

Minutes of the WLCG Monitoring Consolidation meeting, 25/10/2013

Local participants: Andrea, Pablo, Eddie, Valentina, Nicolò, Ivan, Pedro, Marian, Luca, Lionel, Julia, Costin, Alessandra, Simone

Remote participants: David

Pablo presents the first draft of the working group report.

The project consists of two phases: the first is the analysis of the current status (as summarised in the well-known table) and the second is the implementation.

We started by reviewing the applications supported by the group and identified two or three that can be discontinued. Most of the improvements are planned for the site and infrastructure monitoring. Chapter 5 describes our vision on how to do maintenance and development of all monitoring applications.

Then Pablo goes into more detail in describing the proposed system, consisting of eight layers of applications. He points out that aggregation will be either internal or external.

Concerning deployment, the idea is to follow what is done in the Agile infrastructure, that is using Puppet, Foreman, Hiera, etc. All services will be redeployed using these tools and SLC6 will be the reference platform.

Andrea points out that in figure in page 7 it would be better to replace "HammerCloud jobs" with something like "Experiment jobs", as long as HammerCloud jobs are considered simply as a particular subsample of jobs with a value for monitoring. The proposal is approved.

The code will be moved to git and the new documentation will be done using a common tool (and not in Confluence). About ticketing, the different support units in GGUS and SNOW have been already merged and Savannah will be replaced by JIRA.

Concerning probes, Nagios will become optional and job submission probes for Condor-G and CREAM will be developed and maintained.

Pedro suggests to add something about the schema.

Concerning transport, MonALISA, HTTP and messaging are used and will continue being used, as each one works well in its applications. Of course, the latest versions of the corresponding libraries will be used.

About Flume, Julia says we need to understand if it would be relevant for a distributed infrastructure. Pedro comments that having a fourth transport mechanism would not be a big deal.

About storage, the idea is to converge schemas for similar applications. There is already a prototype and we will evaluate more database backends.

Luca suggests to describe the SSB metrics in the document as they are not necessarily known to the reader.

As database backends, ElasticSearch has been evaluated and found very promising for the SSB, while for other use cases it will be necessary to wait for the implementation of multiple grouping.

There are still open questions, for example we didn't yet evaluate HDFS.

Julia adds that we also need to understand how much data will be kept; if it's not too much, ElasticSearch is certainly a good choice. Eddie comments that HDFS is supposed to be good for offline analysis. Julia concludes that an evaluation will be needed.

Pedro asks if there is a data retention policy. In fact, there is and old data are being already cleaned up.

Concerning aggregation, aggregated metrics will be treated like any other metric.

Costin asks if the aggregation will be done of the fly or delayed by some time. Pablo answers that doing it on the fly would require too many changes, but the latency is of the order of 10' anyhow, which is quite acceptable.

Luca argues that if we say that the aggregation is made in the storage, we are ruling out the possibility of using lots of alternative technologies. Pablo agrees to move the aggregation to a different place in the picture, not to have it attached to storage, to leave room for more options.

Concerning visualisation, there will be a clear separation between the server and the client side. Servers would also do some caching, and clients will be based on Javascript, a Model-View-Controller paradigm, etc. In case we go for ElasticSearch, tools like Kibana will be useful.

The next steps are described in detail, in particular the list of applications to be discontinued and those that will still require significant work. Timescales are not yet defined and will be later on.

For the transition period, we will start with a proof-of-concept prototype using a fast cycle of iteration. The goal is to reach production quality by summer. The current infrastructure will be kept as needed.

For the future, the document needs some estimate of the effort to maintain the current and the future system and compare them.

Comments should be sent during the next two weeks; a shared Google Document will be temporarily used to submit comments.

Luca asks who will be the audience of the document. Pablo answers that it is intended for the WLCG MB and GDB. It is of paramount importance that we all agree with the final version.

The next meeting will be on November 8, where the second draft will be presented. Jacobo will have a demo of the prototype of the new SSB-SUM combination.

David asks what are the plans for allowing sites to import metrics into their site Nagios instance. Pablo and Julia answer that everything will be made available via API and sites will be able to publish information in Nagios.