

CPPM News



- Ressources :
**stockage 800 To
CPU 1400 cores
Réseau 10 Gb LHCONE**
- **Électricité: 800 kW en entrée + coffret salle, ajout onduleur 60 kVA**
- **Matériel : achat 2 baies réfrigérées eau froide
CPU : remplacement des wn par Dell C6100/C6200
stockage: ajout de baies MD1000/MD1200 sur serveur existant, derniers achats DELL R720/MD1200**





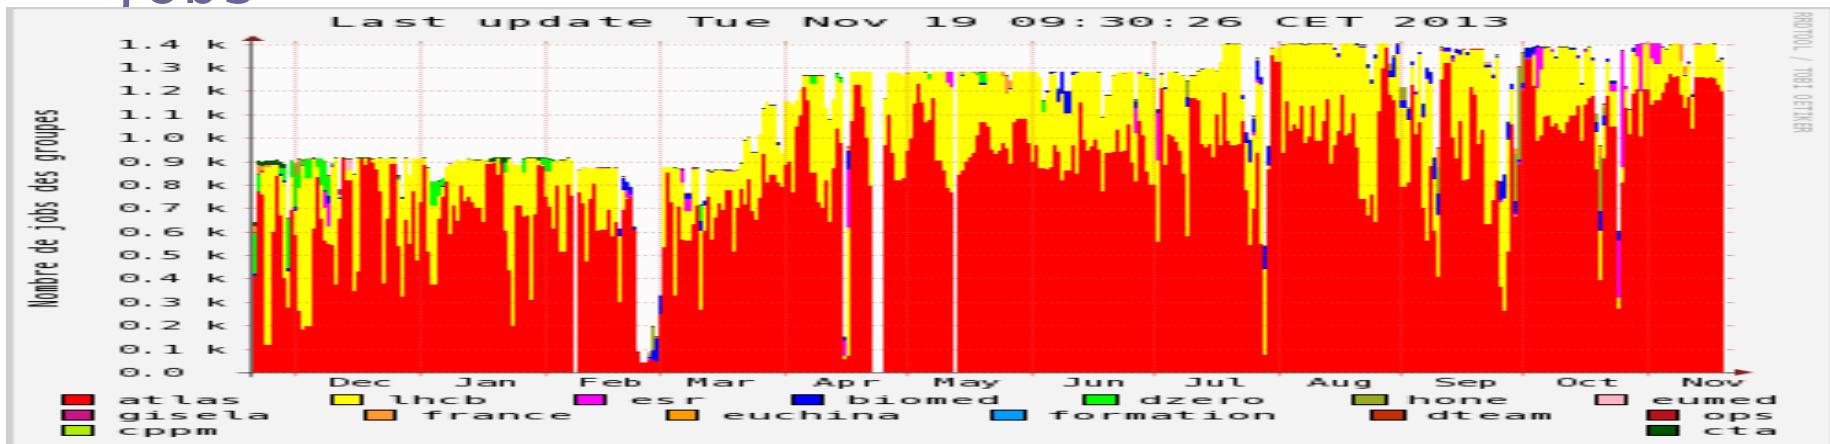
Services

- 2 Cream CE SL5 et EMI2 + 1 serveur Torque et WN in SL640
- 1 SE DPM SL5 avec accès webdav et xroot (ATLAS fed)
- 2 WMS 1 SL5 EMI2 multi VO, 1 SL5 EMI3 nagios regional
- VOMS server et VO CPPM arrêtés
- PERFSOAR (2)
- ARGUS
- SQUID
- QUATTOR 13.1.2 svn. Migration 13.9 github to do
- 1 serveur Nagios

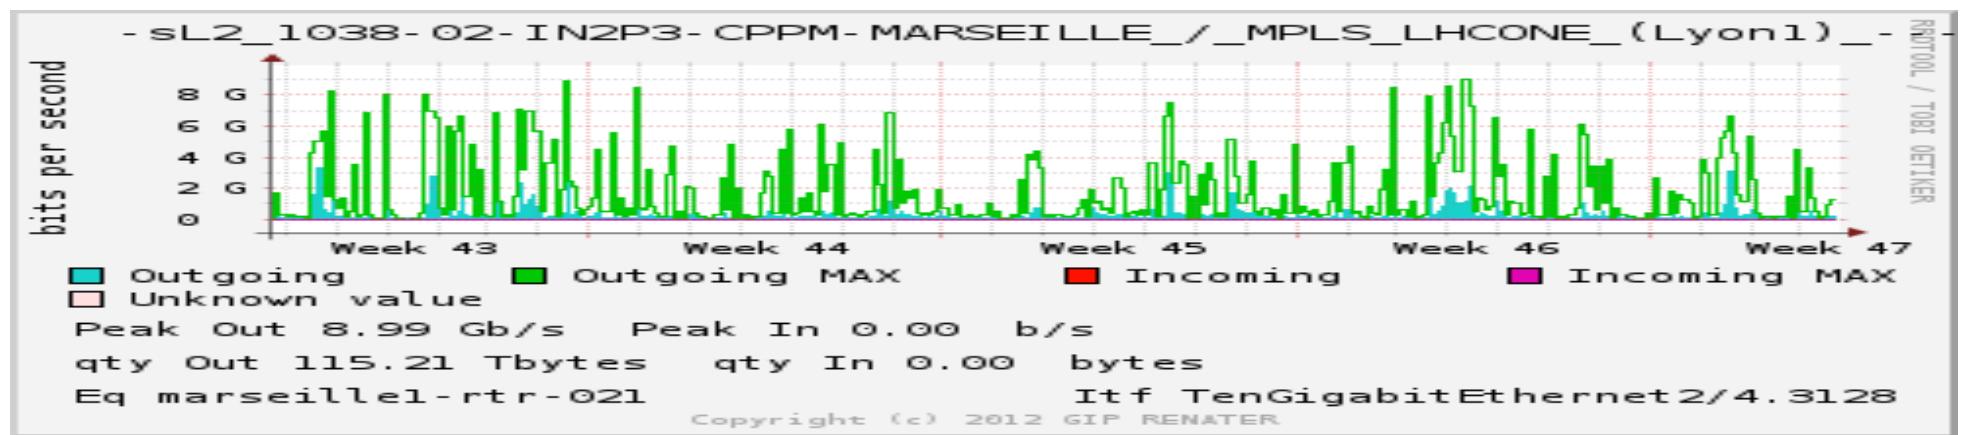


Utilisation

jobs



Network



Monitoring au IN2P3-CPPM



Carlos Carranza
Rencontres France Grilles et LCG France
26-28 novembre 2013

Des sujets à aborder

Supervision du matériel.

- Données de métrologie
- pnp4nagios

Un cas de supervision de A a Z.

Supervision stockage

- Les serveurs de disques.
- Le service DPM et catalogue DPNS

CernVM-FS dans les nœuds de calcul

- Le service
- Gestion du cache

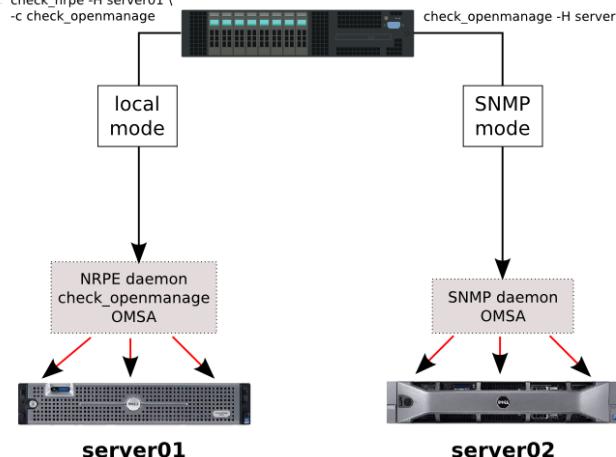
Le Matériel (99% Dell)

OpenManage Server

- Fournit par Dell
- Déployée vers tout les nœuds
- Utilisation multiple
 - Interface graphique
 - Interface en ligne de commande
 - Via SNMP ou TCP



Nagios



Plugin Nagios

- `check_open_manage`
- Dernier version 3.7.11
- Identification des composants matérielles.
- Personnaliser messages alerte, seuils, etc.
- Black listes



OK	1/3	OK - System: 'PowerEdge R710', SN: 'XXXXXXXX', 144 GB ram (18 dimms), 2 lo
CRITICAL	3/3	Memory module 5 [DIMM_A6, 4096 MB] needs attention: Single-bit error loggi
OK	1/3	OK - System: 'PowerEdge 2950', SN: 'XXXXXXXX', 16 GB ram (8 dimms), 1 logic
OK	1/3	OK - System: 'PowerEdge R610', SN: 'XXXXXXXX', 24 GB ram (6 dimms), 2 logic
WARNING	3/3	Physical Disk 0:0:1 [Dell ST146852SS, 146GB] is Replacing [78% complete] Logical Drive '/dev/sda' [RAID-1, 136.12 GB] needs attention: Degraded
OK	1/3	OK - System: 'PowerEdge M600', SN: 'XXXXXXXX', 16 GB ram (8 dimms), 1 logi

Métrologie



Le plugin retourne :

- Un statut global
- Un commentaire contextuel
- Des données de métrologie

Performance Data:

```
T0_FCB_Ambient1=20C;[N/A];50 T2_MLB_TEMP_2=58C;85;87
T3_MLB_TEMP_3=50C;85;87 T4_Processor_1=45C;81;85
T5_MLB_TEMP_1=57C;85;87 T6_Processor_2=48C;81;85
A0_PS_Current=22A;255;255 A1_PSU_2_POUT=49A;0;0
A2_PSU_1_POUT=51A;0;0 V0_PS_12V=11.966V;0;0
V1_PS_5V=5.018V;0;0 V2_STBY_3.3V=3.336V;0;0
V3_VCORE_1=1.029V;0;0 V4_VCORE_2=1.055V;0;0
F0_FCB_FAN1=8100rpm;25500;25500 F1_FCB_FAN2=8000rpm;
25500;25500 F2_FCB_FAN3=8000rpm;25500;25500
F3_FCB_FAN4=8000rpm;25500;25500
```

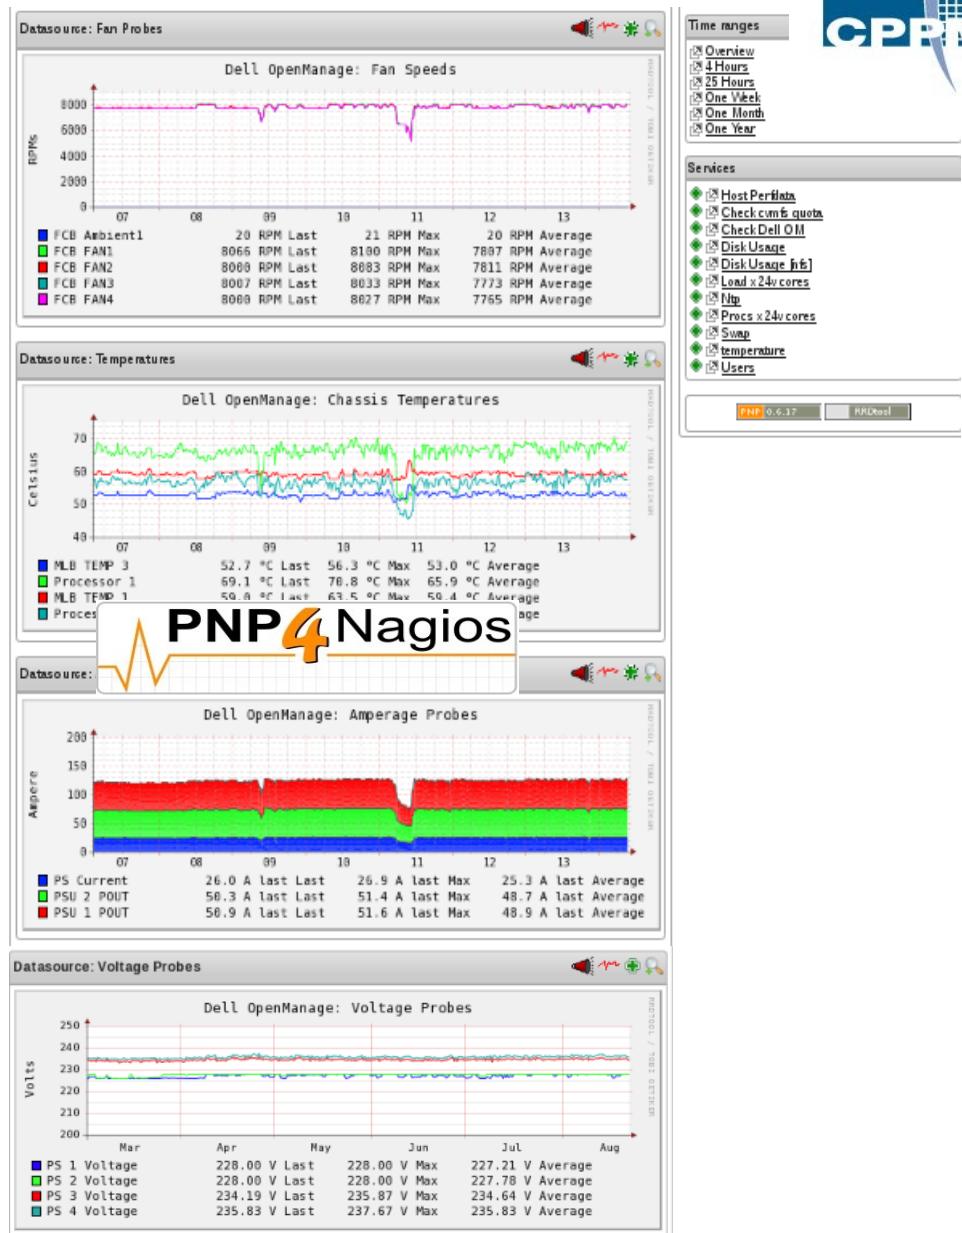


Performance Data



PNP4Nagios

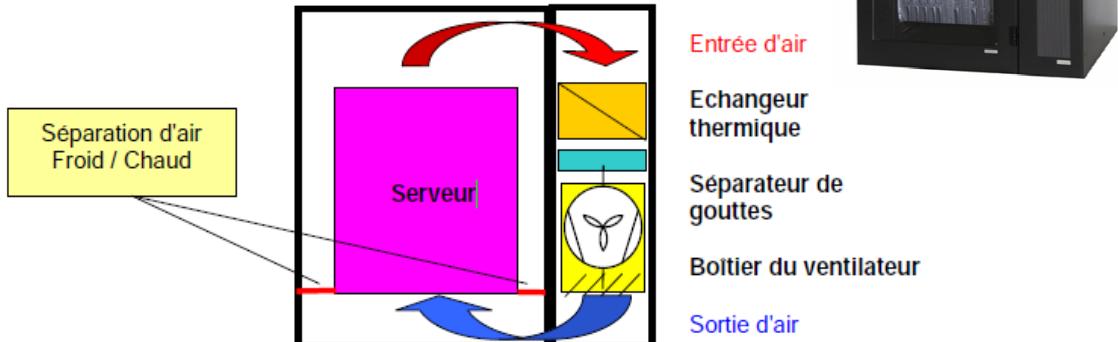
- Remplit une base de données RRD
- Utilise un template php pour rendre :
 - Des graphes.
 - Des historiques
 - Des rapports..
 - La disponibilité du service.



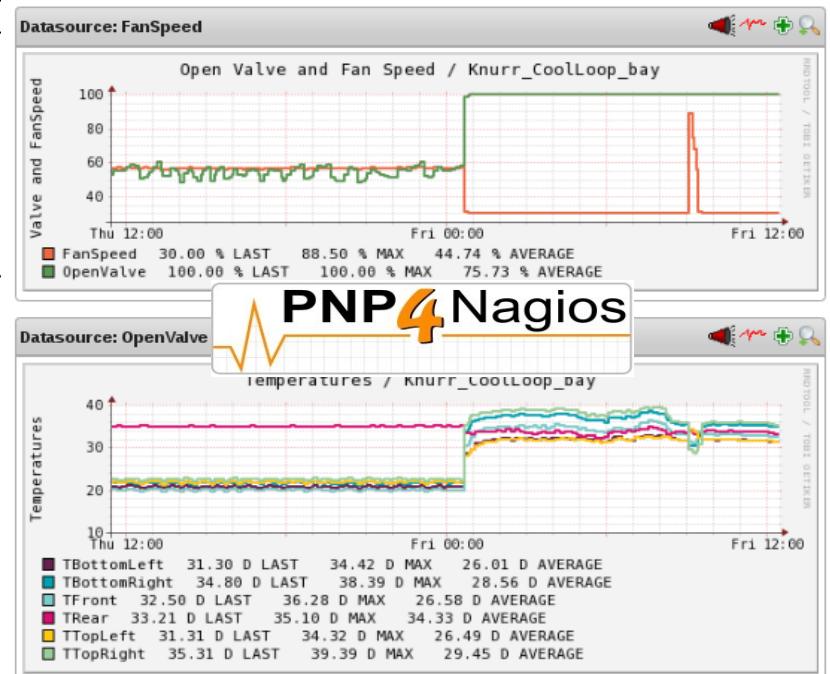
Un cas de A à Z

Supervision d'une baie réfrigérée Knurr Coolloop

- Évacuation des charges thermiques de 10 à 30 kW.
- Le refroidissement par échangeur thermique air/eau.
- Refroidissement adaptée à la charge thermique.



Content	Object 1	Object 2	Object 3	
enterprises.knurr.cooling coolControl 1.3.6.1.4.1.2769.2.1.				
coolEnvironment	1			Analog inputs
temp1 temp1Name temp1CalibrationOffset temp1Current temp1SettingLevel temp1WarningLevel temp1AlarmLevel		1		Temperature server return air CL / CT 1 Warm air 2 Offset temperature 3 Temperature value, return air 4 Normal value 5 Upper warning threshold 6 Upper alarm value



```
$check_coolloop.pl -H 192.168.1.15 -C abcxyz
```

```
check_coolloop.pl OK - FanSpeed<30.0>, OpenValve<59.0>,  
TBottomLeft<26.4>, TBottomRight<22.7>, TFront<23.4>,  
TRear<32.9>, TTopLeft<25.0>, TTopRight<24.3>, |  
FanSpeed=30%;95;98 OpenValve=59%;95;98  
TBottomLeft=26.4D;28;35 TBottomRight=22.7D;28;35  
TFront=23.4D;25;27 TRear=32.9D;43;45 TTopLeft=25D;28;35  
TTopRight=24.3D;28;35
```

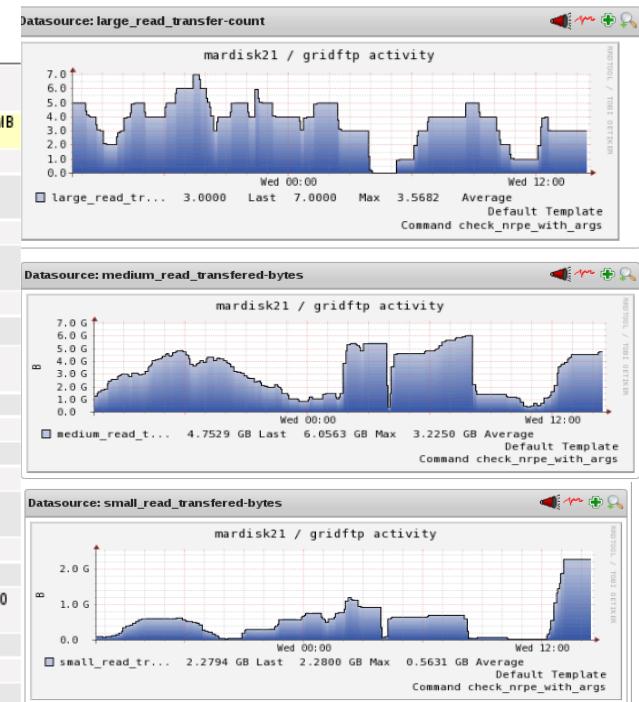
Le stockage

- Plus d'une vingtaine des serveurs de disque derrière un serveur DPM.
- 900 TB de stockage
- Des serveurs Dell PE R510 (~ 20 TB)
- Des baies attachées Powervault MD1200 (~ 30 TB)
- Connectées au réseau x 10 GB
- Des paquetages spécifiques ajoutées
 - Nagios-plugin-dpm-disk

Service	Status	Last Check	Duration	Attempt	Status Information
Check Dell OM	OK	11-13-2013 10:12:18	17d 14h 32m 42s	1/3	OK - System: 'PowerEdge R510 II', SN: 'XXXXXXXX', 24 GB ram (6 dimms), 2 logical drives, 14 physical drives
Dpm disks [datann]	WARNING	11-13-2013 10:16:38	0d 14h 46m 34s	3/3	DISK WARNING - free space: /data01 207085 MB (2% inode=99%): /data02 206722 MB (2% inode=99%):
Host Certificat	OK	11-13-2013 10:11:44	43d 0h 57m 14s	1/3	check_end_cert.pl OK - 215 days to expiry the certificat
Load x 16v cores	OK	11-13-2013 10:10:18	43d 0h 58m 9s	1/3	OK - load average: 0.04, 0.05, 0.01
Ntp	OK	11-13-2013 10:10:18	43d 0h 56m 14s	1/3	NTP OK: Offset -0.0002278089523 secs
Procs x 16v cores	OK	11-13-2013 10:11:15	43d 0h 57m 41s	1/3	PROCS OK: 319 processes
Ssh	OK	11-13-2013 10:11:55	43d 0h 58m 10s	1/3	SSH OK - OpenSSH_4.3 (protocol 2.0)
Swap	OK	11-13-2013 10:13:42	43d 0h 57m 41s	1/3	SWAP OK - 100% free (7820 MB out of 7820 MB)
System disk [sda]	OK	11-13-2013 10:10:34	43d 0h 58m 9s	1/3	DISK OK - free space: / 14562 MB (79% inode=93%): /var 100622 MB (96% inode=99%): /boot 216 MB (93% inode=99%):
Zombie procs	OK	11-13-2013 10:12:46	158d 20h 34m 57s	1/3	PROCS OK: 0 processes with STATE = Z
check network	OK	11-13-2013 10:14:58	6d 20h 39m 21s	1/3	DM-NET OK - eth2 (in: 5.00 KB/s out: 2.00 KB/s err: 0.00%)
dpm disk activity	OK	11-13-2013 10:14:58	6d 20h 39m 21s	1/3	DM-PART OK - sdb1 (r: 0 KB/s, w: 0 KB/s) sdb2 (r: 0 KB/s, w: 0 KB/s)
dpm-gsftd	OK	11-13-2013 10:13:58	43d 0h 58m 8s	1/3	OK - globus-gridftp-server (pid 5479) is running
gridftp activity	OK	11-13-2013 10:11:27	20d 18h 46m 11s	1/3	DM-GRIDFTP-ACTIVITY OK - 60 read operation (3225735.89B/s for 29075271262.00B), 97 write operation (11521173.72B/s for 1097616820.00B), 0
gridftp transfer	OK	11-13-2013 10:10:52	20d 18h 45m 21s	1/3	DM-GRIDFTP-TRANSFER OK - 0 % of the transfers have failed
ncm-cdispd	OK	11-13-2013 10:10:34	43d 0h 58m 8s	1/3	OK - ncm-cdispd (pid 5270) is running
rfio activity	OK	11-13-2013 10:11:55	20d 18h 47m 32s	1/3	DM-RFIO-ACTIVITY OK - 80 read operation (53536056.02B/s for 84801112732.00B), 0 write operation (0.00B/s for 0.00B),
rfiod	OK	11-13-2013 10:13:42	43d 0h 57m 42s	1/3	OK - rfiod (pid 5578) is running
temperature	OK	11-13-2013 10:12:08	43d 0h 56m 58s	1/3	check_temp.pl OK - 24.00 C
xrootd	OK	11-13-2013 10:10:34	43d 0h 58m 8s	1/3	OK - xrootd (pid 5444) is running



Gridftp read activity for long, medium and small transferred bytes



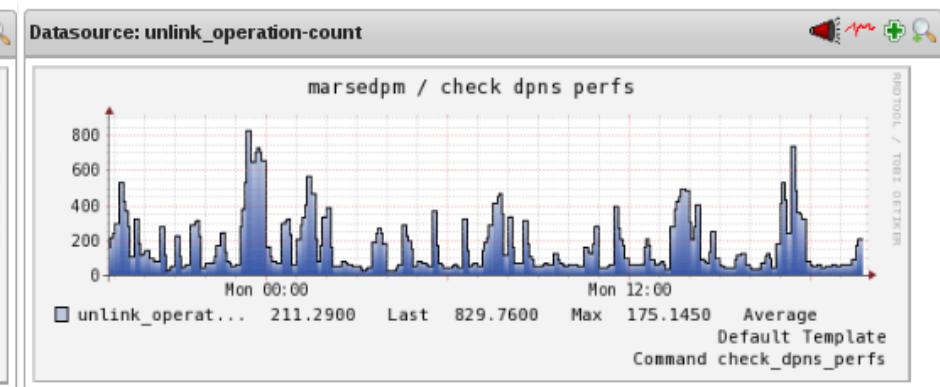
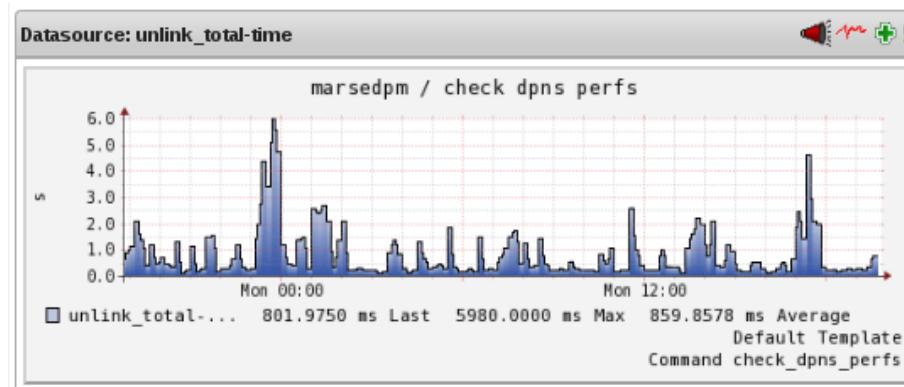
DPM / DPNS

Nagios-plugin-dpm-head

- check_dpm_infosys
- check_dpm_perfs
- check_dpns_perf (create, unlink, mkdir, opendir, rmdir)
- check_dpm_pools
- check_dpm_space_tokens

marsedpm				Check sda	OK	11-18-2013 19:08:05	26d 8h 14m 17s	1/3	SMART overall-health self-assessment test result: PASSED	
				DPM log		CRITICAL	11-18-2013 19:12:48	40d 15h 53m 54s	1/1	CRITICAL: (387): 11/18 19:12:28.862 4469,38 dpm_srv_proc_get: returns 0, status=DPM_FAILED (No such file or directory)
				Disk Usage Server	OK	11-18-2013 19:05:16	26d 8h 52m 39s	1/3	DISK OK - free space: / 13311 MB (72% inode=93%) : /var 35169 MB (30% inode=99%) : /boot 216 MB (93% inode=99%) :	
				check dpm pools	OK	11-18-2013 19:13:11	26d 3h 42m 12s	1/3	DM-POOL OK - All DPM pools are OK	
				check dpm tokens		CRITICAL	11-18-2013 19:13:11	0d 5h 32m 25s	3/3	DM-SPACE-TOKEN CRITICAL - 1 criticals::ATLASLOCALGROUPDISK and 0 warnings
				check dpns perfs	OK	11-18-2013 19:13:12	6d 5h 30m 30s	1/3	DM-DPNS-PERF OK - 2778 functions executed (18.91ms / functions)	
				dpm	OK	11-18-2013 19:09:47	26d 8h 45m 50s	1/3	OK - dpm (pid 4469) is running	
				dpm-qryconf		CRITICAL	11-18-2013 19:08:05	26d 3h 15m 5s	3/3	check_dpm-qryconf.pl CRITICAL - mardisk18.in2p3.fr:/data01 DISABLED ,mardisk18.in2p3.fr:/data02 DISABLED
				dponsdaemon	OK	11-18-2013 19:08:05	26d 8h 52m 16s	1/3	OK - dponsdaemon (pid 4391) is running	
				mysqld	OK	11-18-2013 19:08:05	26d 8h 53m 27s	1/3	OK - mysqld (pid 3905) is running	
				mysqld log	OK	11-18-2013 19:08:06	26d 8h 52m 39s	1/1	OK - No matches found.	
				ncm-cdispd	OK	11-18-2013 19:08:05	26d 8h 51m 50s	1/3	OK - ncm-cdispd (pid 3653) is running	
				rfiod	OK	11-18-2013 19:08:51	26d 8h 51m 1s	1/3	OK - rfiod (pid 4431) is running	
				srmv1	OK	11-18-2013 19:08:51	26d 8h 50m 13s	1/3	OK - srmv1 (pid 4533) is running	

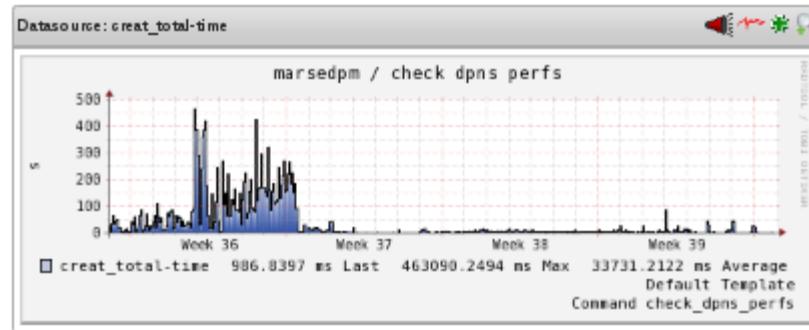
Performances du catalogue DPNS - Cas d'effacement de fichiers.



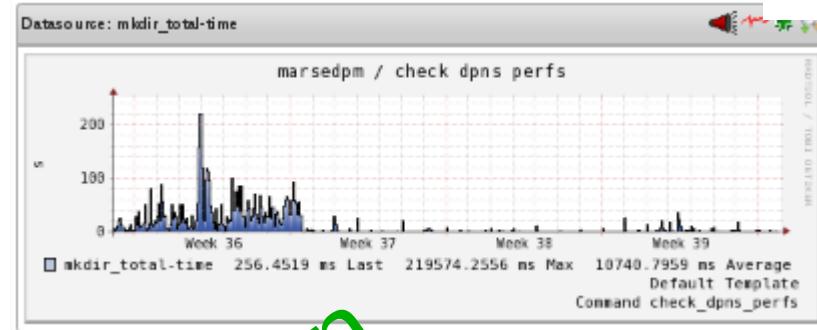
<https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/Monitoring/Probes>

<https://twiki.cern.ch/twiki/bin/view/EMI/EMINagios>

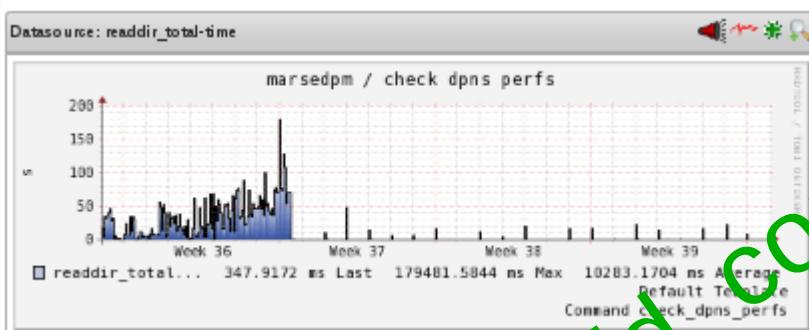
DPNS - mysql configuration



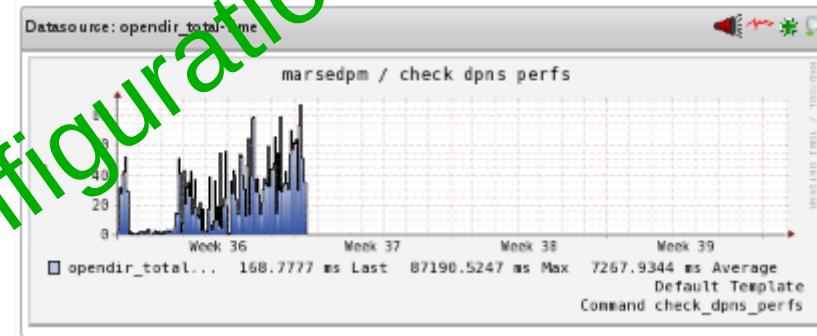
Create total time ms



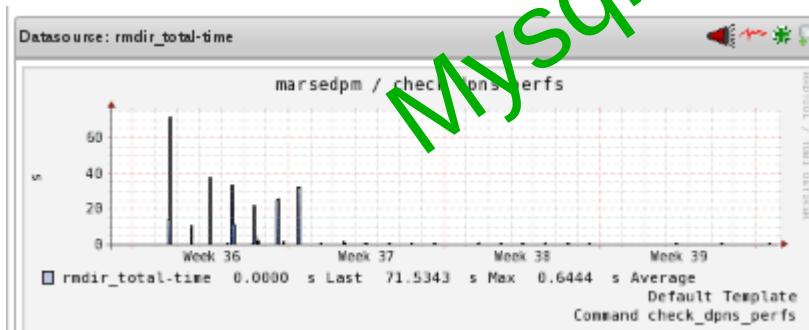
Mkdir total time ms



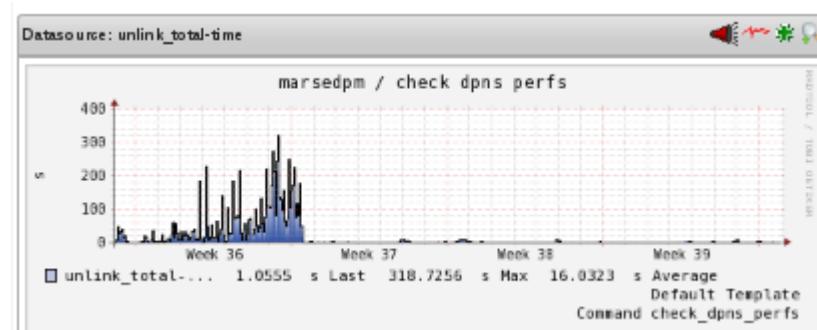
Readdir total time ms



Opendir total time ms



Rmdir total time ms



Unlink total time ms

CernVM-FS – check client



check_cvmfs_1.7.sh

usage :

check_cvmfs <repository> [version]

ex :

[root@marwn32 ~]#check_cvmfs.sh atlas.cern.ch 2.1.14
SERVICE STATUS: OK; repository revision 3647c



<http://cernvm.cern.ch/portal/filesystem/downloads>

cvmfs mounts	OK	10-28-2013 09:25:30	2d 0h 52m 20s	1/3	check_cvmfs.pl OK - atlas-condb.cern.ch : OK: repository revision 863, atlas.cern.ch : OK: repository revision 3536, lhcb.cern.ch : OK: repository revision 16254
--------------	----	---------------------	---------------	-----	--

Status plus courants

repository not available CRITICAL

Current Status: CRITICAL (for 0d 0h 3m 57s)
Status Information: check_cvmfs.pl CRITICAL - lhcb.cern.ch : failed to access lhcb.cern.ch

file .cvmfsdirtab is not readable CRITICAL

host/proxy is unavailable CRITICAL

I/O errors recorded by cvmfs WARNING

open file descriptors is greater than 80% of the max WARNING

resident memory is more than 50M WARNING

Current Status: WARNING (for 0d 0h 2m 25s)
Status Information: check_cvmfs.pl WARNING - atlas.cern.ch : high memory consumption (194m): repository revision 3666
The no. loaded nested catalogs is greater than 750 WARNING

The installed version and the loaded version mismatch WARNING

The installed version and the expected version mismatch WARNING

The cvmfs cache partition is filled more than 95% WARNING

CernVM-FS – Gestion cache



Problème de no libération de ressources

- Restauration de service manuel répétitive.
- Beaucoup de notifications.

Plusieurs solutions ...

Solution choisi :

- Plugin Nagios => seuils critical, warning.
- Event Handler => libération de ressources.
- Une stratégie pour minimiser les notifications.

marwn09		Check Dell OM	OK	11-10-2013 12:17:50	13d 3h 36m 45s	1/3	OK - System: 'C6100', SN: 'XXXXXXXX', 48 GB ram (6 dimms), not checking storage
		Check NFS /home/soft /cream_sandbox [rw]	OK	11-10-2013 12:20:09	1d 7h 24m 25s	1/3	check_mount.pl OK - /home/soft /cream_sandbox
		Check cvmfs quota	OK	11-10-2013 12:22:38	1d 9h 42m 32s	1/3	cvmfs_quota_usage.pl OK - 72% used by atlas.cern.ch
		Checksda	OK	11-10-2013 12:06:30	8d 17h 20m 30s	1/3	SMART overall-health self-assessment test result: PASSED

Host: marwn01 Service: Check cvmfs quota

One Year 10.10.12 14:03 - 25.10.13 14:03

