


DPM Nagios & Puppet


Making life easy(er)

Ricardo Rocha
(ricardo.rocha@cern.ch)

On behalf of the IT/GT/DMS team



Monitoring with NAGIOS



Nagios in DPM

- Why?
 - Widely used in WLCG / EMI
 - Community support
 - Integration with other tools
- How?
 - Support for Nagios v2 or v3, pnp4nagios 0.4.x
 - One additional rpm
 - Multiple configuration options: manual, NCG, Puppet



<https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/Monitoring>

Available Probes

- Initial wishlist
 - Discussed with some of the site admins
 - Some were generic, already available
- Mostly in Python
- Probes
 - Compliant with nagios guidelines (status, output)
 - Database activity, filesystem activity, free and used disk space, service status, ...
 - List is constantly growing

Available Probes

- `check_hostcert`
- `check_oracle_expiration`
- `check_partition_activity`
 - read/write ops
- `check_cpu`
 - no i/o wait (coming soon)
- `check_network`
 - multiple interfaces
- `check_process`
 - memory, cpu, number descriptors, number instances and threads
- `check_dpm_pool`
 - unresponsive disks, filesystems, space usage
- `check_dpns`
- `check_gridftp`
 - get and put
- `check_rfio`
 - get and put

Available Probes

Nagios

General

- Home
- Documentation

Monitoring

- Tactical Overview
- Service Detail
- Host Detail
- Hostgroup Overview
- Hostgroup Summary
- Hostgroup Grid
- Servicegroup Overview
- Servicegroup Summary
- Servicegroup Grid
- Status Map
- 3-D Status Map

Show Host:

- Service Problems
- Host Problems
- Network Outages

Reporting

- Trends

Host ↑↓	Service ↑↓	Status ↑↓	Last Check ↑↓	Duration ↑↓	Attempt ↑↓	Status Information
dpm01_cern.ch	CNS MySQL DB check at: dpm01_cern.ch	OK	05-11-2011 10:24:18	25d 19h 50m 43s	1/5	Uptime: 3521512 Threads: 2 Questions: 899328 Slow queries: 0 Opens: 31 Flush tables: 1 Open tables: 25 Queries per second avg: 0.255
	CPU check at: dpm01_cern.ch	OK	05-11-2011 10:26:31	3d 5h 24m 40s	1/5	DM-CPU OK - Idle=97.31%
	DPM MySQL DB check at: dpm01_cern.ch	OK	05-11-2011 10:28:44	5d 10h 58m 27s	1/5	Uptime: 3521778 Threads: 2 Questions: 899329 Slow queries: 0 Opens: 31 Flush tables: 1 Open tables: 25 Queries per second avg: 0.255
	DPM pools check at: dpm01_cern.ch	OK	05-11-2011 10:30:58	9d 7h 0m 12s	1/5	DM-POOL OK - All DPM disks are OK
	DPM processes check at: dpm01_cern.ch	OK	05-11-2011 10:33:10	16d 8h 30m 14s	1/5	DM-PROC OK - All parameters OK
	DPNS service check listening at: dpm01_cern.ch	OK	05-11-2011 10:25:23	19d 22h 39m 57s	1/5	DM-DPNS OK - 1.8.1-0
	Host certificate check at: dpm01_cern.ch	OK	05-11-2011 10:27:36	19d 22h 42m 54s	1/5	DM-CERT OK - The host certificate expires in 310 days
	Load for host: dpm01_cern.ch	OK	05-11-2011 10:09:49	3d 5h 14m 22s	1/5	OK - load average: 0.03, 0.49, 0.37
	Network check at: dpm01_cern.ch	OK	05-11-2011 10:24:44	1d 10h 32m 19s	1/5	DM-NET OK -
	Ping for host: dpm01_cern.ch	OK	05-11-2011 10:26:57	2d 6h 30m 14s	1/5	PING OK - Packet loss = 0%, RTA = 0.91 ms
	Swap check for host: dpm01_cern.ch	OK	05-11-2011 10:29:10	25d 19h 46m 11s	1/5	SWAP OK - 100% free (3999 MB out of 3999 MB)
dpm02_cern.ch	CPU check at: dpm02_cern.ch	OK	05-11-2011 10:31:24	0d 0h 52m 17s	1/5	DM-CPU OK - Idle=98.65%
	Gridftp service check listening at: dpm02_cern.ch	CRITICAL	05-11-2011 10:23:37	19d 19h 41m 55s	5/5	DM-GFTP CRITICAL - Could not upload the file:
	Host certificate check at: dpm02_cern.ch	OK	05-11-2011 10:25:50	14d 17h 23m 42s	1/5	DM-CERT OK - The host certificate expires in 310 days
	Load for host: dpm02_cern.ch	OK	05-11-2011 10:08:03	3d 5h 16m 8s	1/5	OK - load average: 0.00, 0.00, 0.00
	Network check at: dpm02_cern.ch	OK	05-11-2011 10:30:17	14d 17h 32m 39s	1/5	DM-NET OK -
	Ping for host: dpm02_cern.ch	OK	05-11-2011 10:25:11	4d 13h 32m 0s	1/5	PING OK - Packet loss = 0%, RTA = 1.58 ms
	Swap check for host: dpm02_cern.ch	OK	05-11-2011 10:27:24	26d 17h 30m 52s	1/5	SWAP OK - 100% free (3972 MB out of 3999 MB)

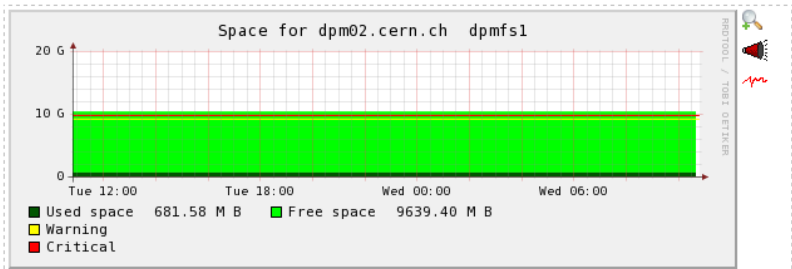
Performance Visualization

- Enabled for all probes
 - Status for some, detailed performance for many
- pnp4nagios (<http://www.pnp4nagios.org/>)
 - There are other tools doing the same, this is the one used today in WLCG / EMI
 - Specific templates provided with our rpm

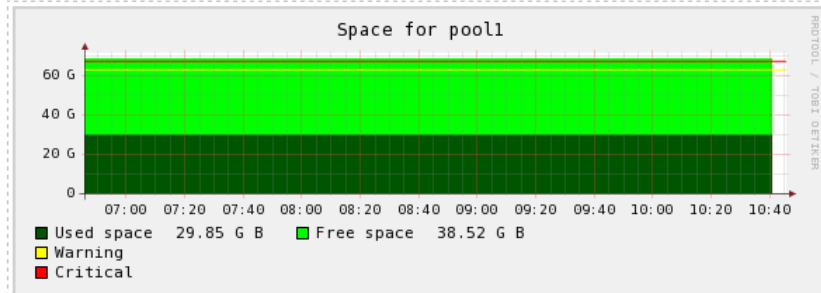
<http://puppet.cern.ch/nagios/pnp4nagios/index.php>

Performance Visualization

Service: DPM_pools_check_at__dpm01.cern.ch
24 Hours (10.05.11 10:48 - 11.05.11 10:48)



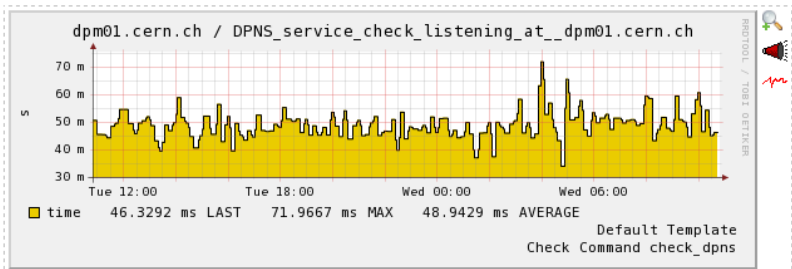
Datasource: pool1



Service: DPM_processes_check_at__dpm01.cern.ch
24 Hours (10.05.11 10:48 - 11.05.11 10:48)

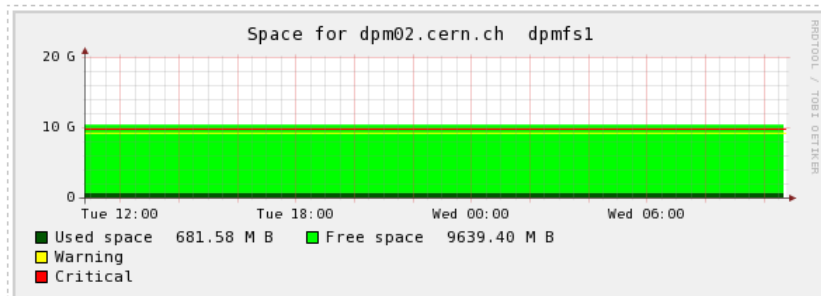


Service: DPNS_service_check_listening_at__dpm01.cern.ch
24 Hours (10.05.11 10:48 - 11.05.11 10:48)



24 Hours (10.05.11 10:45 - 11.05.11 10:45)

Datasource: dpm02.cern.ch dpmfs1




Datasource: dpm02.cern.ch dpmfs2



Future Work

- Release and rollout to sites
 - And collect feedback, suggestions for new probes...
 - RPM available, not yet in any official repository
 - We're happy to fulfill your wishes
- Current requests
 - VO specific usage information (space token might be enough and is much easier)
 - SRM success / fail rate
- Other coming probes
 - Lots of real time performance data
 - Avg access time per operation, avg transfer time, etc
- Extending to other DMS components
 - LFC, FTS

Configuration with Puppet




Puppet in DPM

- Why?
 - Looked for a strong configuration management solution
 - For our internal testbed(s), but reusable at DPM sites
 - Strong momentum, community and commercial support
 - twitter, sun / oracle, rackspace, ...
 - (very) Extensive documentation
- How?
 - Set of specific modules + generic ones
 - Available via puppet forge or github
 - An alternative... not a replacement for YAİM
 - Current support level is « best effort »

<https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/Puppet>

Reusable Modules

- All DPM required areas are covered
- As generic and reusable as possible
 - Go ahead and try them for your service



<https://forge.puppetlabs.com/users/rocha>

ant antivirus apache apache2
application applications
apt archive asadmin auth
authentication authoritative
automation availability aws
backup backuppc bash basic
bigtable bind bugzilla build bzip
caching Cacti capistrano
CentOS cern certificates certs
client cloud cluster cobbler
collectd common compile
compiler concatenate control
controltier cron CVS cyber
dashboard database db
debian deployment developer
development dhcp digester
directory distributed dms dns
dnsmasq dovecot dpm e2
editor enterprise execution
factor facts farm fcgi fcron
FHS file File Server files
fileserver filesystem foo
foreman ftp func gcc gconf
generic Gentoo gix git
glassfish Glider gLite glusterfs
gnome2 graphics greplist
grid groups ha hadoop
hardware hbase hdd headless
heartbeat hg high-availability
hosts htdigest httpasswd HTTP
httpd ignore imagemagick imap
inet6 infrastructure interface
inventory ipaddress iptables
ipv6 ipvs jabber java jsp
keepalived kickstart levm
LANANA languages lcg lcgutil
limits-conf lmsensors load-
lock logging
logical volume logs logwatch
mail master-election
mcollective mediawiki message
messaging meta middleware
module modules mon mongodb
monitor **monitoring**

Cluster Configuration

```
node 'dpm01.cern.ch' inherits dpm-service {
  include dpm::headnode
  include dpm::nfsserver
  include dpm::nagios::headnode

  # setup supported domain/vo(s)
  dpm::headnode::domain { 'cern.ch': require => Service['dps'], }
  dpm::headnode::vo { 'dteam': domain => 'cern.ch', require => Dpm::Headnode::Domain['cern.ch'], }

  # setup pools
  dpm::headnode::pool { 'pool1': require => Service['dpm'] }
}
```

HEAD NODE

```
node 'dpm02.cern.ch', 'dpm03.cern.ch' inherits dpm-service {
  include dpm::disknode
  include dpm::nagios::disknode

  # setup filesystems
  dpm::disknode::filesystem { "dpmfs1": fs => "/dpmfs1", pool => 'pool1', }
  dpm::disknode::filesystem { "dpmfs2": fs => "/dpmfs2", pool => 'pool1', }
}
```

DISK NODES

```
node 'dpm04.cern.ch' inherits dpm-service {
  include dpm::nfsclient
  include dpm::nagios::client
  include voms::client
}
```

CLIENT NODES

Cluster Configuration

```
node 'dpm01.cern.ch' inherits dpm-service {  
  include dpm::headnode  
  include dpm::nfsserver  
  include dpm::nagios::headnode
```

HEAD NODE

GENERIC NODE

```
# set  
dpm::  
dpm::
```

```
# set  
dpm::  
}
```

```
node 'd  
inclu  
inclu
```

```
# set  
dpm::  
dpm::  
}
```

```
node 'd  
inclu  
inclu  
include voms::client  
}
```

```
node cern-service {  
  Package { require => Yumrepo["lcgdm-unstable-etics", "epel"] }  
  
  include dms::unstable  
  include glite::gridmap  
  include cern::slc5  
  include cern::hostcert  
  include cern::afs  
}  
  
node dpm-service inherits cern-service {  
  include voms::atlas  
  include voms::dteam  
  
  cern::afs::user { ["rocha", "aalvarez", "prodrigu", "baud", "okeeble"]: }  
  cern::keytab { "keytab-nfs": serviceType => "nfs" }  
}
```

n.ch'], }

DES

NODES

Puppet Dashboard

puppet dashboard v1.0.4 » Home • Nodes • Groups • Classes • Reports

Nodes 6

- Currently successful 5
- Currently failing 1
- Ever succeeded 6
- Ever failed 4
- Never reported 0
- Not currently reporting 5

Add node

Class

dpm::headnode 0

Add class

Group

Add group

Dashboard

Node failures

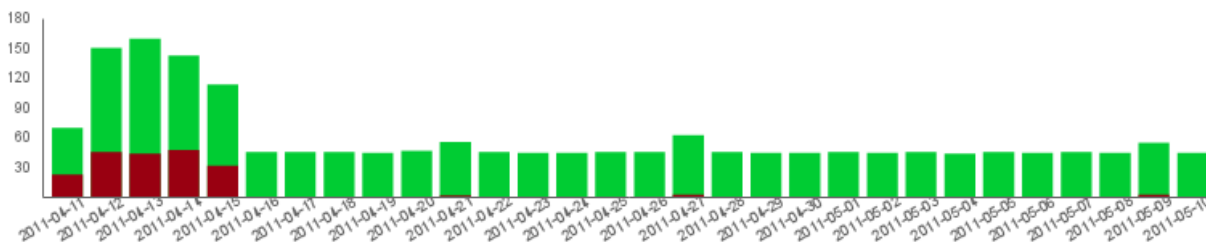
1 node are currently reporting failures: [dpm04.cern.ch](#).

Nodes no longer reporting

5 nodes have not reported in the last about 1 hour: [rocha-slc5.cern.ch](#), [dpm01.cern.ch](#), [dpm03.cern.ch](#), [dpm04.cern.ch](#), and [rocha-emi.cern.ch](#).

Daily run status

Number and status of runs during the last 30 days:



Recently-reported nodes

Hostname	Latest report
dpm02.cern.ch	2011-05-11 09:48 UTC
rocha-emi.cern.ch	2011-05-09 17:52 UTC
dpm01.cern.ch	2011-05-09 15:38 UTC
dpm03.cern.ch	2011-04-27 16:20 UTC
dpm04.cern.ch	2011-04-27 13:12 UTC
rocha-slc5.cern.ch	2011-03-17 13:04 UTC

Puppet & Nagios Integration

- The « cherry on top »
- Puppet has built-in resource types for nagios

```
@@nagios_command {
  "check_ping":
    ensure      => "present",
    command_line => "\$USER1\$ /check_ping -H \$HOSTADDRESS$ -w \$ARG1$ -c \$ARG2$";
  "check_nrpe":
    ensure      => "present",
    command_line => "\$USER1\$ /check_nrpe -H \$HOSTADDRESS$ -c \$ARG1$ -a \$ARG2$";
  "process-service-perfdata":
    ensure      => "present",
    command_line => "/usr/bin/perl /usr/libexec/pnp4nagios/process_perfdata.pl";
}

@@nagios_contact { "nagios":
  ensure      => present,
  alias       => "Nagios Admin",
  host_notification_period => "24x7",
  service_notification_period => "24x7",
  service_notification_options => "w,u,c,r",
  host_notification_options => "d,u,r",
  service_notification_commands => "notify-by-email",
  host_notification_commands => "host-notify-by-email",
  email       => "ricardo.rocha@cern.ch",
}
```


Future Work

- Make it available to interested DPM sites
- Improve DPM modules from their experience
 - Larger setups, bigger problems
- Get other product teams to reuse our common modules
 - And to help us with their maintenance