

Globus Toolkit Use Case: A Monitoring System for the Earth System Grid

15 September 2004

Ben Clifford and Shishir S. Bharathi

Table of Contents

1	Introduction	2
2	Description/Scenario	2
3	Monitoring System Architecture	2
4	Implementation Details.....	5
4.1	Required Software	5
4.2	ESG Implementation.....	5
4.3	Intended Future Implementation	6
5	Postconditions.....	6
5.1	ESG Visualizer	7
5.2	Details Visualizer	8
5.3	Archiver Visualizer.....	8
6	Resources	9

1 Introduction

This use case describes how components provided by the Globus Toolkit's Monitoring and Discovery System (MDS3) were used to create a monitoring system for services deployed in the Earth System Grid (ESG). This system enables ESG administrators and users to determine the status of ESG services, to be notified automatically of changes in service status, and to review historical status information.

2 Description/Scenario

The Earth System Grid (ESG) federates large-scale data and analysis servers to provide community access to large quantities of data produced by climate simulation models. ESG users register with a Web portal (www.earthsystemgrid.org) and can then use that portal to discover, browse, and request (subsets of) datasets, with requested data being returned to the user's workstation.

The ESG infrastructure (Figure 1) comprises a distributed collection of physical devices and software services, including:

- archival storage systems and disk storage systems at different sites
- storage resource managers (SRMs) and GridFTP servers that provide access to those storage systems
- metadata catalog services used to maintain metadata about data
- replica location services used to support the location of dataset replicas
- the Web portal itself.

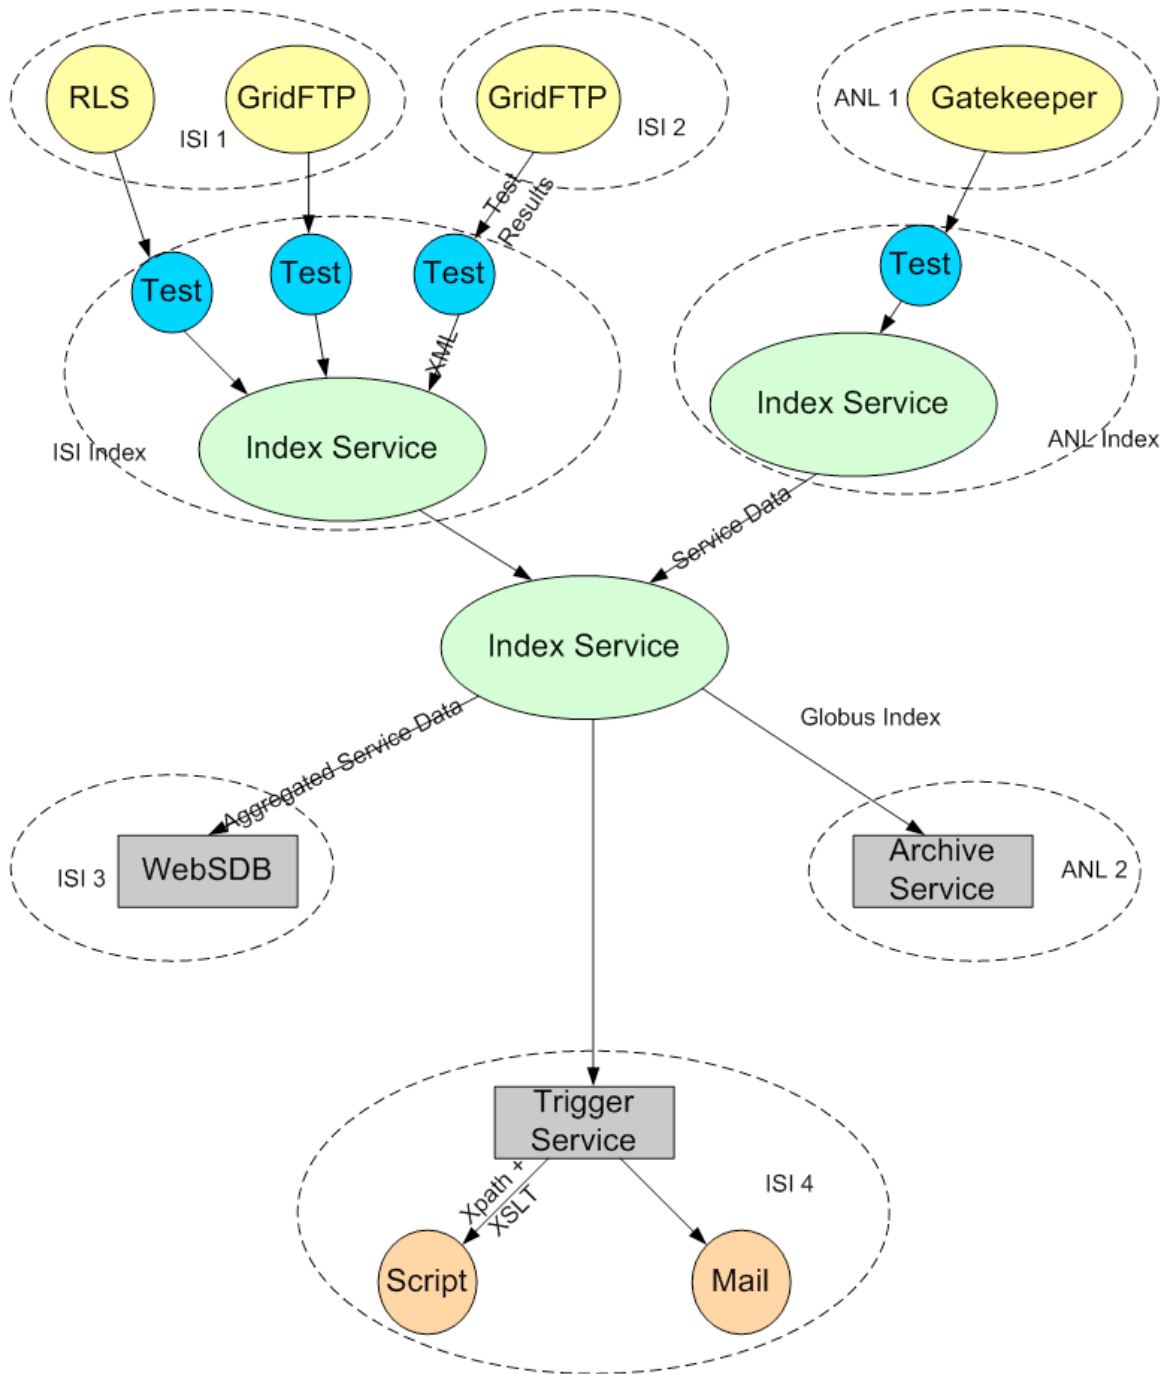
These various components cooperate to allow ESG users to find and access physical files based on descriptive metadata attributes.

ESG needed a way of monitoring the behavior of these various components with a view to detecting and notifying interested parties of failures. We describe here how this requirement was met by using Globus Toolkit components, specifically those provided by the Globus Toolkit's Monitoring and Discovery System (MDS).

3 Monitoring System Architecture

The ESG monitoring system was constructed by integrating four MDS components—Index Service, Web Service Data Browser, Archive Service, and Trigger Service—with ESG-specific information providers and action scripts, as illustrated in the following figure:

Figure 1: MDS components in the monitoring framework



3.1 Index Service

The MDS3 **Index Service** obtains and caches information about other services, and allows clients to query and/or subscribe to that data. The Index Service collects service data from various resources, which it publishes as one big piece of service data. Other

services can query for service data elements or subscribe to be notified when they change. Multiple indexes can be used to obtain scalability to large numbers of services and to configure the system so that it can tolerate failures (i.e. have multiple collectors of data.)

The ESG monitoring system uses the Index Service to collect and publish information about the current status of ESG services.

3.2 Information Providers

In MDS3, **Information Providers** (also referred to in GT3 as *Service Data Providers*) gather and generate data for use by Grid service instances. The data is generated in the form of XML objects, which become the value of a Service Data Element (SDE) emitted as either a Java output stream or a memory-bound Java DOM (Document Object Model) representation.

The ESG monitoring system uses information providers to run a query against a service and, depending on the results, judge whether it's up or down. These are very simple shell scripts that can be written by system administrators and do not involve any configuration. They run periodically on the same machine as the Index Service, wrap the results of these tests in XML, and then publish the resulting XML in the Index Service as *service data*.

3.3 Archive Service

The new **Archive Service** is designed to store and allow time-based queries on Index Service data. The ESG monitoring system uses the Archive Service to enable historical queries about component behavior.

3.4 Web Service Data Browser

The **Web Service Data Browser** (WebSDB) is a customizable display client for service data exposed by grid services. In the ESG monitoring system, the WSDB is used to allow users to query both the Index Service and the Archive Service (see below) to determine the current and past status, respectively, of ESG components.

3.5 Trigger Service

The **Trigger Service** provides functionality such as an email gateway to let system administrators know when services fail (the email IDs of the recipients are generated when the XSLT is performed.) An **action script** may be a simple shell script or a wrapper around other applications such as a grid service clients or mail clients. It is configured with a set of Xpath conditions and, for each condition, an action to be performed when the condition is satisfied plus an optional XSLT stylesheet. The following briefly describes the process:

1. The Trigger Service subscribes to the Index Service.
2. Each time a notification is received, the Trigger Service checks each Xpath condition against the service data in the email.

3. If any condition is met, the Trigger Service runs a specified action script, with the product of the XSLT transformation applied to the matching XML.

4 Implementation Details

The following is a high-level overview of how to set up a monitoring framework using MDS components:

1. Set up an Index Service and configure it to publish data from these tests:
http://www-unix.globus.org/toolkit/docs/3.2/installation/install_config_index.html
2. Set up hierarchy (Index Services aggregating other Index Services.)
3. Set up the Archive Service and subscribe to the correct Index Service.
(link to directions coming soon)
4. Configure the Trigger Service with conditions and action scripts for this service.
5. Write Visualizers (ESG, Archive, Details, explained below), using XSL stylesheets to format and display the data.
6. Set up the WSDB and install Visualizers.
<http://www-unix.globus.org/toolkit/docs/3.2/infosvcs/ws/developer/sbrowser.html>

4.1 Required Software

- Globus Toolkit 3.2
<http://www-unix.globus.org/toolkit/downloads/>
- the Archive Service package
(coming soon)
- the Trigger Service package
(coming soon)
- the Web Service Data Browser
<http://gtr.globus.org/article.php?story=20030723143444410>

4.2 ESG Implementation

The following details are specific to deploying the monitoring framework for ESG:

1. All components are installed on dc-user.isi.edu (CGT machines) using a single Index Service.
2. Status of services is available on this web page:
<http://dc-user.isi.edu:40080/monitor.html>
3. The Trigger Service sends an email when a broken service is detected.
4. The following data is captured:
 - a. Status (whether the server is up or down)
 - b. Percentage of uptime
 - c. Time when the service was last seen alive

- d. Time when the service failed

4.3 Intended Future Implementation

- Displaying archived data in more interesting and meaningful ways
- Deploying on an ESG node
- Monitoring of other services such as SRM, GridFTP, etc.

5 Postconditions

As a result of this deployment, ESG now has the following:

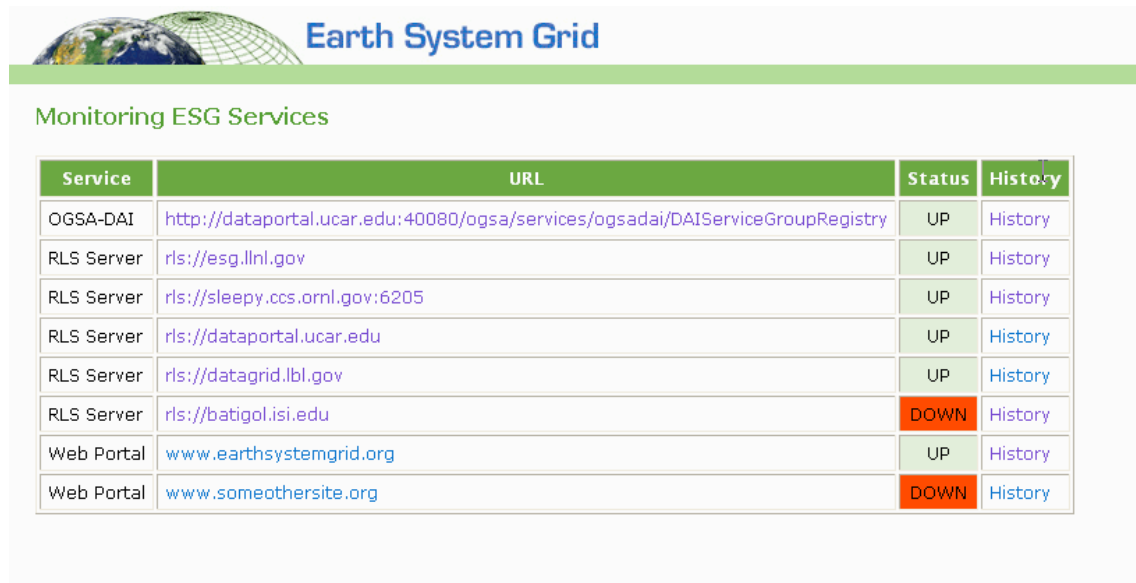
- A scalable system that can easily cope with services being added and removed.
- Status and performance is available for any other service to query via the Index Service and Archive Service
- Test results returned as XML:
 - In our current test schema, we allow for a binary UP|DOWN status value, and also a free form comment
 - the free form comment can contain, for example, RLS stats which include number of mappings in the LRC/RLI, uptime, etc.
 - Complex XPath queries can be run against these results as part of the triggering mechanism
- Web-based visualizers that display the status of its chosen services (see the following topics.)

5.1 ESG Visualizer

This is the 'entry page' for checking out how the various services are doing: status (whether the service is UP or DOWN) and links to more information from the Archiver and Details visualizers.

View the ESG Visualizer at <http://dc-user.isi.edu:40080/monitor.html>

Figure 2: Screenshot of the ESG Visualizer

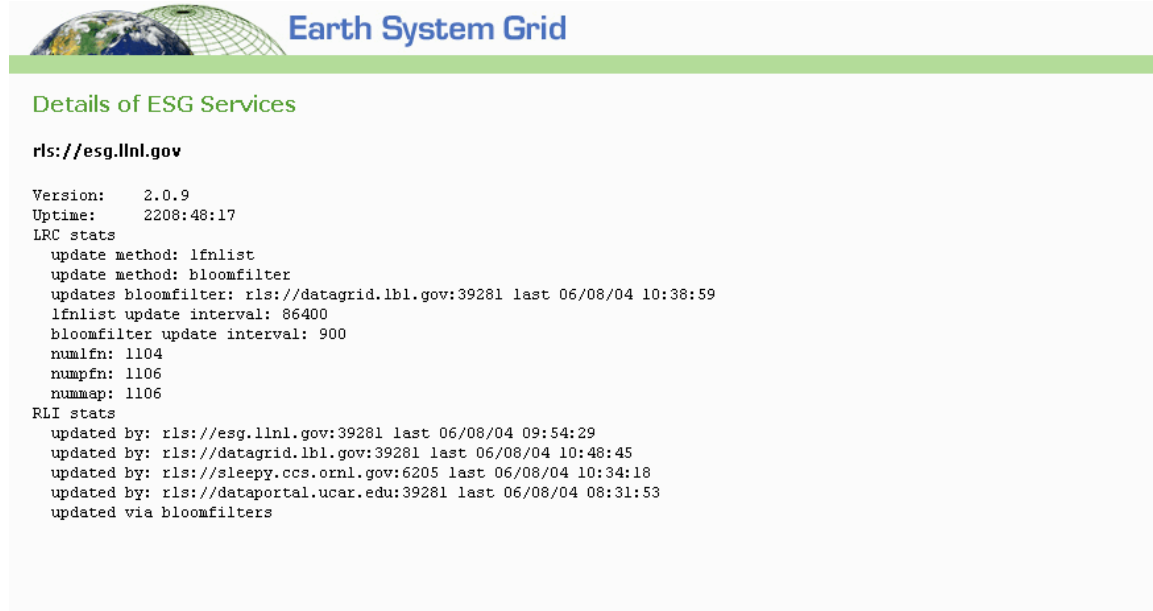


Service	URL	Status	History
OGSA-DAI	http://dataportal.ucar.edu:40080/ogsa/services/ogsadai/DAIServiceGroupRegistry	UP	History
RLS Server	rls://esg.llnl.gov	UP	History
RLS Server	rls://sleepy.ccs.ornl.gov:6205	UP	History
RLS Server	rls://dataportal.ucar.edu	UP	History
RLS Server	rls://datagrid.lbl.gov	UP	History
RLS Server	rls://batigol.isi.edu	DOWN	History
Web Portal	www.earthsystemgrid.org	UP	History
Web Portal	www.someothersite.org	DOWN	History

5.2 Details Visualizer

From the ESG Visualizer, the *URL* link takes you to the Details Visualizer.

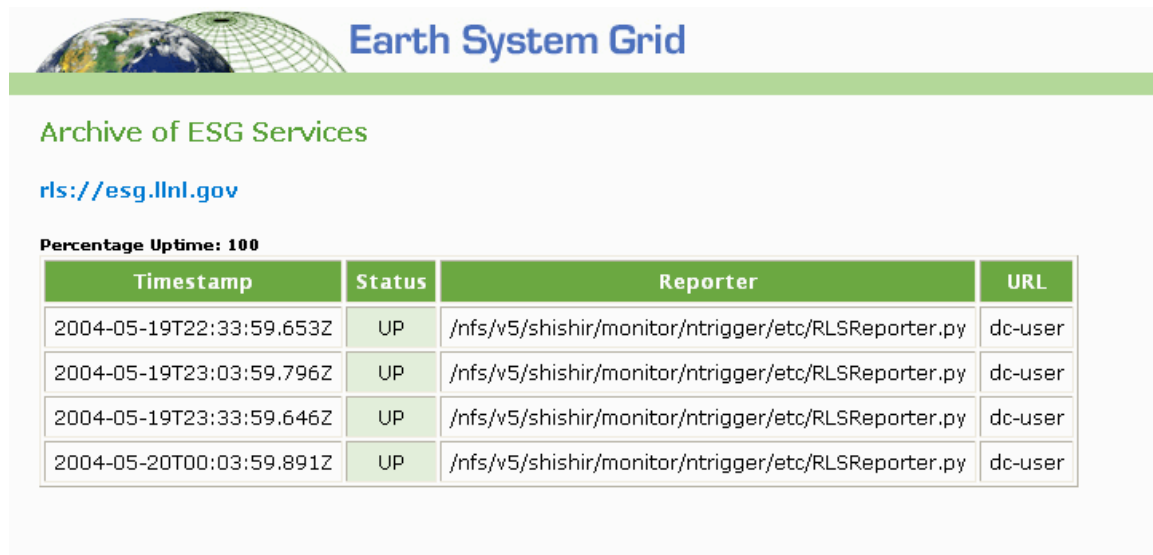
Figure 3: Screenshot of the Details Visualizer



5.3 Archiver Visualizer

From the ESG Visualizer, the *History* link takes you to the Archive Visualizer, which displays the historical data ESG specified (percentage of uptime, the timestamp of the reading, status of the server, reporter (test script) and URL.)

Figure 4: Screenshot of the Archive Visualizer



6 Resources

This page provides information about where and how you can access the tools used in this implementation:

Index Service	The Index Service is currently available as part of GT3.2: http://www-unix.globus.org/toolkit/downloads/3.2/index.html
Web Service Data Browser	Developed by Arslan Javed ajaved@isi.edu Available from the Grid Technology Repository: http://gtr.globus.org/article.php?story=20030723143444410
Archive Service	Developed by June Sup Lee (june@isi.edu) Currently only available through the Globus CVS.
Trigger Service	Developed by Shishir Bharathi (shishir@isi.edu) (coming soon)