Evolution of the Monitoring in the LHCb Online System

Christophe Haen
E.Bonaccorsi and N.Neufeld

LHCb Online team (CERN), GENEVA, CH

10th October 2013

Plan

The LHCb Online System

LHCb

- One of the four large experiments of the LHC.
- Relies on large and heterogeneous IT infrastructure.
- Thousands of servers, different hardware configurations, great variety of tasks

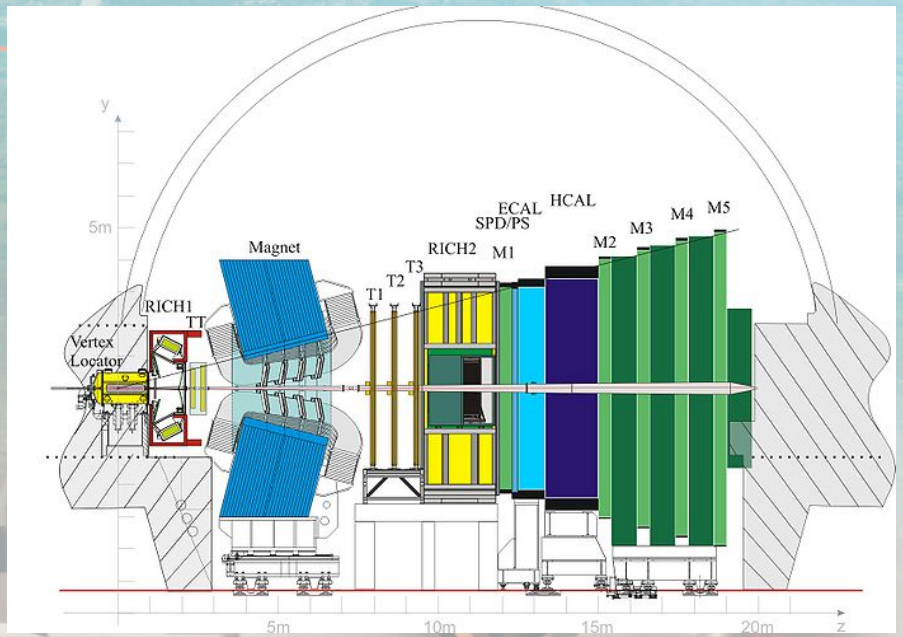

The LHCb Online System

Distributed monitoring infrastructure

- Single Icinga 1.8.4 instance
- ido2db with local MySQL (SSD disks)
- mod_gearman 1.4.2
- NRPE and NSClient++
- Nand for mail aggregation

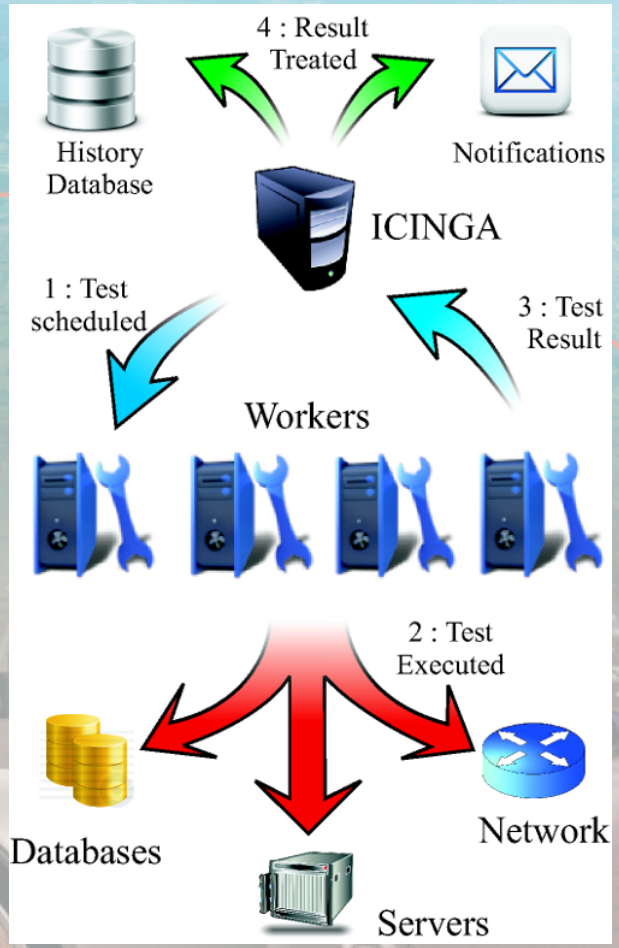

Positive aspects

Pros

- Performance: 40 000 checks in a 5 mn window without latency
- Ease of scalability with mod_gearman
- Group and template functionalities of the configuration: factorization
- New web interface
- Mail aggregation is good and necessary

Negative aspects

Cons

- Icinga instance is a single point of failure
- Dependency system unsatisfactory
- Performance with big environment failures
- Very static: no easy access to live information, no configuration change while running
- (Configuration parsing and loading time in the database)

Nagios4

Nagios4

- Major release, currently in beta version
- Performance improvements
- Better algorithms
- Give up fork system for worker processes (mod_gearman like)
- Claim: -87% iops, -42% CPU, -64% memory
- Configuration logic slightly changed (Beware!)

Shinken

Shinken (part 1)

- Pioneer of the next generation tools
- All in Python
- Extends Nagios' philosophy
- Innovative technical design

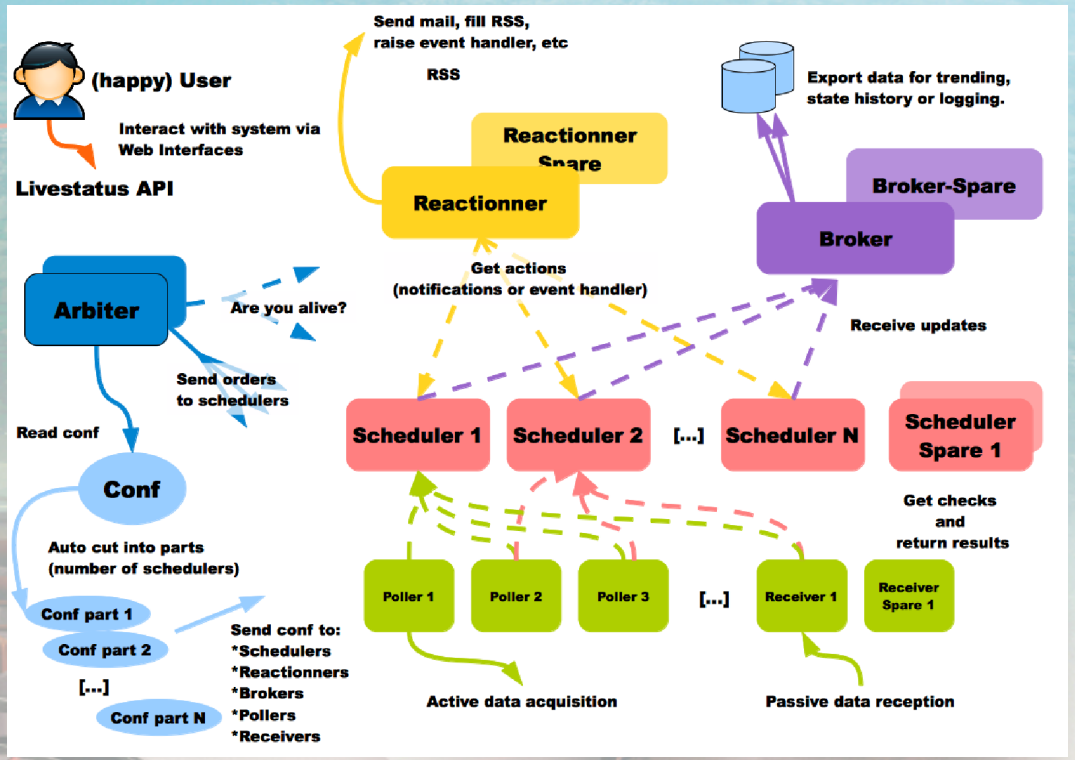

Shinken

Shinken (part 2)

- Dynamic ("calculated" checks, cluster support, virtualization, etc)
- Extends Nagios' configuration (services applied to templates, composition of templates, macro "foreach" etc)
- Automatic configuration generation
- Business oriented

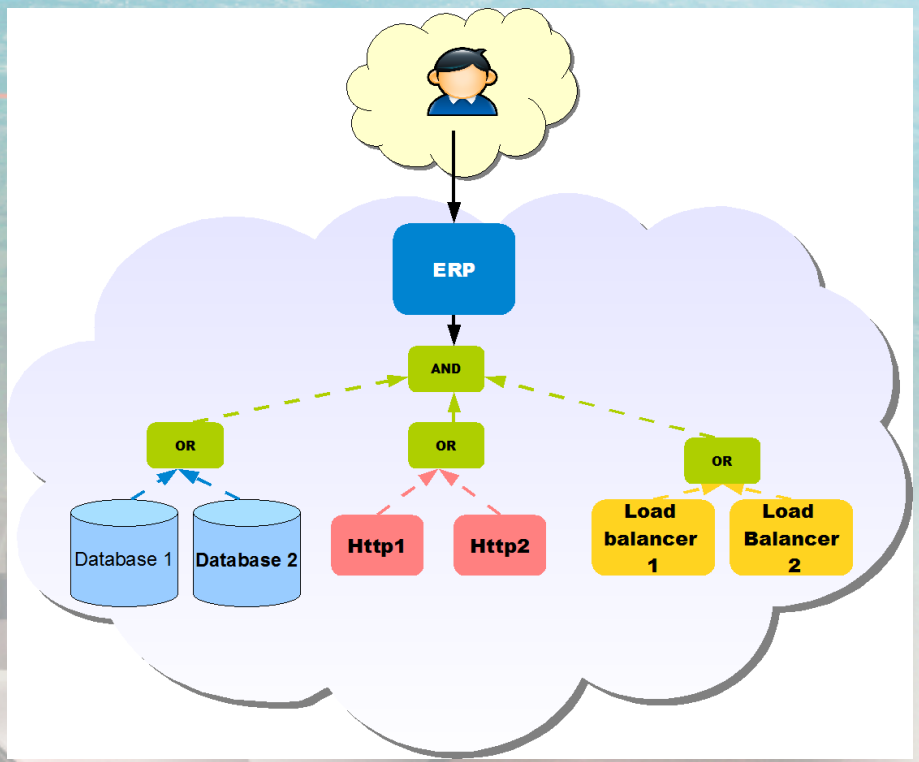

Icinga2

Icinga2

- Early development stages
- Separate branch from Icinga 1.x
- C++ with lots of Boost
- Distributed core
- Totally different configuration
- Remote agent
- Dynamic
- Business oriented

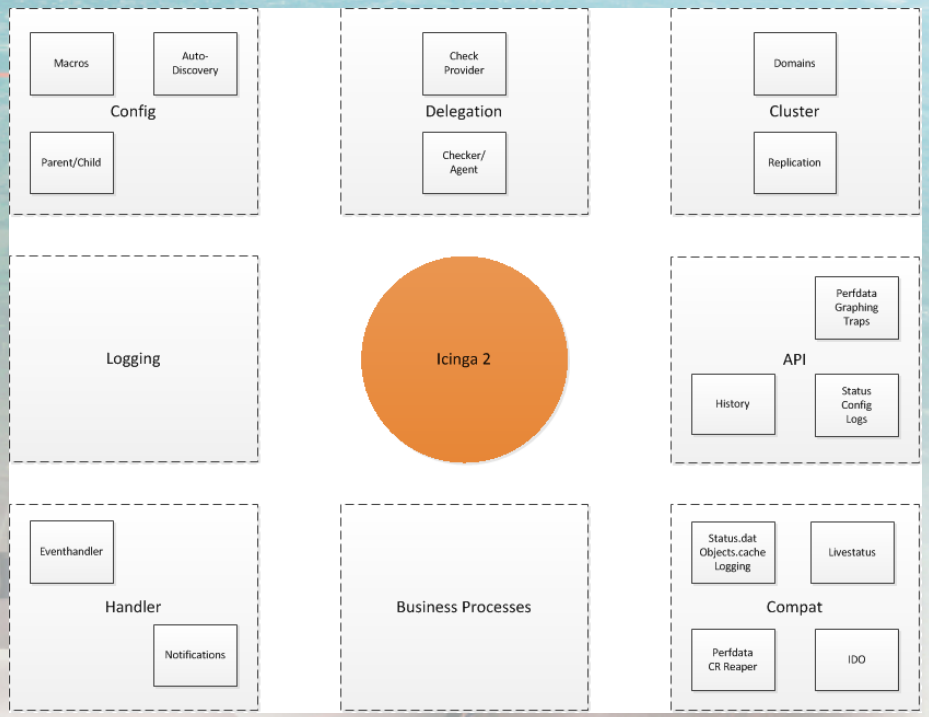

Test bench

Candidates

- Tunned Icinga + 15 remote gearman workers
- Shinken (slightly tuned)
- Out of the box Icinga2
- Out of the box nagios4

Procedure

- 60 000 services on 2 000 hosts
- No historical data
- t=0: everything OK
- t=1000s: 90% services fail
- t=2000s: everything recovers

Latency

Latency

- Icinga_gearman: knock on effect
- Shinken: increase when big failures
- nagios4: flat everywhere
- Icinga2: bump at the beginning

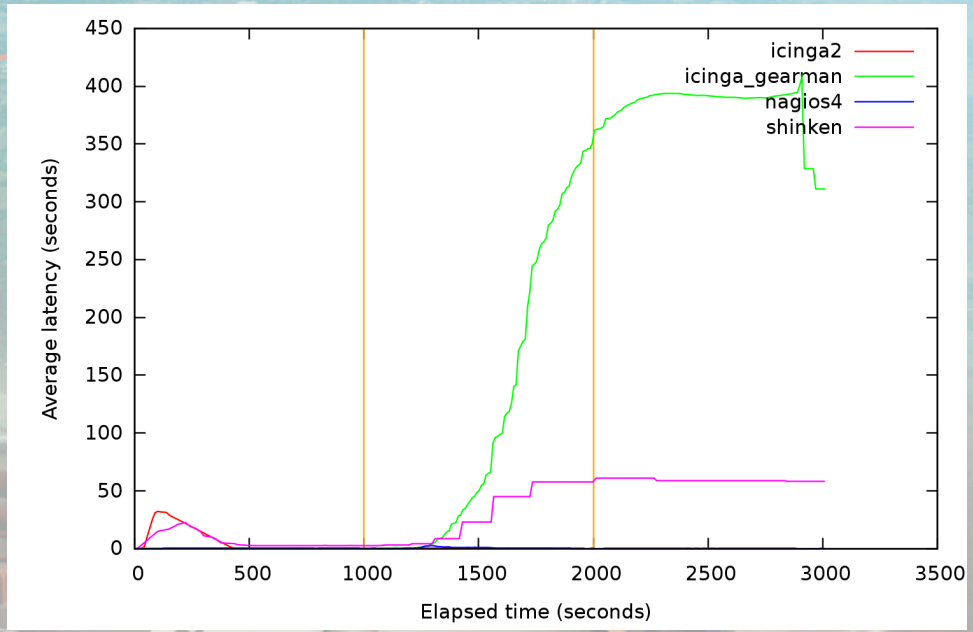

Service checks

Service checks

- Icinga_gearman: slow increase at 1000s
- Shinken: step increase
- nagios4: steep increase after 1000s
- Icinga2: very fast at startup

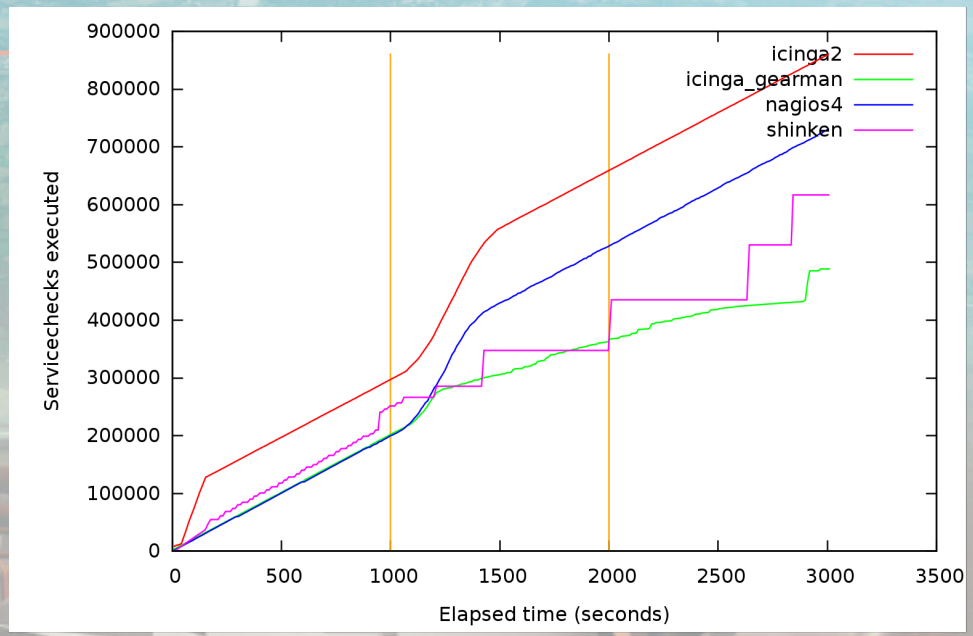

Reaction time

Reaction time

- Icinga_gearman: relatively slow
- Shinken: step function, slow
- nagios4: fast
- Icinga2: very fast

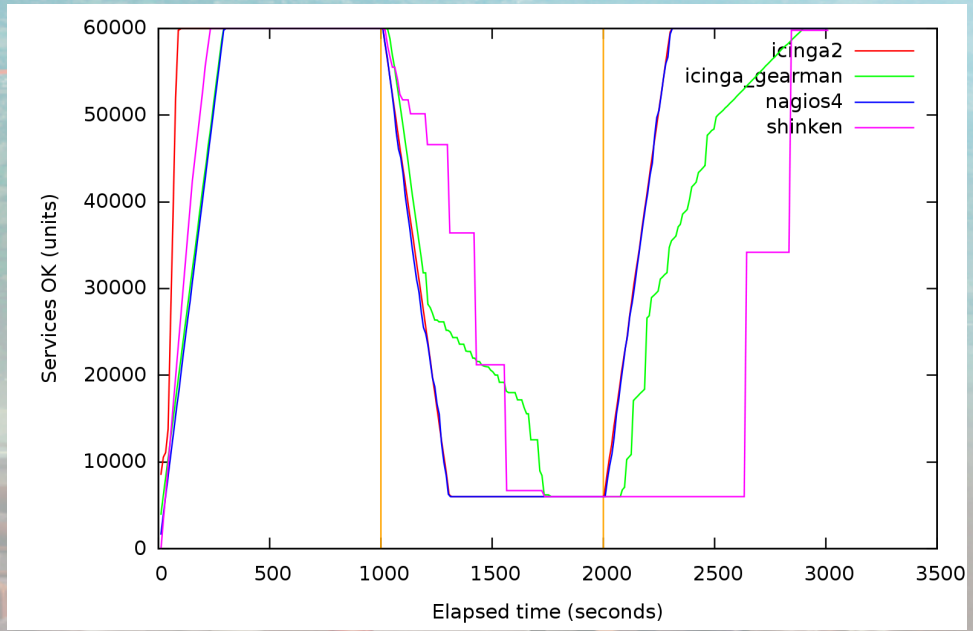

Conclusion

Conclusion

- Icinga with Gearman was a good move
- Still has some weaknesses
- Nagios4 is not an option
- Icinga2 extremely promising performance wise
- Shinken seems slower, but very dynamic and many features
- Further tests to be done when stable version for both