

Editorial Manager(tm) for Journal of Physics: Conference Series  
Manuscript Draft

Manuscript Number:

Title: RSV: OSG Fabric Monitoring and Interoperation with WLCG Monitoring Systems

Article Type: Contributed

Corresponding Author: Mr. Rob Quick, BS Physics

Corresponding Author's Institution: Indiana University

First Author: Rob Quick, BS Physics

Order of Authors: Rob Quick, BS Physics; Arvind Gopu; Soichi Hayashi

# RSV: OSG Fabric Monitoring and Interoperation with WLCG Monitoring Systems

R Quick<sup>1</sup>, A Gopu, S Hayashi

535 W Michigan ITCT 314F Indianapolis IN 46202 USA

E-Mail: [rquick@iu.edu](mailto:rquick@iu.edu)

**Abstract.** The Open Science Grid (OSG) Resource and Service Validation (RSV) project provides solutions to several grid fabric monitoring problems, while at the same time providing a bridge between the OSG operations and monitoring infrastructure of the WLCG (Worldwide LHC Computing Grid) infrastructure. The RSV-based OSG fabric monitoring begins with local resource fabric monitoring, which gives local administrators tools to monitor their status on the OSG using their local monitoring infrastructure. With a set of local grid status probes, the results of which are uploaded to a central collector, a system administrator can monitor and watch their resource locally, while the OSG Operations Center (GOC) can watch from a centralized position. Plug-ins relay RSV results to other popular fabric monitoring software (Nagios) allowing system administrators flexibility to stay aware of their grid status using their chosen status display interface. Additional probes are easily developed and plugged into the RSV structure, and an emphasis is placed on the community to develop additional probes that fit the needs of different categories of users (VO, User, Software Developer) as needed. From the GOC, results are transmitted to a WLCG message broker via a specified format, which can then translate these records into critical statistics to the LHC collaborating projects. RSV has succeeded in meeting these initial goals; future development is centred around usability and extending the project's scope and functionality.

## 1. Introduction

The Open Science Grid (OSG) [1] Resource and Service Validation (RSV) [2] project was conceived in spring 2007 based on a core philosophy of providing grid level service status monitoring of OSG resources that would allow resource administrators to increase the availability of the compute resources and that would allow resource users to gather critical real-time status information before submitting jobs.

The driving concepts behind RSV are that:

- Tests should occur locally and be the responsibility of the resource administrator instead of a central agency.
- Results should be collected and made available to the community in both a human readable and programmatic fashion to allow the best choices for job submission.

---

<sup>1</sup> To whom any correspondence should be addressed.

- Tests should be able to run an existing fabric monitoring framework (e.g. Nagios), or be independently based on the local resource preference.
- Tests should meet a specification that is easily gathered and analyzed by other peering organizations.

The final outcome of the project should produce a more reliable grid structure.

## **2. Timeline and Releases**

### 2.1 RSV Version One

The initial version of RSV was released in the OSG middleware stack in the spring of 2008. The initial version consisted of a set of RSV probes, a Condor [3] scheduling mechanism, a Gratia collector and a messaging service based on ActiveMQ.

The RSV probes are simple perl scripts that can be run locally or remotely and do a specific check of a grid service (e.g. ping, hello world, file transfer, CA date check, etc.) against the middleware stack at an OSG resource. The probes are independent of the scheduling and collection mechanisms and can be run by hand if desired. This allows flexibility of scheduling and collection and ultimately integration with other local fabric monitoring systems. Probe output was based on a specification put forth by the WLCG Grid Monitoring Group [4].

A scheduling mechanism based on Condor's Cron functionality and a collector based on Gratia accounting software, both working pieces already available in the OSG middleware stack, were provided as additional core components to complete the RSV service. The ActiveMQ messaging component that conformed to the WLCG specification completed the goal of allowing easy transfer to collaborating grids.

The missing components of the first version of RSV were a central display to allow easily available near real-time status to be gathered by OSG users and direct plug-ins for common fabric monitoring tools.

### 2.2 RSV Version Two and MyOSG

Phase two of RSV completed and released in Summer 2008 consisted mostly of bug fixes and an updated probe set along with plug-ins that allowed display of RSV probes to be sent to a Nagios server.

In the spring of 2009 the central presentation layer was released and tagged MyOSG [5]. This brought together RSV probe results with existing OSG information services to create a unique grid monitoring experience. MyOSG was not designed as a standard dashboard, consolidating RSV along with accounting, administrative, and other OSG information sources to be presented in a user defined work area using UWA [6] and widget display technologies.

### 2.3 RSV Version Three

Phase three of RSV is scheduled for fall of 2009, RSV will add a still larger and updated probe set, tools for RSV version and configuration management and a central environment to allow bookkeeping against local probes. This will come with a central OSG Nagios instance, which will run a limited critical set of probes against each OSG resource and compare results against what the local probes are reporting. Nagios will also check other critical infrastructure services in the OSG. Management tools will allow easy version updates and re-execution of probes.

## **3. Discussion**

### 3.1 Benefits of Local Monitoring

The decision to bring monitoring to the local level instead of running as a central service served the specific goal of getting the resource administrators involved with the results being published to the OSG. The initial version of RSV met some resistance from resource administrators due to configuration complexity but as the administrators became familiar with the software there was generally a good response from the OSG community. ATLAS and CMS were quick to adopt RSV because the messaging abilities allowed them to publish RSV records to WLCG sources. The overall goal of resource administrators having a vested interest in how and what their resources was reporting paid off. The OSG Operations Center (GOC) regularly communicates with resource administrators about the results, or lack of results, in RSV testing.

### 3.2 Integration with Local Fabric Monitoring

Integrating RSV with site local fabric monitoring has gotten less attention but is an important goal of RSV. RSV now has plug-ins for displaying RSV results in Nagios (figure 1), but at this time there is not a complete Nagios wrapper which will run the RSV probe completely inside of the Nagios environment. Collaboration with WLCG efforts along these lines has been helpful, but at this point this work is still incomplete. When completed, integration with Nagios will further integrate resource providers in the overall monitoring structure of the OSG. Other fabric monitoring systems will be considered based on demand from the resources.

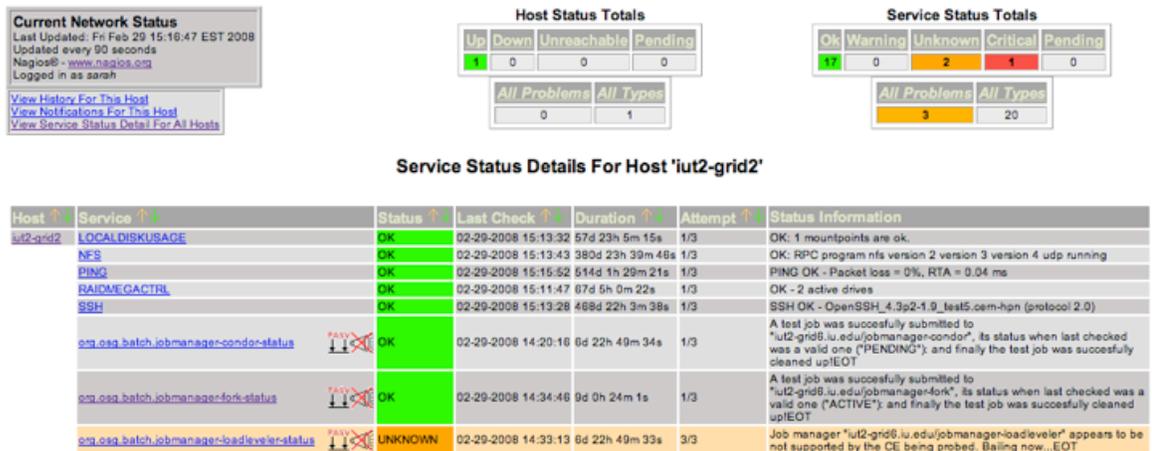


Figure 1 – RSV tests displayed in Nagios [7].

### 3.3 WLCG Information Transmission

OSG personnel serve as members of the WLCG Grid Monitoring Working Group. This allowed that specifications of the monitoring records were agreed upon between OSG and WLCG. These specifications were built into RSV from the beginning of the project and made it possible for monitoring records from OSG and EGEE to be compared side by side in the WLCG SAM and GridView environments. Transmission via Apache ActiveMP was developed to assure records were held upon outage and transmitted when the messaging service became available. Reports doing daily comparisons of OSG versus WLCG are created to assure common results to the OSG stakeholders.

### 3.4 MyOSG Tool Consolidation and Personalized Display

The goal of creating a presentation layer for RSV records intersected with the separate goal of consolidating information about OSG from many websites. MyOSG was developed to aggregate information from many OSG systems: including administrative, accounting, status, monitoring services, and ticket tracking tools. OSG Operations noticed several grid dashboard projects fail over the lifetime of the OSG because they were not flexible to cater to the needs of different user types. Keeping in mind each individual in OSG had a unique view they were interested in, the decision was made to use the UWA specification developed by NetVibes. This allowed generic widgets to be created that individuals could integrate with their personal workflow. Thus MyOSG creates widgets based on user preferences. The information can then be easily exported to a number of tools for display including iGoogle , NetVibes (figure 2), Windows Vista, Apple Dashboard, Live.com, iPhone, Opera, blogs, MySpace, etc.

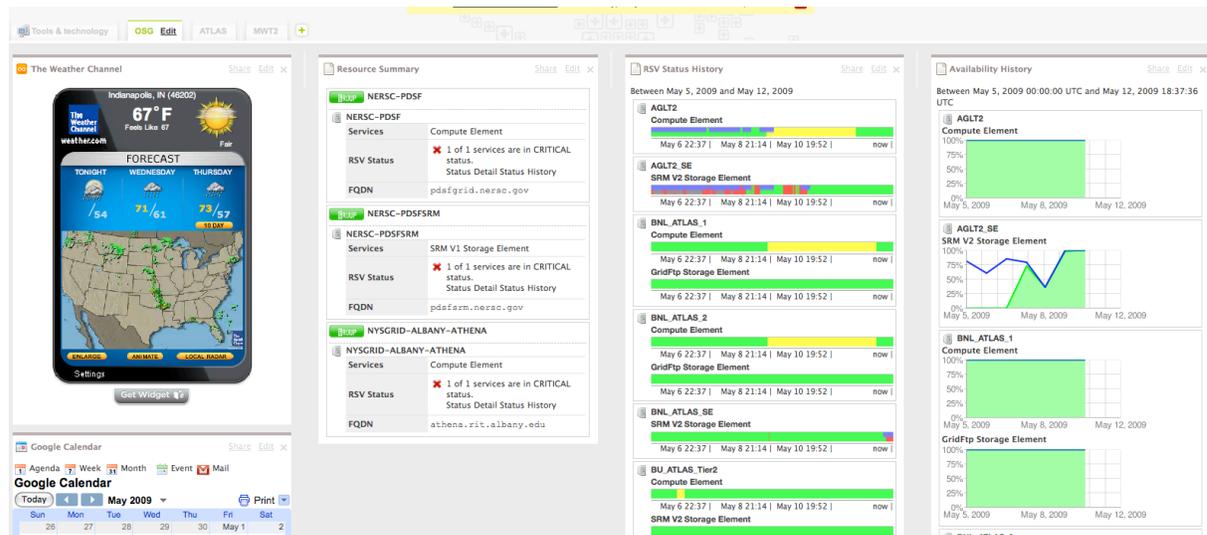


Figure 2 – Various MyOSG status widgets displayed with everyday workflow items (weather and calendar) in NetVibes.

#### 4. Conclusion

The overall goal of RSV and later MyOSG was to make the OSG more reliable overall by making resource status easy to monitor by OSG resource administrators. By creating site-level probes along with an infrastructure to execute and understand the results of the tests, RSV has put grid status monitoring into the hands of those who are closest to the issues it discovers and has improved the reliability of those resources utilizing it.

By listening to the administrators who are now using RSV and MyOSG, the GOC is able to determine that RSV is useful. By making the probe structure and the reporting structure simple, changes can be made modularly and if any component is not working or producing expected results it can easily be removed or updated. There is also flexibility to create application or Virtual Organization level probes. In the end, we find that RSV meets the needs of the OSG and improves performance, and is flexible enough to keep up with an evolving grid structure.

#### References

- [1] “OSG Organization and Model of Operation.” Fermi National Accelerator Laboratory. <[http://www.opensciencegrid.org/About/Learn\\_About\\_Us/OSG\\_Organization](http://www.opensciencegrid.org/About/Learn_About_Us/OSG_Organization)>
- [2] Rob Quick and Arvind Gopu “OSG Resource and Service Validation and WLCG SAM Interoperability.” 2008. Indiana University. 07 Apr. 2008. <<http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=748>>

- [3] Condor Manual. "CronTab Scheduling." 2009. University of Wisconsin-Madison. 14 Apr. 2009  
[http://www.cs.wisc.edu/condor/manual/v7.2/2\\_12Time\\_Scheduling.html#sec:CronTab](http://www.cs.wisc.edu/condor/manual/v7.2/2_12Time_Scheduling.html#sec:CronTab)
- [4] James Casey, Ian Neilson. "Grid Service Monitoring Working Group." 2007. European Organization for Nuclear Research. 19 Nov. 2008.  
<<https://twiki.cern.ch/twiki/bin/view/LCG/GridServiceMonitoringInfo>>
- [5] Soichi Hayashi, Arvind Gopu, Rob Quick. "MyOSG Project Plan" 2008. Indiana University. 27 Feb 2009. <[http://docs.google.com/View?docid=ddtgc5bt\\_44hsps24fw](http://docs.google.com/View?docid=ddtgc5bt_44hsps24fw)>
- [6] Netvibes. "UWA Documentation." 2008. Netvibes. 20 Nov. 2008  
<http://dev.netvibes.com/doc/uwa/documentation>
- [7] Sarah Williams. Indiana University. 14 May 2009. Contributed Nagios Screenshot.