



Enabling Grids for E-scienceE

## Report from LCG monitoring group

*Luigi Zangrando – INFN Padova*

*EGEE JRA1 All-Hands meeting  
Catania, March 09, 2007*

[www.eu-egEE.org](http://www.eu-egEE.org)





- **Main goal: improve the grid reliability**
- **many tools have been developed to improve the understanding and reliability of the grid**
  - SAM, GridView, GridICE, GOCDDB, Gstat, Nagios@CEE ROC, L&B stats, MonaLisa
- **there is a vast amount of monitoring data produced**
- **problems**
  - Interoperability, data exchange formats, replication of information

**some slides taken from the James Casey's presentation "WLCG Monitoring Working Group"**

**OSG Consortium All Hands Meeting, University of California, San Diego, 06-Mar-2007**

<https://indico.fnal.gov/materialDisplay.py?contribId=118&sessionId=1&materialId=slides&confId=468>

*Ian Bird at the Fall 2006 Hepix and at the WLCG Management board proposed to create 3 groups to set up a comprehensive monitoring framework to improve the robustness and the reliability of grid sites.*

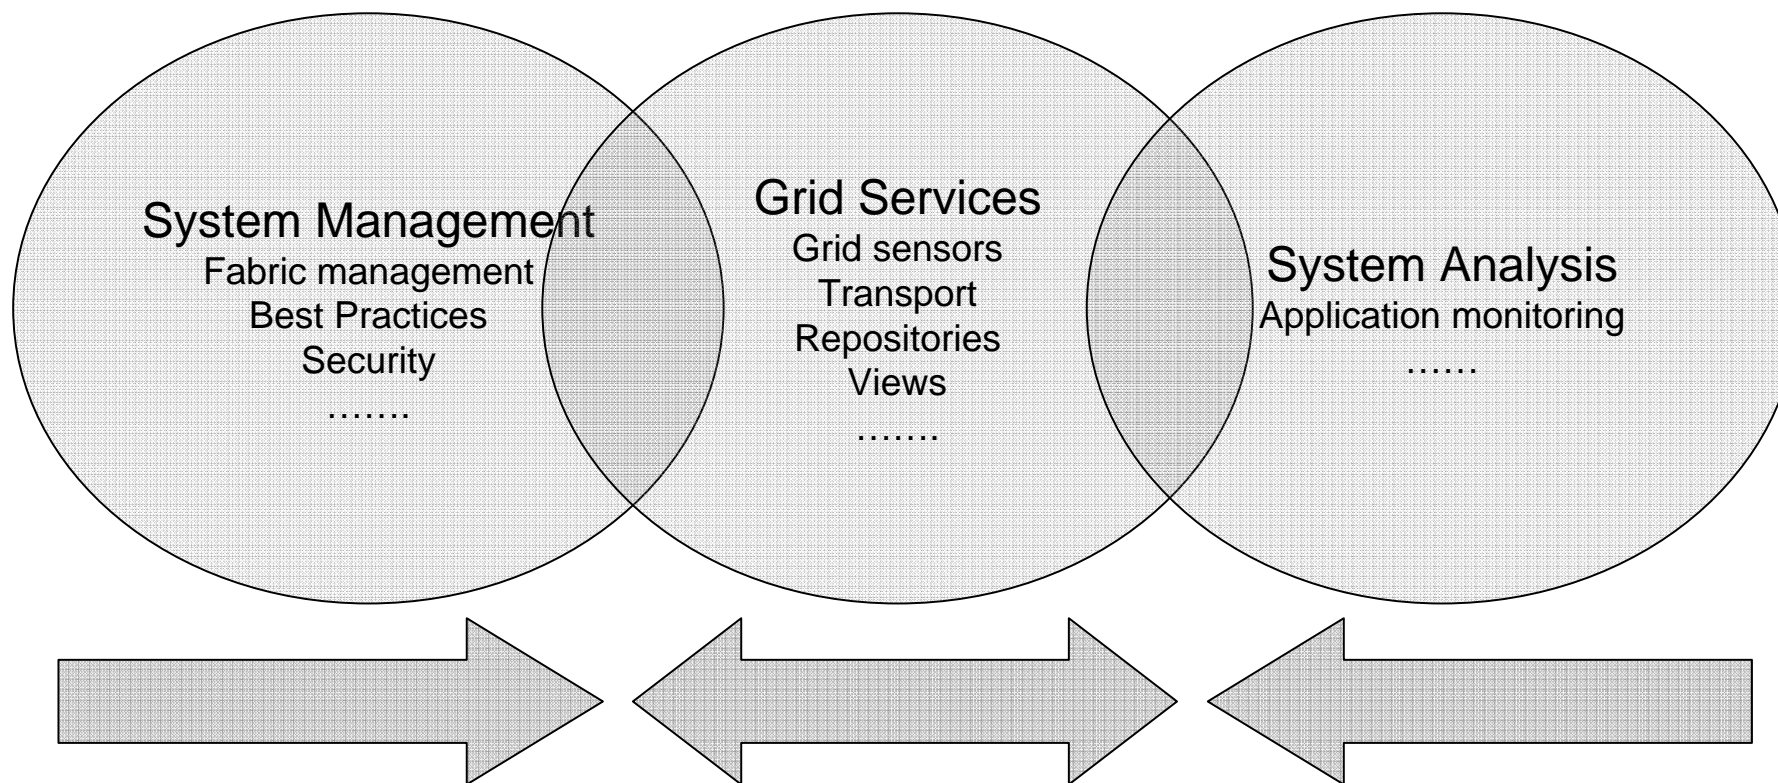
<http://indico.cern.ch/materialDisplay.py?contribId=s0t14&sessionId=s0&materialId=slides&confId=a063271>

<https://indico.fnal.gov/materialDisplay.py?contribId=34&sessionId=8&materialId=slides&confId=384>

### **Three main areas to be addressed:**

- ***System Management WG: system management and fabric monitoring tools and cookbook***
- ***Grid Services Monitoring WG: middleware monitoring and monitoring framework.***
- ***System Analysis WG: monitoring from the application side***

*to improve the reliability of the grid*





- I'll talk about the 'Grid Services Monitoring' group in this talk
- I am a SMWG member and I represent the EGEE-JRA1 group
- the group is coordinated by James Casey, Ian Neilson (CERN)



- **The overall goal of this group is:**
  - *“... 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”*
- **Not in the mandate**
  - *to develop more monitoring tools*
    - *unless a specific need is identified*
  - *to replace existing fabric management systems*
- **Out of scope (for the minute)**
  - *Accounting, Information System (BDII)*

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



- The service monitoring consists of gathering time stamped samples of various metrics from the service.
- There are two types of metrics:
  - Service Status
  - Performance.
- Monitoring of service status consists of checking if the service is running, i.e. available to accept user interaction in a normal method.
- Monitoring of service performance consists of gathering metrics useful in diagnosing the behaviour of a service, e.g. Load, #threads in use, #jobs/transactions processed per minute
- the service monitoring is useful because it can enhance the reliability of the entire grid fabric.



- **At the present time, the grid middleware doesn't usually provides tools for the monitoring of the grid services**
- **Also there is not a description**
  - of what to monitor
  - and how it should be do
- **For this purpose the SMWG decided to create a questionnaire to be completed by software providers**
- **Through the questionnaire the grid service developer provides:**
  - a plain English description of its service
  - an high level information about how to monitor the service.
- **The questionnaire provides a reference for monitoring developers and site administrators creating sensors for integration with local fabric**





- **My first task on the monitoring activity was to request all grid services developers to fill in the questionnaire**
- **<https://twiki.cern.ch/twiki/bin/view/LCG/GridServiceMonitoringInfo>**
- **At the present time 7 service developers have filled in the questionnaire**



- **service name / service context / dependencies**
- **three categories of probes**
  - the metric categorization is based on the type of information that the metric represents and how it is collected):
  - **Category A: “traces in the fabric”**
    - What can we measure from standard fabric monitoring about the service – processes, network ports, files, filesystems, logfiles...
  - **Category B: “interpretation of user-level interaction”**
    - This category of metrics is gathered by running an operation a user might want to do, and parsing the output to interrogate the status of the service. Examples are:
      - 1. *“ping”-type operations which check if the service is listening on a port correctly*
      - 2. *‘lfc-ls /grid’ on an LFC Catalog to check results are being returned from a database correctly*
      - 3. *Submitting a job and monitoring its progress and termination state.*
  - **Category C: “internal metrics presented by the service”**
    - Some services provide a mechanism to obtain more detailed information about the internal state of the service.



- **Service Name: LFC**
- **Contact Details: [hep-service-lfc@cern.ch](mailto:hep-service-lfc@cern.ch)**
- **Context and dependencies:**
  - LFC (LCG File Catalog) is a secure hierarchical file catalog, with a client-server architecture. The LFC server relies on a database backend, either MySQL or Oracle.
- **The dependencies are:**
  - GSI
  - VOMS
  - MySQL or Oracle (LFC server)
  - GRIS (LFC server)



- **Category A - Traces in the fabric**
  - **Process State**
    - Process Name / Number of instances (min/max) / User
    - /opt/lcg/bin/lfcdaemon – 1 process, 20 (default) to 60 threads – lfcmgr
    - /opt/lcg/bin/lfc-dli – 1 process, 20 threads – lfcmgr
  - **Network ports**
    - Port Number / Process Name / Security (GSI)
    - 5010/tcp (incoming access) – lfcdaemon – GSI
    - 8085/tcp incoming – lfc-dli – insecure
  - **Filesystems/Directories**
    - Path / Owner/Group / Permissions / Min. Space
    - /etc/grid-security/lfcmgr/ – lfcmgr:lfcmgr – 755
    - /etc/grid-security/lfcmgr/lfccert.pem – lfcmgr:lfcmgr – 644 (exact copy of /etc/grid-security/hostcert.pem)
    - ....



- **Logfiles**

- Path / Owner/Group / Permissions

For each logfile include any parsing expressions which can be used to identify performance and error conditions together with a description of the meaning and possible actions which could be taken (if know). Include here information written to syslog together with the syslog facility used.

- /var/log/lfc/ – lfcmgr:lfcmgr – 755

- /var/log/lfc/log – lfcmgr:lfcmgr – 644

–The log should be parsed to detect:

- Database error:

- *Look for “error” (MySQL) or “ORA-<number>” (Oracle).*
- *Send an alarm, so that the admin can check.*

- High service load:

- *Check whether too many threads are in use.*
- *Each line contains the thread number used (see below: “,0”):*



- **Category B - Interpretation of a user-level operations**

- Describe (or provide persistent urls to information about) commands which can be used as category B probes including how the status return and output may be parsed to diagnose the performance or status of the service. Include the release packages in which the commands are made available.
  - *Read operation from the database:*
  - *Stat a given directory (ex: /grid/ops). “lfc-ls” or lfc\_statg can be used.*
  - *Write operation to the database:*
  - *Utime on a file, using “lfc\_utime”.*
  - *Note: do not create a file or directory, as it is using a new LFC “inode” each time...*
  - *Performance:*
  - *“lfc-ls” should not take an excessive time to read a given directory (<10s).*



- **Category C - Internal metrics presented by the service**
  - Describe (or provide persistent urls to information about) commands, or otherwise how to invoke mechanisms, to gather category C metrics about this service (see endnote ) including how the status return and output may be parsed to diagnose the performance or status of the service. Include the release packages in which any commands are made available.
    - none



- **7 questionnaires filled in**
- **all questionnaires are more or less equivalent**
- **only 2 out of 7 services provide information about C category**
  - cream and wmpoxy BES (Basic Execution Service) prototypes





- **The service status gathered as described by the category C is an open issue**
- **We (EGEE JRA1) have a proposal for evaluation**
  - xsd schema defining the service status
  - regular snapshots dumped to local filesystem
  - Web-Service interface



- **the information about a site can be gathered:**
  - only locally on the service node (category A)
  - on the server node and/or from a node of a remote site (category B, C)
- **The info gathered locally should be published from the site**
  - Not all information should be published externally
  - Some sort of filtering/aggregation needs to be done
    - gridlce
- **A site should be able to retrieve all info gathered about it by a remote node**
  - SAM/GridView, VO Dashboards, GridPP Real Time Monitor (RB monitoring), FTS (SRM monitoring)



- **Expand work done for SAM Web Service Interface**
  - Will provide one way for a site to query services about data they have regarding a site
- **Provides a data model for talking about sites, nodes, services, grids**
  - And metric value attached to them
- **HTTP-query based**
  - With a schema for the query elements
  - And for the XML response message
- **Designed so it can be added easily onto existing systems**
  - Commitment from many of the WLCG monitoring services to implement it



- **Dec 06**
    - Background research
    - Establish core group
  - **Feb 07**
    - Establish sub-groups (sensors and repositories)
    - Agree interfaces and work plan
  - **April/May 07**
    - Prototype instrumented services to local Fabric Monitoring
    - Remote metrics to local Fabric Monitoring
  - **end-Summer 07**
    - Demonstrated improvement in reliability of grid
- ← they are here



- special thanks to LFC, DPM, LB, wmpoxy, ice, cream and cemon developers for the feedback and collaboration
- questions?

