



The gLite File Transfer Service

Thursday 2 March 2006 16:25 (20 minutes)

In this paper we describe the architecture and implementation of the gLite File Transfer Service (FTS) and list the most basic deployment scenarios. The

FTS is addressing the need to manage massive wide-area data transfers on dedicated network channels while allowing the involved sites and users to manage their policies. The FTS manages the transfers in a robust way, allowing for an optimized high throughput between storage systems.

The FTS can be used to perform the LHC Tier-0 to Tier-1 data transfer as well

as the Tier-1 to Tier-2 data distribution and collection. The storage system

peculiarities can be taken into account by fine-tuning the parameters of the

FTS managing a particular channel. All the manageability related features as

well as the interaction with other components that form part of the overall service are described as well. The FTS is also extensible so that particular

user groups or experiment frameworks can customize its behavior both for pre- and post-transfer tasks.

The FTS has been designed based on the experience gathered from the Radiant service used in Service Challenge 2, as well as the CMS Phedex transfer service. The first implementation of the FTS was put to use in the beginning

of the Summer 2005. We report in detail on the features that have been requested following this initial usage and the needs that the new features address. Most of these have already been implemented or are in the process of

being finalized. There has been a need to improve the manageability aspect of the service in terms of supporting site and VO policies.

Due to different implementations of specific Storage systems, the choice between 3rd party gsiftp transfers and SRM-copy transfers is nontrivial and was requested as a configurable option for selected transfer channels.

The way

the proxy certificates are being delegated to the service and are used to perform the transfer, as well as how proxy renewal is done has been completely

reworked based on experience. A new interface has been added to enable administrators to perform management directly by contacting the FTS, without

the need to restart the service. Another new interface has been added in order to deliver statistics and reports to the sites and VOs interested in useful monitoring information. This is also presented through a web interface using javascript. Stage pool handling for the FTS is being added in order to allow pre-staging of sources without blocking transfer slots on the source and also to allow the implementation of back-off strategies in case the remote staging areas start to fill up.

The reliable transport of data is one of the cornerstones for distributed systems. The transport mechanisms have to be scalable and efficient, making optimal usage of the available network and storage bandwidth. In production grids the most important requirement is robustness, meaning that the service needs to be run over extended periods of time with little supervision. Moreover, the transfer middleware has to be able to apply policies for failure, adapting parameters dynamically or raising alerts where necessary. In large Grids, we have the additional complication of having to support multiple administrative domains while enforcing local site policies. At the same time, the Grid application needs to be given uniform interface semantics independent of site-local policies.

There are several file transfer mechanisms in use today in Data Grids, like http(s), (s)ftp, scp or bbftp, but probably the most commonly used one is GridFTP, providing a highly performant secure transfer service. The Storage Resource Manager SRM interface, which is being standardized through the Global Grid Forum, provides a common way to interact with a Storage Element, as well as a data movement facility, called SRM copy, which in most implementations will again make use of GridFTP to perform the transfer on the user's behalf between two sites.

The File Transfer Service is the low level point to point file movement service provided by the EU-funded Enabling Grids for E-Science (EGEE) project's gLite middleware. It has been designed in order to address the challenging requirements of a reliable file transfer service in production Grid environments. What distinguishes the FTS from other reliable transfer services is its design for policy management. The FTS can also act as the resource manager's policy enforcement tool for a dedicated network link between two sites as it is capable of managing the policies of the resource owner as well as of the users (the VOs). The FTS has dedicated interfaces to manage these policies. The FTS is also extensible; upon certain events user-definable functions can be executed. The VOs may make use of this extensibility point to call upon other services when transfers complete (e.g. register replicas in catalogs) or to change the policies for certain error handling operations (e.g. the retry strategy).

The LHC Computing Project (LCG) is the project that has built and maintains a data storage and analysis infrastructure for the entire high energy physics community of the Large Hadron Collider (LHC), the largest scientific instrument on the planet located at CERN. The data from the LHC experiments will be distributed around the globe, according to a multi-tiered model,

where

CERN is the “Tier-0”, the centre of LCG.

The goal of LCG Service Challenges is to provide a production quality environment where services are run for long periods with 24/7 operational support. These services include the Network and Reliable File Transfer services. In Summer 2005 Service Challenge 3 started with gLite File Transfer

Service and CMS Phedex. The gLite FTS benefited from this collaboration and from the experience of prototype LCG Radiant Service, used in Service Challenge 2. This meant that from the beginning its design took into account

all the requirements imposed by a production Grid infrastructure. The continuous interaction with the experiments was useful in order to react quickly to reported problems, as well as to keep the development focused on real use cases.

Summary

The gLite File Transfer Service

Primary author: MCCANCE, Gavin (CERN)

Presenter: Mr BADINO, Paolo (CERN)

Session Classification: 2b: Data access on the grid

Track Classification: Data access on the grid