

# **Evaluation of Zabbix Monitoring System for CDF Offline**

Gerald Guglielmo

(CD-doc-2654)

## **Abstract**

An evaluation of the Zabbix monitoring system was performed to assess whether it could reasonably meet the requirements needed to improve operational support load of the CDF Offline CAF systems. This represents a stand alone evaluation of the Zabbix system and was not part of a larger comparison of other monitoring systems. The goal of the evaluation was to find a monitoring system that could meet the needs outline, not necessarily find the best one available. Based on an evaluation over the course of several weeks the Zabbix monitoring system sufficiently meets the requirements defined in the Requirements for Evaluation of Monitoring System for CDF Offline (CD-doc-2653).

## Introduction

Early on in the CDF Initiative it was determined that CDF Offline operations lacked sufficient monitoring and issue tracking. In response to this situation the initiative started looking into ways to address these issues. This document describes the evaluation phase performed to address the monitoring situation.

The criteria used for the evaluation of the Zabbix monitoring system for CDF Offline are defined in the *Requirements for Evaluation of Monitoring System for CDF Offline* (CD-doc-2653). As discussed in that document these requirements are limited in scope to designed to address the suitability for improving the current monitoring of CAF operations at CDF. The evaluation is not an exhaustive or highly quantitative evaluation that could be used in a head to head comparison of monitoring infrastructure.

## Zabbix

There are many commercial systems as well as freely available systems available on the web, and Zabbix is one of the freely available systems. Zabbix is an open source monitoring system that can be downloaded at <http://www.zabbix.com/>. The systems consists of agents that can be run on the systems to be monitored, a server that collects data from the agents and stores it in a database, and a web interface for monitoring and configuration based on php.

## Evaluation

The requirements were placed in one of two categories, the first being absolutely necessary and the second being negotiable if not practical. Clearly the most important criteria for the evaluation are the ones that are absolutely necessary. The negotiable ones are not by themselves critical, but if too many of them are not met that could also in aggregate be a serious problem. In this section each an assessment of whether Zabbix can meet requirement will be discussed briefly.

### **Absolutely Necessary**

- *Web based interface for monitoring*

The user interface for administration of the main system, as well as for monitoring is web based. The functionality is provided by PHP scripts which interface to the database that stores the data and monitoring display configuration. The interface does provide configuration of accounts and access control.

- *Ability to easily graph data*

The web interface provides forms for defining graphs of the data. Tests showed creating new graphs was fairly easy as the forms are basically

self-explanatory. During the evaluation several graphs were defined.

- *Ability to store and 10's of Gigabytes of monitoring data*

Based on claimed experience from the CMS Computing Facilities department and on the Zabbix documentation, Zabbix has no trouble storing 10's of Gigabytes of data. There are indications that the CMS Computing Facilities department instance of Zabbix has significantly more data than this requirement, and expects to support an order of magnitude more. The example database sizing in the *Zabbix User's Manual* shows a calculation of about 50 Gigabytes. Zabbix can be configured to use Mysql or Postgres, both of which are known to scale beyond the requirement stated here. Explicit testing of this requirement was not performed as the test environment had modest storage and CPU capabilities.

- *Ability to configure data retention period*

The Zabbix system does support configuring the data retention periods for history data (periodic checks) and trend data (averaging) independently. This allows one to define the retention rules that best fit the needs of the user.

- *Capability of storing monitoring information for thousands of hosts*

This has not been explicitly checked as a test system consisting of thousands of nodes does not exist. However in the *Zabbix User's Manual* there is a table on hardware specifications that outlines requirements for monitoring greater than 10,000 hosts which significantly exceeds the stated requirement.

- *Ability to track 10,000 different items*

The CMS Computing Facilities department claims they are currently performing checks on more than 10,000 and that is sufficient to meet this requirement.

- *Ability to define the frequency of data checks*

Zabbix allows the monitoring frequency of each check to be defined independently. This was verified in tests and was very easy to configure through the web based forms.

- *Ability to define sets of tests that can be applied to hosts*

Zabbix provides a mechanism called templates. A template is a set of check that could be applied to multiple hosts. Once defined a template can be applied to each host desired. The current template implementation in Zabbix allows additions to a template definition to be automatically applied to all hosts that have registered the template.

- *Ability to groups different types of hosts into a logical unit*

Zabbix allows a host to be assigned to multiple groups. This can allow a key piece of hardware to show up as part of several groups. Thus hosts with different checks, even defined through templates, can be assigned

to the same group. For monitoring purposes one can look at all hosts in a given group on one page, hiding other irrelevant hosts. It also means that special hosts can show up on more than one of these pages.

### **Very Important But Negotiable**

- *Web based monitoring by general users around the world*

The user interface for monitoring is through a web server and based on PHP scripts. Individual user accounts or a generic read-only account can be configured to allow access. Network firewalls, ACLs, or other restrictions could block remote access, but in theory these can be addressed as they have for other systems at the laboratory. It is possible that a large number of users could overload the system if they are accessing at the same time, but this does not seem like a big risk.

- *The ability to accept input data from agents in text form to standard output*

Tests of scripts specifically written for the evaluation demonstrate that Zabbix is able to parse text output properly and store it in the database.

- *Ability to leverage existing expertise at the lab*

The CMS Computing Facilities department in the Computing Division, has been running Zabbix for several months and have developed some expertise with the system. During the evaluation period several claims by members of that department concerning Zabbix were confirmed with testing. Furthermore, the Fermilab Experiments Facilities department as well as the Grid Facilities department, both also in the Computing Division, have started their own evaluation programs of Zabbix.

Therefore there is some evidence that local Zabbix expertise exists at the laboratory.

All of the requirements defined for this evaluation have been met by the Zabbix system. While the list of requirements was not an exhaustive list, or one suited for a general monitoring system evaluation, it did outline the needs for the use case intended.

## **Conclusion**

The Zabbix monitoring system meets all of the requirements set out in the *Requirements for Evaluation of Monitoring System for CDF Offline* (CD-doc-2653) document. This includes both the absolutely necessary and the very important but negotiable requirements. Therefore it is recommended that the process of preparing for deployment move forward.