

Distributed Monitoring Infrastructure for Worldwide LHC Computing Grid

P. Andrade¹, M. Babik¹, K. Bhatt², P. Chand², D. Collados¹, V. Duggal², P. Fuente¹, S. Hayashi³, E. Imamagic⁴, P. Joshi², R. Kalmady², U. Karnani², V. Kumar², W. Lapka¹, R. Quick³, J. Tarragon¹, S. Teige³, C. Triantafyllidis⁵

¹ European Organisation for Nuclear Research CERN, CH-1211 Geneva, Switzerland

² Computer Division, BARC, Trombay, Mumbai 400085, India

³ Informatics and Communications Technology Complex, 535 W. Michigan Street, Indianapolis, IN 46202, USA

⁴ University Computing Centre SRCE, Zagreb, Croatia

⁵ AUTH SCC/GRNET S.A, Aristotle University of Thessaloniki, Thessaloniki GR 541 24, Greece

Abstract. The journey of a monitoring probe from its development phase to the moment its execution result is presented in an availability report is a complex process. It goes through multiple phases such as development, testing, integration, release, deployment, execution, data aggregation, computation, and reporting. Further, it involves people with different roles (developers, site managers, VO managers, service managers, management), from different middleware providers (ARC, dCache, gLite, UNICORE and VDT), consortiums (WLCG, EMI, EGI, OSG), and operational teams (GOC, OMB, OTAG, CSIRT). The seamless harmonization of these distributed actors is in daily use for monitoring of the WLCG infrastructure. In this paper we describe the monitoring of the WLCG infrastructure from the operational perspective. We explain the complexity of the journey of a monitoring probe from its execution on a grid node to the visualization on the MyWLCG portal where it is exposed to other clients. This monitoring workflow profits from the interoperability established between the SAM and RSV frameworks. We show how these two distributed structures are capable of uniting technologies and hiding the complexity around them, making them easy to be used by the community. Finally, the different supported deployment strategies, tailored not only for monitoring the entire infrastructure but also for monitoring sites and virtual organizations, are presented and the associated operational benefits highlighted.

1. Introduction

The continuous monitoring of the WLCG infrastructure is a complex task based on a large number of tools and middleware solutions provided by various worldwide projects. At the core of this activity is the Service Availability Monitoring (SAM) framework, which hides this complexity and orchestrates the monitoring workflow execution. The WLCG infrastructure resources are continuously monitored by SAM, which ensures the correct flow of monitoring data from the moment a probe starts to be developed to test a particular service functionality, to the moment where monitoring results for that service are made available. SAM generates, aggregates, and processes monitoring data to provide WLCG management, VOs, and sites with tools tailored for their needs offering an accurate view over the WLCG infrastructure.

2. Software and Projects

To better explain the journey of a monitoring probe in the WLCG infrastructure it is first important to understand the actors involved (projects, tools, services, etc).

2.1. Projects

The tools exploited in the monitoring of WLCG infrastructure are provided by WLCG and various other projects:

- European Middleware Initiative (EMI)¹: collaboration of the major European middleware providers (ARC, dCache, gLite, and UNICORE). EMI aims to deliver a consolidated set of middleware components and extend the interoperability between grids and other computing infrastructures [1]. It is funded by the European Commission;
- Open Science Grid (OSG)²: provides common service and support for resource providers and scientific institutions using a distributed fabric of high throughput computational services based on Globus technology [2]. It is funded by the United States Department of Energy and National Science Foundation;
- European Grid Infrastructure (EGI)³: not-for-profit foundation which coordinates the grid resources provided by National Grid Initiatives (NGIs) and other European International Research Organisations (EIROs) and releases grid technology developed by different EU projects (EMI, ICE).

2.2. Grid Services

The WLCG infrastructure provides many grid services for computing, storage, information systems, etc. These services are developed and maintained by the projects above. Besides the implementation of grid functionality many services also provide monitoring probes to test the service functionality. The most relevant grid services in use today are: gLite (FTS, LFC, BDII, CREAMCE, DPM, WMS, LB, WN), ARC (ARC-CE, ARC-LFC, ARC-SRM, ARC-RLS), UNICORE, dCache, DesktopGrid, and OSG.

2.3. Operational Services

Strictly related to the infrastructure monitoring activity are the operational services. These tools are also provided by the projects listed above and provide functionality to manage and coordinate the operation of the infrastructure. From the monitoring perspective they are useful to consume information about the topology of the infrastructure, scheduled and unscheduled infrastructure downtimes, supported VOs, etc. The most relevant operational services are:

- Configuration Database⁴: provided by EGI, it contains general information about the sites participating in the infrastructure;
- Operations Portal⁵: provided by EGI, it provides information about the infrastructure (dashboards, VO cards, broadcasts, notification) to NGIs and VOs;
- OSG Information Management⁶: provided by OSG, it allows the definition and management of the topology used by OSG systems and services.

¹ <http://www.eu-emi.eu>

² <http://www.opensciencegrid.org>

³ <http://www.egi.eu/>

⁴ <http://goc.egi.eu>

⁵ <http://operations-portal.egi.eu>

⁶ <http://oim.grid.iu.edu>

2.4. Monitoring Services

The monitoring services for WLCG are based on the SAM framework⁷ and the OSG Resource and Service Validation (RSV) service⁸.

SAM is a complete monitoring framework for grid services deployed and managed for distributed environments. It provides two main services: SAM-Nagios, a regional scheduler of monitoring probes supporting local configuration and storage/visualization of probe results; and SAM-Gridmon, a central aggregator of monitoring data and computation engine providing multiple interfaces that exposes curated monitoring results [3].

The RSV software provides a scalable and easy to maintain resource/service monitoring infrastructure for OSG sites. The RSV-SAM Transport component is responsible for forwarding WLCG specific information from RSV to SAM, which is aggregated by the SAM-Gridmon service.

In addition to SAM and RSV, other visualization tools, such as the WLCG experiment dashboards, provide tailored monitoring view for the different WLCG VOs.

3. Monitoring Workflow

When a site administrator or VO administrator checks site availability results of the last month, it is perhaps difficult to realize the length and complexity of the workflow which allowed the computation of such figures. This workflow and its main steps are now explained in detail.

Figure 1 provides an overview of the monitoring workflow in WLCG and depicts a number of details (projects, tools, interactions, etc) of the main phases of this workflow.

3.1. Probe Development

The monitoring process starts with the development of probes. These are simple applications targeted to test concrete functionality and status of a given service. Probes may include multiple metrics to test or verify different functionality of the same service. Due to the need to understand the behavior of the service being tested, each product team developing services for the WLCG infrastructure also implements the probes required to test its services. There are however two main exceptions: probes testing VO specific functionality, which are implemented by VOs, and probes testing operational tools, which are provided by the teams responsible for operations.

A total of 236 metrics are registered today for monitoring the WLCG infrastructure. These metrics are implemented by multiple software providers reflecting the diversity of grid services and operational tools present in the infrastructure. The most relevant providers of metrics are:

- 56 from WLCG experiments testing VO specific services and capabilities;
- 76 from EMI testing gLite, ARC, UNICORE, and dCache services;
- 45 from OSG testing OSG specific services;
- 10 from SAM testing different components of the monitoring services.

The SAM framework plays an important role at this early stage of the monitoring process by defining the policy governing the implementation of new probes⁹. This policy imposes a number of constraints on the following areas: development (programming languages, mandatory arguments), packaging (naming, structure, dependencies, meta-package), integration, and testing. The compliancy to these constraints is mandatory by all probes providers. In a very heterogeneous environment, such as the WLCG infrastructure, the definition of well-defined policies is fundamental for the successful execution of monitoring probes.

⁷ <http://tomtools.cern.ch/confluence/display/SAMWEB/Home>

⁸ <https://www.opensciencegrid.org/bin/view/MonitoringInformation/RSV>

⁹ <http://tomtools.cern.ch/confluence/display/SAMDOC/Probes+Development+Policy>

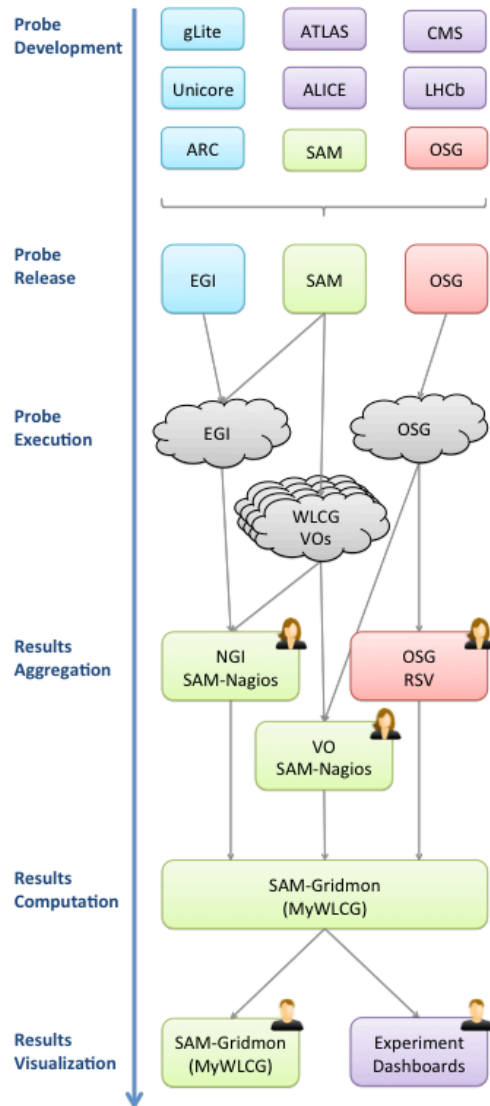


Figure 1. WLCG monitoring workflow

3.2. Probe Release

The second step of the monitoring workflow is the release of probes. Probes are implemented by different software teams and, consequently, there is the need to collect these contributions and prepare new software releases to be deployed in the WLCG infrastructure. The inclusion of probes in integrated software releases guarantees that probes work as expected when testing a given version of a service and when interacting with the SAM framework.

Probes are packaged and released by three entities: (1) EGI for probes related to EGI operations, gLite services, ARC services, UNICORE services, and DesktopGrid services, (2) OSG project for probes related to OSG services, and (3) WLCG for VO-specific probes. A total of 30 probe packages are available today in the WLCG infrastructure.

3.3. Probe Execution

After being released for production, probes are deployed to test all production services available in the WLCG infrastructure. The execution of probes is scheduled by multiple instances of SAM-Nagios and OSG-RSV services. These services are deployed under different scopes (per region, per VO, per site) and schedule the execution of tests for services under that scope. Monitoring of the WLCG infrastructure is a distributed responsibility where each region, VO, or site has the possibility to test its own service instances, using its own set of metrics. To guarantee minimum monitoring results, a number of metrics are always executed against a set of critical services. There are currently more than 40 instances scheduling tests on the infrastructure:

- 38 SAM-Nagios instances testing EGI services;
- 4 SAM-Nagios instances testing WLCG VO services;
- 1 OSG-RSV instance testing OSG services;

The SAM-Nagios instances are based on the Nagios system¹⁰. As such, metrics can be executed in active or passive mode. Active checks are initiated and performed by Nagios, while passive checks are performed by external applications. The interval for metric execution is configurable and usually ranges from 5 minutes to 1 day. For most of the metrics it is set to 30 or 60 minutes. As a consequence, all infrastructure services are constantly being tested. As of May 2012 the WLCG infrastructure was composed of 4,248 services endpoints, deployed at 517 sites, from 66 countries, grouped in 39 regions.

3.4. Results Aggregation

The results of the execution of metrics on the infrastructure nodes are first aggregated at region or VO level by SAM-Nagios and OSG-RSV services, the same that schedule the execution of metrics. This first level of aggregation per region or VO distributes the load and responsibility of monitoring the infrastructure. Metrics results are stored in each SAM-Nagios and different visualization tools are provided to verify the status of metrics execution, including the standard Nagios interface. To allow a central view of all metrics results and compute global availability and reliability figures for the entire infrastructure, a second aggregation of metrics results is done. All SAM-Nagios and OSG-RSV services publish their metrics results to the central SAM-Gridmon service.

The transport of metric results is based on a Messaging infrastructure. This includes the transport of passive tests from the nodes to scoped SAM-Nagios services, and the transport of region and VO metrics results to the central SAM-Gridmon service. Messaging provides a reliable, scalable, and secure transport layer to move metrics results [4].

Metric results can also be aggregated at site level with the deployment of SAM-Nagios configured to support local site services only. This configuration does not provide local storage of results and it only configures Nagios to test a number of local services. This can be used as an extra tool for sites that want to have a complete single view of all their services.

3.5. Results Computation

After aggregated at the central SAM-Gridmon service, metric results are computed to generate status, availability, and reliability statistics for services endpoints, service flavours, and sites. Status represents a discrete value for a set of metric results computed for each service endpoint, e.g. status of a given SRMv2 service is determined by aggregating several SRMv2 metrics. Availability computation gathers the service status information to compute the fraction of time the service was up (in OK state) during the period the service was known. Unlike availability, reliability does not consider scheduled downtime in the service known period [3].

¹⁰ <http://www.nagios.org/>

The results computation process starts with the computation of status, availability, and reliability for each infrastructure service endpoint. After computing the availability and reliability of each service, these results are combined to compute the availability and reliability of service flavours and sites [5]. These final results are then stored in a central database. At this moment results are no longer individual metrics result but instead status, availability, and reliability results for services endpoints, service flavours, and sites.

The SAM-Gridmon computation engine is currently processing more than 20,000 metrics results per hour, around half million results processed per day.

3.6. Results Visualization

The last step of the monitoring workflow is to provide monitoring results to be consumed by others. The most relevant information exposed are status, availability, and reliability statistics for services and sites. In addition, the related monitoring metadata (topology of the infrastructure, probes/metrics, profiles, etc.) is also exposed.

This information is made available by SAM-Gridmon service via two different interfaces: a REST API¹¹ for programmatic access, and a graphic interface, called MyWLCG¹², for end user access. The SAM-Gridmon service is currently serving approximately 15,000 per month visits and more than 1,000,000 hits per month. The SAM-Gridmon graphic interface offers to end users distinct tailored views over the monitoring data: service status, service availability, trends, reports, etc.

Figure 2. SAM-Gridmon MyWLCG home page

Figure 2 shows the MyWLCG home page which gives access to the different monitoring views. Other tools and dashboards use the SAM-Gridmon REST API to provide more specific views

¹¹ <http://grid-monitoring.cern.ch/mywlcg/sam-pi>

¹² <http://grid-monitoring.cern.ch/mywlcg>

over the results computed by SAM. This includes, for example, WLCG experiment dashboards and the EGI operations portal.

4. Improvements

SAM central role in monitoring WLCG sites and services allow a smooth integration of different projects and tools. Nevertheless with so many parties involved and so many interactions at different levels, many improvements are possible in the various phases of monitoring process:

- Probe development: The probes developed to test the grid services deployed in the infrastructure are extremely important from the operations perspective. To ensure the maintenance and future development of these probes it is important that the ownership of these probes is completely moved to the product teams responsible for the development of the related grid services
- Probe execution: From the probes execution perspective two possible improvements can be planned: 1) Currently SAM targets the execution of probes against single service endpoints (hostname plus service flavour). Other granularities of elements to probe, such as workflow of services, could also be considered. 2) VO probes can also be improved by moving away from testing grid service functionality and focus on VO specific characteristics.
- Results aggregation: A very interesting improvement related to the aggregation of results is the integration of external third party monitoring systems. SAM could be able to accept test results produced by external monitoring systems via the messaging infrastructure.
- Results computation: Possible improvements in this area are related to the possibility of correlating final status, availability, and reliability figures with data from other systems, such as job submission, storage elements, and data transfer services.
- Results visualization: From the results visualization point of view it could be studied the possibility of having one common web interface. Currently SAM-Gridmon offers a REST API used by different dashboards to display similar information. This may create confusion and duplication of effort which could be avoided by adopting one single web interface.

5. Summary

The journey of a monitoring probe from its development phase to the moment its execution result is presented in an availability report is a complex process. Thanks to SAM and OSG RSV monitoring services, it is possible to continuously and accurately monitor the WLCG infrastructure resources. The SAM framework periodically executes more than 200 hundred metrics on close to 5000 service endpoints distributed worldwide. This produces around half million results per day first aggregated at different scopes and then processed centrally. These results are finally exposed by different APIs and dashboards to provide official status, availability, and reliability numbers.

References

- [1] P. Fuhrmann, M. Cecchi, L. Field, J. White, B. Konya, J. Nilson, 2012, *Consolidation and development roadmap of the EMI middleware* (CHEP 2012)
- [2] R. Pordes, M. Livny, 2012, *The Open Science Grid Support for Multi-Disciplinary Team Science - the Adolescent Years* (CHEP 2012)
- [3] P. Andrade, M. Babik, K. Bhatt, P. Chand, D. Collados, V. Duggal, P. Fuente, E. Imamagic, P. Joshi, R. Kalmady, U. Karnani, V. Kumar, J. Tarragon, W. Lapka, and C. Triantafyllidis, 2012, *Service Availability Monitoring Framework Based On Commodity Software* (CHEP 2012)
- [4] L. Cons, M. Paladin, 2012, *The WLCG Messaging Service and its Future* (CHEP 2012)
- [5] R. Kalmady, D. Sonvane, P. Chand, K. Bhatt, V. Kumar, P. Nyczyk, and Z. Sekera, 2009, *Monitoring the Availability of Grid Services Using SAM and Gridview* (Grid Computing Part IV, 163-168, DOI: 10.1007/978-0-387-78417-5)