

Open Science Grid

Grid Service Monitoring Working Group

Rob Quick

Open Science Grid Operations
Center - Indiana University


HPDC 2007

Monterey, California - 25 June 2007

Grid Services Monitoring WG

Mandate

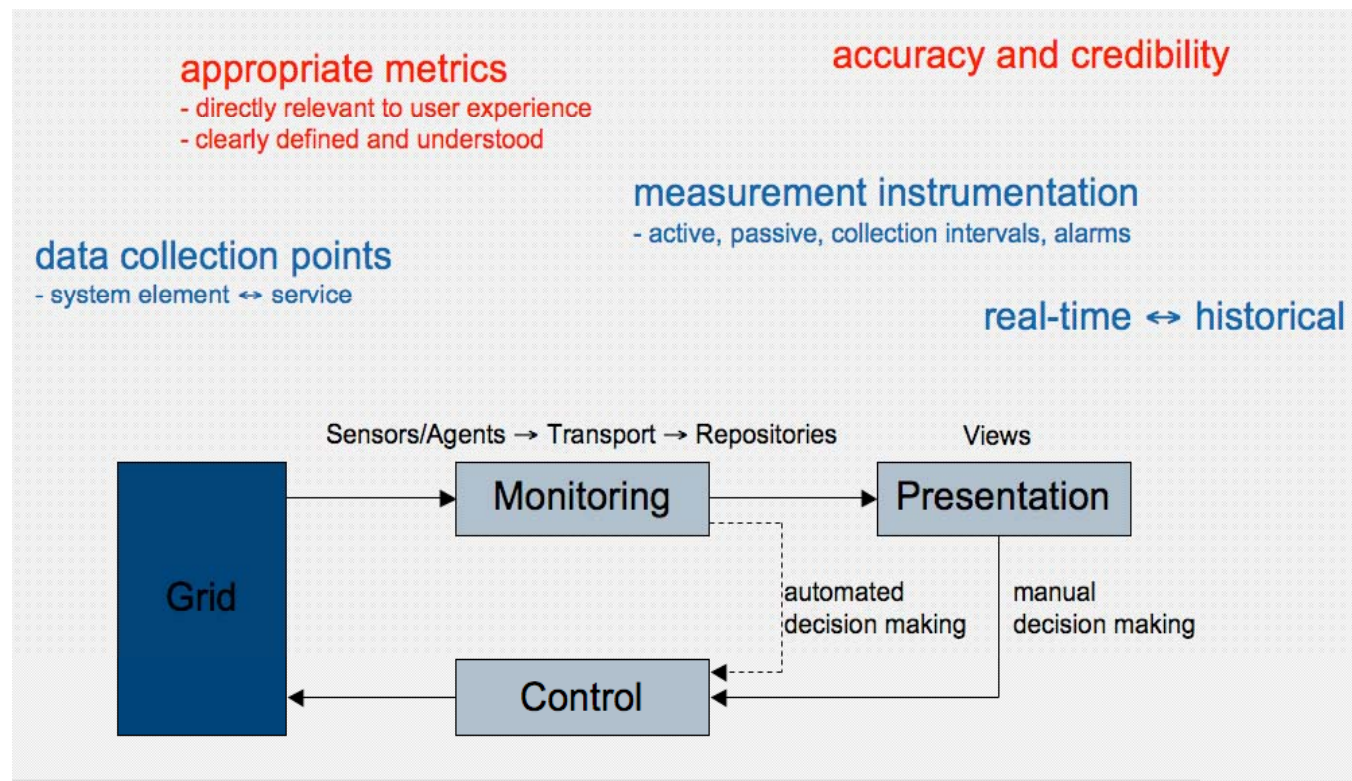
- *“....to help improve the reliability of the grid infrastructure....”*
- *“.... provide stakeholders with views of the infrastructure allowing them to understand the current and historical status of the service. ...”*
- *“... stakeholder are site administrators, grid service managers and operations, VOs, Grid Project management”*



Open Science Grid


Monitoring

You can't manage what you don't measure...



Aims of grid services WG

- Not to provide yet another technical solution
- But,
- Improve reliability of WLCG
 - Consolidate existing solutions
 - Improve communication
 - Reduce overlap
 - Increase sharing




Open Science Grid

How?

- Engage with stakeholders
 - Operations meetings
 - WLCG Workshops
 - Questionnaires to site managers
 - Grid Service providers (EGEE, OSG)
 - Grid Middleware providers (gLite, VDT)
 - Monitoring software providers (SAM, VORS, Gridlce, MonAmi, GridView, LEMON, Nagios, ...)
 - External experts (openlab EDS collaboration)
 - Other Working Groups

Tasks of grid services WG

- Best practice notes
 - How to many grid proxies for monitoring
 - Message-level Security for monitoring
 - What information can/should be passed through site boundary
 - ...
- Create set of ‘standard’ WLCG probes
 - And how to calculate availability based on the metrics produced




Open Science Grid

Direction

- Focus on the interaction points between the different systems
 - Allow for diversity across different grid infrastructures
- *“Specifications, not Standards”*
 - Timescales mean we can't get involved in long and heavyweight standards activities
 - Take best practices from existing systems, and document them
- Implement simple prototypes
 - And mature the bits that work !
- Get something out to the stakeholders
 - Close feedback loop is the key to adoption
 - Plan for a “standards based” solution in the future

High-Level Model



Terminology

- Metric
 - A data value gathered that tells us something about a service
- Probe
 - The actual code which gathers the metric/metrics
- Check & Sensor
 - A 'probe' in Nagios and LEMON respectively

Locality of Probes

- ‘local’ can mean two things ;(
- ‘local’ and ‘remote’ with respect to probing the interface of the service
 - **local** means on the site
 - **remote** means external to the site
- **(host-)local** probes
 - Gathering information from the operating system level
 - Traditional fabric management probes

Specifications

- Probe Specification
 - Defines how a fabric monitoring system can interact with probes that test grid services
 - Simple text-based protocol (lightweight)
 - Decouples grid probes from the specifics of the fabric monitoring system
 - Allows for currently existing probes to be re-used by any monitoring system
 - SAM Tests
 - EGEE CE ROC Nagios testing
 - OSG Tests

Sample Probe Output

org.osg.general.ping-host

```
metricName: org.osg.general.ping-host
timestamp: 2007-06-22T17:42:00Z
metricStatus: OK
serviceType: other
serviceURI: vdt-fc4-ia32.cs.wisc.edu
gatheredAt: vdt-fc4-ia32.cs.wisc.edu
summaryData: OK
detailsData: Host vdt-fc4-ia32.cs.wisc.edu is alive and responding to pings!
EOT
```

Example of exchange format


```
<?xml version="1.0"?>
<root xmlns="http://cern.ch/grid-mon/2007/05/mon-exchange-schema/">
  <Region name="CERN">
    <Site name="CERN-PROD">
      <type>Production</type>
      <status>Certified</status>
      <SiteMetric name="site-daily-avail">
        <measurement>
          <status>ok</status>
          <summary>0.3</summary>
          <timestamp>2007-02-25T00:00:00Z</timestamp>
        </measurement>
      </SiteMetric>
      <Service endpoint="https://ce101.cern.ch:2119/" type="CE">
        <isMonitored>true</isMonitored>
        <inMaintenance>false</inMaintenance>
        ...
      </Service>
    </Site>
  </Region>
</root>
```

Site monitoring

- We can't/won't impose a solution on sites
 - They might/should have something already
- Specification based approach allows our probes fit into any fabric monitoring system
- Data Exchange format allows higher-level services consume the data regardless of fabric monitoring system

Futures and other work

- We focus here on the prototype
 - Since this is what we are delivering now
- Also working on
 - Specifications and example components
 - Security architecture
- Future work includes
 - Probe description database
 - Topology database
 - Messaging architecture for transport layer
- Closely involved with SAM team
 - Looking at how to use Nagios as a submission framework for SAM




Open Science Grid

Summary

- Effort invested to understand the current landscape
- Approach for improvement based on specifications of interfaces between components
- Prototype has been developed and tested on a small scale
- Now looking for early adopters to get feedback

<https://twiki.cern.ch/twiki/bin/view/LCG/GridServiceMonitoringInfo>



Open Science Grid

Links

- **SAM/GridView Monitoring**
Portal: http://gridview.cern.ch/GRIDVIEW/job_index.php
TWiki: <https://twiki.cern.ch/twiki/bin/view/LCG/GridView>
- **SAM OSG Probe Dev**
Homepage: <http://peart.ucs.indiana.edu/docs.osg>
- **(Service Availability Monitor)**
Test Page: <https://lcg-sam.cern.ch:8443/sam/sam.py>
TWiki: <https://twiki.cern.ch/twiki/bin/view/LCG/SamCern>
- **GridICE Monitoring**
Portal: <http://gridice2.cnaf.infn.it:50080/gridice/>
Documentation: <http://gridice.forge.cnaf.infn.it/>
- **Experiment Dashboard**
Portal: <http://dashboard.cern.ch/>
TWiki: <https://twiki.cern.ch/twiki/bin/view/CMS/Dashboard>
- **GridPP Real Time Monitor**
Homepage: <http://gridportal.hep.ph.ic.ac.uk/rtm/> (2D map and 3D globe visualizations)
- **GStat**
Portal: <http://goc.grid.sinica.edu.tw/gstat/>
TWiki: <http://goc.grid.sinica.edu.tw/gocwiki/GstatDocumentation>
- **Lemon**
Portal (CERN Compute Center): <http://cern.ch/lemon-status/>
Documentation: <http://cern.ch/lemon/>



Thank You

- Special to the Thanks James Casey and the GOC Team: John Rosheck, Tim Silvers, Kyle Gross, and Arvind Gopu
- www.opensciencegrid.org
- www.grid.iu.edu