
GT 4.2.1 Using Scheduler Event Generator with GRAM2

Table of Contents

1. Introduction	1
2. Overview of Operation	1
3. Configuration	1
4. Running the globus-job-manager-event-generator	3
5. Troubleshooting	4

<titleabbrev>Using Scheduler Event Generator with GRAM2</titleabbrev>

1. Introduction

Starting with Globus Toolkit 4.1.0, the GRAM2 Job Manager supports using the Scheduler Event Generator (SEG) for obtaining job state information from the scheduler log files in place of running a poll command periodically. This has the effect in most situations of reducing the impact (in terms of CPU usage) of the job manager. This document describes how to configure and use this new feature.

2. Overview of Operation

The WS-GRAM Scheduler Event Generator (SEG) is a program which parses native log files generated by the schedulers supported by GRAM, and uses the information in them to issue events to stdout which are piped back to the WS-GRAM Job Manager service. This avoids the sometimes costly poll operation periodically done by the GRAM scheduler adapters.

For GRAM2, a program called **globus-job-manager-event-generator** runs the SEG and writes job state change records into a log file which all users can read. This log contains the minimal information about jobs to determine when they are queued, become active, and terminate. No user-specific or job-specific data is revealed in this log file. The GRAM2 Job Manager can be configured to use this log file as a source for job state change events. A single instance of the **globus-job-manager-event-generator** will be run for each scheduler on the system.

3. Configuration

3.1. Job Manager Configuration

By default, the job manager uses the GRAM2 script-based polling method. A new command line option (*-seg*) enables SEG-driven job state change notifications.

There are two ways to configure the job manager to use the scheduler event generator: globally, in the `$GLOBUS_LOCATION/etc/globus-job-manager.conf` file, or on a per-service basis in the service entry file in the `$GLOBUS_LOCATION/etc/grid-services` directory.

3.1.1. Global Job Manager Configuration

To enable using the Scheduler Event Generator interface for all Job Managers started from a particular `GLOBUS_LOCATION`, add a line containing the string

-seg

to the file `$GLOBUS_LOCATION/etc/globus-job-manager.conf`.

Example 1. Example `$GLOBUS_LOCATION/etc/globus-job-manager.conf`

```
-home "/opt/globus"  
-globus-gatekeeper-host globus.yourdomain.org  
-globus-gatekeeper-port 2119  
-globus-gatekeeper-subject "/O=Grid/OU=Your Organization/CN=host/globus.yourdomain.org"  
-globus-host-cputype i686  
-globus-host-manufacturer pc  
-globus-host-osname Linux  
-globus-host-osversion 2.6.10  
-save-logfile on_error  
-state-file-dir /opt/globus/tmp/gram_job_state  
-machine-type unknown  
-seg
```

3.1.2. Scheduler-specific Job Manager Configuration

To enable using the Scheduler Event Generator interface for a particular Job Manager, add the string `-seg` to the end of the line in the service's file in the `$GLOBUS_LOCATION/etc/grid-services` directory.

Example 2. Example `$GLOBUS_LOCATION/etc/grid-services/jobmanager-lsf`

```
stderr_log,local_cred - /opt/globus/libexec/globus-job-manager globus-job-manager -conf /o
```



Important

The Job GRAM2 Job ManagerManager does not support using the SEG for the fork scheduler. If the `-seg` option is passed to a fork Job Manager, it will be ignored.

3.2. globus-job-manager-event-generator Configuration

The `globus-job-manager-event-generator` program requires that the `globus_job_manager_event_generator` setup package be installed and run. This setup package creates the `$GLOBUS_LOCATION/etc/globus-job-manager-seg.conf` file and initializes a directory to use for the scheduler logs.

By default, this setup script will create a configuration entry and directory for each scheduler installed on the system. For each scheduler to be handled by the `globus-job-manager-event-generator` program, there must be an entry in the file in the pattern:

```
SCHEDULER_TYPE_log_path=PATH
```

The two variable substitutions for this pattern are

**SCHED-
ULER_TYPE** Must match the name of the scheduler-event-generator module for the scheduler (supported with GT 4.1 are `lsf`, `condor`, and `pbs`).

PATH A path to a directory which must be writable by the account which will run the `globus-job-manager-event-generator` program for the `SCHEDULER_TYPE`, and world-readable (or readable for a group which contains all users which will run jobs via GRAM on that system). Each directory specified in the configuration file must be unique, or behavior is undefined.

Example 3. Example `$GLOBUS_LOCATION/etc/globus-job-manger-seg.conf`

```
lsf_log_path=/opt/globus/var/globus-job-manager-seg-lsf
pbs_log_path=/opt/globus/var/globus-job-manager-seg-pbs
```

In this example, pbs and lsf schedulers are configured to use distinct subdirectories of the `/opt/globus/var/` directory.

Important

For best performance, the log paths should be persistent across system reboots and mounted locally (non-networked).

Important

If a scheduler is added after the configuration step is done, administrator must rerun the setup package's script (`$GLOBUS_LOCATION/setup/globus/setup-seg-job-manager.pl`) or modify the configuration file and create the log directory with appropriate permissions.

4. Running the `globus-job-manager-event-generator`

The `globus-job-manager-event-generator` script creates a log of all scheduler events related to a particular scheduler instance. This script was created for two purposes

- To avoid requiring that all GRAM users have the privileges to read the scheduler's log file. Users may not be allowed read access to the scheduler's log files. Since the Pre-WS GRAM Job Manager runs with the permissions of the user account, it may be unable to access the log files. Instead the **globus-job-manager-event-generator** program will run as a privileged user and then store job state change records in a file which GRAM2 users may access.
- To provide a simple format for the scheduler event generator logs so that the job manager will be able to quickly recover state information if the job manager is terminated and restarted. Some scheduler logs are difficult to parse, or inefficient for seeking to a particular timestamp (as is necessary for recovering job state change information). The data written by this script is easily locatably by date, and it is simple to remove old job information without compromising current job manager execution.

One instance of the **globus-job-manager-event-generator** must be running for each scheduler type to be implemented using the Scheduler Event Generator interface to receive job state changes. This program is located at `$GLOBUS_LOCATION/sbin`. The typical command line for this program is `$GLOBUS_LOCATION/sbin/globus-job-manager-event-generator -s SCHEDULER_TYPE`, where `SCHEDULER_TYPE` is the scheduler name of the SEG module which should be used to generate events (lsf, condor, pbs).

For example, to start the event generator program to monitor an LSF batch system:

```
$GLOBUS_LOCATION/sbin/globus-job-manager-event-generator -s lsf
```

Important

If the **globus-job-manager-event-generator** is not running, no job state changes will be sent from any job manager program which is configured to use the SEG.

5. Troubleshooting

PROBLEM: The `globus-job-manager-event-generator` program terminates immediately with the output:

```
Error: SCHEDULER not configured
```

- Make sure that you specified the correct name for the *SCHEDULER* module on the command line to the **globus-job-manager-event-generator** program.
- There is no entry for *SCHEDULER* in the `$GLOBUS_LOCATION/etc/globus-job-manager-seg.conf` file. See the section on `globus-job-manager-event-generator` Configuration.

PROBLEM: The **globus-job-manager-event-generator** program terminates immediately with the output:

```
Fault: globus_xio: Operation was canceled
```

- The scheduler module selected on the command line could not be loaded by the SEG. Check that the name is correct, the module is installed, and the setup script for that module has been run.

PROBLEM: The Job Manager never receives any events from the scheduler.

- Verify that the directory specified in the `$GLOBUS_LOCATION/etc/globus-job-manager-seg.conf` for the scheduler exists, is writable by the account running the **globus-job-manager-event-generator** and is readable by the user account running the job manager.
- Verify that the **globus-job-manager-event-generator** program is running.
- Verify that the **globus-job-manager-event-generator** program has permissions to read the scheduler logs. To help diagnose this, run (as the account you wish to run the `globus-job-manager-event-generator` as) the command `$GLOBUS_LOCATION/libexec/globus-scheduler-event-generator -s SCHEDULER_TYPE -t 1` You should see events printed to the stdout of that process if it is working correctly.