Don’t you need a fast information service?

iGrid is a novel Grid Information Service developed within the European GridLab project by the ISUFI Center for Advanced Computational Technologies (CACT) of the University of Lecce, Italy.

Among iGrid requirements there are performance, scalability, security, decentralized control, support for dynamic data and the possibility to handle user supplied information.
The iGrid Information Service has been specified and carefully designed to meet these requirements.

Core iGrid features:

- Web Service interface,
- Distributed architecture,
- Based on the PostgreSQL relational DBMS,
- SQL SELECT statement parser,
- Fault tolerance,
- Support for Globus GSI,
- Support for TLS binding to DBMS,
- Support for GridLab Authorization Service (GAS),
- Support for Administrator defined ACL,
- Support for GridLab Mercury logging service,
- Builds on Linux, Mac Os X, tru64 and irix,
- Extensible,
- Extreme performances.

The iGrid distributed architecture is based on two kind of nodes, the iServe and the iStore GSI enabled web services. The iServe collects information related to a specific computational resource, while the iStore gathers information coming from registered iServes. The Information Service is based on a relational DBMS back-end and can handle information extracted from computational resources and also user/service supplied information.

The current architecture allows iStores to register themselves to other iStores, thus creating distributed hierarchies. We are also investigating a possible implementation of a peer to peer overlay network based on one of the current state of the art distributed hash table algorithms in order to improve iGrid scalability. The implementation includes system information providers outputting XML, while user/service information is directly supplied by the user simply calling a web service registration method.
The Web service interface is based on the gSOAP toolkit and on two packages developed within the ISUFI/CACT, the GSI plugin for gSOAP and the GrelC library.

The performances of iGrid are extremely good, as reported in “iGrid, a Novel Grid Information Service”, by Giovanni Aloisio, Massimo Cafaro, Italo Epicoco, Sandro Fiore, Daniele Lezzi, Maria Mirto and Silvia Mocavero, to appear in Proceedings of the first European Grid Conference, Lecture Notes in Computer Science, Springer-Verlag, 2005

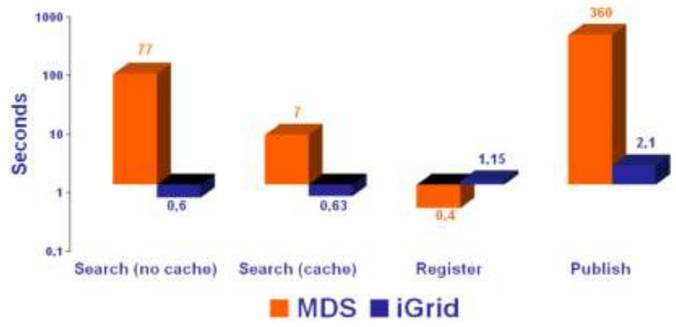

Figure 1. ***The performance of MDS and iGrid***

The performances of the current release are extremely promising, when compared with the GridLab MDS (based on Globus Toolkit MDS v2.4) that has been developed by WorkPackage 10 before iGrid. In what follows, we report about a performance comparison between iGrid and the previous GridLab MDS for a typical query. The clients were in Lecce, Italy, the iGrid and MDS server in Brno,Czech Republic (see more details on Figure 1).
The MDS was queried in order to assess its performance for both cached information and non cached information. iGrid does not need to cache information, since it uses a push mechanism. The Globus MDS instead uses a pull approach, meaning that when information is not available in the GIIS cache, it is pulled from remote GRIS servers. The graph below shows that when information is not cached, the MDS GIIS takes about 77 seconds to return the information requested, while, when the information is found into the cache, this time reduces to 7 seconds. The same query to iGrid returns the information in just half a second. In order to register user supplied information, the MDS enhanced by our WorkPackage to support this functionality (the Globus MDS only provides default information providers for extracting system information)takes about half a second, however this is the time needed to insert the information in a GRIS server. The time for this information to propagate from GRIS to GIIS is about 360 seconds. iGrid needs about one seconds to insert the information on an iServe, but then the time needed to publish the information in the iStore is only about 2 seconds. The following graph summarizes the discussion. On the y axis we used a logarithmic scale in seconds in order to appreciate the performances of iGrid vs the previous GridLab MDS.

iGrid uses a push model for data exchange: information extracted from resources is stored on the local DBMS, and sent periodically to registered iStores, while user supplied information is immediately stored on local DBMS and sent to registered iStores. Thus, an iStore has always fresh, updated information, and does not need to ask iServes for information. Each information is tagged with a time to live that allows iGrid to safely removes stale information from the DBMS as needed. Indeed, on each user lookup data clean-up is performed before returning to the client the information requested. When iGrid starts, the entire DBMS is cleaned up. Thus the user will never see stale information.

The iGrid Information Service provides fault tolerance:
in case of failure of an iStore, iServes remove temporarily the faulty iStore from their registration list. Periodically, the iStore list is updated by adding previously removed iStores when iStores are available again. In this case, the local DBMS is dumped and immediately sent to newly added iStores.

The iGrid Information Service is used by the recently established SPACI consortium, that is building a grid among ISUFI/CACT, University of Naples, University of Calabria and Hewlett-Packard.

For more information please visit the webpage:

www.gridlab.org/igrid