

Introduction

The ATLAS online farm, consisting of nearly 3000 PCs, must be continuously monitored to ensure the correct and optimal working conditions. The monitoring system should be able to check up to 100000 health parameters and provide alerts on a selected subset: the health status of the OS, hardware, selected services and network components. The current monitoring system is based on Nagios v2.5, the most used Open Source monitoring tool. We are evaluating Icinga, that is a fork of Nagios v3, with a very active community supporting it.

Current

A configuration database, complemented by a user interface, pushes the monitoring configuration respectively on

Special server nodes that act as Nagios servers.

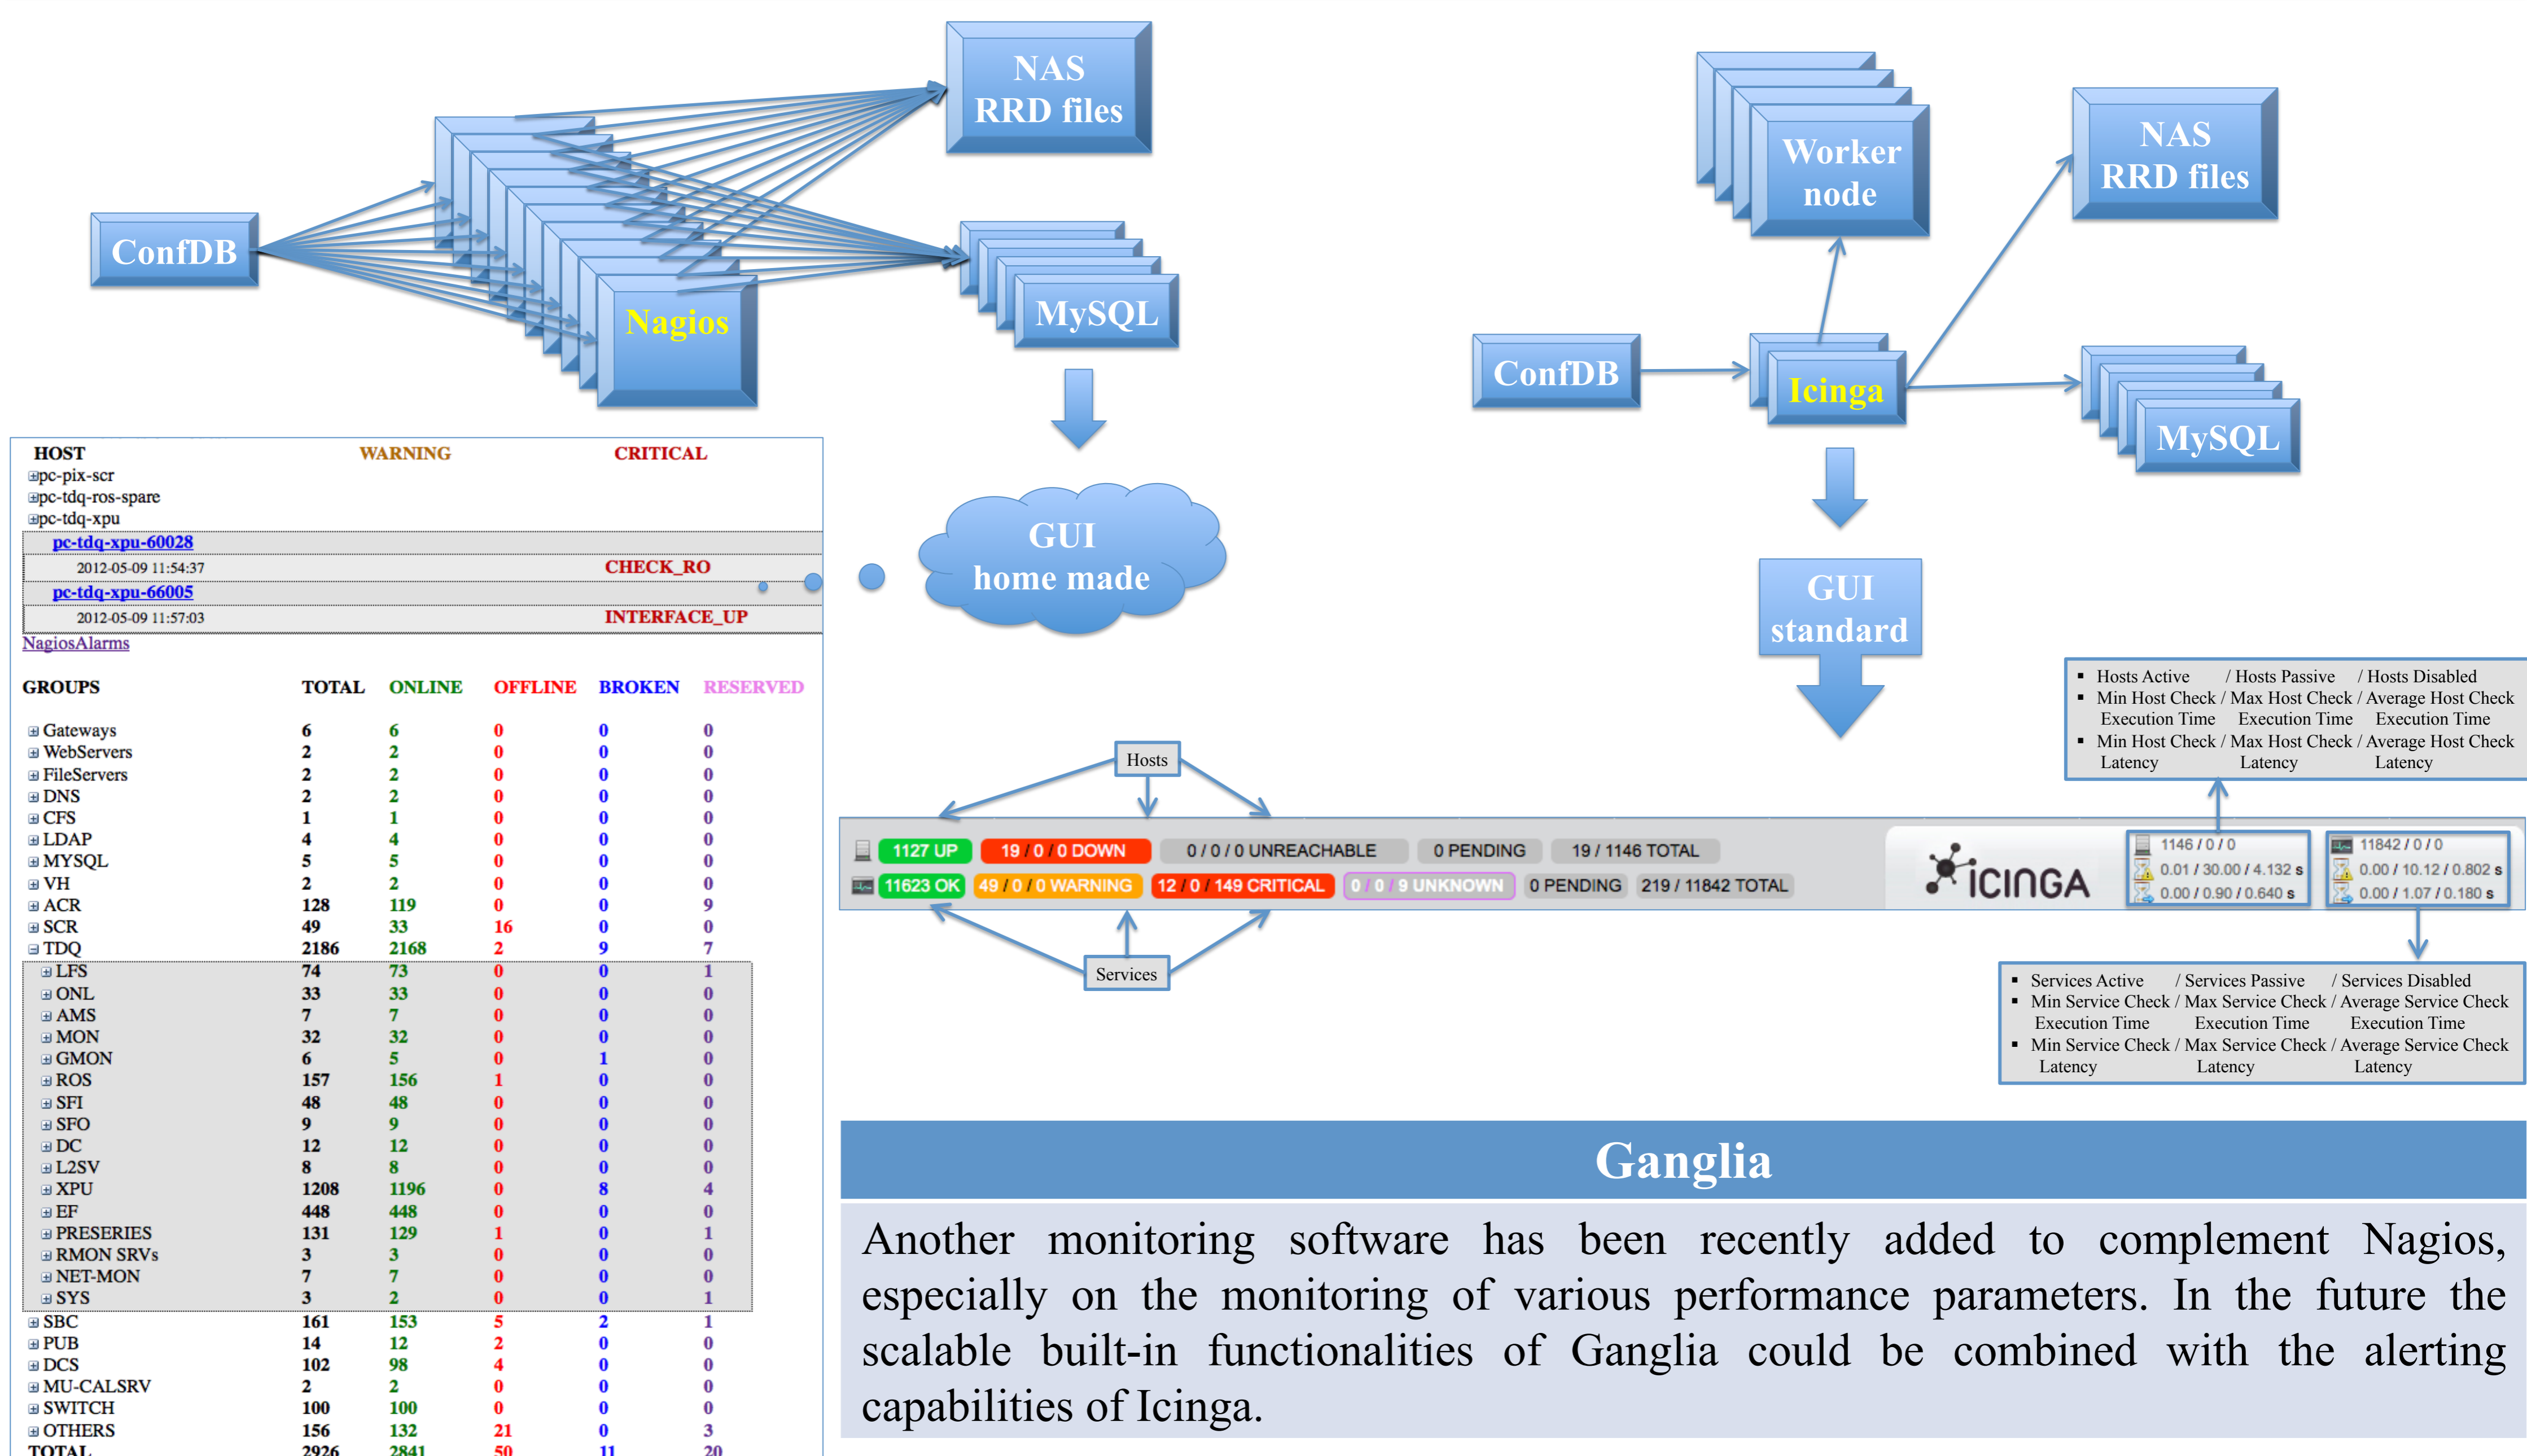
- This was necessary in 2007, when the monitoring system was designed, to distribute the work load of the checks and plotting on the servers
- An in-house GUI has been developed to collect and display summaries of the check results performed by each of the ~80 servers
- ~3000 hosts: ~100000 checks and history

Under evaluation

One (or very few) Icinga server(s).

- In the first tests conducted it seems to cope better than Nagios with the high number of checks performed:
 - running ~12000 checks on ~1100 hosts on a single monitoring server results in an average latency* below 1 second
 - * the delay between the scheduled start time of the check and its effective reporting time

Eventually the data is stored on the MySQL cluster



Ganglia

Another monitoring software has been recently added to complement Nagios, especially on the monitoring of various performance parameters. In the future the scalable built-in functionalities of Ganglia could be combined with the alerting capabilities of Icinga.

Current

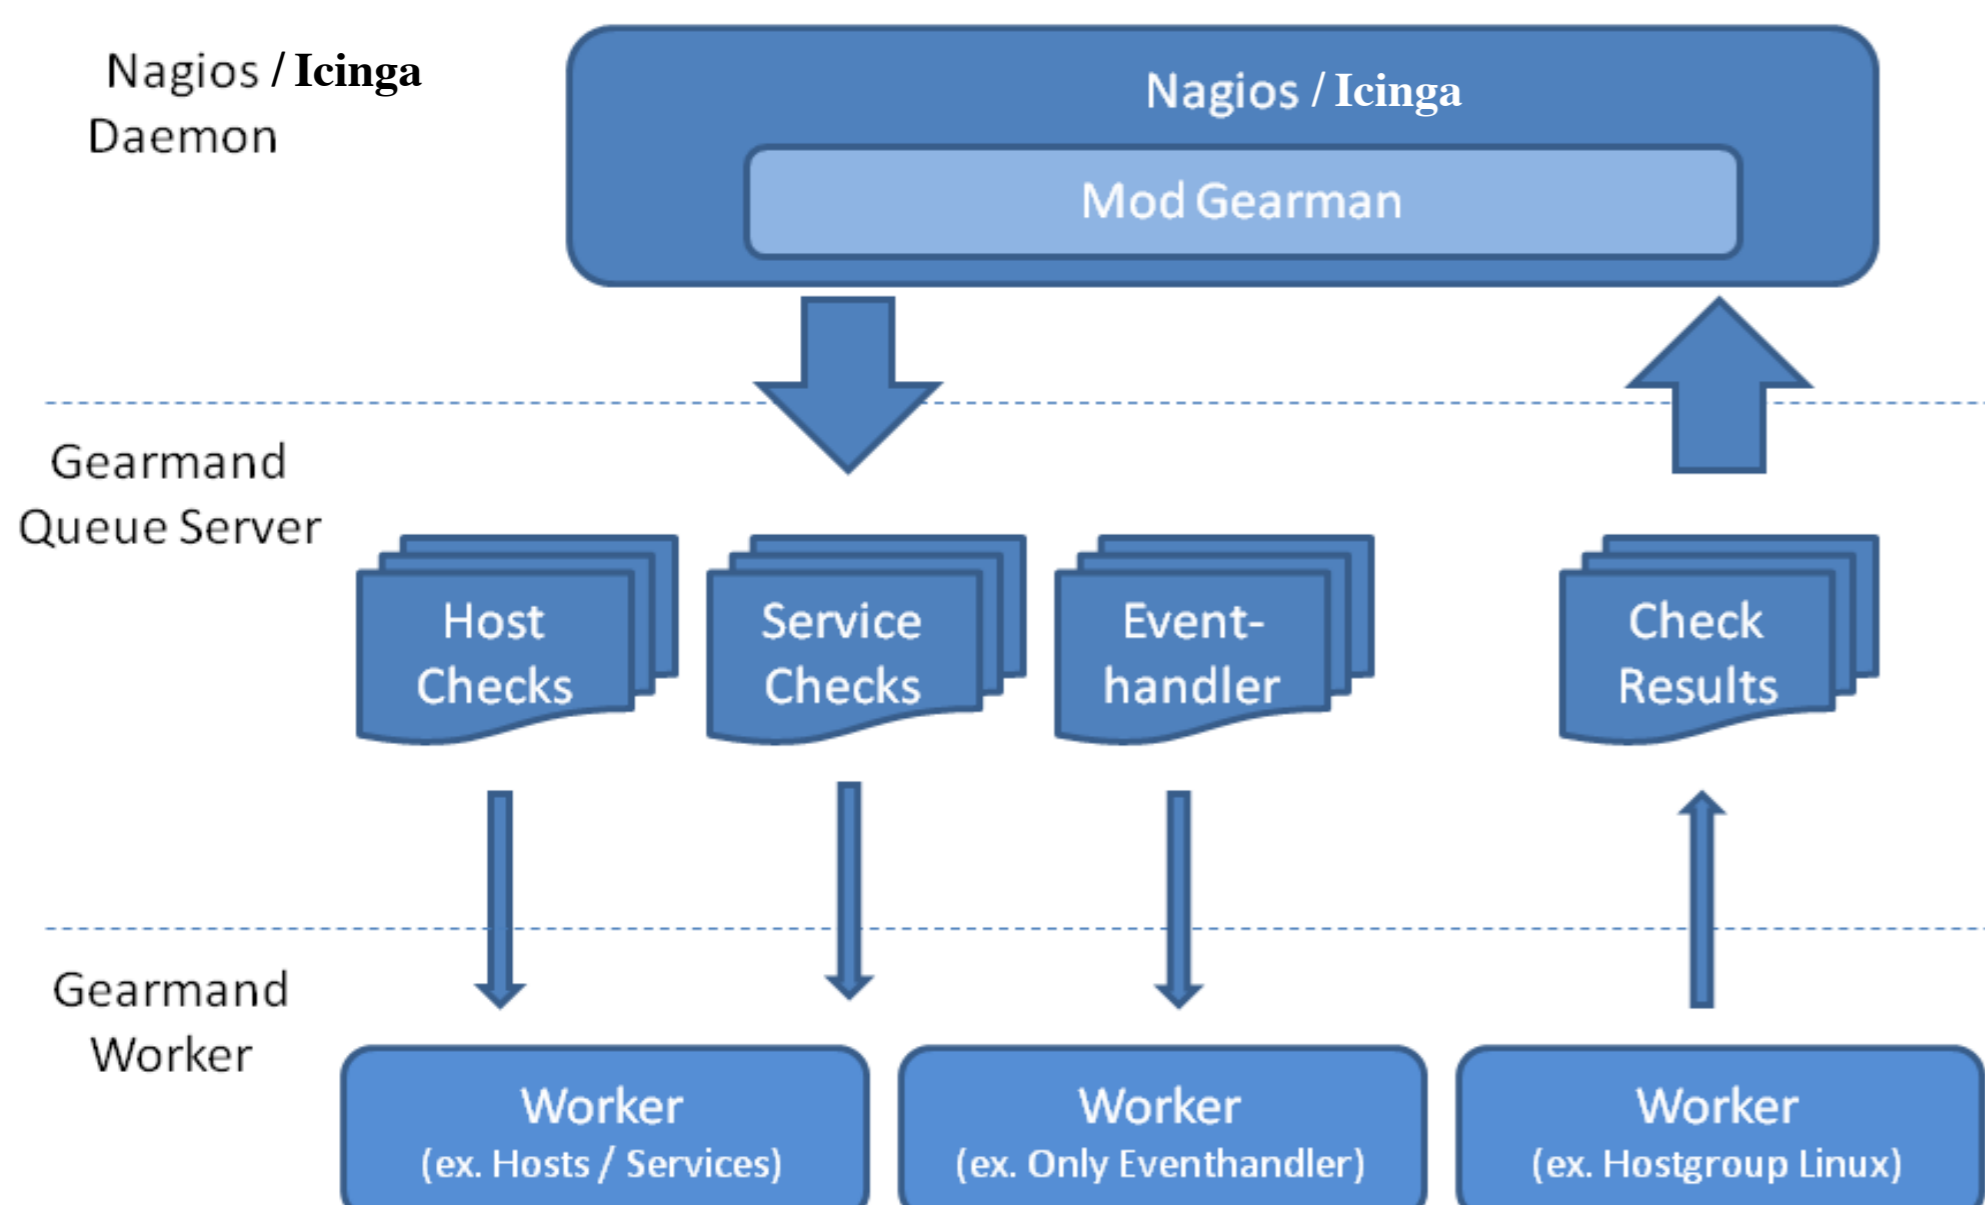
Under evaluation

Advantages

- Flexibility: monitored nodes can be easily moved to another server if needed (e.g. server down)
- Graceful failure mode: losing one server results in the temporary loss of monitoring for a limited number of clients, given the complete independence of the monitoring servers
- Simple schema: only one (or few distributed) servers
- GUI: default one can be used
- Useful options are available to lower the work load
- Proven trivial distribution of work load on other nodes (see gearman/mod_gearman section below)

Disadvantages

- Complicated schema: configuration is scattered on tens of nodes and home-made GUI to be maintained
- Information from many nodes needs to be combined on a single web interface to display summary data
- Many steps are needed to get details about a check
- Single point of failure: losing the central server results in no access to the monitoring information
 - High availability system is possible but complex
 - Not critical for ATLAS data taking



Gearman/mod_gearman

An easy way of distributing active Nagios/Icinga checks across the network, increasing scalability and reducing the load on a central monitoring server. Initial tests have been performed with 2 worker nodes, showing a balanced sharing of the load due to the checks execution.

Queue Name	Worker Available	Jobs Waiting	Jobs Running
check_results	1	0	0
eventhandler	38	0	0
host	38	26	16
service	38	0	22
worker_pc-tdq-lfs-202.cern.ch	1	0	0
worker_pc-tdq-sys-03	1	0	0

Conclusions

The current monitoring system is based on Nagios v2.5 and provides the required information and alert notifications. A new system based on Icinga is being tested to update the current one. The first tests performed using Icinga showed a good behaviour in monitoring ~1100 hosts with ~12000 services; both the execution time and the latency are reasonable, showing that the work load on the server is under control. The additional use of mod_gearman increases the performance as expected.