



# LCGDM Nagios Probes

# Motivation

- In order to facilitate site and node maintenance, monitoring tools can really make life easier to system administrators
- In fact, EMI defines a “standard” monitoring tool to be used: Nagios
  - Which means probes should be part of the standard distribution
  - PT responsibility

# Overview



- A initial “wish” list was used as a base to develop the probes
  - Although some of them already exists outside, so there is no need of reinventing the wheel
    - Database activity, filesystem, free and used disk space...
- 11 probes have been developed
  - Some of them are useful outside LCGDM, as they are generic
- Python is the language
  - One module contains shared functionality, so adding new probes is fast

# Generic probes

- check\_hostcert
  - Expiration time
- check\_oracle\_expiration
  - Password remaining time
- check\_partition\_activity
  - Read/write operations
  - Individual disks can be selected
  - No warning or critical criteria (?)
- check\_cpu
  - Existing probe doesn't care about IO wait, IRQ, etc.

# Generic probes



- check\_network
  - Individual interfaces can be selected
  - No warning or critical criteria (?)
- check\_process
  - Monitors several aspects about a specified set of processes
    - Memory, CPU
    - Number of connections and file descriptors
    - Number of instances and threads
  - Needs sudo

# DPM Specific probes

- `check_dpm_pool`
  - Checks several things
    - Unresponsive DPM disks hosts
    - Unresponsive DPM filesystems
    - Unresponsive DPM heads
    - Free and used space per pool and per subsystem
- `check_dpns`
  - Just calls `dpns_ping`
- `check_gridftp`
  - Check upload/download and data integrity
- `check_rfio`
  - Same as `check_gridftp` for RFIO

# Probes output















- One of the possible Nagios status:
  - Critical, Warning, OK
- Performance data
  - All probes produce it
  - Complies with the “Nagios Developer Guidelines”
    - Which means they can be used with pretty much every plotting extension

# Performance visualization

- Nagios does not keep track of performance data
- An external tools has to be used
  - pnp4nagios (used by EMI Infrastructure Area and EGI)
- Default template is not good enough for most of the data
  - nagios-plugins-lcgdm provides a set of templates ready to use



# Examples – General view


Host 	Service 	Status 	Last Check 	Duration 	Attempt 	Status Information
<a href="#">lxfkra04a04.cern.ch</a>	<a href="#">DM_CERT</a>	 OK	2011-05-10 13:11:10	2d 9h 33m 17s	1/4	DM-CERT OK - The host certificate expires in 72 days
	<a href="#">DM_CPU</a>	 OK	2011-05-10 13:11:45	1d 22h 59m 22s	1/4	DM-CPU OK - Idle=99.90%
	<a href="#">DM_DPNS</a>	 OK	2011-05-10 13:11:21	5d 20h 31m 26s	1/4	DM-DPNS OK - 1.8.0-1
	<a href="#">DM_GSIFTP</a>	 OK	2011-05-10 13:11:36	5d 20h 31m 11s	1/4	DM-GFTP OK - Test passed
	<a href="#">DM_INFOSYS</a>	 OK	2011-05-10 13:11:31	3d 9h 9m 36s	1/4	DM-INFOSYS OK - Published information OK
	<a href="#">DM_NET</a>	 OK	2011-05-10 13:11:47	1d 12h 36m 0s	1/4	DM-NET OK - eth0 (in: 0.00 KB/s out: 0.00 KB/s err: 0.00%)
	<a href="#">DM_ORA_ACCOUNT</a>	 OK	2011-05-10 13:12:42	0d 0h 40m 5s	1/4	DM-ORA OK - The password expires in 174 days
	<a href="#">DM_PART</a>	 OK	2011-05-10 13:12:12	0d 12h 43m 55s	1/4	DM-PART OK - sda3 (r: 0 KB/s, w: 216 KB/s) sdb1 (r: 0 KB/s, w: 0 KB/s)
	<a href="#">DM_POOL</a>	 WARNING	2011-05-10 13:12:08	2d 6h 30m 39s	4/4	DM-POOL WARNING - lxfkra04a04.cern.ch disk free space is under the warning limit 
	<a href="#">DM_PROC</a>	 OK	2011-05-10 13:11:23	5d 20h 31m 24s	1/4	DM-PROC OK - All parameters OK
	<a href="#">DM_RFIO</a>	 OK	2011-05-10 13:11:39	5d 20h 31m 8s	1/4	DM-RFIO OK - Test passed

# Examples – Detailed view

## Service State Information

Current Status:	<b>WARNING</b> (for 2d 6h 31m 14s)
Status Information:	DM-POOL WARNING - lxfsra04a04.cern.ch disk free space is under the warning limit
Performance Data:	'lxfsra04a04.cern.ch:/storage'=2175015717B;5368709120;;1073741824;;0;15600713728 'lxfsra10a01.cern.ch:/storage'=11014920997B;5368709120;;1073741824;;0;15600713728 'lxfsra10a05.cern.ch:/storage'=11012091904B;5368709120;;1073741824;;0;15600713728 'lxfsra12a06.cern.ch:/storage'=11016011776B;5368709120;;1073741824;;0;15600713728 'lxfsra14a03.cern.ch:/storage'=11014709248B;5368709120;;1073741824;;0;15600713728 'the_dpm_pool'=46232749642B;5368709120;;1073741824;;0;78003568640
Current Attempt:	4/4 (HARD state)
Last Check Time:	2011-05-10 13:12:08
Check Type:	ACTIVE
Check Latency / Duration:	0.143 / 0.907 seconds
Next Scheduled Check:	2011-05-10 13:13:48
Last State Change:	2011-05-08 06:42:08
Last Notification:	2011-05-08 06:42:15 (notification 3)
Is This Service Flapping?	<b>NO</b> (0.00% state change)
In Scheduled Downtime?	<b>NO</b>
Last Update:	2011-05-10 13:13:15 (0d 0h 0m 7s ago)

# Examples – Performance graph




# Examples – Performance graph

Datasource: RFIO



RFIO in lxfsra04a04.cern.ch

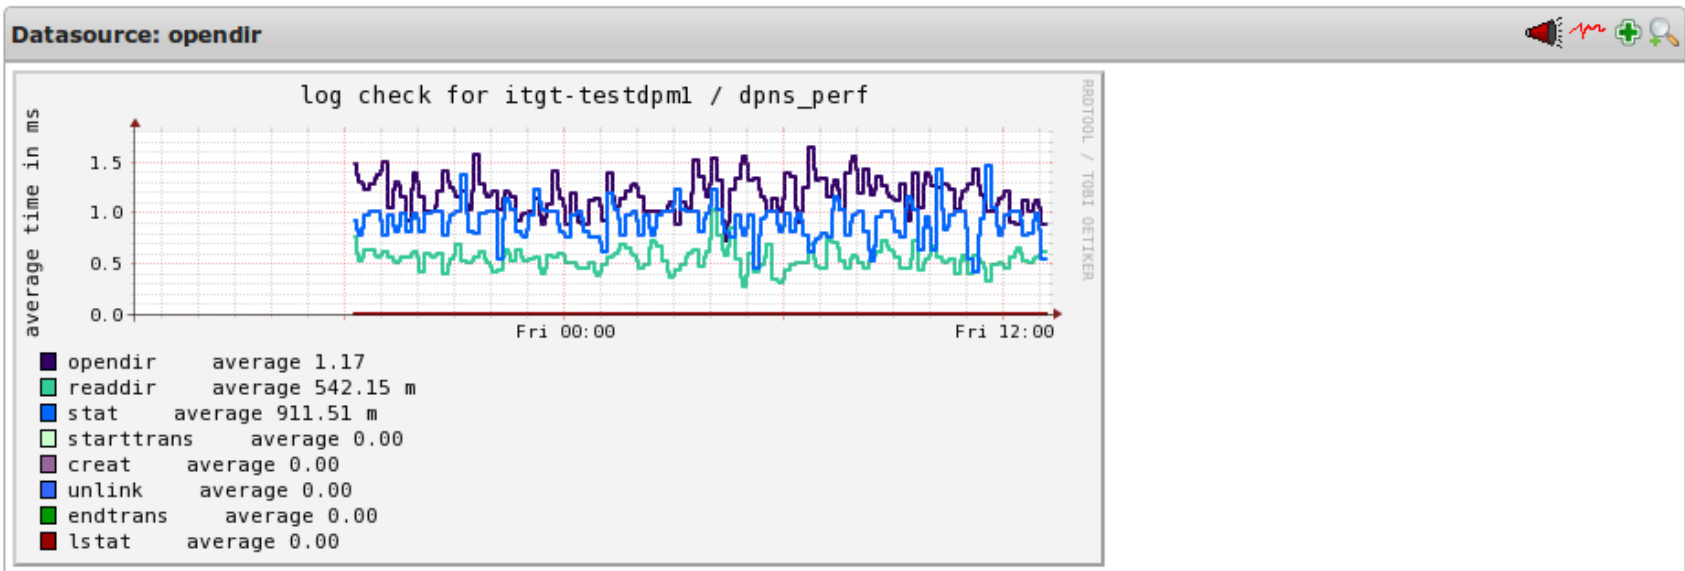


# Performance from logs

## Service State Information

Current Status:	<b>OK</b> (for 0d 19h 5m 47s)
Status Information:	Everything is OK
Performance Data:	opendir=0.857;2.460;3.690:2 readdir=0.625;1.140;1.710:2 stat=0.500;1.860;2.790:2 starttrans=0.000;0.000;0.000:2 creat=0.000;3.040;4.560:2 unlink=0.000;3.120;4.680:2 endtrans=0.000;5.240;7.860:2 lstat=0.000;2.400;3.600:2 mkdir=0.000;6.000;9.000:2 access=0.000;2.000;3.000:2 chdir=0.000;3.000;4.500:2 rmdir=0.000;4.000;6.000:2 getidmap=0.000;4.340;6.510:2 getgrpbygids=0.000;0.000;0.000:2 getgrpbygid=0.000;0.000;0.000:2 address=0.000;3.000;4.500:2 accessr=0.000;2.600;3.900:2 getreplicax=0.000;2.480;3.720:2 delreplica=0.000;1.520;2.280:2 statr=0.000;0.220;0.330:2 getpath=0.000;1.000;1.500:2
Current Attempt:	1/3 (HARD state)
Last Check Time:	05-27-2011 13:09:52
Check Type:	ACTIVE
Check Latency / Duration:	0.170 / 0.072 seconds
Next Scheduled Check:	05-27-2011 13:19:52
Last State Change:	05-26-2011 18:06:53
Last Notification:	N/A (notification 0)
Is This Service Flapping?	<b>NO</b> (0.00% state change)
In Scheduled Downtime?	<b>NO</b>
Last Update:	05-27-2011 13:12:36 ( 0d 0h 0m 4s ago)

# Performance from logs



# Current status

- So far, the consulted sysadmins are happy with the probes
  - Couple of requests pending
    - Requests per VO
    - SRM success/fail rate
    - Number of transfer last day
- The probes are running with no errors
  - In principle they seem to detect failures correctly
  - Also, performance data is being generated with no troubles
- Documentation available in the DPM Trac wiki
  - Probes installation and pnp4nagios configuration

# Future work

- Keep adding probes as system administrators give feedback
  - Again, feedback needed!
- Create specific probes for other components
  - FTS, LFC
  - Always producing performance data
- Pack and distribute together with future DPM releases
- Support and maintenance