

D4Science Infrastructure - Task #2221

Experiment with a new monitoring/metrics tool: influxdata (influxdb)

Feb 10, 2016 12:04 PM - Andrea Dell'Amico

Status:	Closed	Start date:	Feb 10, 2016
Priority:	Normal	Due date:	
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	System Application	Estimated time:	0.00 hour
Target version:	UnSprintable		
Infrastructure:	Development, Pre-Production, Production		

Description

The tool is promising:

- All the components are free software (they sell managed services and consulting)
- Can do alerts, or delegate them to nagios
- It has metrics plugins for almost all the software we commonly use

Drawbacks: It seems that does not natively aggregates logs as logstash does, but in the worst case we could use the same input plugins to route the logs to an ELK instance.

Some references:

<https://influxdata.com/time-series-platform/>

<https://www.elastic.co/guide/en/logstash/current/plugins-outputs-influxdb.html>

<https://www.digitalocean.com/community/tutorials/how-to-analyze-system-metrics-with-influxdb-on-centos-7>

<https://github.com/rochaporto/collectd-ceph>

<https://github.com/pbanaszekiewicz/collectd-ganeti> (Needs modifications because it only supports KVM right now)

History

#1 - Feb 10, 2016 12:25 PM - Andrea Dell'Amico

- Related to Task #1854: Tool for usage statistics on DataMiner added

#2 - Feb 11, 2016 06:16 PM - Andrea Dell'Amico

- Status changed from New to In Progress

#3 - Feb 12, 2016 05:31 PM - Andrea Dell'Amico

- % Done changed from 0 to 20

A preliminar dashboard can be used here: node0-monitoring.research-infrastructures.eu

Beware that there's no authentication and all the dashboards can be edited by anyone. Access is restricted to the ISTI networks.

#4 - Feb 12, 2016 05:38 PM - Andrea Dell'Amico

It seems better suited to collect application and system metrics and maybe send alerts based on them. While ELK seems more suited to aggregate logs.

#5 - Feb 12, 2016 05:58 PM - Andrea Dell'Amico

<http://www.collectd.org> and <http://www.fluentd.org> can talk with influxdb (and other backends) and seem to offer a wide range of input plugins.

#6 - Feb 12, 2016 06:03 PM - Andrea Dell'Amico

And it seems that fluentd can collect the syslog output and feed it to influxdb: <http://www.fluentd.org/guides/recipes/syslog-influxdb>

#7 - Mar 01, 2016 12:28 PM - Andrea Dell'Amico

I see that Kronograf is still very limited and the useful features will not be open sources. It seems that <http://graphana.org> is a better alternative that can also read from elasticsearch, prometheus and other dbs. So it will be maybe possible to have the same graphical interface for all the monitoring sources.

#8 - Mar 16, 2016 03:36 PM - Andrea Dell'Amico

- Related to deleted (Task #1854: Tool for usage statistics on DataMiner)

#9 - Jun 14, 2017 12:30 PM - Andrea Dell'Amico

- Status changed from In Progress to Closed

- % Done changed from 20 to 100

As the interesting tool are commercial, and prometheus seems a better solution, influx will not be deployed.