

DPM in ITALY

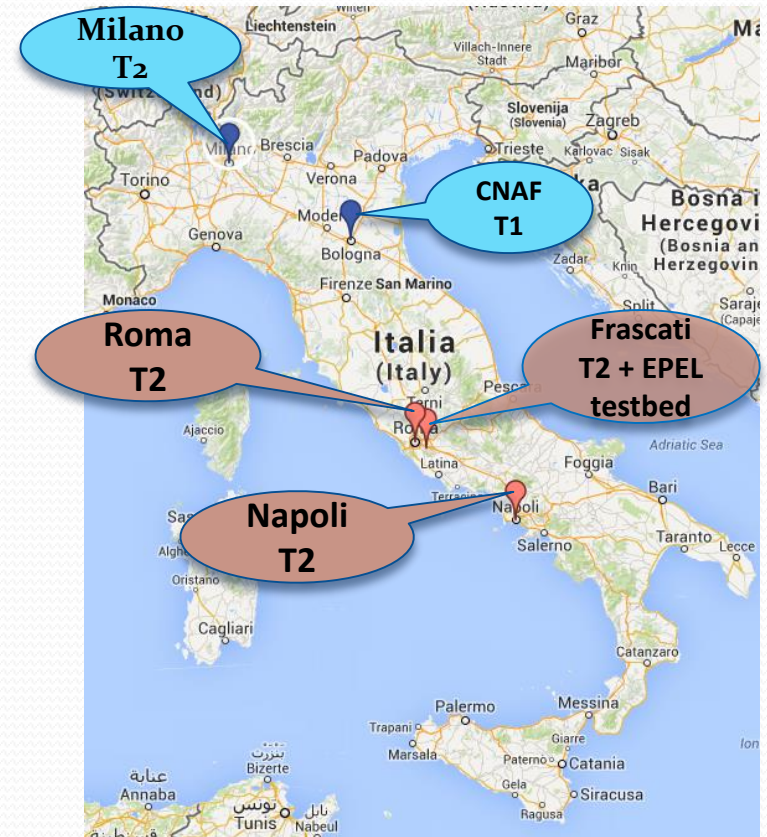
Alessandro De Salvo (INFN, Roma1), Alessandra Doria (INFN, Napoli),
Elisabetta Vilucchi (INFN, Laboratori Nazionali di Frascati)

Outline:

- DPM in Italian sites
- EPEL testbed activity
- Some experiences with consistency checks.
- Tests with DPM in HTTP Dynamic Federation

ATLAS INFN sites in DPM collaboration

- DPM and STORM used in ATLAS italian T1/T2 sites:
 - **DPM at Frascati, Napoli, and Roma Tier2s.**
 - STORM at CNAF Tier1 and at Milano Tier2.
- The Site Managers of the 3 DPM Tier2s are involved in the DPM collaboration .
- Alessandro De Salvo is the ATLAS Italian Computing Coordinator.



DPM setup at ATLAS IT Tier2s

The storage setup is similar, with some peculiarity:

- **INFN FRASCATI:** 1PB, 11 pool nodes
 - Head node on a virtual machine and MySQL DB on a physical server
 - DB replication with MySQL master/slave configuration .
- **INFN NAPOLI:** 2PB, 21 pool nodes
 - Head node and MySQL DB on the same physical server
 - DB daily backups with Percona xtrabackup.
- **INFN ROMA:** 1.5PB , 16 pool nodes
 - MySQL main replica in the DPM head node,
 - Slave replica on a different node, with backup

All the 3 sites use puppet with Foreman for DPM configuration (and for installation and configuration of most/all site servers)

Puppet + Foreman for DPM nodes configuration

Each site has his own Foreman and puppet master, but the testbed uses remotely Roma setup.

Parametrized Puppet modules developed: Puppet manifests released with DPM must be adapted in order to handle parameters from Foreman

Fully generic modules, using Foreman as ENC, can be used to install and configure any server at any site

- Let's see pros and cons of this approach



Edit Atlas SE Nodes

Host Group

Puppet Classes

Network

Operating System

Parameters

Included Classes

tier2recas::dpm::disk
tier2recas::generic



Available Classes

- bdi
- dmlite
- fetchcr
- firewall
- gridftp
- lcgdm
- memcached

- mysql
- stdlib
- t2-frascati
- tier2recas
- voms
- xrootd

In Foreman we can define host groups and reuse classes.
We can pass parameters using Foreman GUI

Hosts

Filter ... Q Search

Name	Operating system	Environment	Model	Host group	Last report	
atlasse00.na.infn.it	Scientific...	production	RS700-E6-ERS4	Atlas SE Nodes	about 3 hours ago	Edit
atlasse01.na.infn.it	Scientific...	production	RS700-E6-ERS4	Atlas SE Nodes	about 2 hours ago	Edit
atlasse02.na.infn.it	Scientific...	production	RS700-E6-ERS4	Atlas SE Nodes	about 1 hour ago	Edit
atlasse03.na.infn.it	Scientific...	production	RS700-E6-ERS4	Atlas SE Nodes	about 2 hours ago	Edit
atlasse12.na.infn.it	Scientific...	production	RS700-E6-RS4	Atlas SE Nodes	about 2 hours ago	Edit
atlasse13.na.infn.it	Scientific...	production	RS700-E6-RS4	Atlas SE Nodes	about 1 hour ago	Edit
atlasse14.na.infn.it	Scientific...	production	RS700-E6-RS4	Atlas SE Nodes	about 1 hour ago	Edit
atlasse15.na.infn.it	Scientific...	production	RS700-E6-RS4	Atlas SE Nodes	about 3 hours ago	Edit
atlasse16.na.infn.it	Scientific...	production	GR360 F1	Atlas SE Nodes	about 2 hours ago	Edit
atlasse17.na.infn.it	Scientific...	production	GR360 F1	Atlas SE Nodes	about 2 hours ago	Edit
atlasse18.na.infn.it	Scientific...	production	GR360 F1	Atlas SE Nodes	39 minutes ago	Edit

Same tool used for host installation and configuration

Puppet + Foreman setup: Pros & Cons

Pros:

- Use Foreman GUI to pass parameters, no file handling.
- Reuse of generic classes , advantage of a hierarchical model
- Same tool used for node installation and configuration

Cons:


- (My) First approach to the tools not so much friendly
- MW configuration “fully puppettized” only for DPM nodes.
- Passwords in clear on Foreman server
- Foreman not yet compatible with Puppet 4

EPEL testbed at INFN-Frascati

Head node + 1 pool node installed in 2013 for the first time.
Validation of DPM 1.8.7, 1.8.8, 1.8.9 as released in epel-testing.


In September 2015 upgraded to DPM 1.8.10.

- HTTP drain not tested yet
 - few data in testbed, better to test in production
- **Space reporting 😊 very useful feature!**
 - Not enabled by default
 - Configure the MySQL DB to accept connections from the disk nodes
 - Modify head and disk node puppet manifests
 - Non enabled in production systems yet.
 - Initialization needed for existing systems
 - how long for large production DBs?



/dpm/lnf.infn.it/home/

Mode	Links	UID	GID	Size	Modified	Name
drwxrwxr-x	1	0	101	0	Tue, 30 Jul 2013 09:03:27 GMT	alice
drwxrwxr-x	13	0	102	3.8G	Fri, 04 Dec 2015 16:21:37 GMT	atlas
drwxrwxr-x	1	0	103	0	Tue, 30 Jul 2013 09:03:30 GMT	cms
drwxrwxr-x	3	0	105	99.6G	Fri, 17 Jul 2015 07:48:18 GMT	dteam
drwxrwxr-x	1	0	106	0	Tue, 30 Jul 2013 09:10:46 GMT	infngrid
drwxrwxr-x	1	0	104	0	Tue, 30 Jul 2013 09:03:32 GMT	lhcb
drwxrwxr-x	1	0	107	0	Tue, 30 Jul 2013 09:10:47 GMT	ops
drwxrwxr-x	1					



/dpm/lnf.infn.it/home/atlas/

Mode	Links	UID	GID	Size	Modified	Name
drwxrwxr-x	919	0	102	0	Fri, 04 Dec 2015 16:20:36 GMT	generated
-rw-rw-r--	1	103	102	760.2M	Thu, 10 Sep 2015 13:33:57 GMT	group.test.hc.NTUP_SMWZ.root.1
-rw-rw-r--	1	103	102	760.2M	Thu, 17 Sep 2015 16:23:00 GMT	group.test.hc.NTUP_SMWZ.root.2
-rw-rw-r--	1	103	102	760.2M	Thu, 17 Sep 2015 16:24:48 GMT	group.test.hc.NTUP_SMWZ.root.3
-rw-rw-r--	1	103	102	760.2M	Thu, 17 Sep 2015 16:28:05 GMT	group.test.hc.NTUP_SMWZ.root.4
-rw-rw-r--	1	103	102	51.7M	Thu, 10 Sep 2015 14:06:52 GMT	log-http.tar
drwxrwxr-x	0	0	0	0	Fri, 04 Dec 2015 16:21:37 GMT	test
-rw-rw-r--	1	103	102	60	Tue, 30 Jul 2013 13:14:41 GMT	test-lcg-cr
-rw-rw-r--	1	103	102	60	Thu, 07 Nov 2013 17:23:36 GMT	test-lcg-cr1
-rw-rw-r--	1	103	102	60	Mon, 28 Apr 2014 16:08:05 GMT	test-lcg-cr2
-rw-rw-r--	1	103	102	60	Mon, 28 Apr 2014 16:08:45 GMT	test-lcg-cr3
-rw-rw-r--	1	103	102	60	Tue, 10 Dec 2013 14:44:12 GMT	test123
-rw-rw-r--	1	101	102	61.5K	Tue, 10 Dec 2013 14:59:04 GMT	test_passwd

Space reporting only for top directories?

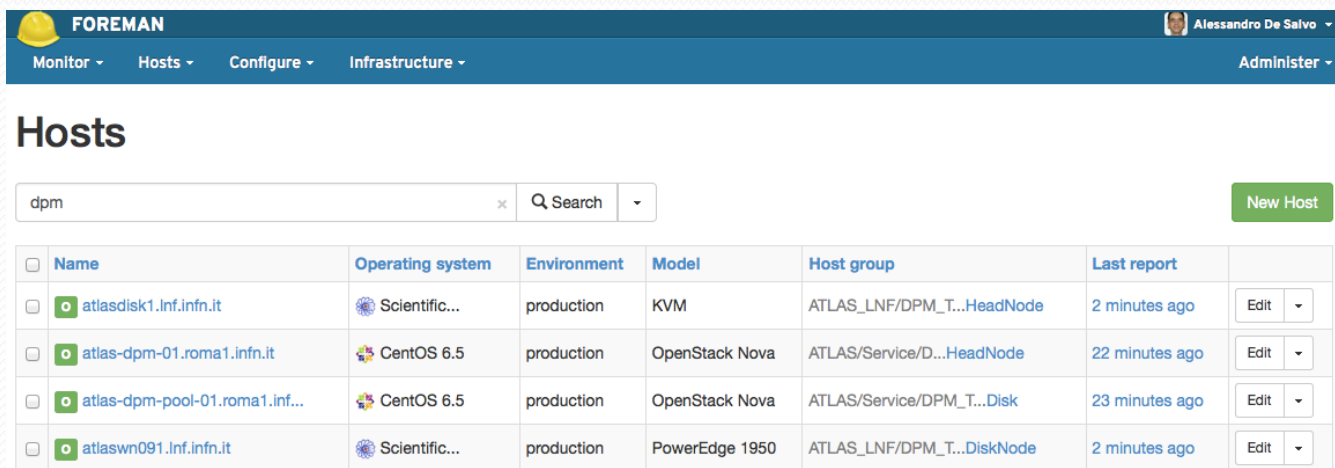
EPEL testbed at INFN-Roma1

Instantiated at the end of April 2014 for the first time, based on OpenStack. Still “private” testbed but fully functional and re-installable very quickly.







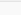

Setup

- DPM head node + 1 Disk Pool running on OpenStack (persistent nodes on cinder volumes)
- MySQL hosted in the headnode VM, but planning to use the existing Percona external DB cluster (8 nodes available)

Same configuration (and same foreman/puppet) used as for the Frascati testbed



The screenshot shows the Foreman web interface. The top navigation bar includes 'FOREMAN', user 'Alessandro De Salvo', and menu items 'Monitor', 'Hosts', 'Configure', 'Infrastructure', and 'Administer'. The main heading is 'Hosts'. Below it is a search bar with 'dpm' entered and a 'New Host' button. A table lists five hosts with columns for Name, Operating system, Environment, Model, Host group, and Last report.

<input type="checkbox"/>	Name	Operating system	Environment	Model	Host group	Last report	
<input type="checkbox"/>	 atlasdisk1.lnf.infn.it	 Scientific...	production	KVM	ATLAS_LNF/DPM_T...HeadNode	2 minutes ago	Edit ▾
<input type="checkbox"/>	 atlas-dpm-01.roma1.infn.it	 CentOS 6.5	production	OpenStack Nova	ATLAS/Service/D...HeadNode	22 minutes ago	Edit ▾
<input type="checkbox"/>	 atlas-dpm-pool-01.roma1.inf...	 CentOS 6.5	production	OpenStack Nova	ATLAS/Service/DPM_T...Disk	23 minutes ago	Edit ▾
<input type="checkbox"/>	 atlaswn091.lnf.infn.it	 Scientific...	production	PowerEdge 1950	ATLAS_LNF/DPM_T...DiskNode	2 minutes ago	Edit ▾

Hot topic for ATLAS sites: deletions and consistency checks, lost files

- Remove ATLASPRODDISK space token
 - Data centrally removed with Rucio, dark data left
 - All data under the same dir `/dpm/na.infn.it/home/atlas/atlasproddisk/`

different ways of performing deletions:

- `rfrm -rf` (slow)
- `dmlite-shell > rmdir -r` (faster, `-r` only in the last release)
- `davix-rm` (not tried 😊)

not found a command to list dpns of files of a space token (only phys names)

- Provide a monthly dump of stored data, to perform consistency checks
 - Ad hoc script provided by ATLAS, reads the DB, creates output files in different formats and writes it in a readable area.
 - Cleanup after cross check with Rucio catalog.

Database checks

dpm-dbck is a powerful tool to find inconsistencies and eventually fix 😊, but it requires to stop the daemons... 😞

It cannot work only on a subtree. This would be useful when there is a known inconsistency in a specific area.

In Frascati a large number of files seemed to be lost:
run a script on a suggestion by Alessandra Forti, based on *dpm-list-disk* and *find*. Quick response (3 mins) to get inconsistencies among MySQL DB and file systems .

Dynamic Federation tests

We made some trials with ATLAS prototype of HTTPS/WebDAV Federation, with PROOF-on-Demand analysis tool.

- we used file replicas located in the DPM endpoints at Frascati and Napoli
- Desy HTTP federator
<https://federation.desy.de/fed>

CHEP 2015: “Tests of PROOF-on-Demand with ATLAS Prodsys2 and first experience with HTTP federation”

The results of the tests show that:

- Being both endpoints up the federator effectively made job connections to the closest storage(Frascati);
- after bringing down the Frascati endpoint, the federator redirected the connections to Napoli;
- the overhead due to the federator is small;
- the support for Rucio naming is in production, can directly open files using the RUCIO syntax.