UPGRADE OF NETWORK INFRASTRUCTURE OF THE KIPT COMPUTING FACILITY FOR CMS DATA PROCESSING

A.A. Kurov

National Science Center “Kharkov Institute of Physics and Technology”, Kharkov, Ukraine

E-mail: aakurov@gmail.com

Based on the analysis of the network infrastructure of the NSC KIPT specialized computer complex for processing of data obtained in the CMS experiment at the Large Hadron Collider, the internal communication channels of the complex were modernized using the link aggregation technology. The upgrade has provided a significant increase of the information transfer rate for the most loaded links. Extension of the external communication channel bandwidth to 2 Gbps and optimization of the routing for the transferred information using the GEANT and LHCONE network infrastructures has led to an improvement in the quality of data transfers by the PhEDEx and XrootD services. The optimal load distribution between the separate external channels and the Internet channel bandwidth subdivision between the PhEDEx and XrootD services have been provided. A set of services have been configured to monitor various parameters of the network.

PACS: 07.05.−t; 29.50.+v

INTRODUCTION

The specialized CMS KIPT computing facility (CF) [1] was built to participate in the distributed processing of data accumulated in the CMS experiment [2] at the Large Hadron Collider (LHC). The CF operates at the Tier 2 (T2) of the hierarchical CMS grid infrastructure, which is a part of the Worldwide LHC Computing Grid (WLCG), and allocates its resources for physics analysis of the CMS data and Monte Carlo (MC) simulation of the experiment.

Since CMS T2 centers are responsible for storing and processing the data, they should follow certain hardware and software requirements. In particular, for the first years of LHC operation, a minimum local (LAN) and wide (WAN) network connectivity of a T2 site should be 1 Gbps [3].

The LHC Run 2 which started in 2015 is characterized by a considerable increase of experimental information being transferred to the WLCG structures. The collider has already reached its nominal luminosity. (In 2017, the LHC peak luminosity of 22 Hz/nb was achieved.) The CMS high-level trigger (HLT) selects events for the offline processing at rate of $\sim 10^3$ Hz, what exceeds the typical HLT rate during LHC Run 1 more than twice.

The new conditions essentially harden the requirements for the CMS grid infrastructure. Accordingly, the CMS computing model undergoes certain evolution [4], with partial smearing functional distinctions between the different infrastructure tiers. In addition, the ‘Any Data, Anytime, Anywhere’ (AAA) project is being implemented to allow an analysis job that runs at a T2 (or T1) site to process a dataset stored at another (remote) center of the infrastructure. Certainly, this implies a proper upgrade of the WAN and LAN parameters for the sites, with a substantial increase of the corresponding bandwidth. In particular, the WAN/LAN connectivity of 10 or even 100 Gbps (for sites having several thousand CPU cores) should be met by CMS T2 centers [5].

Despite the fact that the KIPT CF has been featured with a fair level of stability and reliability of operation as a CMS T2 center since the LHC startup, its hardware parameters were significantly below the CMS requirements for a “nominal” T2 site of the experiment. In particular, its WAN and LAN bandwidth did not exceeded 1 Gbps. So, an upgrade of the CF network infrastructure was needed to mitigate the impact on the system operation of the increased CMS data transfer rate during LHC Run 2.

The details of this upgrade are outlined in the present paper. The result of connection quality monitoring showed.

1. UPGRADE OF THE CF LOCAL NETWORK INFRASTRUCTURE TO MEET INTENSIVE DATA EXCHANGE CONDITIONS

The CF local network is built using three switches, to which computer nodes and disk servers are connected, as well as a router providing external communications. The result of monitoring the quality of communication between the elements of network equipment of the local network has shown that some of the local 1 Gbps channels work with full load. The external 1 Gbps communication channel has been also overloaded for a long time.

The congestion of LAN links emerged as a consequence of information exchange between two network switches, one of which provided connection of the most productive computation nodes – the ones that contribute more than 50% to the total CF computational capacity and, correspondingly, run the greatest number of data analysis and MC jobs. These nodes require the transfer of data stored on both the local CF disk servers and the servers of the remote centers of the CMS grid infrastructure accessible through the external Internet channel. Data for the jobs can be transferred using the AAA architecture, which is based on the XrootD network service [10].

Significant congestion of the channels (Fig. 1) was also observed for the most capacious disk server of the CF’s distributed disk mass storage system (MSS) with a disk space of 247 TB (40% of the total MSS capacity). This server had two 1 Gbps network connections, one of which was used to transfer data to CF local computing nodes using the XrootD service. The second connection was used for the CMS experimental data exchange via the Internet channel through the ‘Physics Experiment Data Export’ (PhEDEx) [11] system, and for data transfers to remote CMS grid centers using the XrootD service. The congestion was caused by a great number of accesses to this server from local and remote computational nodes and a simultaneous intensive data flow governed by the PhEDEx system.

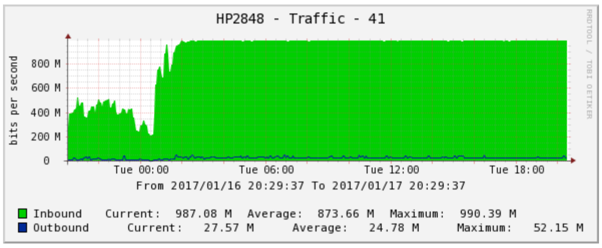

Fig. 1. Daily traffic on the switch port enabling 1 Gbps communication with the most capacious MSS disk server (incoming and outgoing traffic is shown by a shaded diagram and a solid line, respectively)

To solve the problem of channel congestion, we utilized the aggregation technology [12] which allowed us to combine several 1 Gbps Ethernet channels into one logical channel, the maximum throughput of which is the sum of bandwidths of the combined channels. Uniform traffic distribution between the channels, as well as an increased fault tolerance due to the ability to detect a failed channel and redirect the traffic to the working channel is provided by the Link Aggregation Control Protocol (LACP). The LACP protocol is supported by a vast majority of managed switches.

Two new 1 Gbps connections were added to the existing channel between the switches, all combined together into a single logical channel with a total bandwidth of 3 Gbps. Fig. 2 shows the monitoring of traffic through the created aggregated channel during 3 days. It follows from the figure that, firstly, the rate of the information exchange between the switches at separate intervals is really so great that the use of a 1 Gbps channel was a serious factor limiting the overall system performance. Secondly, the threefold increase in the throughput allowed us (at least at this stage) to resolve this problem.

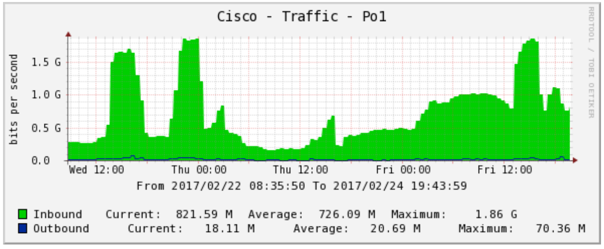

Fig. 2. Traffic through the aggregated communication channel between the switches (the notation is the same as in Fig. 1)

Similarly, with the help of aggregation technology, the network connections of the highest-capacity disk server were extended to 2 Gbps. The corresponding setting was done using the 'bonding' driver [13], which is part of the Linux kernel.

2. OPTIMIZATION OF THE CF EXTERNAL COMMUNICATION CHANNEL

It should be noted that the CF external communication channel is used for a practicalland this requirement y continuous exchange of information with other grid infrastructure centers and is involved in solving several tasks simultaneously. Through the external interface, the CF receives jobs that perform both computer simulation of the CMS experiment (generation of MC events) and processing of real CMS data. In addition, the CF participates in bidirectional transfers of the experimental data through the PhEDEx system and within the framework of the AAA architecture. The intensity of this exchange increased significantly after the LHC Run 2 startup. A typical example of the external communication channel congestion in 2016 is shown in Fig. 3. Note that CMS Tier 2 (T2) centers should meet rather hard requirements. In addition to the availability of necessary computing and disk resources, the "nominal" T2-center of the CMS grid infrastructure should have an Internet uplink bandwidth of ~10 Gbps. Inability to provide such a throughput for the KIPT CF (see Section 1) led to an extensive overload of the external channel and full its saturation with traffic at certain time intervals, as shown in Fig. 3.

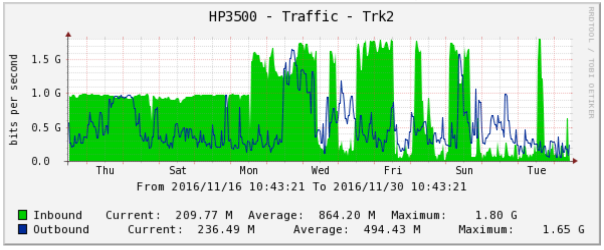

Fig. 3. Traffic through the external CF interface before and after (starting from Monday, 21 November 2016) setting up the aggregated communication channel (the notation is the same as in Figs. 1 and 2)

To solve this problem, a second optical link using a pair of 1 Gbps SFP modules was added to the existing (based on media converters) 1 Gbps connection between the CF router and the provider switch. Analogously to upgrade of the CF local network infrastructure, these connections were combined into an aggregated channel with a total rate of 2 Gbps.

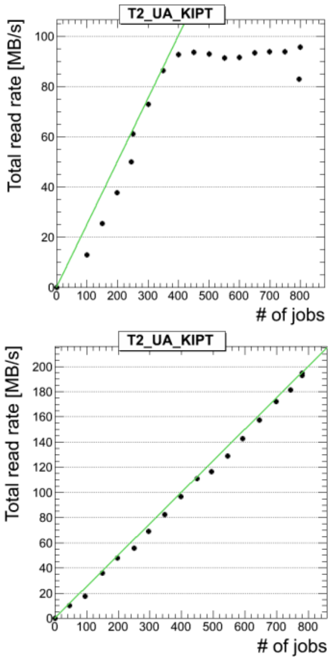

Fig. 4. XrootD data transfer rate for KIPT CF before (top) and after (bottom) upgrade of the external channel (see [14])

The Internet channel bandwidth extension has led to an improvement in the quality of data transfers by the PhEDEx and XrootD services. Fig. 4 shows the dependence of the XrootD data transfer rate on the number of jobs running at remote computing centers of the CMS grid infrastructure and processing experimental data located at the KIPT CF. Previously, the rate had not met the CMS requirement of being at least 150 MB/s for 600 clients. After the extension of the Internet channel, the rate satisfies this condition, as seen in Fig. 4.

The external communication channel upgrade also resulted in a significant improvement in the quality of data transfers governed by the PhEDEx service. This is due to the fact that the file transfer service [15] exploited by the PhEDEx restricts the file transfer time which in case of congestion exceeds this limit. After the upgrade, the transfer time for most files corresponds to this limitation.

3. CONNECTING TO THE LHCONE L3VPN SERVICE

For the exchange of information through Internet communication between scientific and educational organizations of the European continent, the European telecommunication infrastructure GEANT has been created and successfully used [16]. Similar functions on the American continent are performed, in particular, by the Internet2 [17] and ESnet [18] network infrastructures. As a rule, these networks are used by centers of the CMS grid infrastructure to exchange experimental information. To enable connection to these networks is necessary, because it allows one, in particular, to provide communication with centers which, for CMS data transfers, have only blocks of non-routable outside such networks Internet addresses. Note that various scientific institutions have access to these networks, with many of them being not related to the LHC data processing. In the latter case, it is necessary to provide an almost continuous transmission of intensive data flows through Internet channels linking a limited number of centers certified for processing this information. To optimize this flow, the LHCONE L3VPN [19] service was established, with using for these purposes reserved channels of the mentioned above networks.

The access of the specialized KIPT CF that belongs to the NASU telecommunication infrastructure (UARNet) [20] to the capabilities provided by the LHCONE L3VPN service was implemented through a specially configured dedicated local area network (VLAN) which enabled a direct connection with the Ukrainian national GEANT operator URAN [21]. To configure this access, a BGP [22] session with the URAN is established on the KIPT CF router, in which, according to the LHCONE policy, KIPT address blocks directly involved in LHC data processing are announced.

Also, the configured earlier link to the GEANT which provides access to centers operating outside of LHCONE (but connected to GEANT) and the link to the general academic network (UARNet) remain to be used. However, the most intensive information flow is transmitted through the LHCONE, since it is this infrastructure that the majority of the centers, with which the KIPT CF performs CMS data exchange, belongs to (an example of the corresponding traffic is shown in Fig. 5).

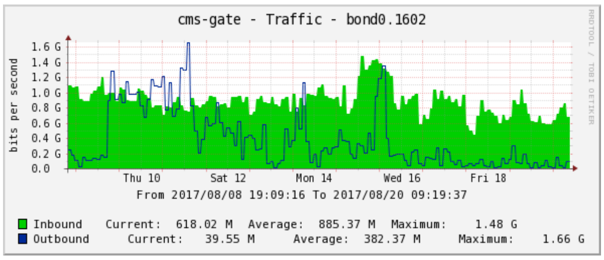

Fig. 5. Traffic of the KIPT CF through the LHCONE link (the notations are the same as in Figs. 1-3)

The LHCONE BGP filtering service provides a possibility to manage the announcement of address blocks to other centers, what enables the most effective use of existing communication channels. For example, using the BGP attribute ‘community 65010: *N*’ allows one to avoid the announcement of network prefixes to a site with autonomous system number *N*.

Another opportunity provided in LHCONE is the use of large-size (“jumbo frame”) packets. The maximum size of the useful data block of one transmitted non-fragmented packet (the so-called “maximum transmission unit” (MTU)) was set to 9000 bytes on all the CF servers and 9100 bytes (as used in the GEANT) on the network interface connected to the URAN. This allowed us to increase the rate of data exchange with the T1 and T2 centers of the CMS grid infrastructure which are configured to use large packets (in particular, T1_ES_PIC, T1_IT_CNAF, T1_US_FNAL, T2_IN_TIFR, T2_US_Caltech, T2_US_Florida, T2_US_Purdue, T2_US_Wisconsin). The most significant increase in the rate occurred during the transfer of experimental information to the centers, the connection with which is characterized by the largest delivery time of single packets.

4. CF TRAFFIC MANAGEMENT

Despite the significant increase in the capacity of external communication channel, it still turned out to be sometimes 100% loaded. For these periods of time, data (from the XrootD service) were transferred from the CF disk servers to jobs running in other CMS grid centers, with the number of XrootD connections that transmit data simultaneously reaching 2000. As a result, the functioning of the PhEDEx service on the CF, which performs independent data exchange (in parallel with the XrootD streams) was disrupted. (Outgoing PhEDEx transfers were interrupted because of exceeding a time limit.)

To solve the problem of outgoing PhEDEx transfers, subdivision of the Internet channel bandwidth between the XrootD and PhEDEx services was configured on the CF router. This restricts the XrootD traffic providing an opportunity to complete the PhEDEx data transfer in the required time.

The bandwidth limit has been set up using the ‘tc’ utility from the ‘iproute2’ Linux package which sets the HTB queue discipline for each of the aggregated interfaces [23] and the classes that determine the maximum traffic rate for the XrootD and PhEDEx services. The traffic for the services was specified using a filter based on packet marking: the XrootD traffic was marked with the ‘iptables’ [24] rule using the TCP port number 1095, whereas for the PhEDEx traffic, we used the range of port numbers configured on the disk servers through the GLOBUS_TCP_SOURCE_RANGE grid-service variable. Due to the difference in the maximum throughput of the external communication channels (VLANs), the marking was also accomplished with respect to the VLAN interface that the packet goes to. Thus, the optimal distribution of the load between the individual channels was carried out.

On the other hand, the performed management of the outgoing traffic has allowed us to solve the problem of packet loss growth, which inevitably arises when communication channels are overloaded.

5. MONITORING OF THE CF NETWORK INFRASTRUCTURE

To provide a necessary quality of grid service operation on the CF, utilities that monitor various parameters of the network functionality have been configured. Monitoring of local traffic on the CF switches has been set up using the Cacti system [25] which exploits the SNMP protocol to obtain traffic data from the ports on the switches and draws diagrams illustrating the information flow for different time periods (see Figs. 1-3, 5).

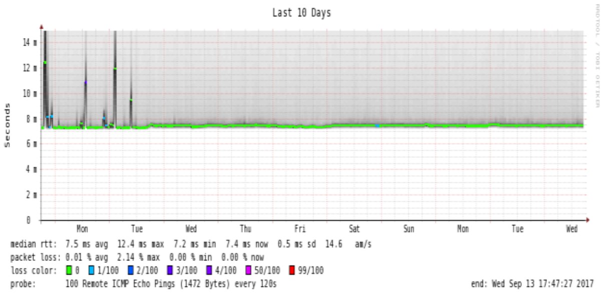

Fig. 6. Visualization of delays and losses for packet exchange with the control point in the UARNet network

The SmokePing system [26] has been also configured on the CF to monitor losses and delays in the delivery of packets with checkpoints both in the local and external networks. Control points in the UARNet and URAN networks were configured for LHCONE channel monitoring. An example of such monitoring for one of the UARNet control points is shown in Fig. 6. In the absence of traffic management (see the previous section), this monitoring identified the problems of significant growth of delays and losses that occurred during channel overloads. In particular, the delays reached 100 ms or even more, and the losses sometimes reached 100%. Fig. 6 shows that, due to the configuration of traffic management, the delay in the delivery of packets increases insignificantly even in case of intensive data transfers from the CF.

A detailed traffic analysis is performed using the NetFlow protocol [27] generated by the router. A dedicated server collects and analyzes the NetFlow data using a set of NFDUMP [28] utilities including 'nfcapd' (the NetFlow data collector) and 'nfdump' (a utility that parses collected data). The NfSen [29] web interface uses the 'nfdump' and 'rrdtool' [30] to visualize the data. Unlike the Cacti, the NfSen allows us to build plots using a number of 'nfdump' filters. These filters were configured to monitor, in particular, the XrootD and GridFTP services and the large-packet ("jumbo frame") traffic – see the example shown in Fig. 7.

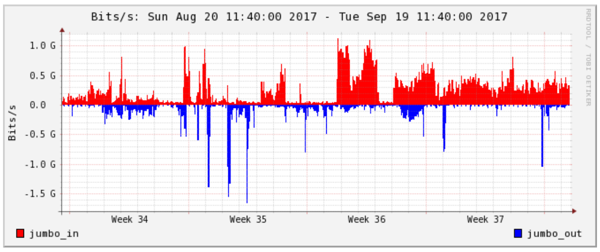

Fig. 7. Incoming (top) and outgoing (bottom) traffic of large-size ("jumbo frame") packets on the CF external interface providing access to the LHCONE network

Monitoring of communication channels between grid sites is carried out using the perfSONAR software [31] consisting, in particular, of the OWAMP service that measures packet delays and losses and the BWCTL service that determines the data transfer rate and the route of packet transmission through the network. The nodes between which the measurements are made must belong to one of the perfSONAR meshes and should be configured using centrally generated configuration files. The perfSONAR software includes a measurement archive and visualization tools for corresponding results. Visualization of these results for a selected perfSONAR mesh is also carried out via the MaDDash system [32]. Also the perfSONAR software includes console utilities that allow one to manually run various quality tests of communication between any given perfSONAR servers, which determine, in particular, the transmission rates and routes for packets transferred both in the forward and backward directions.

On the KIPT CF, two perfSONAR servers have been configured and added to the 'WLCG CMS Bandwidth' and 'WLCG CMS Latency' meshes which, respectively, measure the rate of information exchange between CMS grid sites and test the arising latency and packet losses. To monitor the LHCONE channel on the Kharkov-Kiev network section, tests were set up with the perfSONAR server of the URAN GEANT operator. The results of these tests for one of the two aggregated links are shown in Fig. 8. It can be seen that a stable channel bandwidth of about 1 Gbps is provided, the fraction of losses is within acceptable limits, and delays do not exceed 3.5 ms.

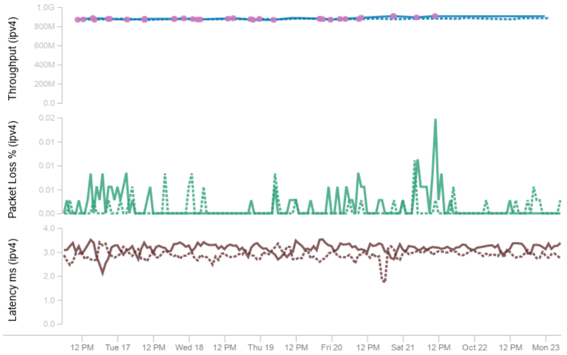

Fig. 8. PerfSONAR tests of the Kharkov-Kiev LHCONE channel section: rate of information exchange between KIPT CF and URAN GEANT operator (top), arising packet losses (middle) and delays (bottom)

CONCLUSIONS

The state of the network infrastructure of the specialized KIPT CF for processing of data accumulated in the CMS experiment at the LHC was analyzed in detail. To solve the problems caused by an inhomogeneous distribution of the computational load on different groups of nodes of the system, as well as by a significantly different intensity of information exchange for these groups, the internal communication channels of the complex have been upgraded. The information exchange between CF nodes has been optimized based on the channel aggregation technology what provided a significant (two-fold) increase of the transfer rate for the most loaded places.

In the same way, an upgrade of the CF external communication channel has been fulfilled. A second optical link using a pair of 1 Gbps SFP modules was added to the existing (utilizing media converters) 1 Gbps connection between the CF router and the provider switch. The achieved Internet channel bandwidth extension to 2 Gbps has led to an improvement in the quality of data transfers by the PhEDEx and XrootD services.

To improve the quality of data exchange between the KIPT CF and other centers of the CMS grid infrastructure, the routing of information transmission was optimized using the GEANT and LHCONE network infrastructures.

To solve the problem of outgoing PhEDEx dataset transfer failures due to an intensive parallel XrootD data flow, subdivision of the Internet channel bandwidth between the XrootD and PhEDEx services has been configured on the CF router. The optimal load distribution between the individual external channels has been performed. The performed management of outgoing traffic also solved the problem of the growth of packet losses during channel overloads.

To provide a necessary quality of functionality of the network infrastructure and grid services, a number of services that monitor various parameters of the network operation is configured in the CF.

The work was supported by grants provided by the National Academy of Sciences of Ukraine (NASU) within the dedicated program “Grid infrastructure and grid technologies for science and scientific applications” and the program of NASU cooperation with CERN and JINR “Nuclear matter in extreme conditions”.

The author thanks L.G. Levchuk for stimulating discussions and helpful comments, as well as A.S. Pristavka, K.A. Klimenko and D.V. Soroka for their help in this work.

REFERENCES

1. S. Zub, L. Levchuk, P. Sorokin, and D. Soroka. Grid middleware configuration at the KIPT CMS Linux cluster // *Nuclear Instruments and Methods in Physics Research*. 2006, A 559, p. 35-37.
2. L.G. Levchuk, D.V. Soroka, M.V. Voronko, S.S. Zub. LCG middleware at the KIPT CMS Linux cluster // *Visnyk Kharkivskogo natsionalnogo universitetu. Series “Matematychne modelyuvannia; informatsiini technologii, avtomatyzovani systemy upravlinnia”*. 2005, № 5 (703), p. 74-86 (in Ukrainian).
3. L.G. Levchuk, P.V. Sorokin, D.V. Soroka, and V.S. Trubnikov. NSC KIPT Linux Cluster for Computing within the CMS Physics Program // *Problems of Atomic Science and Technology. Series “Nuclear Physics Investigations”*. 2002, № 2 (40), p. 49-51.
4. S. Chatrchyan, ..., L. Levchuk, S. Lukyanenko, D. Soroka, P. Sorokin, et al. (CMS Collaboration). The CMS Experiment at the CERN LHC // *Journal of Instrumentation*. 2008, v. 3, № S08004, p. i-334.
5. G. Bayatian et al. (CMS collaboration). The Compact Muon Solenoid Technical Proposal // *CERN/LHCC*. 1994, p. 94-38.
6. H. Newman. A new generation of networks and computing models for high energy physics in the LHC era // *Journal of Physics: Conference. Series 331*. 2011, 012004, p. 1-18.
7. C. Grandi, D. Bonacorsi, D. Colling, I. Fisk, M. Girone. CMS computing model evolution // *Journal of Physics: Conference. Series 513*. 2014, 032039, p. 1-7.
8. K. Bloom. CMS software and computing for LHC Run 2 // *arXiv: 1611.03215v1 [physics. ins-det]*, 2016.
9. Ch. Wissing. CMS networking requirements // *WLCG Workshop*, Manchester, UK, June 2017.
10. CMS Xrootd Architecture // *CERN TWiki*, URL: https://twiki.cern.ch/twiki/bin/view/Main/CmsXrootdArchitecture.
11. PhEDEx – CMS data transfers // *CMSWEB*, URL: https://cmsweb.cern.ch/phedex/
12. H. Frazier, S. Van Doorn, R. Hays, S. Muller, B. Tolley, P. Kolesar, G. Thompson, B. Turner. IEEE 802.3ad Link Aggregation (LAG) // *IEEE 802.3 Higher Speed Study Group Meeting*. 2007, April 17-19, 2007, Ottawa, ON, Canada.
13. Linux ethernet bonding driver // URL: https://www.kernel.org/doc/Documentation/networking/bonding.txt
14. AAA opening and reading tests // URL: http://www.pd.infn.it/~fanzago/TEST/T2_UA_KIPT-xrootd-cms.infn.it.html
15. FTS3 documentation: robust, simplified and high-performance data movement service for WLCG // URL: http://fts3-docs.web.cern.ch/
16. GÉANT // URL: https://www.geant.org/
17. Internet2 // URL: https://www.internet2.edu/
18. Energy Sciences Network ESnet // URL: http://es.net/
19. LHCONE L3VPN Service // *CERN TWiki*, URL: https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome
20. Ukrainian Academic Research Network UARNet // URL: http://www.uar.net/
21. Ukrainian Research and Academic Network URAN // URL: http://uran.ua/
22. A Border Gateway Protocol 4 (BGP-4) // URL: https://tools.ietf.org/html/rfc4271
23. M. Devera. HTB Linux queuing discipline manual // URL:http://luxik.cdi.cz/~devik/qos/htb/manual/userg.htm
24. Netfilter/iptables project // URL: http://www. netfilter. org/projects/iptables/index.html

25. The complete RRDTool-based graphing solution // URL: https://www.cacti.net/
26. SmokePing // URL: https://oss.oetiker.ch/smokeping/
27. R. Hofstede, P. Čeleda, B. Trammell, I. Drago, R. Sadre, A. Sperotto, A. Pras. Flow monitoring explained: from packet capture to data analysis with NetFlow and IPFIX // *IEEE Communications Surveys and Tutorials*. 2014, v. 16, № 4, p. 2037-2064.
28. NFDUMP // URL: http://nfdump.sourceforge.net/
29. NfSen – Netflow Sensor // URL: http://nfsen.sourceforge.net/
30. RRDtool // URL: https://oss.oetiker.ch/rrdtool/
31. perfSONAR // URL: https://www.perfsonar.net/
32. MaDDash – Monitoring and Debugging Dashboard // URL: http://psmad.grid.iu.edu/maddash-webui/index.cgi

Article received 21.02.2018

МОДЕРНИЗАЦИЯ СЕТЕВОЙ ИНФРАСТРУКТУРЫ ВЫЧИСЛИТЕЛЬНОГО КОМПЛЕКСА ННЦ ХФТИ ДЛЯ ОБРАБОТКИ ДАННЫХ ЭКСПЕРИМЕНТА CMS

А.А. Куров

На основе проведенного анализа состояния сетевой инфраструктуры специализированного вычислительного комплекса ННЦ ХФТИ для обработки данных, получаемых в эксперименте CMS на Большом адронном коллайдере, выполнена модернизация каналов внутренней связи комплекса с применением технологии агрегирования, что обеспечило значительное увеличение скорости передачи информации на наиболее загруженных участках. Расширение пропускной способности внешнего канала связи до 2 Гбит/с и оптимизация маршрутизации передаваемой информации с использованием сетевых инфраструктур GEANT и LHCONE привели к улучшению качества передачи данных службами PhEDEx и XrootD. Выполнены оптимальное распределение нагрузки между отдельными внешними каналами и распределение полосы пропускания интернет-канала между службами XrootD и PhEDEx. Сконфигурирован широкий набор служб, осуществляющих мониторинг различных параметров работы сети.

МОДЕРНІЗАЦІЯ МЕРЕЖЕВОЇ ІНФРАСТРУКТУРИ ОБЧИСЛЮВАЛЬНОГО КОМПЛЕКСУ ННЦ ХФТІ ДЛЯ ОБРОБКИ ДАНИХ ЕКСПЕРИМЕНТУ CMS

О.О. Куров

На основі проведеного аналізу стану мережевої інфраструктури спеціалізованого обчислювального комплексу ННЦ ХФТІ для обробки даних, одержуваних в експерименті CMS на Великому адронному колайдері, виконано модернізацію каналів внутрішнього зв'язку комплексу із застосуванням технології агрегування, що забезпечило значне збільшення швидкості передачі інформації на найбільш завантажених ділянках. Розширення пропускної здатності зовнішнього каналу зв'язку до 2 Гбіт/с і оптимізація маршрутизації інформації, що передається з використанням мережевих інфраструктур GEANT і LHCONE, призвели до поліпшення якості передачі даних службами PhEDEx і XrootD. Виконано оптимальний розподіл завантаження між окремими зовнішніми каналами та розподіл смуги пропускання інтернет-каналу між службами XrootD і PhEDEx. Налаштований широкий набір служб, які здійснюють моніторинг різних параметрів роботи мережі.