

JINR (DUBNA) AND PRAGUE TIER2 SITES: COMMON EXPERIENCE IN THE WLCG GRID INFRASTRUCTURE

*J. Chudoba^a, M. Elias^a, L. Fiala^a, J. Horky^a, T. Kouba^a, J. Kundrat^a,
M. Lokajicek^a, J. Schovancova^a, J. Svec^a, S. Belov^b, P. Dmitrienko^b,
A. Dolbilov^b, V. Korenkov^b, V. Mitsyn^b, D. Oleynik^b, A. Petrosyan^b,
N. Rusakovich^b, E. Tikhonenko^b, V. Trofimov^b, A. Uzhinsky^b, V. Zhiltsov^b*

^a Institute of Physics of the Academy of Sciences of the Czech Republic, Prague

^b Joint Institute for Nuclear Research, Dubna

The paper describes the cooperation of JINR specialists and Czech colleagues in the field of deployment and development of modern grid technologies within the framework of the WLCG global infrastructure project. Special attention is paid to the development and usage of monitoring systems.

Статья посвящена опыту совместной работы специалистов ОИЯИ и чешских коллег в области внедрения и развития современных грид-технологий в рамках глобального инфраструктурного проекта WLCG. Особое внимание уделено разработке и внедрению систем мониторинга.

PACS: 02.70.-c

INTRODUCTION

After launching in 2001 the EU Data Grid Project on the creation of grid middleware and carrying out preliminary tests on the initial operational grid infrastructure in Europe, both the Joint Institute for Nuclear Research (JINR) and the Institute of Physics of the Academy of Sciences of the Czech Republic (FZU) got involved in the international grid activities [1, 2]. Since year 2003, JINR and FZU have been taking active part in the large-scale WLCG (Worldwide LHC Computing) Grid Project, the EGEE (Enabling Grids for E-sciencE), and the EGI (European Grid Infrastructure) Projects. During all these years, JINR and FZU have actively cooperated in creating and exploiting their own grid infrastructures. The complexity of the structures of the computing centers at both institutes required the development of local monitoring systems, the quality of which was enhanced by mutually helpful experience exchange and joint meetings held at JINR and FZU.

While work was done in close cooperation, the targets of the monitoring developments were different at the two sites and resulted in specific implementations, tailored to the existing hardware configurations. Both kinds of monitoring solutions secure permanent profiling of the parameters of the physical components of the hardware and do prospective analyses of their expected behaviors in the near future. If an analysis prognosticates expected failure, warnings

are issued (via predefined communication channels) to the subsystem experts charged with the supervision and solution implementation in case of emergency. Automatic measures designed for the system safeguard may also be taken.

The present paper provides parallel descriptions of the most important features of the monitoring systems developed at the two Tier2 sites, emphasizing in-house developments.

JINR GRID SITE AS WLCG TIER2 CENTER

The JINR Central Information and Computing Complex (CICC) is logically constructed as a uniform data-processing resource for all directions of the research activities underway at JINR. All computing and data storage resources are served by unified base software, which allows one to use the CICC resources both in international projects of distributed computing (WLCG, PANDA-GRID, and CBM) and local projects carried out by the JINR users. The system software is optimized to provide the most effective use of the computing resources under as much as possible protected, and at the same time most universal, access to the data. The basic operating system of nearly all the CICC components is OS Linux. The basic system of storing large amounts of data in CICC is the dCache-based hardware-software system.

In 2004, the CICC JINR was fully integrated into the WLCG/EGEE environment as the grid site called «JINR-LCG2». It serves within WLCG infrastructure as a Tier2 level center for the LHC experiments. Apart from maintaining JINR-LCG2 site, part of its servers fulfills important services and functions supporting the Russian grid segment. Currently the share of the computing resources of the JINR grid site within the regional Russian WLCG community makes up 40 to 50 percent and slightly less for data storage resources [3–5].

By year 2012, the JINR grid infrastructure encompassed a farm of 2,102 processor slots, more than 50 specialized servers, 2 disk data storage systems on dCache basis and 3 disk data storage systems on XROOTD basis total volume of nearly 1,234 terabyte, and 5 interactive machines for users' work. The current version of the middleware, the gLite 3.2, as well as Scientific Linux OS v.5 are used.

The internal network of the JINR-LCG2 site is constructed on the 10 GbE basis, and aggregation of links to 8×1 GbE is used to increase network bandwidth to individual services and group switching Ethernet modules. The site's central network switch is connected to the JINR's edge switch at a rate of 10GbE. JINR has a direct channel to MSK-IX (Moscow Internet eXchange) at a rate of 2×10 GbE with 10 GbE devoted to Tier2 connections.

The CICC JINR provides the following basic grid services: Berkeley DB Information Index (top level BDII), site BDII, Computing Element (CREAMCE), Proxy Server (PX), Workload Management System (WMS), Logging&Bookkeeping Service (LB), Accounting Processor for Event Logs (APEL publisher), LCG File Catalog (LFC), Storage Element (SE), User Interface (UI) and specialized grid services (ROCMON; VO boxes for ALICE, CMS, CBM and PANDA; Storage Element of XROOTD type for ALICE; Frontier — caching service for access to CMS and ATLAS central databases). For the LHC virtual organizations, the actual versions of the specialized software are supported: AliROOT, ROOT and GEANT for ALICE, atlas-offline and atlas-production for ATLAS, CMSSW for CMS, and DaVinci and Gauss for LHCb. The VO access is secured through the global file system.

For the time being, CICC JINR, as a grid site of the global grid infrastructure, supports computations of 11 Virtual Organizations (alice, atlas, bes, biomed, cms, dteam, fusion, hone,

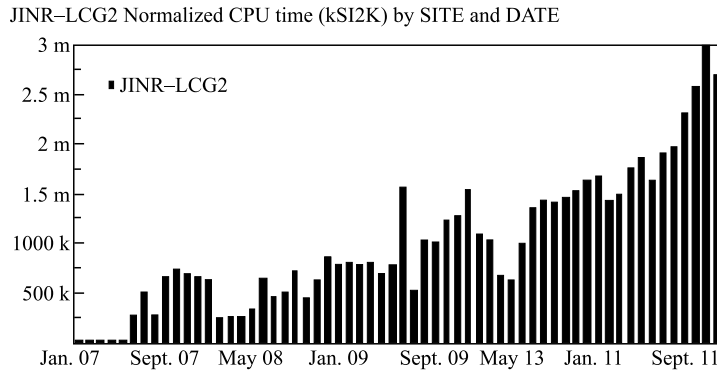


Fig. 1. CPU time (in kSI2K units) provided to virtual organizations at the JINR grid site for the period January 2007 through December 2011

lhcb, rgstest and ops), and opened access to grid resources for the future CBM and PANDA experiments. The basic users of the JINR grid resources are the virtual organizations of the experiments at LHC (ALICE, ATLAS, CMS, and LHCb). The CICC JINR resources were used during the constructing phase of the experiments at the LHC and now, at the running phase of the experiments, they are used for the mass event Monte Carlo production, physics analysis and storage of data replicas of large volumes. Figure 1 illustrates the CPU consumption within 2007–2011 period.

THE CICC LOCAL MONITORING SYSTEM

JINR has rich experience in performing Grid monitoring activities [6–10]. The CICC centralized local monitoring provides clock monitoring of all resources, timely warning of failures and comprehensive analysis of the complex [11]. The system was developed based on the open-source Nagios software product; a number of add-ons and plugins (pluggable software modules) are written specially for the needs of the CICC. It is based on an architecture characterized by universality, scalability, and flexibility for use in the development of various monitoring instruments. Objects of local monitoring are organized in three levels in accordance with particular characteristics of their monitoring and maintenance:

1) At the bottom, there is the hardware level, which includes data collection and display on individual nodes of the network, their hardware and software, their network accessibility check, processor and memory load, state of power supply, temperature control, etc.

2) The network level takes into account devices and services that support the operation of local networks as well as the availability of external networks necessary for the operation of the complex.

3) The top service level system controls the functionality of services provided to end users:

- basic services, such as SMTP, POP, DNS, E-mail (using standard Nagios plugins);
- dCache filesystem (based on scripts running through NRPE to collect necessary metrics);

- gFTP service;

- RAID disk arrays (3Ware, Adaptec), supplying dCache work (with the help of monitoring tools from RAID manufacturers integrated into a single plug-in).

The current web interface of the system was implemented mainly in PHP using XHTML and Javascript for page layout. AJAX (jQuery) was used for building real-time graphics. The main mode represents an object tree, where the upper level nodes are hostgroups, and branches of the tree are single hosts and, in certain cases, individual services. The web interface of the system is available at <http://litmon.jinr.ru>. A special form at the home page allows one to find the desired object using a network name or address.

In case of emergency, the system sends alerts, via email or SMS, to the persons responsible for resolving problems. Data obtained via this monitoring system has repeatedly contributed to identifying, locating and troubleshooting CICC services as well as optimizing its individual elements.

High quality of work of such a system is a common ground for the global grid monitoring, ensuring the correct operation of the site based on controlled infrastructure and providing relevant information on its work at higher levels of monitoring. Data provided by the service has great importance both for the network administrators, who are responsible for the equipment and channels, as well as for the developers and users of the service grid.

Figures 2 and 3 show visualization snapshots in the local monitoring system.

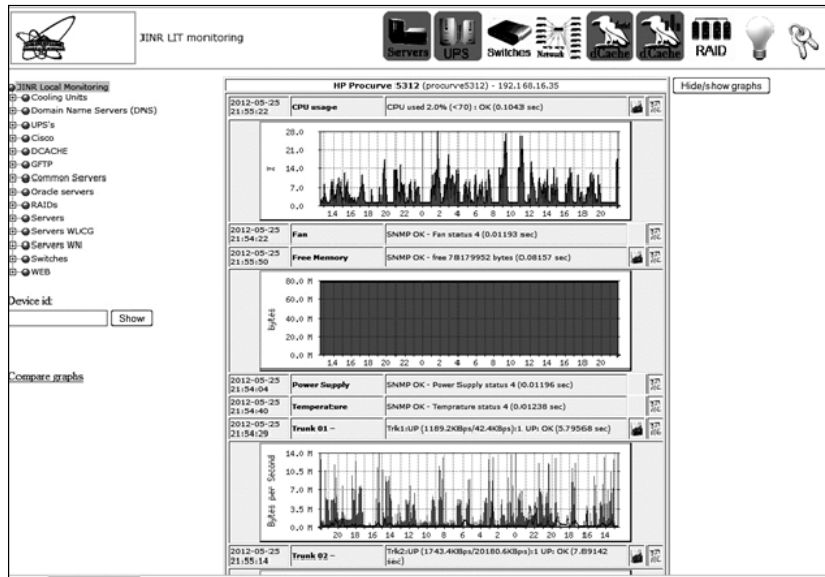


Fig. 2. The main screen of the local monitoring system




Fig. 3. A number of dCache transfer processes (maximum, average and current values) during a 36-hour time period

FZU COMPUTING CENTER BRIEF HISTORY

The Institute of Physics of the Academy of Sciences of the Czech Republic (FZU) has participated in the delivery of HEP computing capacities since 1998. Early computations were done on Unix workstations, and data occupied 0.3 TB. In 2001, our computing center joined EDG, and since 2004, all computing HW has been located in a dedicated fully-featured server room (UPS, air conditioning, raised floor). The new server room helped computing center become a Regional Center for Computing in High Energy Physics and Tier2 for the ATLAS and ALICE experiments. Our center is called **praguelcg2** in central grid databases. The evolution details till 2009 were presented at the NEC 2009 conference [12]. At the beginning of 2010, the server room was equipped with water cooling. This step, the protection it provides and consequences were presented at GRID2010 [13].

Since mid-2012, we have operated about 400 servers with approximately 30,000 HEPSPECs and 2 PB of disk storage space. Our main users are still D0 and two WLCG experiments: ATLAS and ALICE. Other VOs are supported via our participation in the European Grid Initiative, namely Auger and CTA.

Our main network provider is CESNET — Czech Research Network Provider. We are connected to the Internet via 10 Gbps network. A second 10 Gbps link is dedicated to connections to grid participants: BNL USA, FNAL USA, Academia Sinica Taiwan, and several potential Tier3 sites in the Czech Republic. One of the busiest links is a dedicated 1 Gbps link to ATLAS Tier1 center in Karlsruhe. This link is often fully saturated and should be upgraded to 10 Gbps.

PRAGUELCG2 MONITORING

The monitoring at our site is built exclusively on open-source products. There are five main parts that we use on a daily basis:

- Nagios
- Munin
- RRD graphs created by our custom scripts
- MRTG + Weathermap
- Netflow

Nagios is used in two instances. We have one for monitoring hardware health (temperatures, RAID arrays), system health (up-to-date packages, machine load, memory consumption), and basic software health (services running, opened and listening network ports, batch client running, CVMFS status). Most of the checks are active, which means that the Nagios server is the initiator of the checking and waits for the result of the check. However, we also use some passive checking, which means that another software pushes health status to Nagios. This approach is mainly used for getting information from central syslog that monitors system logs for various error messages and for getting status from SNMP traps sent by various hardware (UPS, thermometers, air conditioning, disk arrays). We receive Nagios notifications via email and SMS only for important services. The problems on less important nodes are reported three times a day by a summarizing script.

This instance of Nagios currently monitors 466 hosts and 7,174 services on these hosts.

The second Nagios instance we manage runs SAM (Service Availability Monitoring) for sites participating in NGLCZ (National Grid Initiative for the Czech Republic — part of EGI

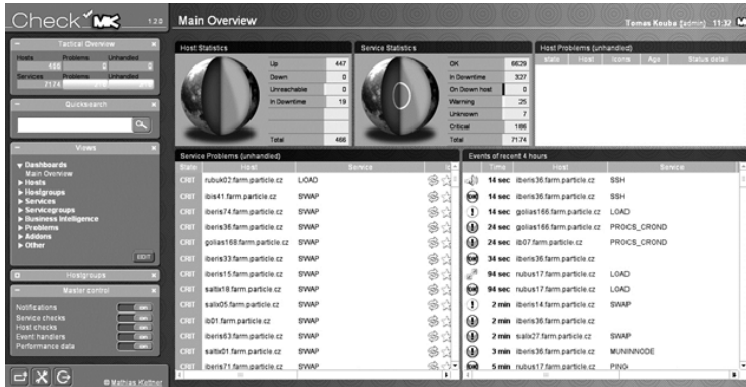


Fig. 4. Multisite dashboard with data from pragueclg2 Nagios

project). This one tells us how our site is doing from the perspective of WLCG data transfers and job execution.

When working with Nagios, we like the alternative user interface called multisite (Fig. 4). It is especially good for massive operations, like putting whole groups of nodes in downtime or rescheduling particular check on many nodes with only few clicks.

For trend graphs, we do not currently use pnp4nagios, which is very popular among Nagios users, but we use Munin. It is mainly because of its simplicity in creating new graphs and availability of many existing sensors (metrics).

Munin runs as a cronjob on the monitoring server and periodically polls all machines for various metrics. The basic ones are part of the Munin package: load, disk occupancy, number of processes, memory consumption. We have developed some other metrics for our specific needs. Some of the interesting ones are DPM tokens usage, accounting of network transfers

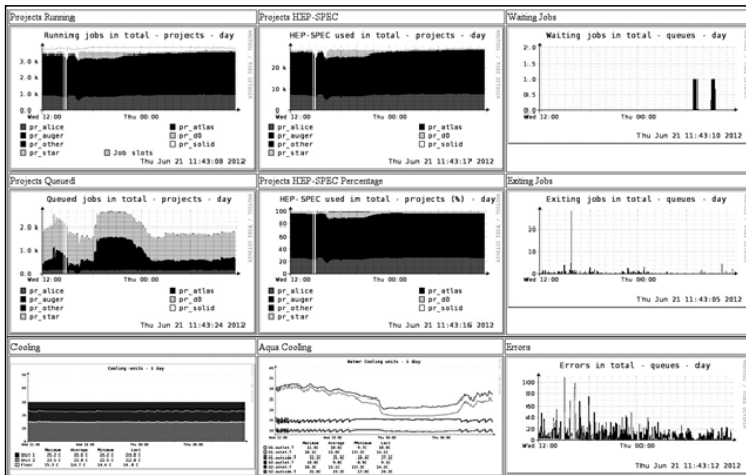


Fig. 5. Set of graphs created by the custom scripts

(using iptables), number of log lines added to DPM/LFC/WMS logs, so we can correlate the load of the machine with the actual service work.

Munin is very good at graphing metrics of running system, but sometimes we need graphs from special facility hardware (UPS, air conditioning, water cooling), or graphs to have special attributes, i.e., size, scale, multiple sources of data, and different sampling period (Munin has 5 minutes). In these cases, we use RRD with custom scripts directly. In Fig. 5, there is a set of our graphs related to running and waiting jobs, grouped by experiments and scaled to HEPSPREC usage. There is also the status of air and water cooling.

We monitor some network aspects by Nagios: switch health and SNMP traps from switches. Other network aspects are monitored by Munin, namely traffic on server's network interfaces, error rate on network interfaces, and number of open connections. However, these are just bits that, on the one hand, do not give an overall picture on how network is utilized and, on the other hand, do not show detailed information needed for debugging deeper problems.

The overall picture of our network (Fig. 6) is created with MRTG and Weathermap. MRTG collects performance data from switches and draws RRD graphs. Weathermap collects the topology of the network and draws a picture of this topology with associated MRTG graph. An example showing our central CISCO router can be seen in Fig. 7. More detailed description of network monitoring was presented at CHEP 2010 [16].

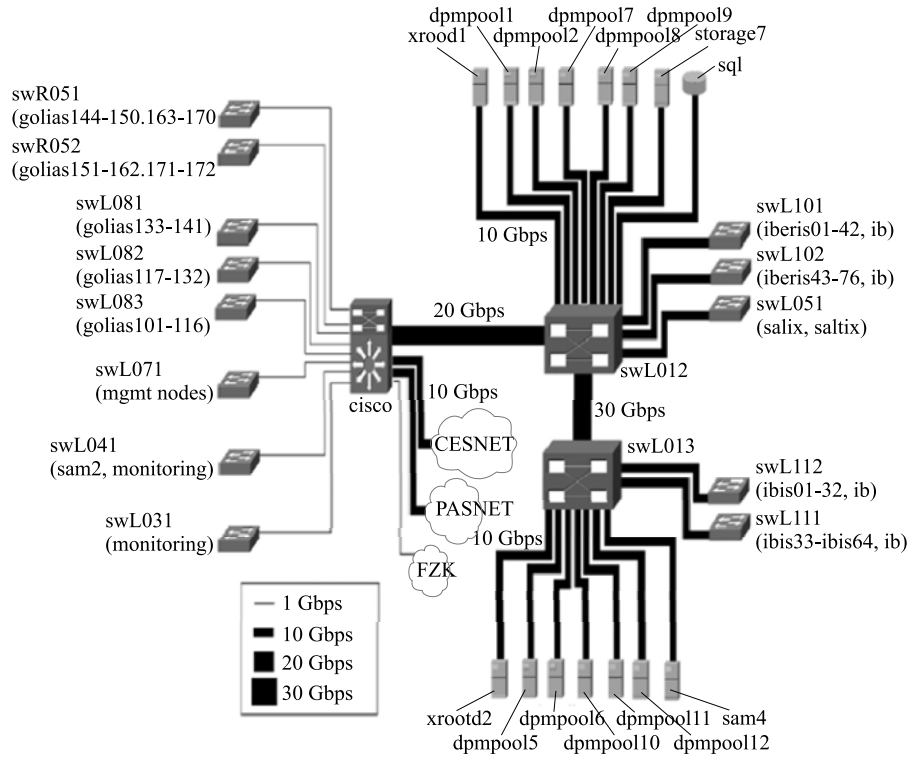


Fig. 6. PragueLcg2 network




Fig. 7. The main CISCO router

For debugging deeper network problems, it is useful to keep record of all transfers between important network parts (routers and main servers). This task is performed by FlowTracker. It is an open-source software for collecting data generated in netflow standard [17]. It creates a database of all transfers so that we can see what machines were communicating with each other, how many packets were transferred, what ports were used, etc. The FlowTracker tool is complemented with FlowGrapher, a graphical tool. An example output can be seen in Fig. 8.

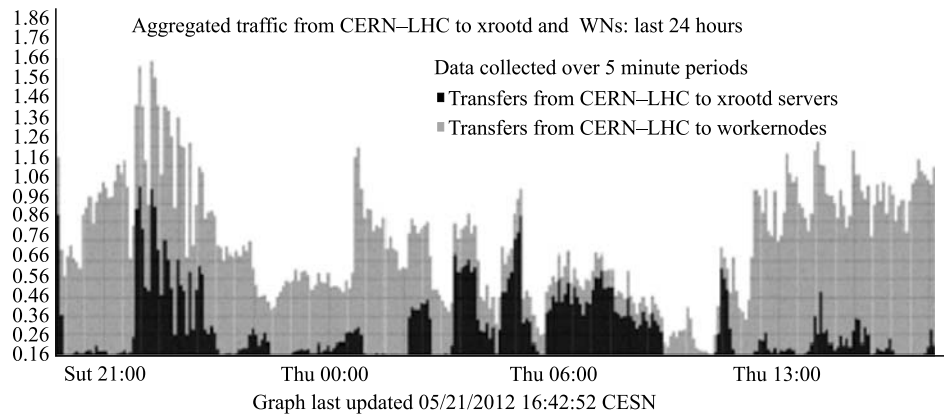


Fig. 8. FlowGrapher showing transfers from CERN to FZU

It shows data coming from CERN to our worker nodes (red area) and data coming to our rootd servers (green area). This is actually the reason why our ALICE jobs have low efficiency — they download too much data from remote sources and spend too much time in I/O waiting state.

Detailed aspects of monitoring at prague1cg2 were presented at NEC 2009 [14].

COMPUTING FACILITY MANAGEMENT

Since 2008, our services have been managed by CFengine version 2 (<http://cfengine.com/>). It is a tool for declarative description of configuration of every node. This system allows the administrator to define what packages should be installed, what services should run, and many other aspects of a running system. The system ensures that all the defined aspects (the so-called promises) are fulfilled or reports difficulties related to the promises.

We try to configure as many services as possible via CFengine as well as keep this configuration in subversion repository. This helps us both in auditing what is configured and in searching history of changes.

The usage of CFengine has two consequences in our monitoring setup:

1. We have to monitor every node if the CFengine client is executed correctly, and all configured steps are really executed on the node.

2. We can use CFengine to fix automatically the problems that Nagios detects.

Point one has been solved with our cfagent Nagios sensor. It is a python script that checks CFengine logs for fresh records (and signals error if the log is too old, telling us that CFengine is probably dead). The sensor also checks the logs for unsatisfied promises and signals warning in such a case.

Point two means that we can use CFengine to fix problems proactively. For example, if srmv2.2 daemon dies on our DPM head node, we get a message from Nagios; however, CFengine knows that this daemon should be running, so it starts the daemon when the CFengine runs again (it runs every half an hour).

The detailed description of our CFengine setup and experience were presented at GRID2010 conference in Dubna [15].

SUMMARY

The carefully tested monitoring infrastructure was further developed and systematically implemented in the JINR and FZU computing centres both to the computing and storage grid infrastructure solutions. The JINR and FZU teams prepared and implemented monitoring environment that was tested in JINR and FZU computing centres. As a result of construction and operation of Tier2 centers at JINR and FZU, the proper conditions have been provided for the physicists of the both institutes for their full-scale participation in the experiments at the LHC running phase. Staff members of JINR and FZU have also been actively involved in the study, use and development of advanced grid technologies. The cooperation of the two institutions in these activities has played a significant and positive role in these activities. The results of common work were published and reported at several conferences and workshops, in particular, at symposiums on Nuclear Electronics and Computing (NEC) in Varna, Bulgaria

(in 2009 and 2011), GRID conferences in Dubna, Russia (in 2008 and 2010), CHEP2007 (Victoria, Canada), CHEP2009 (Prague, Czech), CHEP2010 (Taipei, Taiwan), and HEPIX workshops (Ithaca, NY, USA, 2010; Vancouver, Canada, 2011; GSI Germany, 2011).

REFERENCES

1. *Korenkov V., Tikhonenko E.* The GRID Concept and Computer Technologies during an LHC Era // Part. Nucl. 2001. V. 32. P. 6.
2. *Korenkov V., Mitsin V., Tikhonenko E.* The GRID Concept: On the Way to a Global Informational Society of the XXI Century. JINR Preprint P11-2001-58. Dubna, 2001.
3. *Ilyin V., Korenkov V., Soldatov A.* RDIG (Russian Data Intensive Grid) e-Infrastructure: Status and Plans // Proc. of Conf. «NEC'2009». Dubna, 2010. P. 150.
4. *Korenkov V.* GRID Activities at the Joint Institute for Nuclear Research // Proc. of Conf. «GRID'2010». Dubna, 2010. P. 142.
5. *Belov S. D. et al.* Grid Activities at the Joint Institute for Nuclear Research // Proc. of «NEC'2011». Dubna, 2011. P. 68.
6. *Andreeva J. et al.* Dashboard for the LHC Experiments // J. Phys. Conf. Ser. 2008. V. 119. P. 062088.
7. *Uzhinski A., Korenkov V.* Monitoring System for the FTS Data Transfer Service of the EGEE/WLCG Project // Calc. Meth. Program. 2009. V. 10. P. 96–106 (in Russian).
8. *Belov S. D., Korenkov V. V.* Experience in Development of Grid Monitoring and Accounting Systems in Russia // Proc. of «NEC'2009». Dubna, 2010. P. 75.
9. *Andreeva J. et al.* Job Monitoring on the WLCG Scope: Current Status and New Strategy // J. Phys. Conf. Ser. 2010. V. 219. P. 062002.
10. *Andreeva J. et al.* A Tier3 Monitoring Software Suite (T3MON) Proposal. ATL-SOFT-PUB-2011-001. CERN, 2011.
11. *Dmitrienko P. V. et al.* Development of the Monitoring System for the CICC JINR Resources // Scientific Report of the Laboratory of Information Technologies (LIT): 2010–2011. Dubna, 2011. P. 20.
12. *Lokajicek M. et al.* Prague Tier2 Computing Centre Evolution // Proc. of Conf. «NEC'2009». Dubna, 2010. P. 8.
13. *Kouba T. et al.* Prague WLCG Tier2 Report // Proc. of Conf. «GRID'2010». Dubna, 2010. P. 159.
14. *Kouba T.* Prague Tier2 Monitoring Progress // Proc. of Conf. «NEC'2009». Dubna, 2010. P. 94.
15. *Kouba T. et al.* A Centralized Administration of the Grid Infrastructure Using CFEngine // Ibid. P. 106.
16. *Elias M. et al.* Network Monitoring in the Tier2 Site in Prague // J. Phys. Conf. Ser. 2011. V. 331. P. 052024.
17. *Clais B.* Cisco Systems NetFlow Services Export Version 9. RFC 3954. IETF, 2004. <http://www.ietf.org/rfc/rfc3954.txt>.

Received on September 18, 2012.