

Pegasus FAQ

This page collects frequently asked questions for Pegasus. The goal is to start adding questions here as they come up, and then when we have enough of them, we will add them to the Pegasus User Guide.

How do I implement a custom DAG scheduling algorithm in Pegasus?

Pegasus does not support custom scheduling algorithms in the traditional sense. It is possible to implement a custom site selector in Pegasus, but the scheduling of jobs at runtime is handled by DAGMan and Condor. Site selection allows Pegasus to choose which HPC center or grid site a particular job will be sent to, but it does not enable Pegasus to choose which machine(s) a job will run on. Sites are generally assumed to be geographically separated, and controlled by different administrative domains with different site-level schedulers and parallel file systems.

What does Globus Error 17 mean?

```
GRAM Job failed because the job failed when the job manager attempted to run it (error code 17)
```

This error occurs when GRAM attempts to submit the job to the local scheduler (PBS, SGE, SLURM, etc.). Often it is caused by a bug in the site-specific Perl module that is used by GRAM to interface with the local scheduler. In other cases this error might be caused by invalid job attributes, such as an incorrect queue or project ID. In order to debug this issue, users should log into the remote system and try to submit a job with the same requirements to get a more informative error message.

What does Globus Error 47 mean?

```
GRAM Job submission failed because the job manager failed to open stderr (error 74)
```

This error typically indicates that there is a network problem. The GRAM job manager on the remote system is trying to send back stderr for the job, but it is unable to connect to the client. Typically this is caused by a firewall on either the outgoing side of the server, or, more commonly, the incoming side of the client. To debug this issue the user should set GLOBUS_TCP_PORT_RANGE, check their firewall, and try running a job using the command-line globus-job-run tool.

What does this Globus GRAM error mean?

Open Science Grid has some useful documentation for debugging Globus GRAM errors: <https://wiki.grid.iu.edu/bin/view/Documentation/Release3/TroubleshootingGuide>