

# Capacity and Performance Analysis in Cloud Computing

UDC 681.3.016=111

*S.V. Mescheryakov, D.A. Shchemelinin*

## **ANALYTICAL OVERVIEW OF ZABBIX INTERNATIONAL CONFERENCE 2013**

Zabbix International Conference is a growing annual meeting of professionals from various countries and IT companies where Zabbix enterprise-class automated control system is used. Zabbix Conference doesn't have its own proceedings published; only abstracts and presentation media are available on Zabbix web site [1]. This article is an analytical overview of six top presentations that are the most interesting, innovative and valuable for computer science and business.

**AUTOMATED CONTROL; DATABASE; BIG DATA; CLOUD COMPUTING; DISTRIBUTED ENVIRONMENT.**

*С.В. Мещеряков, Д.А. Щемелинин*

## **АНАЛИТИЧЕСКИЙ ОБЗОР МЕЖДУНАРОДНОЙ КОНФЕРЕНЦИИ ЗАББИКС-2013**

Международная конференция по системе Заббикс является быстро развивающимся ежегодным собранием профессионалов из различных стран и IT компаний, где используется автоматизированная система управления масштаба предприятия Заббикс. На Конференции не издаются тезисы докладов, на интернет-сайте Заббикс доступны только краткие аннотации и медиапрезентации [1]. Данная статья содержит аналитический обзор шести наиболее интересных докладов, имеющих научную новизну и практическую ценность.

**АВТОМАТИЗИРОВАННОЕ УПРАВЛЕНИЕ; БАЗА ДАННЫХ; БОЛЬШИЕ ДАННЫЕ; ОБЛАЧНЫЕ ВЫЧИСЛЕНИЯ; РАСПРЕДЕЛЕННОЕ ОБОРУДОВАНИЕ.**

Zabbix automated control system is the enterprise-class solution based on both hardware and software, which can be used for real-time monitoring, alerting, troubleshooting, computer-aided control, capacity analysis and other business purposes in a large distributed production environment [2]. Zabbix can be effectively used for monitoring performance and automated control of multiple hosts in a cloud computing infrastructure. Object-relational databases from different vendors, for example MS SQL, Oracle, and others, can be used to store big data and analyze historical trends.

Zabbix control system is an open source tool and it is free to install at any enterprise. Nevertheless, every implementation of the Zabbix system meets certain problems, such

as monitoring data delay, out of memory outage, network storm, database performance bottlenecks, events visibility, etc. So, any feedback getting from all Zabbix experienced users and the cooperative discussion of production issues are always helpful for further improvement of Zabbix interface and internal architecture.

In 2013, the Zabbix International Conference took place on September 6-7 in Riga, Latvia. About 150 registered participants arrived from more than 20 countries, including USA, Brazil, Japan, UK, Germany, France, Spain, Italy, Austria, Netherlands, Poland, Russia, Baltic and Nordic countries, and some other Europeans. The most active countries are UK (13 attendees), Russia (12), Poland (8),

Netherlands (6), Lithuania, Italy and Germany (5 per each). The Conference combined 23 presentations in different subjects – database architecture, performance, integration, capacity, and monitoring experience. The sections below are focused on repeatable, recurring and chronic problems of Zabbix enterprise implementation – active control of external objects, database performance, system scalability, integrated interface and dashboards.

**Zabbix: Where We Are.** The current status of the Zabbix production system, the progress since last year and future plans are announced by Alexey Vladishev, the founder and CEO of Zabbix SIA (Riga, Latvia).

Major features are introduced in the new product release Zabbix 2.2 and they are declared as follows:

1. Given modern tendency towards virtualization and a growing number of virtual machines (VM) populated in enterprise data centers, now VMs monitoring is supported in Zabbix system, particularly VMWare, vCenter, vSphere, including auto-discovery of guest VMs, built-in checks and Zabbix pre-defined templates.

2. Zabbix system performance is improved by a factor of 2 to 5, depending on the number



Fig. 1. Zabbix Enterprise Appliance ZS-5200

of monitoring hosts, metrics and database size. This advantage can be reached by means of extending Zabbix cache and less number of update transactions on DB server.

3. Reaction to Zabbix events became faster and the triggers are able to act on the events even when a certain host is disabled and the items are in the unsupported status for some reasons.

4. Zabbix 2.2 enterprise control system has been supported for five years with fixing product bugs, critical, security, and other issues.

The new Zabbix enterprise appliance ZS-5200, made in Japan (Fig. 1), allows for monitoring up to 20K items from external hosts, network devices, and other hardware. That is



Fig. 2. Zabbix Partnership Map

an alternative of using existing Zabbix server with web interface and automatic configuration instead of manual setting up.

The Zabbix enterprise-class automated control system is still an open source product, but not open core, and it is widespread all over the world and it has partnerships with 68 business companies (Fig. 2).

This presentation is available on Zabbix and YouTube web sites [1, 3].

**Perobbix + Zabbix DB Monitoring.** A new approach for database monitoring is presented by Julio Cesar Hegedus, Senior Linux Consultant, Yenlo BV (Amsterdam, the Netherlands).

Perobbix stands for Perl + Oracle + Zabbix. The idea is to monitor DB servers using Perl scripts running SQL queries against relational databases, such as Oracle, MySQL, etc. Perobbix solution is implemented in Yenlo enterprise infrastructure consisting of 800 hosts, 130 DBs, 80K items.

Perobbix architecture looks like Fig. 3. The approach has the following advantages:

1. There is no Zabbix agent required to install on a host under Zabbix monitoring.
2. Not only a DB host itself but the databases, that is DB integrity, DB sizes, performance of DB transactions, etc., can be monitored.
3. Different methods, drivers, and engines can be used for DB connection. A read-only access is enough that is more safety.

4. There are no limitations for the number of queries in a batch to be executed against DB.

5. Performance of the Perobbix system does not depend on the Zabbix data delay if it happened.

A full presentation is available on Zabbix and YouTube web sites [1, 4].

**When It Comes to Scalability.** The new distributed architecture of the Zabbix monitoring system is proposed by Leo Yulnets, Operations Tools Lead, RingCentral (USA) [5]. NoSQL solution based on MongoDB open-source document-oriented data warehouse [6] for storing historical data is described.

RingCentral private cloud infrastructure is the biggest one ever monitored by the Zabbix system, consisting of more than 4K hosts total and about 0.5M metric values per minute. Moreover, RingCentral multi-service for over 300K US customers is a fast growing IP-telecommunication industry, having up to 40 % annual increment. Sooner or later, the performance degradation and, as a result, the Zabbix data delay became a showstopper for further scalability in a rapidly growing enterprise environment.

Reducing the number of monitoring items and/or extending polling time intervals is of a big manual effort, which is not a good idea. Adding more servers, proxies, high-performance storages, or other hardware is a




Fig. 3. Perobbix Architecture

too expensive solution and it does not help a lot because MySQL DB is the main bottleneck due to multiple read-write transactions executed simultaneously.

As a temporary workaround, Zabbix is split into two monitoring systems, each working with a separate database. Special reports and dashboards are created to observe events, alarms, and other control data from both locations on a single monitoring console. As the next step, historical data should be consolidated for more efficient troubleshooting and analysis. But there can be the risk of data delay depending on partitioning between the real-time data amount and history size. Monitoring data delays and reporting gaps for at least 6 min may finally lead to a missed customer service outage that

can be of really high cost for the company business.

The alternative Zabbix architecture, named Octopus, is proposed to reach horizontal scalability with a rather low cost (Fig. 4). The Octopus architectural solution is developed and is currently being implemented at RingCentral. Several Zabbix monitoring systems are consolidated on different levels to meet the enterprise monitoring requirements, including data delivery with no delays and gaps, calculation of items and triggers within a short period of time, keeping long period of history data, providing events and data visibility on all levels.

One day statistics is kept as real time data in MySQL relational database, while all the

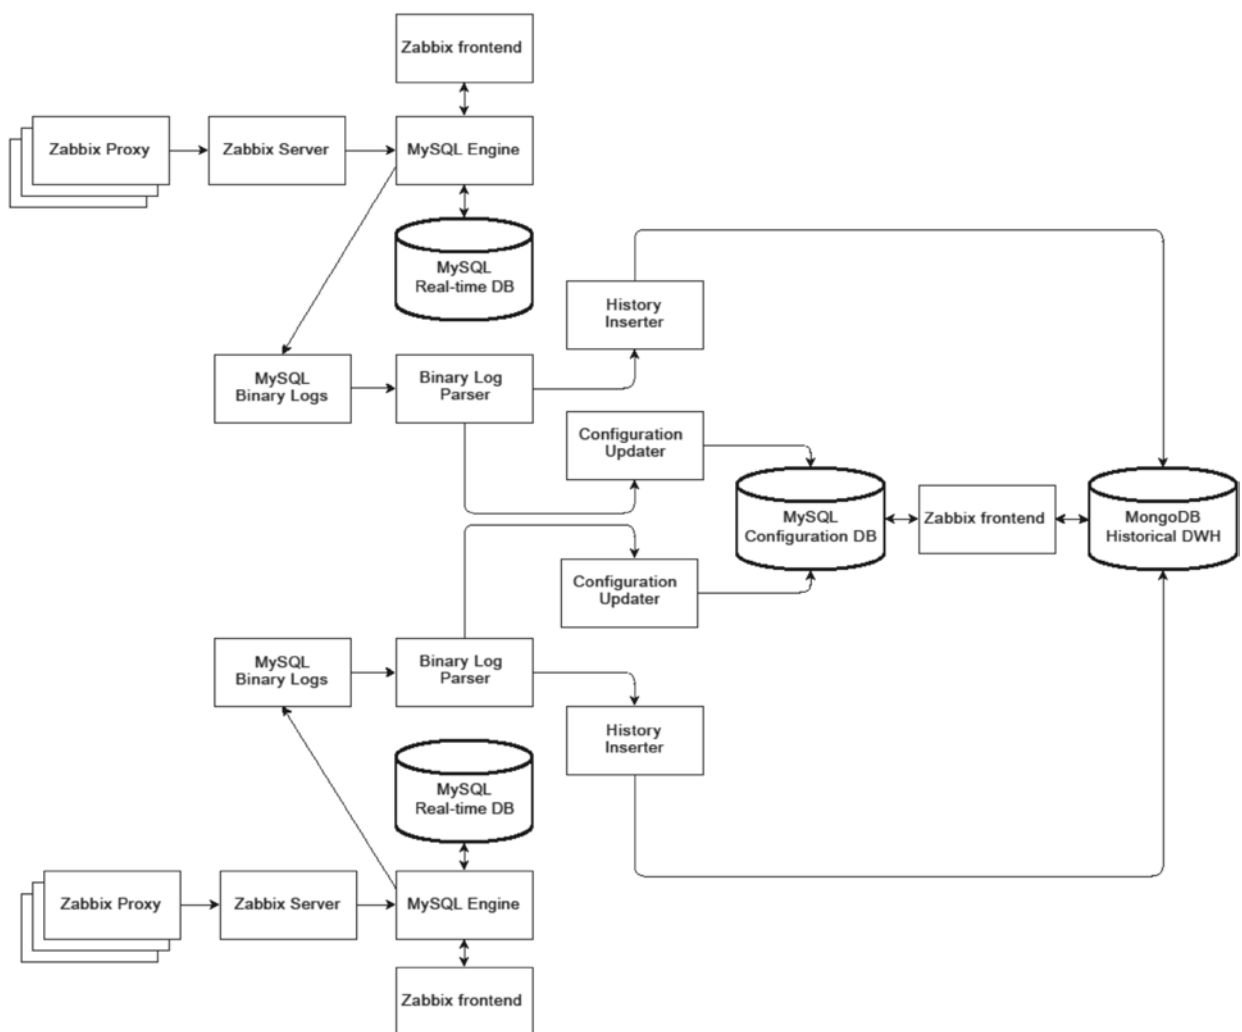


Fig. 4. Octopus Distributed Architecture

other history and the trends are transferred to MongoDB noSQL data warehouse. Real time data does not require separate network storage and can be stored on a local drive, or even in memory. Zabbix system events and alerts are consolidated on a single dashboard using web API.

The following benefits of the new approach and Octopus distributed architecture are expected:

1. Scalability improvement, allowing as many Zabbix monitoring systems as needed.
2. There are no monitoring data delays and no reporting gaps, even in peak time with high workload.
3. The reads from and the writes to history DB can be separated for better performance.
4. Extending historical data retention from 3 to 12 months, or even longer period of time.

The main disadvantage of the Octopus distributed architecture is that programming skills and resources are required to synchronize different Zabbix instances, create consolidated custom dashboards with real time data and the reports with historical data.

Presentation slides and video are available on Zabbix and YouTube official web sites [1, 7].

**Which Database is Better for Zabbix? PostgreSQL vs. MySQL.** The results of measuring performance of the new product release Zabbix 2.2 along with the latest versions of MySQL and PostgreSQL are presented by Yoshiharu Mori, Consultant, SRA OSS (Japan).

The two categories of the performance stress tests, simple and partitioning ones, are carried out in the cloud computing environment, having 600 hosts, 26K monitoring items, and 10K Zabbix triggers. The number of values that Zabbix server was able to process per second is used as the main metric to estimate the performance of both MySQL and PostgreSQL databases.

Some results are shown in Fig. 5. All the others are provided in the presentation that is available on Zabbix web site [1] and YouTube video [8].

Generally, PostgreSQL is more stable under high input/output load and the Zabbix performance is better than MySQL though the difference is not significant. One more conclusion is that DB tuning (buffering, partitioning, transaction logs) is required in each case of certain implementation.

**Integrated Dashboard Design.** A new design of the integrated Zabbix dashboard is presented by Lukasz Lipski, IT Operations Specialist, Nordea Bank Polska S.A. (Poland).

An integrated dashboard provides a real business value for the company management. This presentation demonstrates some ways to display Zabbix reports along with external data, such as JIRA, SharePoint, BMC Remedy, and the third party software, with minimal development effort. Real world scripts, written in Perl, Ruby, Twitter Bootstrap, for getting information from external data sources are provided.




Fig. 5. Zabbix Performance Test

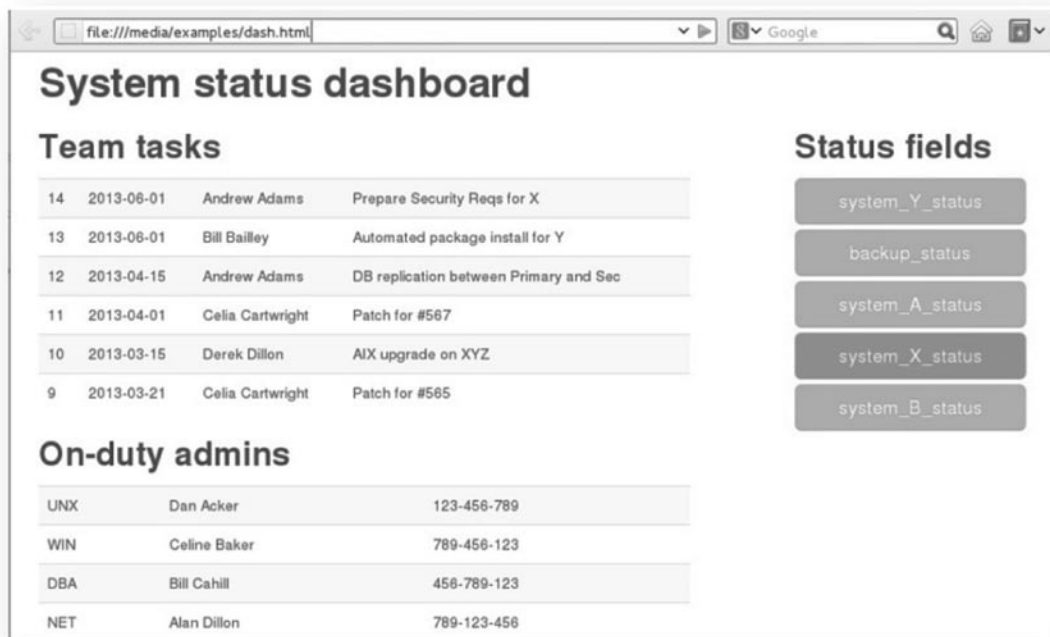


Fig. 6. Example of Integrated Dashboard

One example of the integrated dashboard is shown in Fig. 6. Specific dashboards are implemented in Nordea Bank Polska e-banking system, having 300 hosts and 40K items under Zabbix monitoring:

1. Central monitoring system that shows Zabbix events and alerts from different monitoring systems.
2. Task list, which consolidates Zabbix data with JIRA tasks.
3. Contact list, which combines Zabbix reports with the detailed information from SharePoint.

More examples of integrated dashboards are available in the presentation slides on the Zabbix web site [1] and YouTube video [9].

**Complete Log Infrastructure with Zabbix Alerting.** This is a conceptual presentation of using Zabbix for log capturing and alerting by Pieter Baele, Linux System Architect, ICTRA NMBS-Holding (the ICT provider of the National Railway Company of Belgium).

The big cloud computing infrastructure needs flexible and powerful tools for log analysis. The Zabbix system is a perfect platform for monitoring resources and alerting, but it is not




Fig. 7. Log Parsing



Fig. 8. Splunk Architecture

aimed at storing, transformation, analyzing a huge number of log files. Specific tools are required to capture the logs in a big distributed production environment.

The following concepts of the log infrastructure are proposed depending on real enterprise purposes:

1. Due to popular demand, file content and log file parsing using regular expressions has been implemented in Zabbix 2.2. Returning a part of a string can significantly reduce the workload on the Zabbix database.

2. Rsyslog is an open source logging tool having application level reliability, central log repository, filtering, high-precision timestamps, and configurable outputs.

3. Logstash is easy to deploy and it is used to collect logs, parse and store for later analysis (Fig. 7). Various log formats and filters are supported.

4. Elasticsearch is a distributed search, indexing and analytics technology. It is horizontally scalable and supports archiving on storage. Kibana interface is used as the web frontend.

5. Graylog is a logging tool using Lightweight Directory Access Protocol (LDAP) to distributed information services over the Internet. LDAP integration is good for applications and syslogs.

6. Octopus is an open source log management having LDAP feature and a lot of templates included. It can be integrated with Zabbix alert

senders and it is nice for enterprise usage.

7. ELSA stands for Enterprise Log Search and Archive. It uses LDAP and MySQL though the query language is specific. Email alerts to Zabbix server are possible.

8. Splunk is non-free scalable product (Fig. 8). It is easy to install and can be integrated with Zabbix using Zabbix sender.

9. Fluentd seems to have better performance in the line of other tools, but it is not tested enough. The largest user is currently collecting 5 TB of daily logs from 5000 servers with a rate of 50,000 messages per second.

The above approaches, except the 9th, are tested and have experience of monitoring hundreds of Linux servers. Practical examples of custom scripts for both the OS and web applications are shown on presentation slides [1, 10].

The benefit of annual Zabbix International Conference is to share knowledge and experience of using the Zabbix enterprise-class system between IT professionals in order to improve and automate monitoring the multi-host production environment.

SQL, Perl, or other type of scripts is a universal technique that can be used for any external check of the network environment where Zabbix agent is hard to install.

Database is always a bottleneck in the cloud computing infrastructure. DB tuning is required for better performance with no reporting data

delay in each case of certain implementation.

The distributed Zabbix architecture is valuable when an enterprise system grew up to its scalability threshold in terms of hosts and monitoring metrics. In the case of the Zabbix

database split, the integrated dashboard should be created to consolidate monitoring data from separate sources into one report.

The accuracy of conclusions is also approved by similar investigation [11].

#### REFERENCES / СПИСОК ЛИТЕРАТУРЫ

1. Zabbix International Conference 2013. Agenda and presentation abstracts. Riga, 2013. Available: [http://www.zabbix.com/conf2013\\_agenda.php](http://www.zabbix.com/conf2013_agenda.php)
2. Zabbix official web site. Available: <http://www.zabbix.com/product.php>
3. **Vladishev A.** Zabbix: Where We Are. *Zabbix Internat. Conference*, 2013. Available: <http://youtu.be/c8K83MJS-jg>
4. **Hegedus J.C.** Perrobix + Zabbix DB Monitoring. *Zabbix Internat. Conference*, 2013. Available: [http://youtu.be/pJCV\\_ui0orQ](http://youtu.be/pJCV_ui0orQ)
5. RingCentral official web site. Available: <http://www.ringcentral.com/>
6. MongoDB official web site. Available: <http://www.mongodb.org/>
7. **Yulenets L.** When It Comes to Scalability. *Zabbix Internat. Conference*, 2013. Available: <http://youtu.be/1Eq-9q16AOU>
8. **Mori Y.** Which Database is Better for Zabbix? PostgreSQL vs. MySQL. *Zabbix Internat. Conference*, 2013. Available: [http://youtu.be/LHitd\\_GTC-w](http://youtu.be/LHitd_GTC-w)
9. **Lipski L.** Integrated Dashboard Design. *Zabbix Internat. Conference*, 2013. Available: [http://youtu.be/\\_gy4qzyZf\\_o](http://youtu.be/_gy4qzyZf_o)
10. **Baele P.** A Complete Log Infrastructure with Zabbix Alerting. *Zabbix Internat. Conference*, 2013. Available: [http://youtu.be/f5IfBJdj\\_Xk](http://youtu.be/f5IfBJdj_Xk)
11. **Mescheryakov S.** Zabbix Conference Overview. RingCentral Inc., USA, 2013. Available: <http://wiki.ringcentral.com/display/OPS/Zabbix+Conference+Overview>

---

**MESCHERYAKOV, Sergey V.** *St. Petersburg State Polytechnical University.*  
195251, Politehnicheskaya Str. 29, St. Petersburg, Russia.  
E-mail: [serg-phd@mail.ru](mailto:serg-phd@mail.ru)

**МЕЩЕРЯКОВ Сергей Владимирович** – профессор кафедры автоматов Санкт-Петербургского государственного политехнического университета, доктор технических наук, доцент.  
195251, Россия, Санкт-Петербург, ул. Политехническая, д. 29.  
E-mail: [serg-phd@mail.ru](mailto:serg-phd@mail.ru)

**SHCHEMELININ, Dmitry A.** *RingCentral Inc.*  
1400 Fashion Island Blvd., San Mateo, CA, USA 94404.  
E-mail: [dshchmel@gmail.com](mailto:dshchmel@gmail.com)

**ЩЕМЕЛИНИН Дмитрий Александрович** – руководитель департамента развития и эксплуатации облачных ИТ платформ компании RingCentral, кандидат технических наук.  
1400 Fashion Island Blvd., San Mateo, CA, USA 94404.  
E-mail: [dshchmel@gmail.com](mailto:dshchmel@gmail.com)