The CMS Online Cluster: Setup, Operation and Maintenance of an Evolving Cluster

J.A. Coarasa [1]

CERN
CH-1211 Geneva 23, Switzerland
E-mail: coarasa@cern.ch

Gerry Bauer (6), Ulf Behrens (1), Olivier Bouffet (2), James Branson (4), Sebastian Bukowiec (2), Olivier Chaze (2), Marek Ciganek (2), Sergio Cittolin (4), Christian Deldicque (2), Marc Dobson (2), Aymeric Dupont (2), Samim Erhan (3), Dominique Gigi (2), Frank Glege (2), Robert Gomez-Reino (2), Christian Hartl (2), Andre Holzner (4), Lorenzo Masetti (2), Frans Meijers (2), Emilio Meschi (2), Remigius K. Mommsen (5), Carlos Nunez-Barranco-Fernandez (2), Vivian O'Dell (5), Luciano Orsini (2), Christoph Paus (6), Andrea Petrucci (2), Marco Pieri (4), Giovanni Polese (2), Attila Racz (2), Olivier Raginel (6), Hannes Sakulin (2), Matteo Sani (4), Christoph Schwick (2), Michal Simon (2), Andrei Cristian Spataru (2), Fabian Stoeckli (6), Konstanty Sumorok (6).

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

The CMS online cluster consists of more than 2700 computers running about 15000 application instances. These applications implement the necessary services to run the data acquisition of the CMS experiment. In this paper the IT solutions employed on the cluster are reviewed. Details are given on the adopted solutions which include the following topics: implementation of a redundant and load balanced network and core IT services; deployment and configuration management infrastructure and its customization; a new monitoring infrastructure. Special emphasis will be put on the scalable approach allowing to increase the size of the cluster with no administration overhead. Finally, the lessons learnt from the two years of running will be presented.

The International Symposium on Grids and Clouds (ISGC) 2012
Academia Sinica, Taipei, Taiwan
February 26 – March 2, 2012

---

[1]   Speaker

1.Introduction

The Compact Muon Solenoid (CMS) experiment [1][2] is one of the four experiments at the Large Hadron Collider [3] (LHC) at the European Organization for Nuclear Research (CERN). It is a general-purpose experiment covering a wide range of physics at the TeV scale [4][5][6].

Its Data Acquisition (DAQ) system [7][8] reads data from more than 50 million channels through approximately 650 custom links. At the receiving end, data are fed into commercial networking hardware. Subsequent stages of the data flow are implemented using commercial components.

Event selection in CMS is based on a two level trigger system. The first level, implemented in custom hardware, reduces the event rate to a maximum of 100 kHz. The second trigger level, also called high level trigger (HLT), is implemented in software running on a dedicated computer farm receiving full events from the event-building network.

The total data volume per event depends on the number of underlying minimum bias events in the selected collisions (pile-up), which in turn depends on the luminosity per bunch crossing. LHC was designed to operate at a luminosity of 10^{34} cm^{-2}s^{-1} with 2835 bunch crossings per orbit (88μs). With these parameters approximately 23 pile-up events are expected in every collision resulting in an event size of about 1 Mbyte. Hence the total data flow from the front-end electronics through the event builder to the HLT farm is 100Gbyte/s.

CMS event building (EVB) is done in 2 stages [9]. The first stage is based on a Myrinet [10] switched network whereas the second is based on multigigabit ethernet networks. Both stages are networking intensive, with 100Gbytes/s throughputs.

The HLT is a CPU intensive process that selects 1 out of 1000 events. The selected events are subsequently stored and forwarded to the central data recording center at CERN (Tier 0) by the storage manager infrastructure which can sustain a throughput of around 1.5Gbytes/s (heavy ion runs). About 300 Mbytes/s sustained throughput is sent to Tier 0, totalizing a 20Tbytes/day.

In order to implement the aforementioned functionalities and the control of the experiment a computing cluster has been built. This paper discusses details of its information technologies (IT) implementation and operation over the past years.

2. The CMS online cluster

The CMS online cluster [11] serves the needs of a community of more than 900 users. It is autonomous (i.e. independent from all other networks) and provides uninterrupted operation around the clock on two distant (~200 m) physical locations: A cavern (USC) approximately 80 m below the ground level, where the experimental area, the custom electronics and control computers are located; and a building in the surface (SCX5), where most of the online cluster and the control room are located. The CMS online cluster has been designed with redundant services to avoid wasting accelerator time and to scale easily to accommodate expansions. It features a fast configuration turnaround to cope with the evolving nature of DAQ applications and the large scale of the cluster.

In the following sections, the implementation of the IT infrastructure and the management and operation during the past years is presented.

2.1 The IT infrastructure of the CMS online cluster

The CMS online cluster is composed of more than 2700 computers (mostly running Scientific Linux CERN 5), more than 120 networking switches (Myrinet and Gigabit Ethernet), 2 network-attached storage (NAS) systems and 2 storage area networks (SAN). The cluster has been growing to its actual size in stages to meet the increased performance requirements. In the past year the processing power of the High Level Trigger farm was increased by 50% (288 nodes were added to the initial 720) without disruption to ongoing operations. A new expansion is currently underway.

In the following subsections a detailed description of the computing resources (2.1.1), networking (2.1.2) and IT services (2.1.3 to 2.1.8) will be given.

2.1.1 Computing Resources

A detailed breakdown of the cluster components, according to the different functionality of the computers, follows:

- 640 (2-core) computers (named readout units) equipped with two Myrinet and three independent 1Gbit Ethernet links for data networking. They run applications for both stages of the CMS event building. These event building applications require aggregated high bandwidth networking;
- 720 (8-core) and 288 (12-core) computers (named BUFUs) equipped with two 1Gbit Ethernet lines for data networking. They run applications for the second stage of the event building and the high level trigger (which are CPU intensive);
- 16 computers (storage manager) with access to a Fibre Channel storage area network of 300Tbytes and equipped with four 1Gbit Ethernet lines for data networking and 2 trunked additional 1Gbit Ethernet lines for networking to the Tier 0. They run applications to temporarily buffer the event data and locally handle the transfer to Tier 0;
- More than 400 computers used by the subdetectors and other subsystems;
- 90 computers running Windows for the Detector Control System;
- 12 computers for the experiment Online Database implemented via an ORACLE RAC [12];
- 15 computers running the CMS run control system called RCMS [13];
- 30 computers used for collection of the monitoring and error reporting of the CMS application software, this comprises O(20000) of interdependent XDAQ applications to be monitored in near real-time [14];
- 50 computers as desktop control terminals for the different subdetectors;
- 10 computers hosting the cluster IT services;
- 6 computers to allow access to the online cluster from the CERN campus network;
- 200 computers for commissioning, integration and testing;
- 250 computers as active spare computers.

In contrast to most of the typical computing clusters, the configuration (number of networks, privileged users, running services…) and software deployed in the computers varies substantially depending on their functionality and on the hardware on which this functionality is running. In Section 2.1.5 a description of the configuration management infrastructure allowing this highly specialized configuration is given.
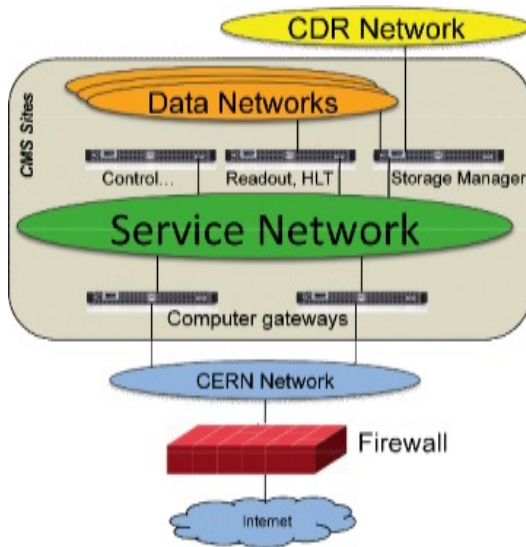

Figure 1. Ethernet networking structure of the CMS online cluster. The LHC Technical network and the small private networks for subsystems and ORACLE RAC are not depicted here.
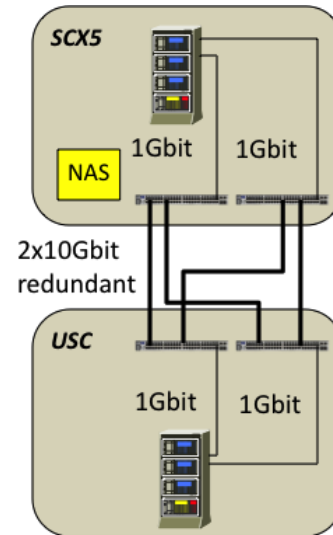

Figure 2. Redundancy of CMS private service network connecting the two locations of the CMS experiment.

2.1.2 Networking in the CMS online cluster

The Ethernet networking structure in the cluster is depicted in figure 1. The CMS online cluster makes use of the following networks:

- The CMS private networks:
  - o The CMS service network. All computers in the CMS cluster are connected to this network. It is built using more than 3000 ports across more than 120 switches. Each of the switch ports support 1 Gb/s. The IT general cluster services are provided via this network, i.e: DNS, Kerberos, LDAP, installation, configuration, etc. This network spans the two physical locations of the CMS computing infrastructure: the underground area (USC, where most of the electronics is located) and the surface building (SCX5, where most of the computers are installed). The redundant interconnection of the two sites is implemented with 4 routers as shown in figure 2. Four 10Gbit lines are used to interconnect the routers. The router pairs in each location are configured in a failover mode. The interconnectivity within each location is also redundant: in each rack, a switch connects to all computers in the rack and to both routers. This cabling tolerates single failures up to the rack switch.

      o The CMS data networks. These are eight networks used for the event building, connectivity to the high level trigger filter farm and to the storage manager. To implement these high bandwidth networks, eight Force 10 [15] Gbit Ethernet switches with 6 high density line cards each are used. One of these networks connects 80 readout units, 126 BUFUs and all 16 storage manager computers through all of their data network interfaces [7][8], totaling more than 4000 ports for the eight switches. In order for the applications to make use of the aggregate bandwidth of all interfaces in each computer the data networks are divided in to Virtual LANs (VLANs) and the computers are configured to use source routing;

      o The subdectector's and ORACLE RAC's private networks (not shown in figure 1);

- The CERN campus public network. This network connects the CMS online cluster to the outside world through the CERN routers and the CMS access servers, that act as computer gateways;

- The LHC technical private network. This network connects all the accelerator related computers (only a few computers are connected to this network. It is not depicted in figure 1);

- The central data recording private network. This network allows transferring data from the storage manager computers to the Tier 0 at CERN for later analysis and archiving. It is serviced by two 10 Gbit lines in redundant failover configuration.

- An additional Myrinet network connects the computers involved in the first stage event building through twelve Myrinet switches. More details about the custom protocols used there can be found in [7][8].

2.1.3 The Network Attached Storage

Important data that do not need to be read by many computers at the same time and need to be kept safe are stored in a *NetApp* network-attached storage filer (NAS) [16]. At the present time, these file servers provide the user home directories, CMS analysis data such as calibration and data quality monitoring, system administration data, software repositories and configuration management data.

The actual configuration of the *NetApp* filer (figure 3) consists of two sets of one head (model FAS3040-R5) and 3 storage shelves (model DS14MkII). Each set is located in a different rack and powered with regular and UPS power for redundancy. The two heads are working in a failover configuration. One of the three shelves in one rack is mirrored to one in the other rack to be able to continue running even if one complete rack fails. The shelves are configured with internal Dual Parity RAID 6. The internal snapshot feature [17] of the Netapp filer is configured to allow user initiated backup restorations in order to minimize the need to restore data from the (independent) backup system. The deduplication option [18] is also active. This saves between 25 to 55% of the available space.

The *NetApp* filer heads are connected to the CMS network routers with redundant 10 Gbit links. Before commissioning the 10 Gbit redundant links, the system was operated with four 1 Gbit trunked lines connecting one of the heads to one of the routers. Measurements carried out

with this configuration were limited by the bandwidth of the trunked link. The bandwidth was measured to be 380 Mbytes/s.
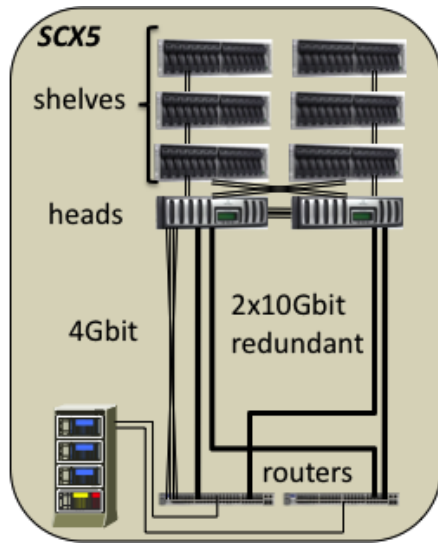

Figure 3. Redundancy in the CMS *NetApp* network-attached storage configuration and its networking.
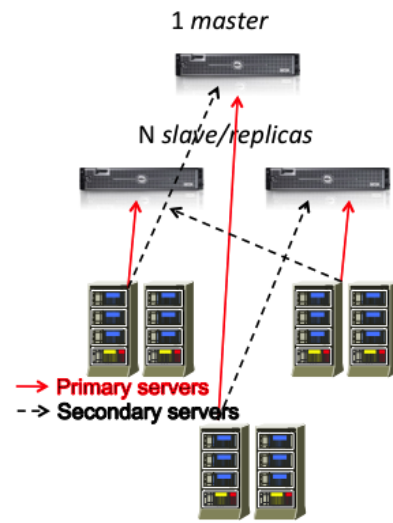

Figure 4. Model for the redundancy and load sharing of the CMS structural services.

Table 1. The CMS IT structural, installation and configuration services.

Service	Replication through	Load Balancing through
DNS	named	Round Robin
DHCP	in-house scripts	No. First who answers
Kerberos	in-house scripts	Explicit segregation
LDAP	slurpd	Explicit segregation
NTP	-	Explicit segregation
syslog	-	None. Single server, only used for stress test purposes
Nagios monitoring	-	Explicit segregation
PXE boot/TFTP		Tied to DHCP answer
Kickstart Installation server		DNS Round Robin
yum repository Installation server	No replication needed. NFS mount from the NAS.	DNS Round Robin
Quattor repository Installation server		DNS Round Robin
Quattor cdb configuration server		Not yet

2.1.4 The CMS IT structural services

The cluster IT structural services, i.e., DHCP, DNS, Kerberos, LDAP, NTP…, have been designed in a redundant easily scalable way that has allowed and still allows expansion.

For the CMS IT structural services (see table 1) a model with one master server and N (where N is now 3) replica servers has been chosen. The administration is done in the master server. It is the single point of administration for the service. This requires that all services' data

are replicated to the replica servers. Some of the services allowed for standard replication at the time of implementation (i.e. slurpd for LDAP or the master/slave mechanism for the DNS data), but some other services did not provide a mechanism to do it (i.e. the Kerberos or the DHCP services). For such services, in-house developments were created to replicate the data from the master servers to the replica servers. For more details refer to table 1.

The services are provided, where possible, by a DNS aliased name. In case of need (administration purposes, recovering from a crashed computer, etc.) the DNS alias can be changed to point to a different computer minimizing the interruption of the service.

Following standard practices in IT, these services are hosted in computers with mirrored system disks and powered with a UPS and non-backed up power. For extended redundancy, the different computers providing the same service are located in different racks, so that if the network or the power fails in one rack the service is still provided.

Having several computers serving the same services allows balancing the load among them. So far, two "load sharing" techniques are used:

- DNS Round Robin;
- Segregating the computers of the cluster in groups that point to different primary servers for a particular service, as depicted in figure 4. This allows the load to be shared in a controlled and reproducible environment while still keeping redundancy for the services. In case of failure of the primary server, in this schema, only the affected group of computers will have to default to their secondary server.

2.1.5 The CMS installation, management and configuration infrastructure

Management of large clusters requires remote control and a configuration management system on top of the infrastructure services described previously. Their implementation for the CMS cluster is detailed below.

Remote control of the CMS cluster is provided using tools for the *Intelligent Platform Management Interface* (IPMI) [19]. The IPMI allows interacting with the remote computer through a dedicated network interface to perform operations like boot up, power off, access the console, etc. In addition it allows reading environmental conditions of the computer like ambient temperature, rotation speed of fans, etc.

As a configuration management system, Quattor (QUattor is an Administration ToolkiT for Optimizing Resources) [20] is used. All linux computers, including the infrastructure servers, or the Quattor servers themselves, are configured using Quattor modules and packages (RPMs) distributed with Quattor. As an example, the BIOS parameters, the IPMI interface parameters, and all the networking parameters (routing tables and IP addresses for the data networks) are configured in the computers at boot time through shell scripts packaged in RPMs. CMS online software (XDAQ [21]) is also packaged in RPMs and distributed through Quattor.

The installation and configuration management infrastructure is based on: PXE [22], TFTP, the anaconda Redhat kickstart [23] and Quattor. All these services follow the same model as the structural services which have been described above (table 1). The services are DNS aliased and they are load shared among 3 servers using DNS Round Robin. In this case, the configuration data do not need to be replicated because they are kept on the NAS, where all

of the servers have access. The master server, defined via a DNS alias, is given write access to the configuration data.

During the installation phase each computer reads it's kickstart file through TFTP, starts the anaconda Redhat installer, installs the Quattor client packages through yum in the postinstall section of the anaconda installer, and applies the configuration information found in the Quattor description templates and reboot.

During the installation or reconfiguration phase of a large number of computers the network traffic to the NAS is limited, due to the fact that the necessary packages are read only once by the repository servers and kept in cache memory while distributed to the clients. This way the quality of the NAS service is guaranteed.

The CMS production Quattor implementation is based on Quattor 1.3: cdb 2.0.4, swrep 2.1.38, PANC 6.0.8. A hierarchy of configuration templates and a set of policies have been developed for the CMS cluster along with tools that ease the administration of the configuration of the computers. In particular, they ease the upgrade of the online software, make the procedure less error prone and accessible to users without knowledge or privileges to the Quattor system. These tools include:

- A tool that lists the groups of computers with the same versions of specific online software. These lists allow automating other administration tasks, such as running a specific command only on a specific group of computers;
- A dropbox for RPMs allows the users to deploy an RPM in a directory to upload it to the Quattor repository for later distribution;
- A template updater can automatically update the templates of the computers. This allows the responsible users (librarians) for a group of computers, without having Quattor knowledge, to update specific software for which they are responsible in these computers;

The CMS Quattor implementation currently manages more than 2500 Scientific Linux CERN 5 computers of more than 90 types, not counting hardware differences. A compilation of all templates takes approximately 12 minutes on the 8 core 2.66 GHz computer hosting the active master Quattor server. Fully reinstalling one computer (with full formatting of the disk) takes between 6 and 30 minutes depending on how much software is loaded. The existing infrastructure has been measured to support reinstalling 1000 computers from scratch in less than an hour and fifteen minutes at the time of commissioning of the infrastructure [11]. A big (O(1Gbyte)) installation in half of the cluster can be done in less than 25 minutes. It is presently limited by the network traffic to the repository servers. Small configuration changes may take only a couple of minutes, even if done cluster-wide.

2.1.6 CMS Monitoring infrastructure

Monitoring is a key element in being able to diagnose and solve a problem. Especially in the case of having so many instances of services, correlating the different failures to the real cause can take quite some time if not automated.

The CMS monitoring infrastructure is based on Nagios [24]. It uses 4 servers in a federated configuration, with the consequent differentiation on the configuration for each server.

Apart from including the standard Nagios tests (ping, ssh accessibility, disk usage...) specific modules were developed to check:

- If the Myrinet links are functional;
- If a configuration and installation finished properly in the computers;
- If the IPMIs are reachable;
- If the storage manager components are operational.

The CMS Nagios monitoring infrastructure performs presently about 25000 service checks in more than 2500 hosts. The latency for the ping test is one minute, while for the rest of the services it is around 5 minutes.

The relatively long latency period of the non-ping services, the lack of a good scaling evolution with the increasing number of services and the difficulty in maintaining the federated configuration triggered the implementation of a new infrastructure based on icinga [25], a fork of Nagios, that allows profiting from plug-ins already developed for nagios. The new implementation uses the following modules:

- mod_gearman [26]. This module allows a master "monitoring" server and as many workers as needed, that share the same configuration, making it easy to scale the system while maintaining ease of configuration;
- check_multi [27]. This module allows grouping the checking of several services so that it will require just one connection to do them, reducing the overhead of the connections. It can cascade providing a way to functionally group the services;
- PNP4nagios [28]. It allows the gathering of performance data. Due to the large number of services being monitored it needed rrdcached [29] in order to work without I/O problems.

The new infrastructure monitors more services, grouped in less checks, with a latency of one minute for all services and provides historical performance data in addition. It allows for easier scaling up, by just adding more workers.

Work is being done to add a system that will allow users to self enroll to receive notification of service failures.

2.1.7 CMS Databases

In CMS, databases are used, among other many uses, to host the conditions under which the data have been taken to allow later analysis in an easier manner. To do this, there are two databases that are hosted in the CMS network, namely: CMSONR (the production online database) and CMSINTR (the integration online database). They are replicated to the CERN site that hosts additionally other CMS databases used mainly by the offline analysis.

The database infrastructure located in the CMS network has been recently entirely rebuilt to be fully redundant. It is implemented using an Oracle 11.2.0.3 RAC architecture [12] with 6 active blades, and additional standby ones, and a *NetApp* filer as storage backend. The total space (including mirroring), at the time of the migration, used by the two CMS online databases is 30 TBytes.

2.1.8 CMS hardware tracking system

The hardware and software tracking system used in CMS is based on open source tools:

- OCS (Open Computer and Software Inventory Next Generation, [30]), that provides the inventory system, giving a detailed picture of software and hardware on the computers by running once a day a client that reports to a centralized server;
- And GLPI (Gestionnaire Libre de Parc Informatique, [31]), that provides the tracking layer, and is synced daily to the OCS inventory.

The inventory and tracking system has proven instrumental in the operation and management of the cluster.

2.2 CMS cluster management and operation

The CMS cluster design, management and operation is carried out by 6 full time equivalents.

As seen from the overall picture of the cluster hardware and services given in the previous section, the complexity of the cluster is not only due to its large scale (around 500 computers/FTE) and number of users but also to the variety of services that are provided. Additionally, the cluster has been growing for more than five years and extended warranty has been purchased where it made sense, however, some hardware is already out of warranty, and some computers connected to the electronics need swift replacement if broken (~1 hour).

The challenge has been overcome partly by investing in automation. Scripts are designed to automate what is time consuming, like the scripts that allow librarians to update the software on their own with no additional overhead on the administration, or to handle critical circumstances for the operation, for example, some scripts allow the computers to monitor the temperature and other services on their own and take corrective actions if possible.

Spare computers are kept in standby and a fast turnaround in reinstallation and configuration has been setup to accommodate hardware failure.

The manpower usage is rationalized by trying to interact with users through a ticketing system and acting only as a 2^{nd} level support when in 24/7 operation.

A special emphasis is put on continued running. To achieve this effort has been put into:

- Learning from previous failures by trying to understand the causes, the consequences and trying to avoid its repetition.
- Trying to prevent service failures before they happen by:
  - Adding redundancy where possible:
    - A new fully redundant database is being commissioned;
    - The Detector Control System is also adding redundancy;
    - Virtualization is being explored to add redundancy where it was not in the original design;
  - Carrying out preventive campaigns:
    - Old BIOS batteries where replaced on ageing computers;
    - Ageing SAS controllers, still under warranty, were exchanged;
    - A full clean up and reinstallation of the whole cluster was carried out in 2011 and a selective one is being carried out this year;
- Identifying the failures before they affect the running:

- o   Tests are done to test the redundancy of the services;
- o   A new monitoring system has been implemented;
- o   Scripts have been created to report problems (such as the SAS controller one) before they affect the running;
- •   And testing before deployment. There are two testing and integration areas and two more are being built for specific purposes.

Inevitably failures happen. The rate of failure for the past year of operation, enumerating in decreasing order of largest effort needed to recover, is summarized in table 2. The computer failures include the SAS ageing problem that was responsible for most of them until the replacement campaign was carried out.

Table 2. Failure rates for the past year of operation.

Failure type	Failure rate
Unscheduled power cuts	~1 every 2 months
Unscheduled cooling cuts	~1 every 3 months
Failure of network switches	~1 every 3 months
Computer failures	~3 every week

3.Summary

The CMS online cluster has been working successfully since 2007 during the commissioning phases and runs of the CMS detector.

It consists presently of more than 2700 computers, it will be expanded in the next months to 3000. These computers run applications implementing the control of the experiment, the event building, the high level trigger, the online database and the control of the buffering and transferring of data to Tier 0 at CERN.

The design of the cluster and networking allows the CMS DAQ system to handle large data volumes up to 100 Gbytes/s.

The chosen IT infrastructure and services provide high availability and allow autonomous and independent running of the cluster.

The design allows for easy scalability in order to accommodate further expansions.

Acknowledgements

References

[1] The CMS Collaboration (Adolphi R et al.) "The CMS Experiment at CERN LHC", *JINST* **3** S08004 361, 2008.

[2] The CMS Collaboration, "The Compact Muon Solenoid Technical Proposal", CERN/LHCC 94-38, 1994.

[3] The LHC Study Group, "The Large Hadron Collider Conceptual Design Report", CERN/AC 95-05, 1995.

[4] The CMS Collaboration, "The CMS Physics Technical Design Report, Volume 1", CERN/LHCC 2006-001, 2006. CMS TDR 8.1.

[5] The CMS Collaboration, "The CMS Physics Technical Design Report, Volume 2", CERN/LHCC 2006-021, 2006. CMS TDR 8.2.

[6] The CMS Collaboration, "High Density QCD with Heavy Ions; Physics Technical Design Report, Addedum 1", CERN/LHCC 2007-009, 2007. CMS TDR 8.2-Add1

[7] The CMS Collaboration, "Trigger and Data Acquisition Project: Level-1 Trigger,", CERN/LHCC 2000-38, Vol. I, 2000.

[8] The CMS Collaboration, "The Trigger and Data Aquisition project", CERN/LHCC 2002-26, Vol II, 2002.

[9] G Bauer et al., "The CMS event builder and storage system", 2009 Prepared for International Conference on Computing in High Energy and Nuclear Physics (CHEP 2009), Prague, Czech Republic, 21-27 Mar 2009.

[10] http://www.myricom.com/myrinet/overview/

[11] G Bauer et al., "The CMS Online Cluster: IT for a Large Data Acquisition and Control Cluster". International Conference on Computing in High Energy and Nuclear Physics, Prague, Czech Republic, 21-27 Mar 2009. J.Phys.Conf.Ser.219:022002,2010.

[12] http://www.oracle.com/technology/products/database/clustering/index.html

[13] Bauer G et al. 2008, "The run control and monitoring system of the CMS experiment", *J. Phys. Conf. Ser.* (CHEP 07, Victoria, BC, Canada, 2-7 Sep 2007) 119 022010.

[14] G Bauer et al., "Monitoring the CMS data acquisition system", International Conference on Computing in High Energy and Nuclear Physics (CHEP 2009), Prague, Czech Republic, 21-27 Mar 2009. doi:10.1088/1742-6596/219/2/022042.

[15] http://www.force10networks.com/

[16] http://www.netapp.com/

[17] http://www.netapp.com/us/products/platform-os/snapshot.html

[18] http://www.netapp.com/us/products/platform-os/dedupe.html

[19] http://www.intel.com/design/servers/ipmi/index.htm

[20] Quattor: Tools and Techniques for the Configuration, Installation and Management of Large-Scale Grid Computing Fabrics. Journal of Grid Computing Vol. 2/4. Springer Verlag, December 2004.

[21] G Bauer et al., "The CMS data acquisition system software", International Conference on Computing in High Energy and Nuclear Physics (CHEP 2009), Prague, Czech Republic, 21-27 Mar 2009 doi:10.1088/1742-6596/219/2/022011.

[22] "The Preboot Execution Environment specification v2.1" published by Intel & Systemsoft. http://download.intel.com/design/archives/wfm/downloads/pxespec.pdf

[23] http://www.redhat.com/docs/manuals/linux/RHL-7.3-Manual/custom-guide/ch-kickstart2.html

PoS(ISGC 2012)023

[24] http://www.nagios.org/

[25] https://www.icinga.org/

[26] http://labs.consol.de/nagios/mod-gearman/

[27] http://my-plugin.de/wiki/projects/check_multi/start

[28] http://www.pnp4nagios.org/

[29] http://oss.oetiker.ch/rrdtool/doc/rrdcached.en.html

[30] http://www.ocsinventory-ng.org/

[31] http://www.glpi-project.org/

PoS(ISGC 2012)023