GRID4ALL

Scope

Imagine you are a secondary school teacher and you have an ambitious plan for a scientific project for your students along with some other schools with which you have contacts. The project will include lectures by external experts and collaborative work between students guided by the teachers. Students will use geophysics simulation software with real data. They will analyse the results with the teachers. How to set up an IT organisation to implement your project? Will you have all computational resources to run applications, store data, and incorporate content provided by experts and research institutes? Will you need to pay specialist system administrators to maintain your IT infrastructure? Will students be able to participate from wherever they are?

Or imagine you are returning home from work and have invited family and friends to watch a holiday film. You need to transcode it before 21:30 and with your PC at home you will not make it in time. You use your operator's online service to transcode on the net so as to have your film back by 21:30, indicating that this is urgent for you.

Grid4all embraces the visions of a "Democratic Grid" as a ubiquitous utility whereby domestic users, small organisations (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.

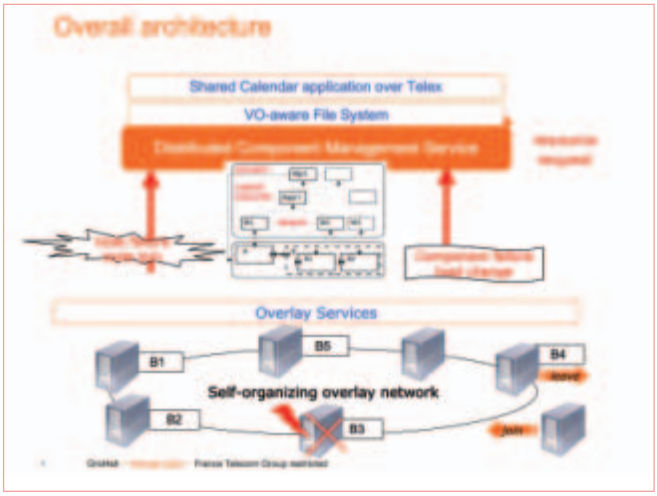

Advances

The main Grid4all results consist of software frameworks, middleware, and tools. The project aims at building proof-of-concept prototypes (middleware, a component framework, services) that will contribute to building a Democratic Grid and at developing a number of prototype self-managing applications.

Support for self-management in applications, middleware, and services is required to ease administration and use by inexperienced users. **The Distributed Component Management Service (DCMS)** is a middleware framework that provides an expressive Application Programming Interface to program self-managing component-based distributed applications. DCMS aims to achieve predictable behaviour in uncertain environments with churn and volatility. Management tasks can be programmed independently of the deployment configuration. DCMS automatically moves the management components when necessary, e.g. due to failures or churn. Deployment configurations can be expressed declaratively and are managed automatically by the middleware. The DCMS prototype relies on the scalability and robustness of the Niche structured peer-to-peer overlay network. A distributed file storage service to illustrate and evaluate our framework has been prototyped. DCMS uses the Fractal component model, extends the Fractal Architecture Description Language, and provides primitives to configure sensors, monitors and actuators that implement the management tasks.

The project has released **VO-aware distributed file system (VOFS) software** to build federative peer-to-peer file systems that allow members of a virtual community to share data. Virtual Organisations (VOs) pose a number of issues in file-system design. For example, VOs tend to appear and disappear dynamically, yet documents must remain persistent. A user may choose to export only a view, i.e. a subset of his files, to a given VO. He might export different, possibly overlapping views to different VOs; despite this, VOs must remain isolated from one another. VOFS is user-oriented and allows participants to access networked data (represented as files), to create new work spaces, or to join shared workspaces. Users may be mobile and even disconnected. VOFS administrators may set up collaborative environments controlling the access to the different users. The VOFS architecture provides a separation between meta-data, access control, and the consistent management of the data itself. It also supports distributed applications by providing file system based communication.

Telex is a middleware tool that helps design collaborative applications. Contrary to tools such as Wikis, Telex is totally distributed across the participants and hence there is no single point of failure. Telex increases mobility and its users may work off-line. They execute tentative operations on their local replica and reconcile later when reconnecting. Telex leverages on VOFS for persistency. It supports multiple applications; a shared calendar, a distributed dictionary, and a distributed database have been developed using Telex.

P2P- Logging and Timestamping for Reconciliation (P2P-LTR) is a middleware service for reconciling updates on replicated data in a P2P system. P2P-LTR supports a simpler, yet not as powerful, alternative to the Action Constraint Formalism (ACF) implemented in Telex. Besides validating an alternative implementation to ACF, this is useful to study the feasibility of extending an existing collaborative application with Grid4all functionality. The existing application is Xwiki, a wiki system which is used intensively for collaboration among small enterprises or within large organizations.

Positioning in global context

Currently there are not many alternatives to DCMS to program distributed applications that may execute on resources within a volatile environment. It provides end-to-end support; from programming, via deployment, to run-time management and reconfiguration.

Telex middleware has successfully been used to build a wide variety of collaborative applications. It is based on the Action Constraint Formalism which has enabled reasoning about consistency properties. It has also enabled designing a new generalised peer-to-peer consistency protocol.

Contribution to standardization and interoperability issues

The project uses the Fractal component model to design and develop its software. The Fractal component model is used in the Grid Component Model that is being promoted within the Coregrid NOE.

Target users / sectors in business and society

Potential users of Grid4all software and frameworks are mainly application developers, and technology and service providers. End users may use the produced middleware, services and applications to setup virtual organisations and to use the tools for collaboration.

Overall benefits for business and society

Grid4all provides an opportunity for citizens to democratise the global network by opening up participation by lowering costs, facilitating usage, and supporting collaboration and data sharing.
The DCMS framework simplifies programming of robust self-managing applications that could execute on unpredictable environments. This is one step towards truly global computing, where idle Internet resources can be reliably used by society and small businesses.
The Telex middleware simplifies programming of distributed collaborative applications. Its replication and consistency service allows developers to focus on application logic and promotes a clean separation of concerns. With Telex a wide category of collaborative applications may be used in volatile environments.

Examples of use

How could Grid4all results be used in the scenario of the secondary school project above?
The school project will require computational and storage resources on which simulation software may be executed. Sufficient quantity of resources may not be available within the schools themselves. Resources could be obtained from 'outside' the schools; for example, resources could be volunteered by parents, etc. Hence the underlying infrastructure on which the virtual organisation is built is formed by some resources of the schools themselves and extended dynamically by resources obtained from outside. This environment is subject to volatility and churn. The simulation application (and other applications) should be self-managing and should not require excessive management by teachers, students or other participants. Volunteered resources may also be used as part of the storage system required to host the project documentation.
Data management is required by a collaborative school project that involves various participants, such as students, teachers, and expert researchers. The data provided by the participants need to be accessible by all the other participants of the virtual organisation that is created for the project. With some help from teachers, students will produce documents describing the project and its execution. Students follow specific courses given by the teachers and experts and collaboratively – in small, remote groups, reply to tests assessing what they have learnt. The VOFS may be used to set up a virtual file system allowing the different participants to export their storage and data and control the access to these data. Applications developed using Telex, allow students to work on documents even when they are disconnected from the rest since it allows the different contributions to be reconciled.
Software failures and the complexity of their management are major problems. The DCMS middleware has been used to develop a self-managing storage application that aggregates disk space contributed by desktops and portable computers. The application is robust to churn, that is, it automatically recovers when one of the computers forming the storage system leaves or fails. Applications (such as the simulation software) may also be adapted to use DCMS. In this case, the middleware automatically handles churn and failures.
The Telex middleware has been used to develop a shared calendar. Telex supports a wide category of collaborative applications. Currently, it has been used to develop a shared calendar, a distributed dictionary and a distributed database. The well-known collaborative file sharing application CFS is also being adapted to use Telex. This extends FireFox with concurrent data sharing to allow decentralised collaboration between groups of people.

Achievements

- **The Distributed Component Management Service (DCMS): Middleware and API**, alpha release
  Contact: Konstantin Popov, kost@sics.se
  https://gforge.inria.fr/projects/grid4all/ under wp1/Jade and wp1/YASS
- **The VO-aware distributed file system (VOFS) : Software package**
  Contact: Georgios Tsoukalis gtsouk@cslab.ece.ntua.gr
  http://danaos.cslab.ece.ntua.gr/g4a-svn
- **The Telex: Middleware and APIs**
  Contact: Marc Shapiro marc.shapiro@acm.org ;
  Jean-Michel Busca jean-michel.busca@lip6.fr
  http://telex2.gforge.inria.fr (alpha release)
- **The P2P- Logging and Timestamping for Reconciliation (P2P-LTR): Middleware**
  Contact Patrick Valduriez (Patrick.valduriez@inria.fr)
  http://gforge.inria.fr/projects/appa/ (alpha release).

title
Self-Grid: dynamic virtual organizations for schools, families, and all

contract number
034567

type of project
Specific Targeted Research Project

contact point
Ruby Krishnaswamy
FRANCE TELECOM SA, FR
E-mail: ruby.krishnaswamy@francetelecom.com

project website and partner list
http://grid4all.elibel.tm.fr/

EC contribution
2 900 000 €

start date
01/06/2006

duration
36