
GT 5.0.0 Release Notes: GRAM5

Table of Contents

1. Component Overview	1
2. Feature summary	1
3. Summary of Changes in GRAM5	2
4. Bug Fixes	2
5. Known Problems	2
6. Technology dependencies	4
7. Tested platforms	4
8. Backward compatibility summary	5
9. Associated Standards	5
10. For More Information	5
Glossary	5

<titleabbrev>Release Notes</titleabbrev>

1. Component Overview

The Grid Resource Allocation and Management (GRAM5) component is used to locate, submit, monitor, and cancel jobs on Grid computing resources. GRAM5 is not a *job scheduler*, but rather a set of services and clients for communicating with a range of different batch/cluster job schedulers using a common protocol. GRAM5 is meant to address a range of jobs where reliable operation, stateful monitoring, credential management, and file staging are important.

2. Feature summary

New Features new since 4.2.x

- Server-side architectural changes to improve scalability, performance, and reliability
- Improved error notification protocol compared to GRAM2
- Teragrid Gateway Identity support for job auditing
- Usage stats messages
- Added support for Sun Grid Engine (SGE)

Other Standard Supported Features

- Remote job execution and management
- Uniform and flexible interface to local resource managers
- File staging before and after job execution
- File and directory clean up after job termination
- Service auditing for each submitted

Removed Features

- The GRAM5 client tools have dropped support for the Duroc API for task coallocation
- The GRAM5 service no longer streams output and error during job execution; instead this data is send after the job terminates
- The GRAM5 service no longer provides intra-job communication via the DUCT API
- The GRAM5 does not rely on XML schemas and WSDL service definitions

3. Summary of Changes in GRAM5

GRAM5 represents a significant improvement from GRAM2 and GRAM4 service implementations. GRAM2's limitation is scalability. GRAM4's is reliability. GRAM5 is both reliable AND scalable. It is important to note that GRAM5 is GRAM2 compatible. There are other improvements as well, like completely rewritten service logging based on the [CEDPS logging best practices](#)¹, Teragrid Gateway Identity support for job auditing, support for job exit codes, and [usage stat support](#)².

We have been very encouraged by our [performance results](#)³, which shows greater than 10x scalability than GRAM2 and roughly 10x reduction in resource consumption on the service host. We welcome your feedback as you integrate GRAM5 into your production grids.

4. Bug Fixes

- [All Fixes](#):⁴ This link lists all 67 of the improvements and bug fixes that were done for the GT 5.0.0 release

5. Known Problems

The following problems and limitations are known to exist for GRAM5 at the time of the 5.0.0 release:

5.1. Limitations

- [list limitations]

5.2. Outstanding bugs

- [Bug 108](#):⁵ Fork perl zombies
- [Bug 106](#):⁶ Fix test failures with SGE LRM adapter
- [Bug 105](#):⁷ Held Condor jobs should be reported as SUSPENDED
- [Bug 104](#):⁸ globus-job-manager-event-generator loads all historical events the first time run

¹ <http://www.cedps.net/index.php/LoggingBestPractices>

² <http://www.globus.org/toolkit/docs/5.0/5.0.0/execution/gram5/user/#gram5-usage-stats>

³ <http://www.globus.org/toolkit/docs/5.0/5.0.0/execution/gram5/qp/#gram5-reports-throughput>

⁴ <http://tinyurl.com/yagefou>

⁵ <http://jira.globus.org/browse/GRAM-108>

⁶ <http://jira.globus.org/browse/GRAM-106>

⁷ <http://jira.globus.org/browse/GRAM-105>

⁸ <http://jira.globus.org/browse/GRAM-104>

- [Bug 103:](#)⁹ Ease two phase end commit timeout
- [Bug 102:](#)¹⁰ Fix Two Phase Commit Semantics for Failed Jobs
- [Bug 100:](#)¹¹ one of the RSL parameters is not supported error doesn't indicate which it is
- [Bug 99:](#)¹² Add a high-level diagram for the approach doc
- [Bug 98:](#)¹³ Add Condor-G doc for using GRAM 2 and 5
- [Bug 96:](#)¹⁴ GRAM-106 SGE LRM mishandles invalid environment definition
- [Bug 95:](#)¹⁵ GRAM-106 SGE LRM doesn't check for executable permissions
- [Bug 94:](#)¹⁶ GRAM-106 SGE LRM doesn't check for executable existence
- [Bug 93:](#)¹⁷ GRAM-106 SGE LRM script doesn't handle environment vars with whitespace
- [Bug 92:](#)¹⁸ stdout to local file doesn't work if count >1
- [Bug 88:](#)¹⁹ Missed two phase commit causes job to not be destroyed
- [Bug 86:](#)²⁰ Prioritize script invocations to improve throughput
- [Bug 80:](#)²¹ GRAM 5 beta2 release
- [Bug 79:](#)²² Add support for OSG's "NFS Lite" concept
- [Bug 77:](#)²³ GRAM zombie
- [Bug 71:](#)²⁴ GRAM protocol test package contains expired test certificate
- [Bug 70:](#)²⁵ globus-job-status acts strange for completed jobs in GRAM5
- [Bug 69:](#)²⁶ globus-job-get-output -f doesn't work in GRAM5
- [Bug 68:](#)²⁷ Bad error when proxy is too short-lived
- [Bug 54:](#)²⁸ make globus-job-manager-event-generator not require configuration by default

⁹ <http://jira.globus.org/browse/GRAM-103>

¹⁰ <http://jira.globus.org/browse/GRAM-102>

¹¹ <http://jira.globus.org/browse/GRAM-100>

¹² <http://jira.globus.org/browse/GRAM-99>

¹³ <http://jira.globus.org/browse/GRAM-98>

¹⁴ <http://jira.globus.org/browse/GRAM-96>

¹⁵ <http://jira.globus.org/browse/GRAM-95>

¹⁶ <http://jira.globus.org/browse/GRAM-94>

¹⁷ <http://jira.globus.org/browse/GRAM-93>

¹⁸ <http://jira.globus.org/browse/GRAM-92>

¹⁹ <http://jira.globus.org/browse/GRAM-88>

²⁰ <http://jira.globus.org/browse/GRAM-86>

²¹ <http://jira.globus.org/browse/GRAM-80>

²² <http://jira.globus.org/browse/GRAM-79>

²³ <http://jira.globus.org/browse/GRAM-77>

²⁴ <http://jira.globus.org/browse/GRAM-71>

²⁵ <http://jira.globus.org/browse/GRAM-70>

²⁶ <http://jira.globus.org/browse/GRAM-69>

²⁷ <http://jira.globus.org/browse/GRAM-68>

²⁸ <http://jira.globus.org/browse/GRAM-54>

- [Bug 53](#):²⁹ Generalize log path configuration
- [Bug 51](#):³⁰ configurable control of number of perl scripts that can run simultaneously
- [Bug 47](#):³¹ simplify the throughput tester program and use improved version as doc
- [Bug 24](#):³² Debug/verbose flags for globusrun, globus-job-run
- [Bug 23](#):³³ Improved error codes and error reporting for users
- [Bug 22](#):³⁴ client connections can't be timed out
- [Bug 15](#):³⁵ transition from httpg to https
- [Bug 14](#):³⁶ increase availability of GRAM in linux distributions
- [Bug 12](#):³⁷ Gatekeeper's syslog output cannot be controlled
- [Bug 5](#):³⁸ Add gram-level prologue and epilogue script execution for mpi jobs
- [Bug 4](#):³⁹ Add support for a "managed fork" service
- [Bug 2](#):⁴⁰ Investigate how to setup GRAM5 services in a HA setup

6. Technology dependencies

GRAM depends on the following GT components:

- Globus Common
- GSI C
- GridFTP server

7. Tested platforms

Tested platforms for GRAM5:

- Linux
 - CentOS 5.3 x86_64
 - Debain 4.0 x86_64

²⁹ <http://jira.globus.org/browse/GRAM-53>

³⁰ <http://jira.globus.org/browse/GRAM-51>

³¹ <http://jira.globus.org/browse/GRAM-47>

³² <http://jira.globus.org/browse/GRAM-24>

³³ <http://jira.globus.org/browse/GRAM-23>

³⁴ <http://jira.globus.org/browse/GRAM-22>

³⁵ <http://jira.globus.org/browse/GRAM-15>

³⁶ <http://jira.globus.org/browse/GRAM-14>

³⁷ <http://jira.globus.org/browse/GRAM-12>

³⁸ <http://jira.globus.org/browse/GRAM-5>

³⁹ <http://jira.globus.org/browse/GRAM-4>

⁴⁰ <http://jira.globus.org/browse/GRAM-2>

- Mac OS X
 - Mac OS X 10.5.8

8. Backward compatibility summary

Protocol changes in GRAM since GT4.2.x series:

- The GRAM5 service uses a superset of the GRAM2 protocol for communication between the client and service. The extensions supported in GRAM5 are implemented in such a way that they are ignored by GRAM2 services or clients. These extensions provide improved error messages and version detection.
- GRAM5 does not support task coallocation using DUROC and its related protocols. Jobs submitted using DUROC directives will fail.
- GRAM5 does not support file streaming. The standard output and standard error streams are sent after the job completes instead of during execution.

9. Associated Standards

None

10. For More Information

See [GRAM5](#) for more information about this component.

Glossary

J

job scheduler

See the term [scheduler](#)¹.

¹ #scheduler