

Open Science Grid Dependence on Globus/CDIGS components, related to the supplement request for Globus support

Open Science Grid Executive Team, December 10th 2010

DRAFT

The Open Science Grid depends on a sustained, supported set of middleware technologies that allow its users to make productive use of the distributed computing infrastructure. The Globus software components in the November 2010 supplement request (V4) are included in this middleware:

Component used in OSG	Description
Grid Security Infrastructure (GSI)	APIs and tools for authentication, authorization and certificate management
GRAM. Only GRAM2. No date set to deploy GRAM5.	Technologies which enable users to locate, submit, monitor and cancel remote jobs on grid-based compute resources
GridFTP	High-performance, secure, reliable data transfer technologies optimized for high-bandwidth wide-area networks
jGlobus	Java libraries for interfacing with Globus GSI, GRAM and GridFTP.

Of the components listed above, OSG deploys the latest versions of the software in a timely fashion with the exception of the GRAM component. OSG continues to use and support GRAM2. The major European Infrastructure EGI (previously EGEE) continues to use GRAM2. Part of the reason OSG has been able to continue to use GRAM2 is due to the increased stakeholder adoption of Condor based “pilot-job” technologies. These technologies reduce the load on the GRAM service thus allowing GRAM2 to be used for much higher job throughputs.

The LHC experiments and EGI in Europe are now deploying CREAM, a new European developed remote job management component. The US LHC collaborations have requested OSG evaluate CREAM and GRAM5. To date, there is no decision to deploy either one into production on OSG.

Support:

The OSG VDT team provides a first level of support and problem triage for OSG users of the OSG software stack, including the Globus components. In the earlier years there were a significant number of support and enhancement tickets placed to the Globus team. In the past two years, we have seen a reduction in the number of such tickets:

2009: 5 Bugzilla tickets submitted by OSG;
2010: 2 Bugzilla tickets + 1 Jira ticket submitted by OSG.

There is continuing support for the OSG communities outside of this list. For example, the bugzilla report for LIGO¹ is 2009: 5 and 2010: 1 ticket submitted.

Below is a more detailed report on the outcome of the OSG tickets submitted:

ID	Problem Report	Creation & if Resolved Dates	Status	Comments
6688	Need fix to LSF job manager to use bacct	2009-03-11	RESOLVED FIXED	Resolved with help from OSG site manager with LSF.

¹ http://bugzilla.globus.org/bugzilla/buglist.cgi?content=LIGO&field-1-0-0=product&field-1-1-0=content&product=&query_format=specific&reanction=&type-1-0-0=anyexact&type-1-1-0=matches&value-1-0-0=&value-1-1-0=LIGO&query_based_on=&columnlist=opendate%2Cchangeddate%2Cbug_status%2Cresolution%2Ccomponent%2Cshort_desc

ID	Problem Report	Creation & if Resolved Dates	Status	Comments
	when bhist doesn't work	2009-05-26		
1934	Gatekeeper's syslog output cannot be controlled.	2004-09-22	NEW	Updated request in 2009 to point to GRAM5. "Unscheduled at present"
6840	The PBS job manager doesn't handle large environments well	2009-08-19	NEW	No information since creation date.
6874	Need help with stat not failing properly in a GridFTP plugin	2009-10-13	NEW	No information posted since 2009-10-14.
6883	Cog Jglobus code fails to process RFC 3820 compliant proxy created with voms-proxy-init tool	2009-11-06 2009-12-09	RESOLVED FIXED	Removed the check that imposes the key usage. Fix committed to globus_4_0_branch and trunk. Will be part of CoG JGlobus 1.8 release
6977	Bad name for syslog in gatekeeper	2010-03-08 2010-03-09	RESOLVED FIXED	EGEE reported a bug.... It's a trivial Fix --just a single character.
5820	Improve Condor Logfile Processing in GRAM	2008-01-30 2010-07-06	RESOLVED FIXED	GRAM5 in 5.0.2 will process each job in a separate log file, see http://jira.globus.org/browse/GRAM-130 for details.
GRAM-189	PBS jobmanager needs to check for additional exit code	30/Sep/10 16/Nov/10	Resolved:	We (the VDT) got a bug report about the PBS jobmanager. It boils down to one line in pbs.in that says: if(\$exit_code == 153) It should apparently say: if(\$exit_code == 153 \$exit_code == 35)

We are asked to estimate the effort needed within OSG to take on support for the Globus components for the OSG stakeholders:

Component	Initial Resources to learn the code	Resources for Continuing support	Use
GSI	3 FTE months	1 FTE month/year	Basic security library for all OSG software
GRAM2	4 FTE months	2 FTE months/year	Client-Server that enables jobs submitted remotely from a pilot-factory or user client interface to be executed at a remote site.

GridFTP	8 FTE months	?3 FTE months/year ²	Client used for transport of all files transferred between nodes in OSG. Service manages transfer of files between sites, between nodes within a site when using BestMan or Hadoop storage implementations.
jGlobus	3 FTE months	2 FTE month/year.	Java Globus interface library for use by java based software in the OSG software stack.

Enhancements:

Of the enhancements in the extension proposal, several have been in the queue for more than a year. In some cases OSG has workarounds deployed in the field that have proven adequate for our current and projected needs. Other comments on the proposed enhancements:

GridFTP: All OSG stakeholders depend on the continued support and maintenance of GridFTP.

GRAM: There have been no enhancements to GRAM2 for several years. They are incorporated into GRAM5 and our understanding is that all enhancements are targeted for that version of the software.

GSI: GSI is a well-used and stable software component. It is unlikely that any major enhancements could be adopted in a short time frame.

Native Packaging: At the request of our stakeholders, especially LIGO, OSG is transitioning the packaging of the VDT software components from the existing Pacman (initially developed by the US ATLAS experiment) to “native packaging” (e.g. RPMs and Debian packages) used by the major LINUX distributions and most large open source software projects. We have released the Globus components to LIGO as native packages already.

Error Reporting: Globus has improved the error reporting in the past couple of years. Consistent and complete error reporting and easy diagnosis remains a challenge for all software in the complex, and non-deterministic, end-to-end distributed system environments. This is a difficult and never-ending task.

Documentation and Training: OSG has been pleased with the level and quality of the Globus documentation over the past several years. We are not aware of our stakeholders making use of any of the GLOBUS education and training activities.

Finally we list the enhancements proposed, giving their priority for OSG stakeholders:

GRAM		
Proposed Enhancement	OSG Priority and Comments	Target Communities
Provide a mechanism for preventing overload when GRAM is used as a fork job manager.	Not Needed. This request for GRAM2 has been in the queue for several years. (This has been discussed in person previously and documented in the 2009 document from OSG.) OSG currently works around this with the managed fork software and would not regard this as needed at this time.	<ul style="list-style-type: none"> • OSG • IGE
Allow clients to supply GRAM with	Not Needed.	<ul style="list-style-type: none"> • OSG

² Will be able to give an estimate for the out-years following initial learning of the code and needs.

tracking information for more detailed logging and usage statistics.	Long term request in the queue. Patch to globus applied by VDT. ~1 month of development. Would not expect this to get done for GRAM2.	<ul style="list-style-type: none"> • TeraGrid • IGE
Support load balancing and failover in GRAM.	No Longer Needed. ATLAS is requesting the evaluation of GRAM5 and CREAM to obviate this need. This is a deeply non-trivial implementation. In Condor failover took 1-2 FTE years of effort	<ul style="list-style-type: none"> • FNAL • ATLAS • CMS
Support commonly used environment management tools (e.g. modules, VDT).	Not Needed. For past 2 years VDT patches Globus to set environment variables for jobs. Would be few weeks work to implement solution.	<ul style="list-style-type: none"> • TeraGrid • OSG
Support third party file transfers for stdin and stdout to and from arbitrary GridFTP resources	Priority not known at this time. Increases data staging flexibility. Enables clients to move stdin and stdout data without requiring GASS.	<ul style="list-style-type: none"> • LIGO
GridFTP		
Proposed Enhancement	OSG Priority and Comments	Target Communities
Enable GridFTP to perform on-the-fly message digest computation.	Medium. Allows communities to detect any inadvertent content changes and to vouch for the integrity of data and results.	<ul style="list-style-type: none"> • ESG • BIRN • OSG
Enable configuration of the GridFTP server to support a maximum bandwidth policy against a user-defined ³ set of endpoints.	Medium. Improves administrator control over network resources and allocation of GridFTP server resources to their users.	<ul style="list-style-type: none"> • LIGO, APS • TeraGrid • OSG • IGE
Enable a configurable timeout in GridFTP servers that allows fread() to abort a transfer and notify client.	Medium Improves quality of service by allowing a transfer to fail gracefully in the event that it cannot be completed. LIGO uses some large tape backends and it can sometimes take a long time to load a tape. LIGO is using tape less and for advanced LIGO expect all data to spin and tape will only be used for true archival and backup and not access by users and codes. If this doesn't appear in a year or less will probably not actually use it.	<ul style="list-style-type: none"> • LIGO • Communities using HPSS
Security (GSI)		
Proposed Enhancement	OSG Priority and Comments	Target Communities

³ Administrator-defined

Reconcile the Java and C GSI behavior in managing expired Certificate Revocation Lists (CRLs).	Not Needed. OSG works round this and will “not care”. Discussed for a couple of years in conversation.	• OSG
Provide alternative to httpg delegation that uses standard SSL	Low. Was agreed to for FY10 delivery for dCache. Was delayed, now unknown when will be needed. Waiting for dCache SRM task force to reconvene.	• OSG • dCache • BestMan
Package, distribute, and support the jGlobus Java GSI implementation as an independent library.	Not Needed. Managing to work around now. May be obviated by rewrite of authorization component by dCache team.	• OSG • FermiGrid • LBNL
Support more sophisticated authorization callouts (e.g., LCAS, LCMAP) for C toolkit services.	Medium OSG is working round this and continuing to support existing VDT module PRIMA until it is available or will do the coding ourselves to enable use of new authorization components from Europe.	• OSG/VDT • TeraGrid • IGE
Performance and Usability		
Provide improved diagnostic interfaces for all services by returning detailed error descriptions in addition to basic exit codes.	Low. We have been coping to date. Everyone hates the situation and would like it to be better. This is a lot of work. Between 1month and 1 year of work. Would be useful to understand the scope.	• All
LIGO request: globus-gridftp-server natively packaged on Solaris.	Medium. Note that it is only globus-gridftp-server and things it depends on...not the rest of GT. Only globus-gridftp-server used on Solaris.	• LIGO
Release rest of Globus components as independent products using native packaging.	Not Needed. VDT already releases Globus components in native packaging. If Globus does it then VDT needs to coordinate with them. LIGO currently takes Globus distribution from VDT.	• LIGO • OSG • IGE