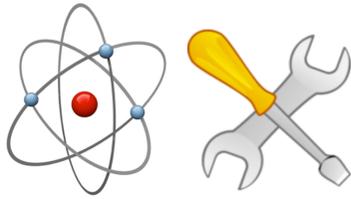


Development and deployment of a T2-oriented monitoring infrastructure

SPINOSO Vincenzo, DONVITO Giacinto (INFN Bari)
MAGGI Giorgio Pietro (INFN e Politecnico di Bari)



LHC physics user groups will need a simple and intuitive *user-level summary of their T2 services' status, showing available, busy, unavailable resources...*

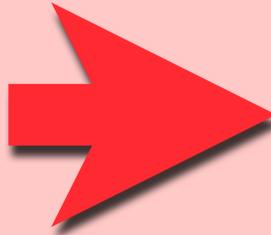


Sysadmins need to monitor (and to be automatically notified about) a huge amount of *technical-level aspects of their farms, in order to guarantee available/reliable services...*

We need a cluster monitoring tool which is able to rule all of those needs

From a "raw solution" ...

Often we discover something new (and interesting) to monitor, and often no existing monitoring tool is able to monitor it (manufacturer specific monitoring tools, new parameter/feature to monitor...) Writing Bash scripts by hand to monitor events and send back notification via mail/web is our "patch" to this problem.



... to a "new monitoring framework"!

Why not collect all those scripts under a common framework, simple to manage and to deploy?
Why not use SSH (already used for administrative purposes) to transport messages from sensors to the information collector?
Why not use databases to store information and build statistics?

Mon2 is a monitoring framework, based on Perl/Bash scripts and a DB

- every host is monitored by daemonized Bash scripts, downloaded from the Mon2 server
- every 3 minutes each host collects information about itself and uploads them on the central server;
- every 2 minutes the Mon2 server gathers all **host information** with summary information about the status of high level services of the farm, then publishes them on the Web
- all interesting information are stored into a **local database**, so all needed history is available for future investigations
- sends event-driven **alerts** to sysadmins through emails, RSS, SMS
- Mon2 can monitor **what both users and sysadmins need**: temperatures, PBS job monitoring, central availability tests (SAM tests, Job Robot, HLR...), OS (CPU/filesystems/network), RAID devices health status, system clock offsets
- Mon2 is **flexible**: all you need to create/modify sensors is Bash, all you need to create/modify Web pages and alerts is Perl
- Mon2 is **portable**: it is based on standard Unix/Linux tools (bash, perl, rsync, ssh, cron)
- Mon2 is **secure**: all host-to-server connections are SSH-based, and no need to login as root using root-to-host connections; no server-side scripting (HTML pages are generated in batch mode using Perl and gnuplot); no DB server is used (we use SQLite)

PBS monitoring



We need weekly/monthly/yearly views of PBS jobs...

Storage

Hostname	Partition	Used(MB)	Total(MB)	Used(%)
gridse03.ba.infn.it	/data2	2721	2793	97%
gridse03.ba.infn.it	/data1	1325	1367	96%
gridse02.ba.infn.it	/data3	1294	1353	95%
gridse03.ba.infn.it	/data3	1348	1426	94%
gridse01.ba.infn.it	/data2	1030	1117	92%
gridse01.ba.infn.it	/data3	1031	1117	92%
gridse02.ba.infn.it	/data2	1272	1377	92%
gridse01.ba.infn.it	/data1	1008	1117	90%
gridse02.ba.infn.it	/data1	1233	1375	89%
cmshome				
gridse05				
gridba6.ba.infn.it	/data/1	751	929	80%
gridse04.ba.infn.it	/sdb1	7780	9778	79%
fermidisk-01.ba.infn.it	/data1	1536	1955	78%

How many TB are available now?

gridse04.ba.infn.it RAID devices

Ctl	Model	(F)Ports	Drives	Units	RotOpt	EBate	VBate	BBU
c0	8006-2LP	2	2	1	1	2	-	-
c1	9650SE-16HL	16	16	1	0	4	5	-

Unit	UnitType	Status	KB/Cpl	%I/M	Stripe	Size(GB)	Cache	ATrify
u0	RAID	DEGRADED	-	-	-	232.885	OFF	-

Port	Status	Unit	Size	Blocks	Serial
p0	DEGRADED	u0	232.885 GB	488397168	9RY00ELH
p1	OK				

Unit	UnitType	Status	Size	Blocks	Serial	ATrify	
u0	RAID-6	OK	-	64K	9778.74	UM	OFF

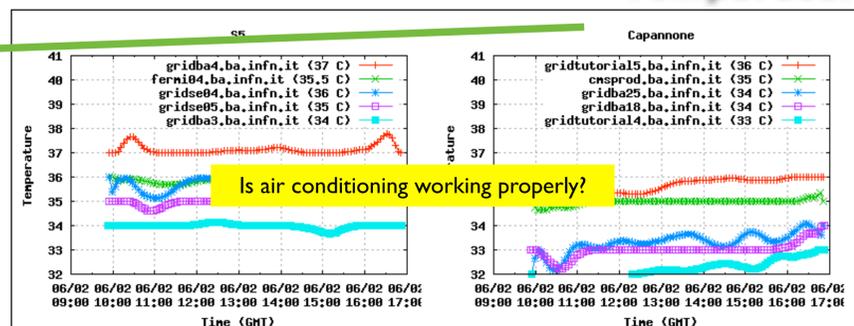
Port	Status	Unit	Size	Blocks	Serial
p0	OK	u0	698.63 GB	1465149168	SQD0HEJW
p1	OK	u0	698.63 GB	1465149168	SQD0M4DY
p2	OK	u0	698.63 GB	1465149168	SQD039TC

this RAID5 is degraded and needs sysadmin intervention!

Alerts

Item	Description	Where	Creation
Timeout	alicegrid15.ba.infn.it providing no fresh info. Is it alive?	alicegrid15.ba.infn.it	2009-2-6 5:27
Timeout	gridba13.ba.infn.it providing no fresh info. Is it alive?	gridba13.ba.infn.it	2009-2-1 4:54
Timeout	cream-ce-4.ba.infn.it providing no fresh info. Is it alive?	cream-ce-4.ba.infn.it	2009-1-30 10:35
Timeout	cream-ce-3.ba.infn.it providing no fresh info. Is it alive?	cream-ce-3.ba.infn.it	2009-1-30 10:29
Device errors	Some devices with problems on gridse04.ba.infn.it, please read http://cmsprod.ba.infn.it/mon2/devices.html#gridse04.ba.infn.it	gridse04.ba.infn.it	2009-1-2 12:14

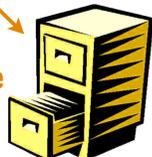
Temperature



Send notifications



Store to database



CPU/fs/net

Hostname	IP	IP (hex)	MAC	Ping?	SSH?
alicegrid08.ba.infn.it	212.189.205.10	D4BDCDDA			
alicegrid09.ba.infn.it	212.189.205.11	D4BDCD08			
alicegrid10.ba.infn.it	212.189.205.12	D4BDCDOC			
alicegrid11.ba.infn.it	212.189.205.80	D4BDCD50	00:30:48:79:B7:6A		
alicegrid12.ba.infn.it	212.189.205.81	D4BDCD51	00:30:48:79:DD:3A		

What's wrong with host.ba.infn.it???

Hostname	Load/Core	Uptime	ExpCert	MAC	eth
alicegrid11.ba.infn.it	1.3%	2040457	0	00:30:48:79:B7:6A	eth0
alicegrid12.ba.infn.it	15%	2040468	0	00:30:48:79:DD:3A	eth0
alicegrid13.ba.infn.it	12%	2040452	0	00:30:48:79:B9:06	eth0
alicegrid14.ba.infn.it	14%	2040446	0	00:30:48:79:B7:92	eth0

Coming soon:

- complete inventory of the farm to fulfil other goals:
- how many **CPU/cores/TB** are **supposed to be running** on our farm (pledged resources) ?
 - how many **VM** do we have? Where are they?
 - where is "host.ba.infn.it"? Can you show me it using a **graphical representation**?