



Contribution ID: 76

Type: **Oral Presentation**

A new scheme for efficient distribution of information in Grid infrastructures

Tuesday, 12 April 2011 16:30 (30 minutes)

Conclusions

In this work we present the architecture of the system and show the quantitative results to support its advantages.

Our system design and implementation are sufficiently generic, thus allowing applications in other areas, where the distribution of information follows the same patterns.

Overview

The gLite grid middleware utilizes the so-called BDII as an information system in a layered approach. In our work we developed an alternative path for distribution of the information, which enables considerable decrease in the amount of data that is exchanged. We use the fact that much of the information that is exchanged does not change over time by introducing a versioning scheme. In this way most of the traffic consists of data in a compressed diff-like format.

Impact

The new scheme can be deployed in combination with the current BDII solution. It reduces drastically (by orders of magnitude) the amount of network bandwidth being used, which will be important to accommodate for further increases in the number of sites in the infrastructure and the complexity of information about them. It also decreases the latency in the movement of data between the layers.

Description of the work

We implemented an alternative path for distribution of the information that is currently available via the gLite BDII service. The new lightweight middleware components are based on the well known properties of the information that is currently transferred between the various layers in the production Grids based on gLite and follow a simple protocol in order to ensure consistency.

Since the information that a site-level information system displays changes at once, the problem is how to distribute in an efficient way the changes between successive “versions” towards the nodes that require them (mostly top-level BDIIs). Since the changes are usually limited to certain attributes and the appearance or disappearance of whole sections of information is related to rare events, like failing or new nodes, we use the routines from the “diffutils” package to compute the difference between successive “versions” of the information that a site displays. Moreover, we compress the result, which decreases even further the amount of information that needs to be transferred. A client, which can be co-located with a top-level BDII, can query our service over the http protocol, supplying a hash of the information that it has at a given point. If this information is sufficiently recent, within the last 1-2 minutes, then only the gzipped difference is transferred.

The client proceeds to uncompress the patch and applies it to its copy of the information, obtaining the same result as when all the information is downloaded. We have implemented a testbed deployment of the new system, where we mimic the infosystems of the European Grid Initiative sites in real time, which enabled us to perform a series of benchmarks and to analyse the performance of the new scheme.

Primary author: Dr ATANASSOV, Emanouil (IICT-BAS)

Co-authors: Mr DIMITROV, Dimitar (IICT-BAS); Dr GUROV, Todor (IICT-BAS)

Presenter: Dr ATANASSOV, Emanouil (IICT-BAS)

Session Classification: Technologies for Distributed Computing

Track Classification: DCI - Implementation