Grid4All
democratize the grid

NEWSLetter

issue n°2, February 2008

Events

UPRC took part on the *European Services, Software and Grid Technology Days FP6 Projects Concertation Event* celebrated on Wednesday 26th & Thursday 27th September 2007 at the Diamant Centre, Bd. A. Reyerslaan 80, 1030 Brussels, Belgium. - Day 1 - Parallel Sessions Track 1: Ontologies

UPC took part on the *European Services, Software and Grid Technology Days FP6 Projects Concertation Event* celebrated on Wednesday 26th & Thursday 27th September 2007 at the Diamant Centre, Bd. A. Reyerslaan 80, 1030 Brussels, Belgium. - Day 1 - Parallel Sessions Track: Distributed Collaborations & Workflows

THE GRID4ALL PROJECT: One year investment

The first year has mainly focused on gathering use scenarios, defining requirements, analysing the state of the art and defining the global architecture. From a DHT-based overlay called DKS (for Distributed K-ary System), brought in as background, SICS has released a prototype of new self-managing overlay services implementing multicast primitives. INRIA has integrated its component based management architecture on the overlay and this framework provides three key Grid-management functionalities: deployment, monitoring and discovery. These services are built taking into account churn and volatility.

A first prototype of the management framework for virtual organisations (VO) that can deploy and manage application components has been released. A VO is mapped on an overlay. Application components are described using an extended version of the Fractal Architecture Description Language. These extensions facilitate design of self-managing applications. This framework will be improved in order to provide policy-based management, enabling a rule-based decision control without requiring development of new code.

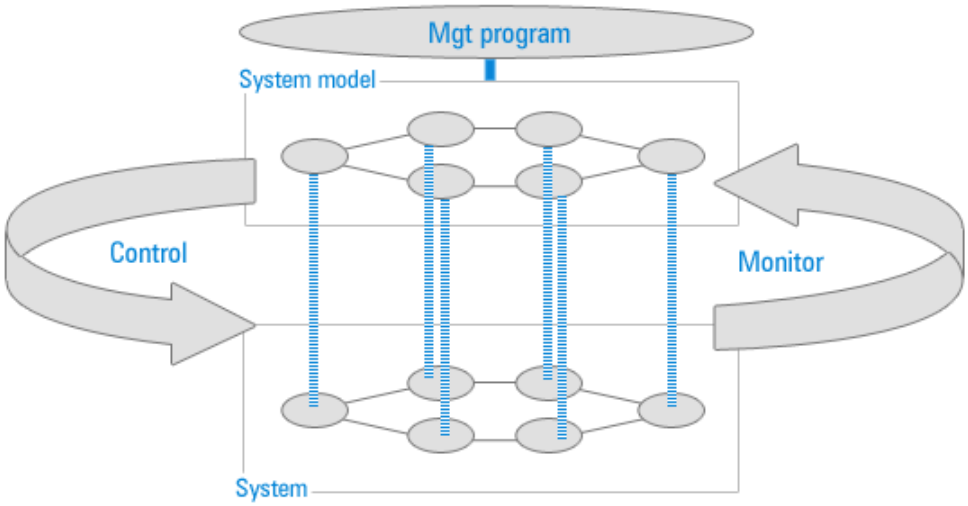

Figure 1: Feedback-based management

To address collaboration, the project has released a prototype of a VO-Aware distributed file system (DFS) that aggregates different types of storage providers through a common layer, the "Virtual Block Storage". This DFS implements a hierarchical and decentralized meta-data (directories and file systems) database providing information on the file objects that have been created or exported. It is built as a layer linking the file namespaces pooled into the virtual community. The advantages of this architecture are (a) separation between meta-data, access control and the consistent management of data itself, and (b) application-specific data-coherency control even in the face of volatility and disconnected operations.

P2P collaborative work is addressed through a novel middleware that handles replication, consistency and conflicts. This takes application semantics into account and leverages on DFS for persistency.

TECHNICAL RESULTS

Semantic Information Service

We have designed and developed ontology for the representation of traded resources in e-marketplaces. The ontology has been decided in a way that orders (offers or requests) can be automatically matched using OWL (Web Ontology Language) subsumption mechanism: offers of traded resource configurations matched with specific requests placed in specific e-markets.

A Semantic Information System (SIS) designed and implemented, as a component of the Grid4All Market-Based Resource Brokering System (RBS). This provides a registry of the e-markets available, together with a retrieval and ranking service for those markets. Peers may query the SIS for orders (i.e. requests and offers) that match certain attributes and criteria, locating orders that concern matching resources, traded in forthcoming (initiated but not yet started) or ongoing markets. Results of queries shall be ranked according to the preferences and intentions of providers and consumers, as well as according to the characteristics of resources and markets. SIS acts as a services' registry for the following types of services: Markets, Application Services and Services exposing (traded) resources. The SIS may

be queried by software agents as well as by human users to select advertised services and resources: Matchmaking happens through different criteria concerning resources and other application specific traded domain entities, as well as services' profiles. The returned query results are ranked according to resources/services matching characteristics and providers'/consumers' features by exploiting the Grid4All resources ontology, as well as OWL-S services profile specifications.

SIS functionality may be accessed by two types of interfaces: A user interface, accessible through a web-based portal and an API. Though the portal, users can submit advertisements and queries. The web forms for both advertisements and queries are automatically created based on the underlying ontology schema. An Application Programming Interface has been also designed for accessing the SIS. This API will be available as a Web Service. The API will be available as a WSDL description as well as in the form of a set of Java stub classes for facilitating web service calls. The SIS API provides a set of operations for performing advertisements and queries of services. It also provides a set of data types which define an abstraction over the ontology schema so that programmers need not be aware of the ontology in order to publish advertisements of available services. The API will provide the functionality mentioned above, namely advertisement and querying of services.

Grid4All Programming Platform

WP1 produced the first version of the Grid4All programming platform together with a demo application that demonstrates it. The Java-based platform allows to design distributed applications consisting of the component-based functional part, and the self-management part containing objects that asynchronously communicate with each other and the platform. The latter can subscribe to resource monitoring (sensing) and component and node failure events, and can use the API provided by the platform to repair or change the architecture accordingly. The demo application is a functionally simple distributed file storage service that at the moment can repair itself, but soon will also perform self-configuration and self-tuning using the same technical framework. The platform itself provides at the moment for designing and testing applications, and later will provide for reliable execution of self-management code using overlay services.

Middleware for collaborative applications

This middleware addresses collaboration. In order to collaborate, users must be able to write to the same data, be it a file, a database, or more generally a set of documents. Within Grid4All, we have designed and developed a virtual organisation aware file-system, the VOFS, which however provides only best-effort consistency.

In a collaborative application e.g. a shared calendar, the actions of different users may conflict. Understanding and resolving these conflicts requires taking application semantics into account. The Semantic Store of the Grid4All architecture, called Telex, is a common middleware enabling collaborative work modes for any application. The middleware enables remote sharing of documents and remote collaboration. It facilitates the design of collaborative applications by taking care of complex application-independent aspects such as replication, conflict repair and ensures eventual commitment. Telex allows an application to access a local replica without synchronizing with peer sites. The application makes progress, executing uncommitted operations, even while peers are disconnected. Telex supports the whole spectrum of collaboration modes, from on-line sharing (``what you see is what I see'') off-line, check-out/check-in isolated work. Telex integrates seamlessly into Grid4All's VOFS. A collaborative document is persistent, can be named just like ordinary files, and is accessible by all members of any VO it is exposed to. Access to collaborative documents is subject to access control through the security infrastructure.

We have demonstrated the middleware by implementing a shared calendar using this middleware. The SC is a simple prototype of a joint decision-making application. Authorised users manage their calendars collaboratively. One user may create and manage a meeting, e.g., changing the time or the list of attendees; this may cause conflicts with other meetings, which are detected by Telex; Telex proposes possible resolutions and users vote for their preference; finally Telex identifies and commits the winning resolution.

Collaboration

The consortium finds essential the exchanging within other investment projects and/or institutions in order to share knowledge and results and to collect perception from end-users to achieve objectives of Grid4All Project. Partners are carrying on collaborative actions which are briefly detailed above.

France Telecom is collaborating with SORMA (FP6 Project) in the design and development of a combinatorial auction model to trade CPU and storage resources and the evaluation of implementation based on CPLEX. It is also participating within collaborative task group TG7 and TG8 that took place during the European Services, Software and Grid Technology Days in Brussels on the 27th of September 2007

The different **INRIA** teams collaborate with leading European research groups and projects. INRIA-Atlas and ECOO team at LORIA, Nancy to combine a prototype implementation of reconciliation based on operational transforms) with distributed time-stamping algorithm developed by the Atlas team to demonstrate reconciliation for the P2P Xwiki application. INRIA-Regal collaborates with the Universita Nova de Lisboa on commutative replicated data types for collaborative editing. INRIA-Grandlarge collaborates with SZTAKI of Budapest for sandboxing technologies for Desktop Grids.

KTH is collaborating with the SELFMAN project (Self-management for large-scale distributed systems based on structured overlay networks and components) through the exchange on know-how and software (DKS/NICHE structured P2P middleware), including joint development of software, technical meetings on support for self-management in distributed systems and a component management framework based on Fractal; support for overlay services. KTS collaborates with researchers from the NextGrid project on Grid security architecture. It also consults and discusses with researchers from the TRUSTCOM project (completed in 2007) on policy-based access control, authorization mechanisms and software.

ICCS collaboration activities target the area of data management. Since ICCS also participates in FP6 project GREDIA, a few meetings have been organized internally in ICCS between the two research teams. GREDIA has designed and developed a prototype data transfer protocol (GridTorrent) which has the ability of simultaneously transferring different parts of the same collection of data from multiple sources. ICCS considers adding this protocol to the data transfer protocols that are already supported by DFS (VBS internal, FTP).

Moreover, ICCS has participated in the Collaboration Task 1 Exploitation of synergies/technical concertation-Data Management (TG5) meeting that took place during the European Services, Software and Grid Technology Days in Brussels on the 27th of September 2007. After this meeting, it was decided that Grid4All initiates the procedure of becoming a regular participant in TG5, sharing several aspects of the technology developed in Grid4All Data Storage (WP3) with the projects that already participate in these collaboration activities. The next target is to amend TG5 white paper with the development and experience resulted from the work on data management that takes place in Grid4All.

UPC participates in a Spanish project with the Universidad de Navarra. The objective is the decentralization of Grids and social networks and how social networks can inform and optimize decentralized grids. Along with the Universidad Rovira i Virgili in Spain, they are planning research on simulation and optimization of decentralized grid systems that will be developed during 2008.

FT

After the publication of a reference article about Grid4All Project on the Internal FT communication the 10th of December, 3 projects (for Orange services) have contacted the Grid4All group with objective to explore further on how Grid4All results can be used within these projects.

ANTARES

One of Antares objectives on the Grid4All Project is to develop a version of its e-tutor application getting advantage of a decentralized environment such as Grid4All.
The E-Tutor application (on its stand alone version which has just been released) is currently being used by institutions such us VAESA (Volkswagen Audi Spain) or Linguactive (an online language learning platform). FIBA (International Basketball Federation) is also going to integrate E-Tutor application on its refereeing LMS shortly. Some e-learning consultant corporations have shown their interest in the application.
Antares is starting to disseminate the concept of grid4all among these e-learning consultant corporations which are interested in getting acquainted with its results.

Prospects

At this stage the overall project goals are to realise proof-of-concept implementation in order to evaluate the system focusing on user and scenario requirements. With respect to the overall Grid4All vision, this particularly concerns schools; they are the primary targets as eventual end users. Our current objective is to develop/adapt applications using the Grid4All middleware to render them self-managing. Then we will focus on evaluation of the platform.
Security and trust are crucial aspects before going towards pilot studies. It is enormously important that the providers of resources are assured their security is not being compromised, and correspondingly consumers that their applications execute in a secure environment. We will capitalise on work on the potential of virtual machine technology to preserve this isolation.

Glossary and Acronyms used:

VO: Virtual Organisation
DHT: Distributed Hash Table
DKS: Distributed K-ary System, a peer-to-peer middleware by SICS that is based on a DHT
DFS: Distributed File System
Fractal: Extensible and hierarchical language-agnostic component model