

070119-WG PhoneCon

19 January 2007

14:59

Grid Monitoring WG Meeting

Date and Location -----

19 January 2007
15:00 - 16:30

Attendees -----

Dave Colling
Alessandra
Luigi
Sergio
James
German
Piotr
Sonvane
Julia

Subject -----

Alessandra: Request : Minutes to be in plaintext (or at least not only Word !)

Sergio: DPM session overlap at WLCG meeting. Will work out in co-ord with Sophie Lemaitre.
Dave Colling : Collab Board clash on Thursday next week

Dave: Do we want the data from RTM?

James: Perhaps separate from the pure sensor systems. Fits better with gridview, and will allocate a session to it. Should have a separate discussion with Judit Novak next week on volumes etc...

Tool survey:

LEMON:

Agent runs on a node. Communicates with sensors running on the node via push or pull.
Sensor extracts **metrics**, which can be organized in **metric classes**. E.G Metric, is 'checking if a process exists'.

Metric has an **instance** e.g. 'checking for the httpd process'.

Checking a metric generates a **sample**.

Samples are sent to a **repository**. There are then interfaces for querying. SOAP, CLI, visualization (similar to Ganglia)

Protocol for talking to repository is either TCP or UDP. Simple - circular buffer for sample. Also can be encrypted.

James: Is it assumed the repository is local to agents?

German: Not assumed, but no setup exists. But designed for high reliability link between agent and repository.

Luigi: Is it possible to change the metrics at runtime, or have a new metric at runtime.

German: Implies the agent on the machine needs to be reloaded. But you need to register the metric on the server (for the oracle based server)

Luigi: Is there secure communication from agents? And for the web service.

German: agent + sensor on same machine (sensor is forked child) so no security implications. You can also have a secure channel between agent and repository - auth exists, and also

encryption is in progress. No X509, since overhead is too heavy. You can use https for the visualization system, and use an authentication callout (e.g. NICE)

German: At CERN, we monitor over 4000 nodes.

James: Do you know metrics/node, samples per node.

German: 10 sensors, 70 metric classes, 350 metrics. 2G per day.

Julia: how do you guarantee sync of schema between server and the agent.

German: Assume it is done outside of the system. And we, at CERN, assume a fabric management system is used.

Julia: As a user, can I change the things on the host and collect them.

German: If you want a new metric, reconfig is needed. If you want just a metric that is already collected on other hosts, then you simply add it to your host, and restart the agent on the node (no DB op needed).

Sonvane: how do you decide the DB schema for a metric.

German: For LEMON it's quite simple - timestamp, metric ID, variable number of fields (e.g. string, integer, standard SQL types) - no complex data type.

Sonvane: Are there separate tables per metric

German: for the oracle implementation, separate tables per metric. A new table needs created (by the app) when you add one.

Luigi: Is it possible to access the service via a WS interface.

German: currently on the repository, there is a SOAP interface. Want to discontinue it. Replace with HTTP/XML gateway.

James: If we want to query to forward to the grid, what's the best way.

German: HTTP/XML will be the standard way.

German: LEMON does not provide a concentrator of data. It stores the data as received by all and every node. For the visualization, this is stored in RRD graphs. Not a correlation engine.

James: Is there a concept of cluster in LEMON DB?

German: No - this is outside, you need to know the hostnames.

Sonvane: It needs to be named metrics, not logs?

German: Yes, it's not a logfile collecting mechanism - it needs to be specific metrics.

James: Does the visualization system use XML/HTTP or direct DB?

German: Direct DB.

James: Is there a direct interface to the RRD database?

German: No. This isn't exposed. Users usually don't go the repository directly, except in a few general cases.

Sonvane: Generic visualization of metric for text data?

German: No.

Piotr: Why then store text data

German: We try and ask users not to do this, but sometimes you want to make a correlation between a number of metrics.

German: Also **exceptions**. An sensors on the node looks at values from sensors and compares against conditions, and can do an action. This can be launching a recovery action, or mark as a exception - just another metric.

GridICE

LEMON was a natural choice when starting their tool.. Developed grid sensors and a way to publish for a set of sites. Works for ROC managers or site managers. Added a concentrator - which is tailored for different communities - VO,

Sergio: maybe next week also good to show the information we provide in the UI, which is aggregated.

Also, there is a notification part.

James: Is the concentrator only GUI, or is it also interface

Sergio: All information published via graphical interface is also exposed via XML. Also an API to have some information directly from the DB, but only for known people.

James: And AAS on the interface?

Sergio: all information not relevant to privacy is open to all. We have an integration that takes care of the privacy. Based on the DN.

Julia: How can we influence the deployment of this.

James: We do have the influence. We need to work out components are needed, and get them deployed. We can't change the fabric monitoring at a site, but we can propose different solutions - e.g. have Nagios sites be able to publish into the aggregator. Most always focus that we're trying to improve reliability.

Sergio: We are integrated into gLite, and should be deployed. When started, we have in mind an aggregator for the entire grid, but then sites said it could be useful for fabric monitoring too, so it got extended for that.

Sonvane: Important to focus on how tools can integrate the remote data.

Luigi: We have a draft proposal on the service status description. Service status is status at a well-defined time, t. Can be a single attribute, or a tuple. We need a list of attributes that define the status. Different service have different attributes. Tried to write this as an XML schema. This is important for the service.

German: Agree this is good, and also SLS is complementary.. It deals at the service level. We have 150 services doing this - the concepts are very generic and useful.

German: Also there is the SAM approach, of emulating the users. Always easier to mimic the user rather than try and calculate out of the metrics.

Sonvane: At the start it should be the metric info and then we need to description for the probes.

Sergio: We need to consider OSG and other services e.g. dCache.

James: Yes, we need to consider other services too, outside of JRA1.

ADMIN: We'll try to show the visualizations next week. James to find time - either another slot, or extending the current one.

TODO: Availability calculation for a 'service'.

AIMS: Subset of people to start writing a 'Sensor Description Questionnaire'.

AIMS: Subset to look at integrating remote probe data into fabric (what data, and what action)

Next Meeting: Thursday afternoon during WLCG Workshop - IT Aud.
