

dCache sensors & monitoring

A proposal to share sensors

Gerard.Bernabeu@pic.es

Functional check

- We rely on puppet for all server's setup
 - but PoolManager.conf, for that we use IN2P3 XML config generator
- Functional check always before/after updates
 - Minimalistic but very useful
- dCache update and basic verification in < 15 minutes (~80 servers, 5.7PB on disk)
 - Unless something goes wrong!
 - Still have to wait for pool initialization

Functional check config

Same script to verify 3 different instances
I believe it's easily adaptable to any dCache installation (improvements very welcome)

```
vo=`voms-proxy-info -vo`  
  
case $1 in  
    "prod")  
        srm=srm.pic.es  
        dcap=dcap.pic.es  
        gsidcap=gsidcap.pic.es  
        gridftp=gridftp.pic.es  
        pnfsDestFile=pnfs/pic.es/data/$vo/FunctionalTest`date +%F-%H%M``$$.txt  
        ;;  
    disk)  
        srm=srm-disk.pic.es  
        dcap=gridftp-disk.pic.es  
        gsidcap=gridftp-disk.pic.es  
        gridftp=gridftp-disk.pic.es  
        pnfsDestFile=pnfs-disk/pic.es/$vo/FunctionalTest`date +%F-%H%M``$$.txt  
        ;;  
    *)  
        srm=srm-test.pic.es  
        dcap=dccgftp01.pic.es  
        gsidcap=dccgftp01.pic.es  
        gridftp=dccgftp01.pic.es  
        pnfsDestFile=pnfs/pic.es/data/$vo/FunctionalTest`date +%F-%H%M``$$.txt  
        ;;  
esac
```

Functional check at work

```
[bernabeu@ui02 ~]$ bash ./FunctionalTests/dCacheFunctionalTest.sh prod
Logging to /nfs/pic.es/user/b/bernabeu/logs/FunctionalTest2012-04-16-1426.txt.log
globus-url-copy -dbg file:///etc/group gsiftp://193.109.172.147:2811/pnfs/pic.es/data/dteam/FunctionalTest2012-04-16-1426.17233.txt.gftp3
globus-url-copy -dbg gsiftp://193.109.172.147:2811/pnfs/pic.es/data/dteam/FunctionalTest2012-04-16-1426.17233.txt.gftp3 file:///tmp/FunctionalTest2012-04-16-1426.txt.gftp3
Result (1s): 0
uberftp 193.109.172.147 rm pnfs/pic.es/data/dteam/FunctionalTest2012-04-16-1426.17233.txt.gftp3
....
....
....
srmls -2 srm://srm.pic.es:8443/pnfs/pic.es/data/dteam
Result (5s): 0
srm-advisory-delete --debug=true -2 srm://srm.pic.es:8443/pnfs/pic.es/data/dteam/FunctionalTest2012-04-16-1426.17233.txt.srmv2t1d0
Result (4s): 0
Everything is OK. 77 seconds elapsed.
[bernabeu@ui02 ~]$
```

dCache generic sensor

For each cell check status on the web interface (if exists) + listening ports + connection to main server +java procs

Service	Status	Last Check	Duration	Attempt	Status Information
BONDING_STATUS	OK	04-16-2012 15:55:05	2d 21h 13m 7s	1/1	[OK]: bond0 IEEE 802.3ad Dynamic link aggregation under Ethernet Channel Bonding Driver: v3.6.0 (September 26, 2009)
CLEAN_ENCP	OK	04-16-2012 16:57:14	3d 0h 40m 58s	1/3	OK: no zombie encp detected
DCACHEPIC	CRITICAL	04-16-2012 16:59:13	3d 1h 12m 59s	3/3	[CRITICAL] according to http://dccore-disk.pic.es:2288/cellinfo pool is Offline. [OK] Connection with dccore established: . java version "1.6.0_26"
DCACHEPOOLLOGPARSE	OK	04-16-2012 17:05:33	3d 0h 47m 39s	1/3	[OK] No dead error found in dCache's log: 13 Apr 2012 16:20:11 (dc026_1) [] Pool mode changed to enabled [OK] No dead error found in dCache's log: 13 Apr 2012 16:14:45 (dc026_2) [] Pool mode changed to enabled [OK] No dead error found in dCache's log: [OK] No dead error found in dCache's log: 13 Apr 2012 16:14:45 (dc026_4) [] Pool mode changed to enabled
DISK_FREE_TMP	OK	04-16-2012 17:03:42	3d 0h 4m 30s	1/3	[OK]: Free space on /tmp = 186849896 Kbytes
DISK_FREE_VAR	OK	04-16-2012 17:05:52	3d 0h 2m 20s	1/3	[OK]: Free space on /var = 186849892 Kbytes
ENSTORE_CONFIG	OK	04-16-2012 15:37:53	13d 1h 45m 53s	1/2	Instance=disk No Enstore setup
HARDWARE_COMMON	CRITICAL	04-16-2012 16:55:02	3d 1h 23m 10s	3/3	[CRITICAL] detected 16/36 in cntrI2 : [CRITICAL] Some logical drives degraded for Adaptec controller 2: [CRITICAL] Less than 3 logical drives detected for Adaptec controller 2: [WARNING]: Adaptec Controller 2 with some cache Disabled: : [OK] PS: [OK] System Temp at 32C : [OK] md4 (2/2/2) : [OK] md3 (2/2/2) : [OK] md2 (2/2/2) : [OK] md5 (2/2/2) : [OK] md1 (2/2/2) : [OK] md0 (2/2/2) : [OK] md6 (2/2/2):[OK] Adaptec6445-1 (5.2-0(18301)) 45 disks:[OK] Adaptec6445-2 (5.2-0(18301)) 16 disks:[OK] JBOD PS (2/2)

dCache generic sensor config

Same (dynamic) sensor for different server profiles
(SRM, pool, etc.).

```
instance=`grep instance /opt/PICnodeInfo | cut -d= -f2`  
node_type=`grep node_type /opt/PICnodeInfo | cut -d= -f2`  
dCacheVersion=`grep dCacheVersion /opt/PICnodeInfo | cut -d= -f2`  
dcachelayout=/opt/d-cache/etc/layouts/`grep dcache.layout /opt/d-cache/etc/dcache.conf | cut -d= -f2`.conf  
  
case $instance in  
    prod )  
        dcip=dcip.pic.es #dCache information provider server  
        dccore="193.109.174.39"  
        ;;  
    test )  
        dcip=dcip-test.pic.es #dCache information provider server  
        dccore="193.109.174.41"  
        ;;  
    disk )  
        dcip=dccore-disk.pic.es #dCache information provider server  
        dccore="193.109.175.249"  
        ;;  
    *)  
        echo "Unknown instance $instance. Check puppet maintained /opt/PICnodeInfo file"  
        exit 2  
        ;;  
esac
```

More specific sensors

- On pools: parse specific pool log errors, mounted PNFS, enstore config, zombie encp
- On doors: parse gridFTP logs for errors, certs&CA

Host ▾	Service ▾	Status ▾	Last Check ▾	Duration ▾	Attempt ▾	Status Information		
dc103.pic.es	BONDING_STATUS	OK	04-16-2012 15:02:35	0d 2h 15m 28s	1/1	[OK]: bond0 fault-tolerance (active-backup) under Ethernet Channel Bonding Driver: v3.6.0 (September 26, 2009)	<input type="checkbox"/>	
	CLEAN_ENCP	OK	04-16-2012 17:13:12	4d 22h 4m 51s	1/3	OK: no zombie encp detected	<input type="checkbox"/>	
	DCACHEPIC	OK	04-16-2012 17:17:35	4d 6h 15m 28s	1/3	[OK] pool. [OK] Connection with dccore established: . java version "1.6.0_26"	<input type="checkbox"/>	
	DCACHEPOOLLOGPARSE	OK	04-16-2012 17:09:31	0d 2h 23m 32s	1/3	[OK] No dead error found in dCache's log: 16 Apr 2012 14:41:00 (dc103_1) [] Pool mode changed to enabled	<input type="checkbox"/>	
	DISK_FREE_TMP	OK	04-16-2012 17:09:41	5d 0h 8m 22s	1/3	[OK]: Free space on /tmp = 449710696 Kbytes	<input type="checkbox"/>	
	DISK_FREE_VAR	OK	04-16-2012 16:51:35	4d 21h 26m 28s	1/3	[OK]: Free space on /var = 449710768 Kbytes	<input type="checkbox"/>	
	ENSTORE_CONFIG	OK	04-16-2012 16:43:45	4d 20h 34m 18s	1/2	[OK] Enstore configuration in status OK. v3_6d. y-pic2	<input type="checkbox"/>	
	HARDWARE_COMMON	OK	04-16-2012 16:55:54	4d 21h 22m 9s	1/3	[WARNING] FANs not monitored: [OK] PS: [OK] System Temp at 24C : [OK] md1 (2/2/2) : [OK] md0 (2/2/2) : [OK] MegaRAID Battery 56% : [OK] MegaRAID 36 online disks : [OK] MegaRAID 36 online disks	<input type="checkbox"/>	
	LOAD_AVERAGE	OK	04-16-2012 17:03:04	4d 22h 44m 59s	1/3	OK - load average: 0.00, 0.03, 0.00	<input type="checkbox"/>	
	PING	OK	04-16-2012 17:14:14	0d 2h 33m 49s	1/3	PING OK - Packet loss = 0%, RTA = 0.41 ms	<input type="checkbox"/>	
	PNFS_MOUNTED	OK	04-16-2012 16:01:30	0d 1h 16m 33s	1/2	[OK] PNFS dcns.pic.es mounted 1 time at this pool	<input type="checkbox"/>	
	PUPPET_LOG		OK	04-16-2012 16:59:07	1d 3h 18m 56s	1/3	Mon Apr 16 16:17:31 +0200 2012 Puppet (notice): Finished catalog run in 18.88 seconds	<input type="checkbox"/>
	SSH	OK	04-16-2012 17:13:43	4d 22h 19m 20s	1/3	SSH OK - OpenSSH_5.3 (protocol 2.0)	<input type="checkbox"/>	
	SWAP	OK	04-16-2012 17:08:53	0d 2h 9m 10s	1/3	[OK]: Swap 0% [Swap-io:0 cpu-user:100]	<input type="checkbox"/>	
	VAR_LOG	OK	04-16-2012 17:06:03	4d 22h 12m 0s	1/3	/var/log is OK	<input type="checkbox"/>	

Misc monitoring

Check enough freespace, files properly landing to Enstore, gridftp functional check, queued movers

Host ▾	Service ▾	Status ▾	Last Check ▾	Duration ▾	Attempt ▾	Status Information	☒	
dcmon01.pic.es	BACULA-FD	OK	04-16-2012 12:05:10	27d 11h 13m 4s	1/5	SNMP OK - 0 "[OK]: Bacula file daemon OK"	<input type="checkbox"/>	
	DCACHE_CHECKMIGRATED	OK	04-16-2012 08:00:00	13d 1h 47m 7s	1/3	SNMP OK - 0 "[OK] no pool migration problem detected"	<input type="checkbox"/>	
	DCACHE_FREESPACE	OK	04-16-2012 16:39:21	13d 2h 50m 48s	1/3	SNMP OK - 0 "[OK] Freespace is ,linkgroup-magic:56474957941734,linkgroup-pau:32970090886610,linkgroup-turbo:858993459200,linkgroup-cms:1628605229277308,linkgroup-atlas:978793632180700,linkgroup-lcb:182473937854184,linkgroup-t2k:1305829858588,linkgroup-crg:36005022175232,linkgroup-dteam:23164723485217,linkgroup-cta:6793187039101,linkgroup-archive:2104637713909"	<input type="checkbox"/>	
	DCACHE_QUEUEDMOVERS	OK	04-16-2012 16:41:30	5d 9h 18m 2s	1/3	SNMP OK - 0 "[OK] No movers queued at dCache instance=prod."	<input type="checkbox"/>	
	DISK_FREE_TMP	OK	04-16-2012 17:11:20	5d 5h 21m 27s	1/3	SNMP OK - 0 "[OK]: Used space on /tmp = 51 %"	<input type="checkbox"/>	
	DISK_FREE_VAR	OK	04-16-2012 17:13:29	5d 5h 19m 18s	1/3	SNMP OK - 0 "[OK]: Used space on /var = 51 %"	<input type="checkbox"/>	
	GRIDFTP_SRM_MONITORING	UNKNOWN	04-16-2012 17:03:39	3d 21h 9m 8s	3/3	(Service Check Timed Out)	<input type="checkbox"/>	
	LOAD_AVERAGE	OK	04-16-2012 17:30:37	5d 5h 17m 10s	1/3	SNMP OK - 0 0.12	<input type="checkbox"/>	
	PING	OK	04-16-2012 17:27:46	5d 5h 15m 1s	1/3	PING OK - Packet loss = 0%, RTA = 0.45 ms	<input type="checkbox"/>	
	PUPPET_LOG		OK	04-16-2012 16:59:54	1d 5h 32m 53s	1/3	SNMP OK - 0 "Mon Apr 16 16:11:10 +0200 2012 Puppet (notice): Finished catalog run in 54.28 seconds"	<input type="checkbox"/>
	SLS_RUNNING	OK	04-16-2012 16:45:48	13d 2h 44m 21s	1/3	SNMP OK - 0 "[OK] SLS metrics is running OK"	<input type="checkbox"/>	
	SSH	OK	04-16-2012 17:22:03	5d 5h 10m 44s	1/3	SSH OK - OpenSSH_4.3 (protocol 2.0)	<input type="checkbox"/>	
	SWAP	OK	04-16-2012 17:09:12	3d 20h 53m 35s	1/3	SNMP OK - 0 "[OK]: Swap 0% [Swap-io:0 cpu-user:100]"	<input type="checkbox"/>	
	TAPE_TOKENS_SPACE	OK	04-16-2012 15:47:58	23d 15h 49m 54s	1/2	SNMP OK - 0 "[OK] Freespace is OK"	<input type="checkbox"/>	

What about sharing nagios sensors?

Anyone interested?

I'm interested in your sensors :)

dCache sensors in a common repository?

<https://github.com/gerardba/dCacheProbes>

Should be easy to separate site-dependant config in a file...

We also have some ganglia ad-hoc graphs (ie: each pool plotting their mover queues, JVM metrics) which rely on dCache web interface.