Grid4All
democratize the grid
NEWSLetter
issue n°3, September 2008

Events

Decentralized self management for Grids, P2P an user communties *Oct. 20-21, 2008, Isola di San Servolo (Venice), Italy*

Workshop held in conjunction with SASO 2008, Oct. 20-24, 2008

The Internet is a fantastic tool for information and resource sharing. User communities such as families, friends, schools, clubs, etc., can pool their resources and their knowledge: hardware, computation time, file space, photos, data, annotations, pointers, opinions, etc.

However, infrastructures and tools for supporting such activities are still relatively primitive. Existing P2P networks enable world-wide file sharing but are limited to read-only data and provide no security or confidentiality guarantees. Grids support closed-membership virtual organizations (VOs), but their management remains largely manual. Web 2.0 social networks, blogs, and wikis remain centralized and have limited functionality.

This workshop examines issues of decentralized self management as they relate to these areas. The workshop is co-located with SASO 2008, the Second International Conference on Self-Adaptive and Self-Organizing Systems.

It is sponsored by the European projects Grid4All and SelfMan and with corporate sponsorship from France Télécom R&D.

Organizing committee:

- Peter Van Roy, Université catholique de Louvain, Belgium
- Marc Shapiro, INRIA & LIP6, Paris, France
- Seif Haridi, SICS & KTH, Stockholm, Sweden

INRIA Saclay presented the grid4all project at the third **NEGST workshop**, June 24th 2008 in Toulouse, France. NEGST's (NExt Grid Systems and Techniques) objective is to promote the collaborations of Japan and France on grid computing technology. Workshop's participants come from top-level universities and premium research institutes from both Japan and France.

INRIA Atlas organized and participated to the **Second Int. Workshhop on High Performance Management on Grids** (HPDGRID 2008) in Toulouse on June 13, 2008, with one keynote conference and one paper presentation related to Grid4All.

The Distributed Component Management System was demonstrated by means of a distributed storage service, at the **7th edition of Orange**

Editorial

The Grid4All project embraces the vision of a “Democratic Grid” as a ubiquitous utility, whereby domestic users, small organizations (e.g. schools) and enterprises can aggregate own resources and draw on resources on the Internet without having to individually invest and manage computing and IT resources. The Democratic Grid should allow its users and resource providers to form Virtual Organizations (VOs) on demand for specific purposes.

The major achievements of the project in the second year include refining the user scenarios and sharpening user requirements; definition of the overall Grid4All (Democratic Grid) architecture and APIs; development and implementation of prototypes of Grid4All components and demonstrator applications. In the second year, the project also spent efforts on defining the evaluation plan for the Grid4All prototype to be evaluated in the third year.

The overall Grid4All architecture was specified and presented at the M18 review meeting in February 2008. The architecture includes the following layers:

- Fabric: computing, networking and storage resource provided by VO members;
- Core VO support: overlay services, VO management (including VO membership service and security infrastructure), distributed components;
- Collaborative & federative services: execution service, VO file system and Semantic Store;
- Inter-VO services: resource brokerage, Grid4All marketplace, information services
- Applications.

In the second year of the project, a number of Grid4All components, services, frameworks and self-managing demonstrator applications were implemented and integrated with each other. Some of the implemented software is described in this newspaper. Most of the components, APIs and alfa quality end-user oriented applications were released by June 2008. Integrated parts of the Grid4All environment are expected to be released by January 2009. After further evaluation, the project will deliver the final prototype systems by June 2009.

A number of end-user oriented applications are being developed that use in different ways the APIs provided by the components, services and frameworks developed within the project. Note that this is still ongoing work and further integration is to be done in the third year. The set of demonstrator applications includes

- Collaborative File Sharing (CFS) that allows users to collaborate, interact and share information;
- Collaborative Network Simulator Environment (CNSE) that aims at supporting collaborative learning scenarios within the context of networking education;
- eMeeting, which is an on-line synchronous collaborative tool that allows sharing not only voice and video but also text, and other forms of sharing such as a group annotation and drawing tool and a polling tool;
- gMovie, which is a service accessible via a web interface, allowing distributed video transcoding on top of VO resources;
- Shared Calendar (SC) which is an application that allows multiple participants to schedule meetings handling conflicts among concurrent and overlapping proposals for meetings or tasks.

There are also some software results including:

- Distributed Component Management System (DCMS) which is a framework for building self-managing services and applications. DCMS is based an extended Fractal component model and allows building management code (feedback control loops) of an application using three types of Management Elements: watchers, aggregators and managers. The first two types of MEs are used to monitor a managed application, whereas managers should perform analysis, planning and actuation stages of the management loop. In order to demonstrate abilities of DCMS, Grid4All has developed a self-

Labs research fair held on the 10th and 11th of June 2008. These research fairs are organized at the end of each semester and services as a window to the R&D activities. Visitors from the entire France Telecom group, come to discover the demonstrations representing main research results. This research fair is important since deciders at the group level inform themselves of the latest R&D innovation.

Grid4All project was one of 16 EU-projects selected for presentation at **the European Project Showcase event** that took place for the first time alongside the 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2008), May 19-22, 2008, Ecole Normale Superieure de Lyon, Lyon, France. The European Project Showcase event highlighted European and nationally funded Grid projects in Europe. The Grid4All project presented the project objectives, the vision of the Democratic Grid; the Grird4All architecture; core VO support, and collaborative/federative services developed in the project.

Participants of the Grid4All project met participants of the XreemOS project at **the joint XtreemOS/Grid4All meeting** held on July 1-2, 2008, LIP6, Paris. The purpose of the meeting was to exchange information about project objectives and results. Researches from both projects made technical presentations on different topics of common interests, e.g. VO management, file systems for Grids, self-management.

Grid4All project was presented at **Gecon2008**: The 5th International Workshop on Grid Economics and Business Models, in Las Palmas, Spain; August 25th-26th, 2008; co-located with Euro-Par 2008. It presented the Grid4All vision as it relates to the VO-oriented marketplace. The presentation was entitled "Grid4All: Open Market Places for Democratic Grids" and it was presented in a session on "Research Projects on Market Mechanisms", with other two related projects. The supporting paper will be published in the LNCS series.

managing service called YASS (Yet-Another-Storage-System) for storing files. YASS can self-heal, self-organize and self-tune in a Grid4All dynamic environment.
- Grid4All Security Infrastructure (SI), which includes a set of security components that can be used to protect Grid4all resources and services from unauthorized access. Grid4All SI supports policy-based authorization. We have used the Sun's XACML implementation in order to develop and to implement a prototype of the Grid4All security framework. Grid4All SI has been integrated with VO Membership Service in order to provide authentication of VO members.

At the current stage, the ongoing work includes further integration of Grid4All software components, frameworks, services and subsystems, as well as further integration and improving of integration of demonstrator applications with Grid4All components, e.g. DCMS, Telex, VOFS, VO Membership Service and security components. Evaluation of the Grid4All prototype is our major future work to be done in the third year.

Telex: Principled System Support for Write-Sharing in Collaborative Applications

Telex is an innovative middleware to help designing collaborative distributed applications. Contrary to existing tools, for instance wikis, Telex is totally distributed across the participants.

The benefits of Telex are multiple:

- Telex increases mobility: Telex users may work off-line, executing tentative operations on their local replicas, and reconcile later when reconnecting.
- Telex increases the availability and the responsiveness of the system: there is no single point of failure.
- Telex decreases the infrastructure cost since the service is partially provided by the users.

Handling accesses to shared resources is a non-trivial issue when designing distributed collaborative applications. Telex helps developers by clarifying building blocks, and providing consistency control and persistence management facilities.

The Telex system is designed for sharing mutable data in a distributed environment, particularly for collaborative applications. Users operate on their local, persistent replica of shared documents; they can work disconnected and suffer no network latency.

The Telex approach to detect and correct conflicts is application independent, based on an action-constraint graph (ACG) that summarises the concurrency semantics of applications. The ACG is stored efficiently in a multilog structure that eliminates contention and is optimised for locality.

Telex supports multiple applications and multi-document updates. The Telex system clearly separates system logic (which includes replication, views, undo, security, consistency, conflicts, and commitment) from application logic. An example application is a shared calendar for managing multi-user meetings; the system detects meeting conflicts and resolves them consistently.

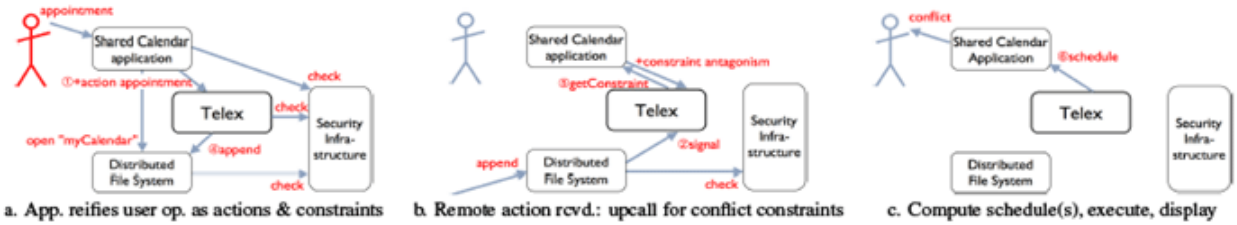

a. App. reifies user op. as actions & constraints b. Remote action rcvd.: upcall for conflict constraints c. Compute schedule(s), execute, display

A demonstration application for Telex is a Shared Calendar application (SCA) supporting multi-user meetings. As Telex is parameterised by the semantics of the application (the shared calendar in this case) it is aware of conflicts such as double bookings. Telex combines the meeting requests from all users, in a peer-to-peer way, and proposes optimal, conflict-free schedules. Eventually all sites converge to the same schedule.

Another application above Telex is the Cooperative File System (CFS). CFS enables groups of users to manage a library of files and file versions in a distributed, peer-to-peer, eventually-consistent way.

The distribution logic is encapsulated into Telex, which takes care of recording and propagating updates, replaying them, detecting conflicts, and reconciliation. The developers of SCA and CFS were able to concentrate on their application logic.

Telex is described in a technical report available here: http://fr.arxiv.org/abs/0805.4680. Telex is open source, and is maintained at http://gforge.inria.fr/projects/telex2/

Some Concrete Results of the Grid4All

VOFS: Virtual Organization File System

VOFS is a federative peer-to-peer file system for use with Grid4All VOs. It makes it easy for VO members to expose their files in shared workspaces and collaborate. It also provides applications with convenient functions for general communication and for distributed file access.

A VOFS prototype will be available as a product of Grid4All.

In contrast to common filesystem usage, VOFS focuses on the user, being a personal tool rather than a system tool. Users command their own peers and are free to explore a network of files creating or joining shared workspaces. Structured activities are easily organised by authorities, such as instructors, who create a suitable environment that students may join.

Behind the apparent simplicity, flexible mechanisms allow VOFS to integrate with administrative services such as VO-wide policy-based access control and self-healing storage resource management.

VOFS supports user mobility through disconnected operation and can support distributed applications by helping them communicate through the filesystem. An important example of such an application is the Telex middleware which provides a powerful framework for collaborative applications.

VOFS is to be used firstly by end users as part of their collaborative VO environment. Secondly, VOFS is used by administrators to create collaborative environments and interact with users that participate. Thirdly, VOFS is used by application writers that wish to exploit the VOFS environment in their applications.

DCMS: Distributed Component Management System

Distributed Component Management System (DCMS) is a framework that provides an expressive API for programming self-managing component-based distributed applications. Application's functional and self-management code are programmed separately and independently of deployment configuration on the network. The framework extends the Fractal component model by the component group abstraction and one-to-any and one-to-all communication patterns. DCMS can automatically move application components responsible for self-management when necessary due to resource churn. The framework supports a network-transparent view of system architecture simplifying designing application self-* code.

Programming application self-* behaviours with DCMS requires just a few dozens lines per application component. The implementation of the framework relies on scalability and robustness of the Niche structured p2p overlay network. It has been also developed a distributed file storage service to illustrate and evaluate the framework.

What is the innovation?

- DCMS is leveraging the latest component models with state-of-art in overlay networks to benefit from their self-managing capabilities.
- DCMS is introducing a network transparent view of system architecture which simplifies the design of self-managing applications.
- Robustness and reliability of management components are handled automatically by the system. DCMS maintains reliably the description of the application architecture and automatically reconfigures the architecture in case of failures, adaptation to load.
- DCMS provides network transparency simplifying the design of self-managing applications. Applications are developed once and deployed anywhere.

For whom?

- Software developers and architects to improve the application development cycle. The principled APIS and reliable management allow applications to focus on functional concerns.
- Administrators and system managers can rely on self-configuring and self-optimizing capabilities and focus on specifying high-level policies.
- Service providers who would like to reduce total cost of owning and operating IT infrastructure.

Which usages?

- Allows separation of application functional and non-functional parts concerning both development and execution.
- Allows applications to be written once, and reused for many different environments, ranging from heterogeneous community Grids to high performance clusters with minimal or no changes to the application functional parts.
- Allows applications to be deployed in volatile target environments where the cost of manual management would be prohibitive. Enabler to deploy services and applications on off-the-shelf hardware, global and desk-top computing.

APPA P2P-LTR (Log and timestamping for reconciliation)

P2P-LTR is a service of the APPA P2P platform (http://gforge.inria.fr/projects/appa/) developed in the INRIA Atlas team. P2P-LTR provides two major functions: logging of user actions in a DHT and continuous, distributed timestamping of these actions. This is useful to perform reconciliation of replicated data.

P2P-LTR is used for an implementation of the Semantic Store based on the APPA P2P system as promised in the Grid4All project proposal by ATLAS. The main implementation and validation of SS is the Telex middleware of the Regal team, which provides applications with replication and consistency services based on the Action Constraint Formalism (ACF). P2P-LTR is useful to support a simpler, yet not as powerful, alternative to the Action Constraint Framework (also developed in Grid4All). Besides validating an alternative implementation to ACF, this is useful to study the feasibility of extending an existing collaborative application with SS functionality which is a hard problem.

The existing application is Xwiki, a client-server wiki system which is used intensively for collaboration among small enterprises or within large organizations.

SIS: Semantic Information System

In a market-based environment as envisaged by the Grid4All project, grid resources and services are made available through peer-initiated markets in a distributed manner. The Semantic Information System (SIS) within Grid4All provides a matching and selection service between peer negotiators willing to offer or use resources and services. Grid4all adopts a distributed market model where (resource) consumers and providers negotiate on certain traded entities in auctions initiated by providers, consumers or by third party entities. In order for an auction to take place and given the distributed nature of market creation and management, markets must be discovered by potential participants. The SIS provides a directory of resources, services and markets, that is, a registry of descriptions of resources and services. By searching though this registry, peers can allocate available markets negotiating resources and services. Provisional provider/consumer software agents and human users are able to query the SIS registry for available service descriptions that match certain attributes and criteria concerning price, time of availability, quality of a service, etc. The query results are ranked according to the capacity of resources, the preferences and intentions of providers and consumers in a selection process.

The SIS uses Semantic Web technologies to facilitate the advertisement, discovery, matching and selection of services and resources. The system provides an API for both advertisement and discovery of market and application services.

SIS is addressed to both simple users and application programmers which are agnostic of the low-level details for managing the resources and services offered and requested.

APPA selection service

The selection service is part of the Semantic Information System (SIS) and is a service of the APPA P2P platform (http://gforge.inria.fr/projects/appa/) developed in the INRIA Atlas team. It is responsible of ranking a set of services (requests or offers) according to the buyers' and sellers' preferences, which are mainly based on the queries type (i.e. which resources or services are asked by the query). The particularity of this service is that it balances both preferences based on the buyers' and sellers' satisfaction. Satisfaction in this context means that the preferences of a buyer or a seller are met, in the long run, by the system. The purpose of doing so is to allow buyers and sellers to accomplish their objectives by avoiding that some of them monopolize services, but also to preserve the control over their resources or services.

Grid4All Security Infrastructure

The goal of the Grid4All Security Infrastructure (SI) developed by KTH is to ensure that VO-members can only access resources that they have the right to access. Authorization is policy-based; policies are expressed in XACML (eXtensible Access Control Markup Language). SI includes a number of security components (security points) developed using the Sun's XACML Java packages. A VO policy administrator uses a Policy Administration Point to set and store policies in the Policy Repository. Before accessing any resource, the user logs in to the VO membership service (VOMS) that returns a memberId used to identify the user. When accessing a resource, a user's application sends the user name and memberId to the Policy Enforcement Point (PEP), which protects the resource. Upon receiving a request, PEP interacts with the Policy Decision Point (PDP) that evaluates the access request against VO policies. PDP contacts VOMS to check user's identity, gets the user's roles from the Policy Information Point (PIP), and, finally, evaluates the request. If access was granted, the requestor accesses the resource. A prototype of the Grid4All Security infrastructure integrated with VOMS has been implemented, and it can be can be used in Grid4All components and applications.

Self-Grid: Dynamic virtual organizations for schools, families, and all (GRID4ALL)
Web site: http://www.grid4all.eu/
FP6-IST-Grid Technologies: http://cordis.europa.eu/ist/grids/projects.htm
Project Reference: 034567, Start Date: 2006-06-01, Duration: 36 months.