM Configuration evolution

(as presented during the previous Workshops)



# YAIM



Sep '13

Jan '14

Sep '14



Cont. puppet module development

# DPM Configuration evolution



Sep '14

Now



Enhancements and Bug fixing

# Puppet modules

- “Our” modules:  
dmlite, gridftp, lcgdm, xrootd
- Other CERN modules used:  
bdii, fetch-crl, voms
- External modules:  
firewall, mysql, stdlib, memcached, limits

# Contributions to the modules

- Our modules are available in github
  - <https://github.com/cern-it-sdc-id>
- Got some enhancements and bug fixes during the year from the site managers
  - THANKS!!
- Part of HEP-Puppet organization:
  - <https://github.com/HEP-Puppet>

# Twiki configuration page

[Wiki](#) | [Timeline](#) | [Roadmap](#) | [Browse Source](#) | [View Tickets](#) | [New Ticket](#) | [Search](#) | [Rev Tree](#) | [Blog](#) | [Admin](#)

[Dpm](#) / [Admin](#) / [InstallationConfigurationPuppetSimple](#)

[Up](#) | [Start Page](#) | [Index](#) | [History](#) | [Last Change](#)

## DPM (Disk Pool Manager)

The Disk Pool Manager (DPM) is a lightweight solution for disk storage management. It is part of LCGDM ( <https://svnweb.cern.ch/trac/lcgdm> )

- DPM home ( <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm> )
- Support: [✉ DPM Users Forum](#), [✉ GGUS Helpdesk](#), [✉ Dev List](#)
- User Docs: ( [Tutorial](#), [Cookbook / FAQ?](#), [API Overview](#), [Command Line Tools](#) )
- Admin Docs: ( [Installation](#), [Configuration](#), [Maintenance](#), [Performance](#), [Monitoring](#), [Reference Card](#), [EMI Reference Card](#) )
- Development: ( [Savannah](#), [Certification](#), [Release Procedure](#), [Internal Testbed](#), [Components](#), [Coordination](#), [Recipes](#) )

## Installation and Configuration Guide

This guide explains the installation and configuration procedure of DPM with puppet. It will introduce the puppet components to be installed, but not give an introduction into writing puppet manifests. For that, you can look at the fantastic documentation at [puppetlabs](#). Please also have a look at this, if you want to find out the difference between modules and manifests. This guide also applies to other components, e.g. the LFC, for which puppet modules exist.

If you want to install DPM without using puppet, the following pages describe the configuration manually or with YAIM:

- [Installation Guide](#)
- [Configuration with YAIM](#)
- [Manual Configuration](#)

If you want to know more details about the puppet modules, follow this tutorial. It shows an alternative way to use the puppet modules showing exposing more configuration parameters and modules:

- [Detailed Puppet Installation Guide](#)

## Open Issues

- **GridFTP** The puppet installation instructions are valid for the gridftp frontend based on dmlite, i.e. dpm-dsi >= 1.9.0. To configure the DPM-based gridftp frontend, please use the yaim configuration.

## Overview

Puppet is a configuration management system that can not only preserve and update the configuration of your components, but also manage the installation. That is why this guide is an integrated installation/configuration guide. As puppet modules are developed by a larger community, you can also manage other parts of your system, e.g. the yum configuration or the firewall settings, with puppet. In the following, we describe the settings and how they can be applied with puppet.

## The Sections

<http://lcgdm.web.cern.ch/dpm-configuration>





# Releases

- Several releases during the last year:
  - 11 puppet-dmlite releases
  - 7 puppet-lcgdm releases
  - 1 puppet-xrootd
    - to configure xrootd-redirectors in US
- Tried to extend the documentation as much as possible to cover all possible installation scenarios

# Examples

- <https://github.com/cern-it-sdc-id/lcgdm-puppet-examples>
- Manifests:
  - DPM-Head
  - DPM-Disk
  - DPM-Head+ DPM-Disk
  - LFC
  - repos
- new installation VS upgrade
- GridFTP redirection **NEW!**
- Space reporting **NEW!**

# FIXED issues: service restarts

- In some cases rpms upgrades/conf changes did not trigger DPM services and frontends restarts even after a puppet run
  - Yaim was configured to always restart the service at each execution
- In order to fix we have added the missing service restarts upon file config changes + more dependencies between classes
  - E.g if mysql plugin conf file changes, all frontends are restarted
- Need also to add some restarts at packaging level that was missing
  - Dmlite update triggers frontends restart

# FIXED issues: external DB

- Failed installation in case of Head node with an External DB (started when moving to puppet-mysql > 3.x)
  - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/InstallationConfigurationPuppetSimple#ExternalInstallation>
- We now configure the HEAD node with a ~/.my.cnf conf file with the mysql root pass

```
class{"lcgdm":  
    ..  
    dbhost => "<db_host_name>",  
    mysqlrootpass => "${mysql_root_pass}", # new in lcgdm-lcgdm 0.2.9  
    ...  
}
```

# FIXED issues: MySQL 5.1 on SL5

- In case of SL5 installation the default MySQL version installed using the puppetlabs-mysql is the 5.0 for which the security support has ended.
  - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/InstallationConfigurationPuppetSimple#SL5installation>
- For this reason a small configuration change is requested on the manifest to enable the installation and configuration of Mysql 5.1

```
class{"mysql::server":  
  ...  
  package_name => "mysql51",  
  service_name => "mysql51-mysqld",  
  ...  
}
```

# FIXED issues: Memcached access

- SVG Advisory sent to sites because the default configuration for Memcached was vulnerable
  - [https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/InstallationConfigurationPuppetS\\_imple#Installationandconfigurationofthedmlite-memcacheplugin](https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/InstallationConfigurationPuppetS_imple#Installationandconfigurationofthedmlite-memcacheplugin)

```
class{"memcached":  
  max_memory => <MaxMem>,  
  listen_ip => "0.0.0.0", ← NO (everyone can read if the firewall is open)  
}
```

Fixed with:

```
class{"memcached":  
  max_memory => <MaxMem>,  
  listen_ip => "127.0.0.1", ← only local  
}
```

# What's new

- Tuning Hints
- Xrootd4
- Gridftp redirection
- Space reporting
- HDFS plugin configuration
- CentOS7
- puppet-dpm module

# Tuning Hints in puppet

- Tuning Hints page prepared by Fabrizio in order to speed up DPM performances:
  - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/TuningHints>
- All tuning hints are now available via puppet as well
  - The suggested parameters are default in puppet installations



# Xrootd4 support

- Starting from puppet-dmlite 0.3.0 module the modification for the new Xrootd4 are the default ones
  - Vomsxrd 0.3.0 installation
  - Different type of log rotating
  - Small changes in federations parameters

# GridFTP redirection

- It's now possible to enable GridFTP redirection via puppet:
  - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/GridFTP#GridFTPRedirectionpuppetconfiguration>
- Manifest Templates are available with the specific configuration

Head:

```
class{"dmlite::gridftp":  
  dpmhost => "${::fqdn}",  
  remote_nodes => '<disknode>:2811',  
}
```

```
lcgdm::shift::protocol_head{"GRIDFTP":  
  component => "DPM",  
  protohead => "FTPHEAD", Tells SRM to return gridftp TURLs pointing to the HEAD  
  host => "${::fqdn}",  
}
```

Disk:

```
class{"dmlite::gridftp":  
  dpmhost => "${headnode_fqdn}",  
  data_node => 1,  
}
```

# Space reporting

- New in DPM 1.8.10
  - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/TopdirSpaceReporting#Puppetsetup>
- Puppet modules and manifests have been modified in order to:
  - configure the MySQL DB to accept connections from the disknodes ( so firewall rules, mysql conf and user grants)
  - install and configure the dmlite-mysql-plugin on the disknodes

# Space reporting [2]

- New Manifest templates with the changes

- Head + DB

```
class{"dmlite::head":
```

```
...
```

```
enable_space_reporting => true,
```

```
}
```

- Grants, open DB remote connections

# Space reporting [3]

- Disk

```
class{"dmlite::disk":  
  ...  
  enable_space_reporting => true,  
}
```

- 3 new vars to configure MySQL parameters ( db\_user, db\_pass, db\_host)

```
class{"dmlite::disk":  
  ....  
  mysql_username => "${db_user}",  
  mysql_password => "${db_pass}",  
  mysql_host     => "${db_host}",  
}
```

TODO -> same configuration for HDFS

# HDFS plugin configuration

- New Manifest and extensions to the module have been implemented for an DPM + HDFS deployment
  - HEAD node
    - HDFSplugin configuration
    - Info provider conf changes to publish HDFS pool
    - Gridftp redirection
  - HDFS Gateway node ( ~disk node)
    - HDFS plugin configuration
    - Frontends specific conf
      - No RFN check in gridftp + redirection
      - Xrootd use dmlite I/O

# CentOS7 support

- Some changes needed in order to add support for CentOS7
  - Enable DPM legacy daemon systemd unit file
  - Apache 2.2.4 conf
- Nothing to do DB side as MariaDB is transparently installed and configured if CentOS7 is detected by the puppet-mysql module

# Puppet-dpm

- Puppet-dpm is a “meta-module” + wrapper to existing developed module
- developed by A. Sartirana with small modifications on our side
  - Includes all other module dependencies
  - Hides the modules inter-dependencies
  - Parameters configurable via hiera



# Puppet-dpm

```
class{"dpm::disknode":  
  configure_vos      => false,  
  configure_gridmap  => true,  
  
  #cluster options  
  headnode_fqdn      => "${::fqdn} ",  
  disk_nodes         => "dpm-puppet02.cern.ch",  
  localdomain        => "cern.ch",  
  webdav_enabled     => true,  
}
```

- Same for headnode and head+disknode
- <https://github.com/cern-it-sdc-id/puppet-dpm>

# Puppet-dpm

- Pros
  - Resolution of the modules deps
  - Parameters can be specified like in yaim
- Cons
  - wrapping of all dmlite/lcgdm parameters needed if hiera is not used..
- To use in case of testing installations
  - See demo
- In production it's easier to customize the existing manifests to your needs

# Future work around puppet

- Puppet 4
  - Test compatibility of our modules
- Enhancements
  - Some changes for dpm-xrootd federations configuration are foreseen
- Follow the DPM development directions
  - DPM REST configuration
  - New plugins configurations

# Ansible

- Some site managers ( in particular Raul Cardoso from Brunel) are pushing for a Ansible configuration for DPM
- Ansible is now under RedHat
  - <http://www.ansible.com/blog/red-hat>
- ..but we have no experience and effort at CERN to work on it

# Ansible

- If Raul and others would like to work on this we are more than happy 😊
- but we need to complete some work on our side in order for people to start..
  - Manual installation guide is obsolete for HEAD nodes
  - Disk nodes guide is fine
- Anyone interested just have to contact us

# Docker

- We started to play with it ( as everyone 😊)
- The first ( and only ) test done is with a docker container with puppet installed, following then the DPM puppet configuration
  - <https://hub.docker.com/r/amanzi/>
- `docker run -h head.cern.ch --name=head --privileged --dns-search=cern.ch --dns=137.138.16.5 -p=80:80 -p=443:443 -i -t amanzi/puppet-dpm-head:1.8.10`

# Docker

- I personally had some issues with Networking as we need intercommunication btw Head and Disk containers
- Didn't spend so much time to solve it
- Plan to invest some more time in the future
  - Docker DPM for testing purposes
- Help is welcome 😊

# Hands on

# Questions?





# Steps

- Set up simple DPM Head + Disk
- VM already configured with
  - Certificates
  - Repos
  - Puppet 3.8
- puppet module install lcgdm-dpm

# CERN Puppet infra

- Foreman
  - <https://judy.cern.ch/hosts>
- Several Load balanced Puppet Masters
- Module and manifests in separate git repos
  - <https://git.cern.ch/index.html>
- Tbag for secrets

# Our testbeds

- Open to internet
- Trunk
  - <https://dpmhead-trunk.cern.ch/dpm/cern.ch/>
  - Gridftp redirection + space reporting
- RC
  - <https://dpmhead-rc.cern.ch/dpm/cern.ch/>