


The Compact Muon Solenoid Experiment
Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland



20 June 2012 (v5, 17 July 2012)

Health And Performance Monitoring Of The Online Computer Cluster Of CMS

G Bauer⁶, U Behrens¹, O Bouffet², M Bowen², J Branson⁴, S Bukowiec², M Ciganek², S Cittolin⁴, J A Coarasa², C Deldicque², M Dobson², A Dupont², S Erhan³, A Flossdorf¹, D Gigi², F Glege², R Gomez-Reino², C Hartl², J Hegeman^{2,a}, A Holzner⁴, Y L Hwong², L Masetti², F Meijers², E Meschi², R K Mommsen⁵, V O'Dell⁵, L Orsini², C Paus⁶, A Petrucci², M Pieri⁴, G Polese², A Racz², O Reginel⁶, H Sakulin², M Sani⁴, C Schwick², D Shpakov⁵, M Simon², A C Spataru², K Sumorok⁶

Abstract

The CMS experiment's online cluster consists of 2300 computers and 170 switches or routers operating on a 24-hour basis. This huge infrastructure must be monitored in a way that the administrators are pro-actively warned of any failures or degradation in the system, in order to avoid or minimize downtime of the system which can lead to loss of data taking. The number of metrics monitored per host varies from 20 to 40 and covers basic host checks (disk, network, load) to application specific checks (service running) in addition to hardware monitoring. The sheer number of hosts and checks per host in the system stretches the limits of many monitoring tools and requires careful usage of various configuration optimizations to work reliably. The initial monitoring system used in the CMS online cluster was based on Nagios, but suffered from various drawbacks and did not work reliably in the expanded cluster. The CMS cluster administrators investigated the different open source tools available and chose to use a fork of Nagios called Icinga, with several plugin modules to enhance its scalability. The Gearman module provides a queuing system for all checks and their results allowing easy load balancing across worker nodes. Supported modules allow the grouping of checks in one single request thereby significantly reducing the network overhead for doing a set of checks on a group of nodes. The PNP4nagios module provides the graphing capability to Icinga, which uses files as round robin databases (RRD). Additional software (rrdcached) optimizes access to the RRD files and

¹) DESY, Hamburg, Germany

²) CERN, Geneva, Switzerland

³) University of California, Los Angeles, Los Angeles, California, USA

⁴) University of California, San Diego, San Diego, California, USA

⁵) FNAL, Chicago, Illinois, USA

⁶) Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

^a) Now at Princeton University, Princeton University, New Jersey, USA

is vital in order to support the required number of operations. Furthermore, to make best use of the monitoring information to notify the appropriate communities of any issues with their systems, much work was put into the grouping of the checks according to, for example, the function of the machine, the services running, the sub-detectors to which they belong, and the criticality of the computer. An automated system to generate the configuration of the monitoring system has been produced to facilitate its evolution and maintenance. The use of these performance enhancing modules and the work on grouping the checks has yielded impressive performance improvements over the previous Nagios infrastructure, allowing for the monitoring of many more metrics per second compared to the previous system. Furthermore the design allows the easy growth of the infrastructure without the need to rethink the monitoring system as a whole.

Presented at *CHEP 2012: International Conference on Computing in High Energy and Nuclear Physics*

Health And Performance Monitoring Of The Online Computer Cluster Of CMS

G. Bauer⁶, U. Behrens¹, O. Bouffet², M. Bowen², J. Branson⁴, S. Bukowiec², M. Ciganek², S. Cittolin⁴, J. A. Coarasa Perez², C. Deldicque², M. Dobson², A. Dupont², S. Erhan³, A. Flossdorf¹, D. Gigi², F. Glege², R. Gomez-Reino², C. Hartl², J. Hegeman^{2a}, A. Holzner⁴, Y. L. Hwong², L. Masetti², F. Meijers², E. Meschi², R. K. Mommsen⁵, V. O'Dell⁵, L. Orsini², C. Paus⁶, A. Petrucci², M. Pieri⁴, G. Polese², A. Racz², O. Raginel⁶, H. Sakulin², M. Sani⁴, C. Schwick², D. Shpakov⁵, M. Simon², A. C. Spataru², K. Sumorok⁶

¹ DESY, Hamburg, Germany; ² CERN, Geneva, Switzerland; ³ University of California, Los Angeles, Los Angeles, California, USA; ⁴ University of California, San Diego, San Diego, California, USA; ⁵ FNAL, Chicago, Illinois, USA; ⁶ Massachusetts Institute of Technology, Cambridge, Massachusetts, USA ^a Now at Princeton University

Olivier.Raginel@cern.ch

Abstract. The CMS experiment's online cluster consists of 2300 computers and 170 switches or routers operating on a 24-hour basis. This huge infrastructure must be monitored in a way that the administrators are pro-actively warned of any failures or degradation in the system, in order to avoid or minimize downtime of the system which can lead to loss of data taking. The number of metrics monitored per host varies from 20 to 40 and covers basic host checks (disk, network, load) to application specific checks (service running) in addition to hardware monitoring. The sheer number of hosts and checks per host in the system stretches the limits of many monitoring tools and requires careful usage of various configuration optimizations to work reliably. The initial monitoring system used in the CMS online cluster was based on Nagios, but suffered from various drawbacks and did not work reliably in the expanded cluster. The CMS cluster administrators investigated the different open source tools available and chose to use a fork of Nagios called Icinga, with several plugin modules to enhance its scalability. The Gearman module provides a queuing system for all checks and their results allowing easy load balancing across worker nodes. Supported modules allow the grouping of checks in one single request thereby significantly reducing the network overhead for doing a set of checks on a group of nodes. The PNP4nagios module provides the graphing capability to Icinga, which uses files as round robin databases (RRD). Additional software (rrdcached) optimizes access to the RRD files and is vital in order to support the required number of operations. Furthermore, to make best use of the monitoring information to notify the appropriate communities of any issues with their systems, much work was put into the grouping of the checks according to, for example, the function of the machine, the services running, the sub-detectors to which they belong, and the criticality of the computer. An automated system to generate the configuration of the monitoring system has been produced to facilitate its evolution and maintenance. The use of these performance enhancing modules and the work on grouping the checks has yielded impressive performance improvements over the previous Nagios infrastructure, allowing for the monitoring of many more metrics per second compared to the previous system. Furthermore the design allows the easy growth of the infrastructure without the need to rethink the monitoring system as a whole.

1. Introduction

The CMS experiment's [1][2] online cluster at the LHC consists of 2300 computers and 170 switches or routers operating on a 24-hour basis. This large infrastructure must be monitored in a way that the administrators are pro-actively warned of any failures or degradation in the system, in order to avoid or minimize downtime of the system which can lead to loss of data taking. Each minute of colliding beam in the machine is precious, and therefore fast and reliable monitoring and alerting is crucial. Different groups working on diverse parts of the detector use most of these computers, and the administration and monitoring of all of them is the responsibility of the DAQ (Data AcQuisition) group.

The original monitoring solution, based on Nagios, was found to be unsatisfactory, and was replaced by Icinga, with several useful plugins. The motivations behind the migration, and the selection of the plugins are discussed.

2. Issues and solutions

2.1. Upgrading from a Nagios installation

The overall experience with Nagios in CMS was satisfactory, but there were problems with scaling in the CMS growing cluster. The simplicity of its plugins and their extensibility are its key assets, however the difficult-to-maintain process of splitting the configuration in order to scale the monitoring with the growing infrastructure and the old standards used for the web interface forced a reconsideration of what to use to re-architecture the monitoring system.

While evaluating the different alternatives possible, mostly three products were considered:

- The latest version of Nagios [3] (3.2.3, as 3.3.1 had serious bugs)
- The latest version of Icinga [4] (1.3 at the time, now 1.71)
- Shinken [5]

All these products were able to read the previous configuration with very little modifications, and they could reuse the existing plugins, made the migration easier.

The latest version of Nagios did not have the additional features or enhancements required ,i.e, it still used a CGI web interface written in C, which needed a recompilation for each modification.

The first evaluated version of Icinga was not available in RPM format. An effort was put into packaging it for version 1.3.1, and giving it back to the community, with the benefit that nowadays there are available RPMs in the DAG [6] repository. Icinga web interface was following modern standards. The latest version of Icinga has an improved web interface, which can be easily extended through widgets (called "cronks"), which are XML files. Icinga also offers JSON output for most of the pages (since version 1.4.1), which allows very easy integration with other tools. Icinga had also a very active community behind it.

Shinken looked promising, but seemed still in its beta stage at the time.

Icinga was selected considering the points aforementioned.

2.2. Scalability with respect to a growing cluster

It was difficult to maintain the process of changing the configuration during cluster growth. The load on each monitoring server is proportional to the number of checks and the number of hosts monitored, and once the maximum load on a single server was reached, the configuration had to be split across several servers to handle the load. In order to show a unified view of the whole cluster, the monitoring information had to be aggregated onto a central dashboard, which added an additional layer of complexity. The splitting was a manual task and required both guessing as to how the cluster would grow and cautiousness in writing the configuration so that it could be re-used across servers.

Gearman provides a generic application framework for farming out tasks to other machines or processes that are better suited to do the work. It allows work to be executed in parallel, to balance the processing load on several machines. It can be used in a variety of applications, from high-availability

web sites to the transport of database replication events. It is the nervous system through which distributed processing communicates. [7]

Icinga can leverage the power of Gearman by using a module, which is installed as a broker directly inside the main executable and farms out all checks needed to be run to the Gearman server and checks for results to be processed. This allows easy scaling; the central server is only a scheduler, there can be as many worker nodes as needed, and they each have the same configuration. Gearman can have multiple queues, which Icinga uses for storing (See Figure 1):

Services to be checked, consumed by the workers, injected by the scheduler;

Hosts to be checked, same as for services;

Results to be processed, consumed by the scheduler, injected by the workers and check_multi;

Performance data to be processed, consumed by PNP4nagios, injected by the scheduler.

Gearman quickly became the central piece for the scaling of our monitoring infrastructure, because of its inherent efficiency. It is very stable and reliable, and has proved to be a good choice.

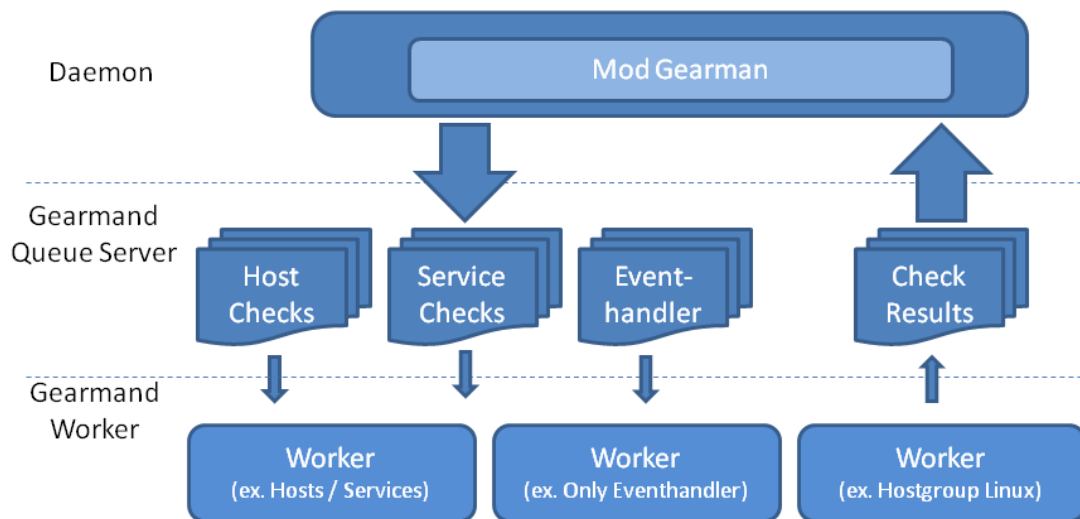


Figure 1 Scalability with respect to the number of checks

A standard check can either be run locally on the monitoring server or remotely through some mechanisms. The most common ones are SSH, NRPE [8] and SNMP. SSH (Secure SHell) has more overhead and is potentially more secure. SNMP is light and fast but difficult to extend and complicated to setup to have good security (SNMPv3). NRPE (Nagios Remote Plugin Execution) is the best of both worlds: easy to setup, secure enough, fast and easy to extend. It has a drawback that required patching: the limit of the output on a 32-bit operating system was limited to 4096 characters.

Regardless of the mechanism to execute remote checks, for each check run from the monitoring server a connection to the remote server needs to be open. Therefore, the load on the monitoring server scales linearly with the number of checks done on a single machine, which was already problematic when having over 20 checks on all hosts.

To reduce this overhead, a plugin called check_multi [9] is used. It can run on the monitoring server one single remote check that can cascade on the remote monitored client and check several plugins in a row. Currently the CMS monitoring server runs two remote check_multi checks per host instead of 27 remote checks. The problem posed by this solution is that all results appear as one single service in Icinga. This causes issues because the alerting is done on a service, which contains multiple checks, some of which requiring alerting and some not. Furthermore a failure on one check causes the whole service to be critical, and will therefore hide any additional check failures in this service. To remedy this, the tests are then split as passive ones, which allow fine-tuned notifications, and proper alerting.

In summary, for the standard 27 tests run on all the monitored hosts, two active tests are defined in Icinga, which make use of `check_multi`. These `check_multi` checks output XML data, which is piped into `send_multi` (part of `mod_gearman`), which then takes care of parsing and splitting the 27 results and sending them back through the Gearman queue. The master Icinga scheduler processes the 27 passive tests. The configuration is a little bit more complex than in plain mode, but `check_multi` provides tools for generating the passive checks from the local/remote `check_multi` configuration, reducing the effort required in configuration and maintenance.

As most of the checks take as long to run as the overhead of the NRPE connection, tunneling the 27 passive checks through `check_multi` reduces the load on the monitoring server roughly by a factor of 15 and cuts the time needed to run them all in half.

2.3. PNP4nagios

Most of the executed checks generate performance data, which are simple metrics used to monitor the health of the machine and its variation. For example the checks for the disk occupancy, the memory usage, the ping response time, and other such variables generate performance data. They are very useful to graph, in order to see trends and correlations in the data. PNP4nagios [10] graphs any plugin's performance data that follows the Nagios plugin development guidelines [11]. This also allows developing custom plugins and having the results graphed.

PNP4nagios can also merge plots, in order to show the metrics from several hosts on a single graph. One example for this is the monitoring of the temperature of the hosts (See Figure 2) (local checks are in place to ensure hosts don't overheat, and turn themselves off in case they reach some thresholds). It is useful to be able to correlate the temperature within a building, a room or a rack, to pinpoint a cooling issue for example, or simply to observe the heat impact of data taking.

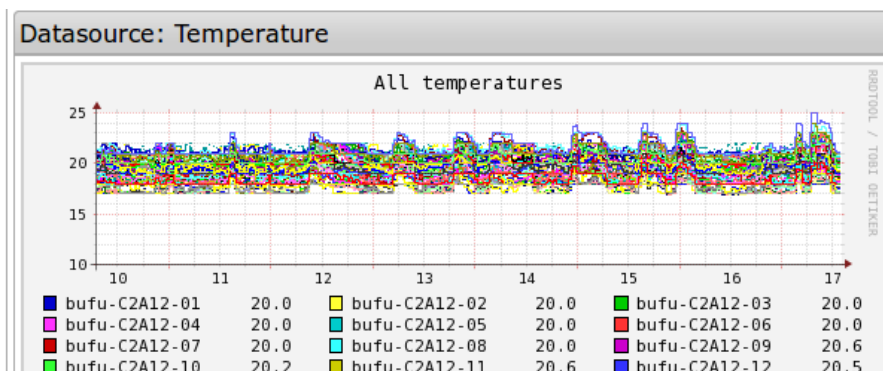
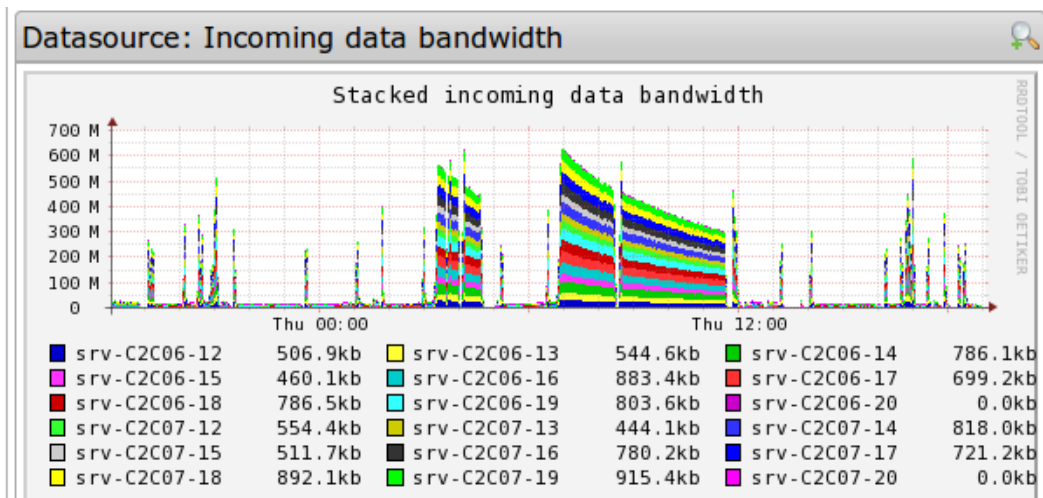


Figure 2 Overall temperature display of 1028 computers from the High Level Trigger farm

PNP4nagios also allows the merging of plots. As all the data are stored in RRDs [12], it is easy to collect metrics from several hosts/checks, and graph them together. For example, in CMS the Storage Manager machines, which are used to store the data selected by the High Level Trigger on disk before shipping it to CERN's Tier0 for storage and processing, have four network interfaces that can all receive data during taking data. In order to graph the total bandwidth usage, all four network cards for all 16 storage managers need to be stacked, which is very easy to do with PNP4nagios (See Figure 3).



Rysunek 1 Stacked incoming data bandwidth on all 16 Storage Managers

Thanks to caching (see next section), the load on the disks isn't too high. The only issue that needed tuning was that PNP4nagios writes a summary file in XML that describes what metrics are available for each hosts. Because of the number of metrics in each file, these XML files are big, and the default refresh time of 5 minutes was too low to handle the graphing of all the network interfaces of the entire cluster. Therefore the frequency of writing these XML files to disk has been decreased to once per day, a compromise against an acceptable delay when new tests are added.

2.4. RRDcached

PNP4nagios generates many RRD files, which needs a lot of I/O to disk. In order to leverage the load, RRDcached is used. It simply caches all updates to a RRD file. Changes are synced to disk only every 5 minutes by default. RRDcached daemon has a feature that the client can use to tell it to force the file to be synced to disk right away. This feature can be used in order to serve the latest data. As trends are more used than live statistics, this feature is not used. RRDcached is part of RRDtool since version 1.4.

2.5. Integration with external tools

Some sub-detectors need to integrate Icinga's results within their own monitoring architecture, in order to show a coherent overall status. This is done using JSON [13] standard outputs (available from version 1.4.1 in Icinga). Most of the monitoring pages in our infrastructure can output JSON data, which can then easily be parsed and displayed inside another application.

2.6. Configuration management

The CMS cluster is managed by a configuration management tool called Quattor [14]. To avoid duplication information, a perl script that extracts information out of Quattor and generates a configuration suitable for Icinga has been developed. The script also generates the network topology, in order to hide all failures located behind a failed switch or router.

3. Conclusion

The new Icinga based system is able to monitor over 2'300 computers and 170 switches, with a granularity of two minutes. Most of the configuration for the hosts and switches are generated from Quattor, getting the hosts and groups from Quattor profiles. On all those 2'500 hosts, 27 standard checks are run through check_multi, with additional checks added on a host-by-host basis depending on its functionality. Alerting is based on host groups, generated from a wiki page, and alerts can be sent either by email or by SMS.

Although the decision to migrate from Nagios to Icinga was mostly motivated by the philosophy behind the Icinga project, and its newer interface and wider feature list, it has exposed many very useful plugins that have been shared across the LHC experiments, and benefited other experiments at CERN [15].

Thanks to the new Icinga based system, all sub-detectors monitor their production critical machines, and some have implemented their own checks. The system is now up and running, and provides enough scalability to suffice also after the upcoming long shutdown.

The use of these performance enhancing modules and the work on grouping the checks has yielded impressive performance improvements over the previous Nagios infrastructure allowing for the monitoring of 65'000 checks run every 2 minutes on average, compared to 20'000 every 5 minutes on the previous system. It also allowed us to scale down the hardware requirements, as we moved from four 8 cores machines to one 4 cores machine running the Icinga server and generating all the RRDs, and one 8 cores machine running all the checks. Furthermore the design allows the easy growth of the infrastructure without the need to rethink the monitoring system as a whole.

Acknowledgement:

This work was supported in part by the DOE and NSF (USA) and the Marie Curie Program.

References

- [1] The CMS Collaboration, *The Compact Muon Solenoid Technical Proposal*, CERN LHCC 94-38, 1994.
- [2] The CMS Collaboration (Adolphi R et al.) “The CMS Experiment at CERN LHC”, JINST 3 S08004 361, 2008.
- [3] Nagios Is The Industry Standard In Infrastructure Monitoring, Nagios.org, <http://nagios.org>
- [4] Icinga takes open source monitoring to the next level, icinga, <http://icinga.org>
- [5] Shinken, The next Industry Standard in IT Monitoring, <http://shinken-monitoring.org>
- [6] DAG, RPM packages for Red Hat, RHEL, CentOS and Fedora, <http://dag.wieers.com/rpm>
- [7] Gearman provides a generic application framework to farm out work to other machines or processes that are better suited to do the work, <http://gearman.org/#introduction>
- [8] NRPE, Nagios Remote Plugin Executor, <http://docs.icinga.org/1.6/en/nrpe.html>
- [9] check_multi is a multipurpose wrapper plugin which takes benefit of the Nagios 3.x capability to display multiple lines of plugin output, Matthias Flacke, http://my-plugin.de/wiki/projects/check_multi/discussion
- [10] PNP is an addon to Nagios which analyzes performance data provided by plugins and stores them automatically into RRD-databases, <http://docs.pnp4nagios.org/pnp-0.6/start>
- [11] Nagios Plug-in Development Guidelines, Nagios Plugins Team, <http://nagiosplug.sourceforge.net/developer-guidelines.html#AEN201>
- [12] RRDtool is the OpenSource industry standard, high performance data logging and graphing system for time series data, Tobias Oetiker, OETIKER+PARTNER AG, <http://oss.oetiker.ch/rrdtool>
- [13] JSON (JavaScript Object Notation) is a lightweight data-interchange format., <http://json.org>
- [14] Quattor is a system administration toolkit, <https://trac.lal.in2p3.fr/Quattor/wiki/Web>
- [15] Tools and strategies to monitor the ATLAS online computing farm, Scannicchio & all, *International Conference on Computing in High Energy and Nuclear Physics (CHEP), NY, May 2012*