

DAQ LHC Workshop Monitoring

Christophe Haen
&

Sergio Ballestrero, Olivier Chaze, Lavinia Darlea, Olivier
Raginel, Diana Scannicchio, Adriana Telesca

14th March 2013

Monitoring?

Why?


- To make sure that everything is working
- To see how the performances change over time
- To correlate problems

What?

- Data collection (and its distribution/load balancing/storage)
- Visualization of collected performance / health data
- Alert triggering on collected data

Good bye

Tools that will disappear





Lemon

- Developed at CERN
- Provides data collection, alerting and performances visualization
- Currently used by ALICE

Why replacing it?

- I.T. will drop the support
- ALICE made a lot of custom changes

Nagios

- Quasi open source industry standard
- Main purposes : collecting & alerting
- Was used by CMS and LHCb as a single instance. ATLAS still uses it as an aggregation of many instances

Why replacing it?

- Satisfying in many features but...
- Lack of performances
- Slow development, because not so open to the community
- Some features are only in commercial version
- a lot of in-house improvements (e.g. done by ATLAS) are now available through new dedicated tools

The new tools

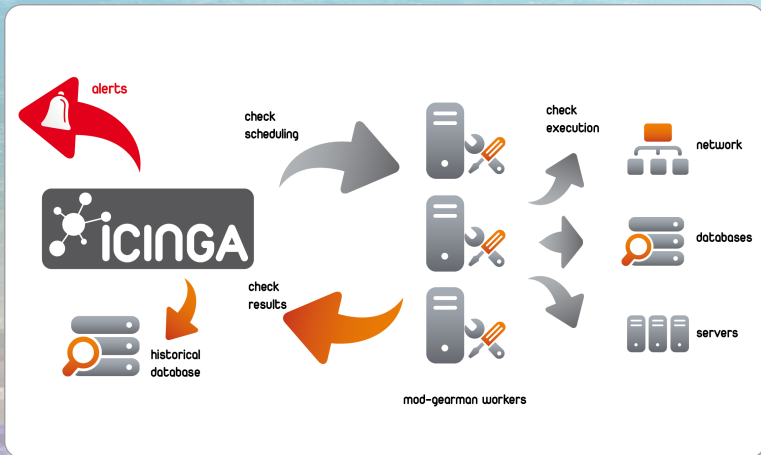
Welcome

Icinga

- A fork of Nagios
- Very strong support and community
- Very modular and many plugins available

Who?

- CMS and LHCb already for 2 years
- ATLAS in a near future to replace Nagios
- CMS uses a plugin for performance graphs (PnP4Nagios)



Ganglia

- Collects and plots graphs (RRDFFiles)
- No alerting
- Very scalable because of a 'tree-like' structure
- Some redundancy possibilities thanks to multicast addressing
- Customizable web interface with advanced comparison features

Who?

- ATLAS has made long duration tests over 300 hosts. They will use it as data collector and graphing also for Icinga
- LHCb has tested it for a shorter time but over 1500 hosts
- Both are happy and will use it

Zabbix

- All in one solution
- Collection, presentation, performances graphs, reporting, discovery...
- Very scalable
- Very extendable

Who

- ALICE
- Has been chosen after careful evaluation of many alternatives by Adriana. (see backup slides, or even better, her :-))
- Only used for performance data collection and visualization

Orthos

- Developed for and by ALICE
- Alarm/triggering and issues follow-up
- Notifying the expert and/or opening a JIRA ticket
- Zabbix will feed Orthos.

Maybe

Will be investigated during LS1

Shinken

- Fairly new but impressively growing community
- Uses and extends the philosophy of Nagios/Icinga...
- ... but with a completely new technical design
- Icinga being reshaped according to similar design, Nagios follows the ideas

Why?

Addresses some of the flexibility problems of Icinga/Nagios => LHCb will have a look

Technical considerations



How do we get the information?

Fetching the information

- SNMP (query or trap)
- NRPE (Nagios/Icinga)
- IPMI (we are all fairly unhappy with this)
- Ping
- Local agents (Ganglia, Zabbix)
- Push data to passive listener (Ganglia gmetrics, Icinga NSCA)
- Usage of 'check aggregator' like check_multi

=> Many options for many situations

Configuration management

How do we generate configuration?

- ALICE : Zabbix API used to change the configuration according to the changes in the configuration database
- ATLAS : custom tool ConfDb
- CMS : twiki page description + quattor profiles + perl scripts
- LHCb : clever configuration schema + set of scripts

=> We did not yet converge on that part because...

- The externally available config tools are limited
- We need to integrate with other custom tools / data sources

Conclusion

Tools exist...

- Do not reinvent the wheel!
- Tools now exist outside, and at bigger scale
- HEP has less and less specificities regarding monitoring


... BUT

- No "turnkey" solution
- Monitoring still requires considerable efforts for customising and integrating

Share!

Keep sharing between experiments, it works!

Questions





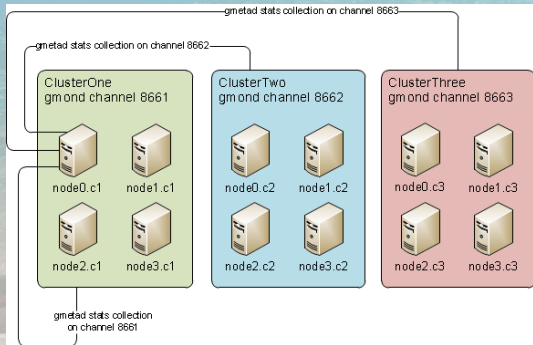
Backup

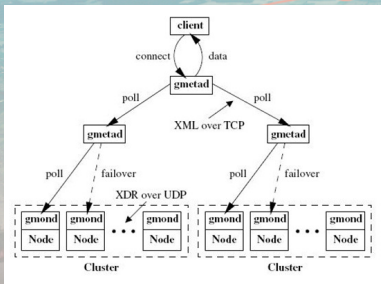
Comparison Adriana

Name	Data gathering	Graphing (0-2)	Alerting (0-1)	Triggering (0-1)	Distributed monitoring (0-1)	Scalability (0-2)	Data Storage	Flexibility / Extensibility (0-2)
Icinga	Agent	0	1	1	1	1 - up to 1000 hosts	DB	2
Cacti	Server	2	1	0	1	1 - up to 1000 hosts	RRDtool - DB	2
Zenoss	Server	1	1	1	1	2 - 1000+ hosts	RRDtool - DB	1
Zabbix	Agent or Server	2	1	1	1	2 - 1000+ hosts	DB	2
Splunk	Agent	2	1	1	1	2 - 1000+ hosts	raw files	2
MonALISA	Agent	2	1	1	1	2 - 1000+ hosts	DB	2

Comparison Adriana

Name	SNMP (0-2)	Web Interface (0-2)	Documentation / User Community (0-2)	Access control (0-1)	Maximum Granularity	Auto discovery (0-2)	Open source	Total
Icinga	1	2	2	1	1 minute / metric	1	1	14
Cacti	2	2	2	1	1 minute / metric	1	1	16
Zenoss	1	2	1	1	1 minute / collector	2	1	14
Zabbix	2	2	2	1	No limit / metric	2	1	19
Splunk	2	2	2	1	No limit / metric	2	0	18
MonALISA	2	2	1	1	1 minute / metric	2	1	18





Enter host regular expression:

