Gridification Task Overlaps

Subject: **Functionality overlaps between existing Globus components and the subsystems of the WP4 Gridification task**

Author: **D. Groep**
Partner: **NIKHEF**

Diffusion: **unlimited**
Information:

1. INTRODUCTION

The gridification subsystem (GS) connects the local fabric to the grid, by providing the interface for grid-wide services to access the local resources and by publishing information about these resources to external parties. The GS is a complete architecture of the components on this interface, and includes all components necessary to make a fabric accessible to external parties via the appropriate protocols.

This area has traditionally also been the focus of the Globus Toolkit [Glo2001], and therefore a complete design will have area's over overlap with existing and planned components of the Globus Toolkit. In this document, an inventory of these overlap areas is made, starting from the subsystem decomposition defined in the Architecture Document [Can2001]. The relative organization of these subsystem is depicted in Fig. 1.

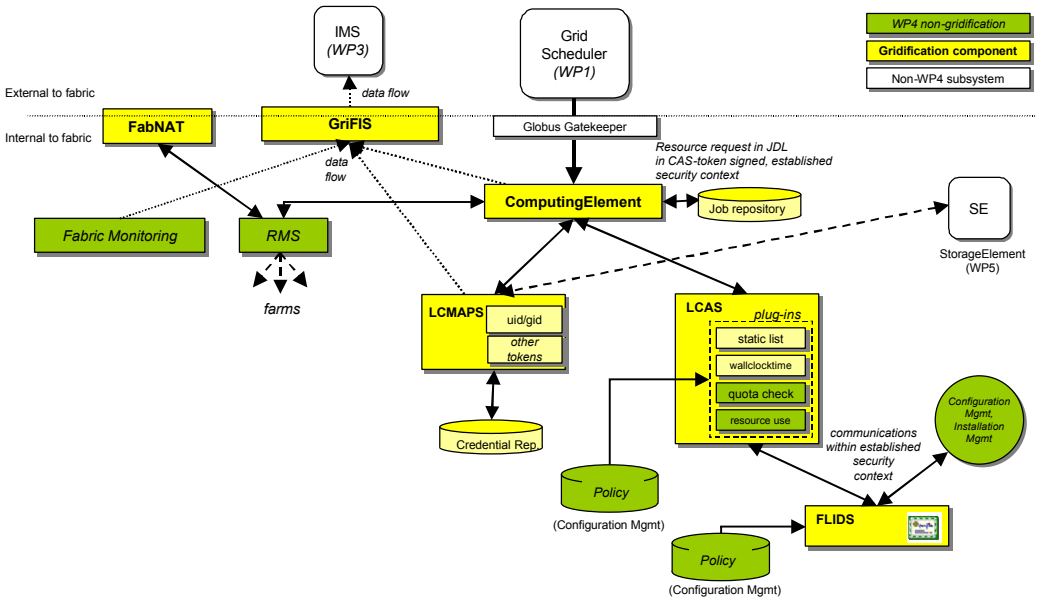

Figure 1: gridification components in WP4 and their relationships.

2. GATEKEEPER AND COMPUTE ELEMENT

The protocol for resource access currently implemented in the Globus Toolkit is GRAM version 2 [Cza1998]. This protocol is based on GSI-secured HTTP, exchanging attribute-value pairs describing the jobs to a resource and returning job status from the resources. The protocol driver is implemented in the Globus gatekeeper. Additions to the protocol added as part of the Globus 2.0 release implement robust job submission using code from the UW Madison Condor group[1].

The GS gatekeeper component can be identified with the Globus gatekeeper. Part of the Globus gatekeeper is a protocol driver for the GRAM protocol. This part will be used almost as-is: it act as the authentication agent and basic protocol driver for the resource access mechanism.

What may change is the authorization component of the underlying GSI, in accordance with new developments in this area. Developments include those coordinated by the EDG Security Group, as well as new GGF efforts on credential delegation and concepts from a generic AAA architecture [RFC2903]. The current authorization model cannot take `dynamic' authorization additions into account, since at this point in the protocol negotiations, the job request data is not yet available.

It may therefore be that authorization will be delegated to dedicated subsystems by the gatekeeper. In particular, a mapping to a local user ID should be postponed until an authorization decision is made; generating or assigning local credentials may be an expensive process (contrary to the Globus 1.1.3 situation where the mapping is only to static UNIX uid/gid pairs). Giving up root privileges at this point would be too early.

The current Globus Gatekeeper will subsequently start a job manager (defined in the contact string and specified as the path to the initial POST request to the gatekeeper). The current Globus job manager will evaluate the RSL, connect the input and output channels to a remote GASS server and start the job using a local submission system (batch queue or fork). The job manager lives as long as the job, and establishes a contact point (secure HTTP listener) for job control.

This component thus offers a subset of the functionality of the GS/CE component. In order to make a fully functional GS/CE component, it is expected that the current job manager (being a backend for the gatekeeper) has to be extended. Current job manager code (jobmanager or jobmanager-condor) can be used as the interface to the RMS.

The GS/CE components that effectively make up an AAA server will have to be designed from scratch (perusing concepts from GGF and IRTF workgroups). These components will be closely integrated with the new-style gatekeeper.

3. LCAS

There is currently no component in the Globus Toolkit that covers the functionality of the LCAS subsystem. Although one might suspect the static-user-list plug-in to be equivalent to the map file used by the Globus gatekeeper in the current Globus versions, the specification this component requires it to give a yes-or-no answer. The current GSI map file approach combines the static user authorization with a local credential mapping (uid/gid). This functionality is split in the GS design to enable a modular AAA approach.

Other authorization components are either absent or currently implemented only in the local batch systems. Implementing authorization that late in the job submission process would make a distributed decision AAA impossible.

4. LCMAPS

The current GSI grid map file approach provides a static implementation of the LCMAPS component, although not at the right stage of the authorization process. Some of the dynamic functionality foreseen for LCMAPS (dynamic account creation or leasing) has been implemented in the grid-mapdir-patch [Nab2001]. This approach has proven useful in actual test bed deployments in the UK and elsewhere, but being tied closely to the current GSI authorization mechanism, the functionality of this patch will have to be re-implemented as part of the LCMAPS system.

Acquisition of other user tokens is partly available in the current GSI. Kerberos tokens are currently obtained through a wrapper program (globus_gram_k5), that will get the relevant Krb5 credentials via either kinit or sslk5. This method can be re-used in the LCMAPS component, but the interface used to obtain these credentials should be identical to the one used for obtaining uid/gid mappings. Some re-implementing may therefore be needed.

5. FLIDS

No equivalent functionality exists in the Globus Toolkit. This is due to the fact that Globus is not concerned with local fabric management (systems install and configuration) and that the current test beds in which Globus is deployed are sufficiently well controlled that "hiding" the complex world of authentication and authorization is not yet needed. However, the mechanisms used by FLIDS are those currently used by GSI (i.e. X.509 and TLS).

6. GRIFIS

The GriFIS component consists of two complementary functions: information providers (also called `producers') and a publishing framework. The first function is partly covered by the current Globus information providers, but in the future much more information, both from the RMS and from the Configuration Database needs to be published. This part of the GriFIS is hardly covered by Globus.

The latter GriFIS component (the publishing framework), is currently available as a part of Globus (the MDS), but will in the future be part of the WP3 information services framework (G-RMA) [Fis2001] or be replaced by new development in either Globus or the GGF (RMA). The role of the Gridification Task in the development of this framework is expected to be largely in the requirements area. The information providers will then have to adapt to this framework.

7. FABNAT

This component is entirely new. The functionality provided by this component has not yet been considered as part of the Grid middleware and is completely absent in the Globus Toolkit. The developments will, however, be heavily influences by development in the networking area (IPv6, IPsec, masquerading/forwarding, QoS) and in the AAA area (end-to-end bandwidth brokering, tunnelling).

In zero-st order, the functionality of FabNAT can be covered with the null component (for existing transparent networks) or with masquerading forwarders (for private-to-public network connectivity).

8. REFERENCES

[Can2001] G. Cancio et al., *Architectural design and evaluation criteria WP4-Fabric Management*, DataGrid-04-D4.2- 0119-2_1

[Cza1998] K. Czajkowski et al., *A Resource Management Architecture for Metacomputing Systems*, Proc. IPPS/SPDP '98 Workshop on Job Scheduling Strategies for Parallel Processing, pg. 62-82, 1998 [Architectural overview, the protocol is not yet publicly documented]

[Fis2001] S. Fisher et al., *WP3 GMA releases*, http://hepunx.rl.ac.uk/grid/wp3/releases.html

[Glo2001] Globus Toolkit, The Globus Project, 2001.

[Nab2001] A. McNab, *gridmapdir patch for Globus 1.1.3*, http://www.gridpp.ac.uk/gridmapdir/, Manchaster University 2001.

[RFC2903] C. de Laat et al., *Generic AAA Architecture*, RFC 2903, ISOC, 2000.

[1] see announcement on the Globus 2.0 beta release at <http://www.globus.org/gt2/release-notes.html#gram>