Grid Survey

Håvard K. F. Bjerke (havarbj *at* idi *dot* ntnu *dot* no)

August 9, 2004

Contents

1 Introduction

The purpose of this document is to give estimates of the likely technical and scientifical impact of some of the Grid and Grid-related projects in various areas of distributed computing. The emphasis is on technical developments and services that are more advanced than provided by simple deployment of the basic toolkits. These toolkits may be used as a basis for comparison where applicable.

In cases of projects intending to provide a service in a form of a production Grid, an effort to evaluate the funding situation and other organizational aspects affecting the feasibility of the plans is made. Similarly, in cases of projects intending to produce new kinds of software, the characteristics of the user and developer communities is taken into account.

As a background work for selecting the projects that are of most interest, a survey of a larger group of projects is made before a smaller group is selected, while the rest is included in an appendix. Due to the rapidly changing nature of the Grid landscape, reviewing these projects at some later point can be useful.

1.1 Executive Summary

The world of grid computing is continuously growing, and new projects are founded in an increasing rate. These projects range from dedicated grid computing infrastructures to public infrastructures, from academic to commercial, from pilot projects to production systems and from proof-of-concept to traditional science applications.

The two organizations Global Grid Forum (GGF) and The Globus Alliance have a high influence in grid computing developments. GGF is an organization which is lead by academic and commercial Grid communities and coordinates the setting of standards for best practices in Grid computing. The Globus Alliance's influence comes mostly through the development of the Globus Toolkit (GT), which is used in many Grid middlewares. Some projects, however, develop their own solutions for middleware, and some base them on the GT, such as NorduGrid, who are developing the Advanced Resource Connector (ARC) middleware on top of this.

Also, different solutions exist for fabrics, from dedicated infrastructures with high speed interconnections and high-end computing nodes to public computing infrastructures, such as BOINC. There exists a trade-off between these two types of computational resources in that the public resources are free and potentially vast but less reliable compared to dedicated resources.

Security mechanisms are important for the sharing of computing resources across VOs. Globus' GSI is an extension to the X.509 security infrastructure. It adapts the security mechanisms of X.509 to Grids with the concept of proxy certificates. The Liberty framework brings security mechanisms to a higher level and allows for authentication across Grid organizations.

The parties involved in the projects, through funding and applications, are important for the successes of the projects, although the amount of funding is not necessarily critical. Also, the job efficiency rates are indicators of the capabilities of the Grids.

2 Background

The development of grid computing is controlled by many different factors and by many parties from different levels in the Grid world; and there exists many types of Grid projects with different research areas. This section provides some background information about some of the important parties and some basic categories for grid computing research areas.

2.1 Research Areas

The projects in this survey are diverse in their intentions. Some projects are created in the interest of seeing a functional computational grid as the end result, while others are created in the interest of evaluating the feasiblity and applicability of computational grids for their computational problems. The typical computational grid or related projects fall into one or more of the following categories.

- *Grid technology development* - Grid projects that focus on developing technologies to be used in Grids.
- *Grid testbeds* - projects that employ Grid infrastructure in order to test and research technology.
- *Field specific* - Grid projects that apply to specific fields, that have production or development status.
- *Proto-Grid* - projects developing competing or complementary solutions.

2.2 Organizations

The following organizations are relevant to the world of grid computing, as they are established as familiar organizations in coordinating grid computing efforts.

2.2.1 Global Grid Forum

The Global Grid Forum (GGF) is a forum initiated by the computational grid community in industry and research, that sets the agenda for making open standards for best practices in grid computing. GGF's process of making standards is modeled after the Internet Engineering Task Force's (IETF) Internet Standards Process, and documents are published similarly to RFC (Request For Comments) documents. GGF is also a forum for the exhanging of experiences from different levels in grid computing, such as research, software, deployment and usage. The participants of GGF come from over 400 organizations in over 50 countries. Funding and support come from sponsoring members in technology production and use, and academic research institutions. [21]

2.2.2 The Globus Alliance

Globus develop the Globus Toolkit (GT), which is one of the most used grid computing software toolkits in grid computing today, either directly or in some derived form. The GT includes most of the services that are needed in a Grid, such as

- security - through the Grid Security Infrastructure (GSI),
- data management - through GridFTP, the Reliable File Transfer Service (RFT) and the Replica Location Service (RLS),
- resource management - through the Grid Resource Allocation and Management (GRAM),
- information services - through the Monitoring and Discovery System (MDS).

The GT software is licensed after the "Globus Toolkit Public License" (GTPL) Version 2. This is a restrictionless open source license which must be inherited by derived software. [31]

3 Research

The information that is presented in this survey, is mostly gathered from the projects' websites, where the projects often publish results of tests, data about their Grid infrastructure and information about their applications. This data is often publically available through presentation slides, meeting proceedings, articles and publications. Some projects also have publical live demos or even open portals to their systems. This gives an opportunity to experience the properties of these systems first hand. Grid-related news articles also provide a useful and possibly more neutral source of information. Finally, some information is gathered from interviews with people who are involved in the projects.

Allover, most of the information is based on documents found on the respective projects' web sites. However, the projects Grid2003, EGEE, GRACE and NorduGrid have been contacted directly, and representatives have presented useful documents as well as information not on print. Information about BOINC has generously been given by office mate Karl Chen, who has been working with the BOINC project. Information about the TeraGrid project has primarily been gathered from its web site and documents of similar projects, while information about Xgrid has been collected from recent articles and posts on related fora.

3.1 Project categories

The definition of a Grid or grid computing suffers from the fact that the term "Grid" is overly hyped. Without intending to define what is a Grid project, and rather including border cases, a range of projects are included in this survey, from those that are related to grid computing, to those that maintain a production grid computing environment. The projects are divided into categories that signify their intensions and status in regard to that:

- *Proof-of-Concept* projects include those that utilize grid technology and demonstrate novel or curious concepts in relation to grid computing.
- *Prototype* projects include those that are prototyping Grid or Grid-related technology, whether or not they intend to produce production level technology.
- *Beyond prototype* projects include those that have produced grid computing environments or technologies that have progressed beyond the state of prototyping and into the state of production.

3.2 Academic vs Commercial

Academic projects often distinguish themselves from commercial projects in many ways that affect the processes and products of the projects. This gives reason to separate these projects into academic and commercial. Differences between academic and commercial projects are, among others:

- While commercial projects often have a tight centralized management, the academic projects more often have a decentralized management.
- Commercial projects have applications that are, to a higher extent, mission critical.
- Commercial projects to a higher extent employ databases, while academic projects use files, for data storage and management.

3.3 Framework

In order to provide a near complete view of the projects' attributes and working areas and provide a background for the analysis of the projects, each project is reviewed in respect to a few attributes that are relevant to the particular project. The following list provides a framework of attributes that are relevant to the development of the projects.

- URL
- Clients and Application Areas
- Goals and Objectives
- Project Status
- Community and Work Force
- Licensing
- Application Environment and Tools
- Mobile Access, Portals
- Security
- Resource Management
- Data Management
- Problem Characteristics
- Monitoring and Performance
- Fabric

4 Analysis

This section analyses aspects of and developments in areas that are relevant to grid computing.

4.1 Stability

Distributed computing projects that use the spare time of any internet-connected computers running their software, are often called public computing or cycle scavenging. Such projects have a potentially vast amount of computing resources. The number of computers connected to Internet was in September 2002, 605.60 million; the number of users contributing to SETI@Home through BOINC (see 7.1.2) is higher than 25000. [25], [3] The amount of processing power that these projects receive from the computers running their software is limited by the number of computers connected to Internet, the popularity of the cause of the project or the project itself, the amount of similar projects, and the amount of CPU time the contributors give to the project.

While such projects may have very large resource pools, the probability that a single computational work unit will succeed, i.e. deliver the right results, if any, is lower than in a fully dedicated pool of processors, such as in a cluster. The main factor that affects the stability of computational jobs running on public computing infrastructures, is the presence of human interaction with the platforms that the jobs run on. Such interactions may include the running of software that makes the jobs crash, turning off the computer in the middle of a computation, and maliciously returning false results.

Distributed and parallel computing projects with large numbers of components also face the problem of component failure, either in software or hardware. The failure of a single computing node will cause the failure of at least one computational work unit, unless preventive actions, if possible, are taken.

4.1.1 Redundant Computing

A way to remedy these problems is to make redundant computations, i.e. to run a single computation on more than one node in parallel, and then compare the results. If the probability of a computation failing on any node is X, when the computation is run redundantly on n nodes, the probability of the computation failing becomes X^{-n}. The trade-off to this procedure is that if the entire resource pool is saturated, this reduces the utilization of the pool to the X-th of the available resources.

A BOINC (see 7.1.2 computation has a chance of failure in the order of 1%, due to the fact that the processors of the BOINC infrastructure are not dedicated. According to a poll, [16] , performed by the main BOINC user, SETI@home, most contributors of computation power are males in the age range 20 to 39, committing a single computer less than 24 hours a day. Even though many contributors of computation power are committed to the cause of SETI's objective for the SETI@home project, most do not surrender all their computers' time to the project. As contributors may interrupt or crash running jobs while starting computer games or similar, this poses a significant cause for job failure.

Due to the popularity of the SETI@home project, the project has abundant computing resources. This allows them to exploit the method of redundancy without underutilizing the available resources. By running three copies of a single job simultaneously on three nodes, the failure rate of jobs running on the BOINC infrastructure is reduced to 10^{-6}.

4.2 Security

One of the core concepts of grid computing is the Virtual Organization (VO). The VO describes the mapping of a single set of credentials to several grid computing resources. This allows for authorization to specific resources based on digital identities that are tied to one or more users.

A simple manner of authenticating members of a VO is to tie their identities to a username and password pair. However, the process of entering a password each time a resource is needed is troublesome and can be insecure, as users make their own mechanisms in order to avoid the process, that are insecure. Other schemes, such as Kerberos and X.509, are deviced to make the process of authentication transparent.

4.2.1 GSI

In a grid computing environment different VOs need differentiated access to files, CPU resources, software and other resources. These resources may be needed sporadically at different times during a job, each time requiring a new authentication process. However, jobs running on a grid may be running for several days and should be able to do so without user intervention - what is often called a ”single sign-on” system. For this the traditional Internet security mechanisms do not suffice.

Globus' GSI is an extension to X.509 and provides mechanisms for a uniform authentication for different grid resources. Furthermore, by delegating the user's authorization to a ”user proxy” and giving it the right to act on the user's behalf for a limited amount of time, the GSI automates job-intermediate authentication processes, eliminating the need for users to supervise jobs.

4.2.2 Liberty

Still, problems remain as each different service providers requires its users to have a identity specifically for their service. For a user who uses services such as banking and other Internet-based services and need access to many different resources, maintaining all the different identities is a challenge. Liberty provides a framework for a user to access several services, across different organizations, with a single identity. This allows for a single sign-on system in the sense that the user only has to log in once and then has access to resources across different web sites. The liberty framework specifies a security framework which is based on standard web technologies such as SSL 3.0.

Liberty presents a potential framework for VOs accross different sites or grid computing projects. As Liberty is limited to traditional web browser technology, the syndication of these services are limited to a higher level than site-internal resources. Such services may be based on standard web browser technology or SOAP and may include portals and mobile access, for example. For single sign-on possibilities to remain in a Grid-internal resource level, lower level authentication mechanisms such as GSI, are needed. [30] suggests integrating GSI and Liberty to Grid security in order to allow single decentralized VO identities across different grid resources.

5 Proof-of-Concept Projects

This section includes the survey of the project GRACE, which is considered proof-of-concept since it uses grid computing in a novel and unique application.

5.1 GRACE

GRACE (GRid seArch and Categorization Engine) is an EU funded grid search engine. The concept of the project is a distributed search engine for structured and unstructured information, with indexing on a grid architecture as opposed to a limiting centralized architecture. In contrast to traditional search engines such as Google, this search engine also acts as a knowledge management tool and introduces the concepts of a Categorization Engine and knowledge domains.

URL http://www.grace-ist.org/

Clients and Application Areas The service may be used, for instance, in university libraries for finding documents in large scientifical document bases.

Goals and Objectives By the project's end in February 2005, a final prototype with all GRACE functionalities is planned to be finished, and validation and evaluation performed by user communities is to be completed. The immediate goals of the project include [12]

- integrating physics, computer science and engineering as knowledge domains,
- finalizing the Grid set-up by the middle of July,
- completing the integration tests by the middle of August,
- completing a final evaluation and validation phase involving end users between September and November, through workshops and free evaluation the GRACE web portal.

Project Status The project has been running since September 2002 and continues until February 2005.[20] Its GRACE is currently being tested on the testbed GILDA and is, after the testing, still considered unstable and needs more testing, optimization and integration of other functionalities. [12] A live protoype demo is running at http://grace.gl2006europe.com:8080/.

Application Environment and Tools GRACE provides a configurable interface between the distributed local search engines and GRACE, named GRACE Toolkit. A local search engine may be such as Google, document servers, etc., and the GRACE local search engine may also be installed. [12]

The GUI of GRACE is a web interface for submitting queries and where search results are integrated; the middleware of GRACE is based on the middleware of LCG-2.

Security In order to submit a query, a certificate is needed, which can be a hassle for the ordinary university library user. The certification will be made either on user, service, or application specific level. [12]

Problem Characteristics The most computationally intensive processes of GRACE come from the Content Engine, which normalizes and indexes documents resulting from the query, and the Categorization Engine, which produces a concept map of the normalized results and renders them to the UI. These two problems are run as Grid jobs.

The total execution time of a search query is determined by the following factors:

- Grid overhead - in the order of two to five minutes.
- GRACE initialization - in the order of five to fifteen minutes.
- Processing by Content and Categorization Engines - variable execution time.

An example job of reading 20 documents amounting to about three megabytes has taken 50 to 60 minutes for the Content Engine, and 20 minutes for the Categorization Engine. [12]

Fabric The initial intention for GRACE was for the data to be stored on the LCG-2 grid; however, problems in the operation of the LCG-2 have delayed the move towards storage on the LCG-2. Currently, the plan is that data is stored locally at each document server, in databases, and GRACE provides an interface for connecting these databases to GRACE. Results from each of these nodes are returned through the interface in a normalized form. The testbed GILDA, which GRACE currently is being tested on, runs EGEE middleware (LCG-2) installed on two sites in Turin and Milan in Italy.

6 Prototyping Projects

The following projects are involved in developing large scale computational grids and grid technology, yet have not exercised production quality stability in general.

6.1 Grid2003

The Grid2003 project has deployed an international computational grid laboratory, Grid3, over sites that are operated by US Grid projects iVDGL, GriPhyN, PPDG and US participants in LHC experiment ATLAS and CMS, jointly.

URL http://www.ivdgl.org/grid2003/

Clients and Application Areas The resources that Grid2003 provide are used in seven different scientific applications, including HEP, bio-chemistry, astrophysics and astronomy. Experiments that have been using services of the Grid3 include

- the LHC ATLAS and CMS,
- the Sloan Digital Sky Survey project,
- the gravitational wave search experiment LIGO,
- the BTeV experiment at Fermilab,

and experiments in molecular structure and genome analysis and computer science research projects. The Grid2003 project claims to have 100 people, divided into 5 VOs, currently with access to the Grid. 11 applications occupy the parts of the larger sites' resources. The registration policy of Grid2003 is similar to that of LCG, and more than 102 users are authorized to use Grid3 resources. [23], [14], [28]

Goals and Objectives The purpose of the project is to build a computational grid environment, "Grid3", that will provide [14], [28]

- an infrastructure for the next phase of the iVDGL laboratory,
- the infrastructure and services needed for running LHC production and analysis applications,
- a platform for computer science technology demonstrators and researchers,
- a common grid computing infrastructure and services for SDSS, LIGO and other scientific projects.

The project intends to do efforts for being consistent with other Grid projects such as LCG.

Project Status The grid infrastructure of the project is deployed and integrated across the participating institutions and has been in operation since November 2003. [14] Grid2003 claims [28]

- 10 applications running,
- over 17 sites,
- a peak number of 1300 concurrent jobs,
- a job completion efficiency of up to 90% in well run sites,
- 4 TB of data transfer per day,
- 40% - 70% resource usage.

Application Environment and Tools The middleware of Grid3 is based on the Virtual Data Toolkit (VDT).

Security Grid2003 uses VOMS (Virtual Organization Membership Service) for maintaining VOs. VOMS is a result of the DataGrid project and is open-source software. VOMS is based on GSI. [14]

Resource Management For job queueing, sheduling and prioritizing, and for resource monitoring, Grid2003 uses Condor. [14]

Data Management Grid2003 uses the Virtual Data System (VDS) for storing representations of the computational procedures used to generate data and the datasets produced by them, which allows for auditing and automatic on-demand re-derivation of the data. VDS is based on a combination of Chimera, Pegasus and DAGMan and is open-source software.

For data storage, Grid2003 uses the Storage Resource Manager (SRM) for storage allocation across heterogenous mass storage systems. [14]

Monitoring and Performance Grid2003 uses a distributed monitoring system, Ganglia, for collection and presentation of system status of Grid, clusters and hosts. Ganglia presents status information about CPUs, memory, processes, load, and some higher level attributes. Also, a monitoring service, MonaLisa, is used for monitoring regional clusters and for dynamical configuration of site and network elements. [14]

Fabric The fabric of the Grid2003 computing grid consists of the processors of sites participating in the project. 27 sites distributed over USA and South Korea collectively provide about 2700 CPUs or a peak of 2800 CPUs. [23], [14], [28]

6.2 TeraGrid

TeraGrid is a field specific American grid infrastructure for scientific research. The project was launched by NSF in August 2001. The project's activities include training and support across its sites. [27]

URL http://www.teragrid.org/

Clients and Application Areas The TeraGrid computing infrastructure is dedicated to open scientific research. Users apply through the National Science Foundations (NSF's) Partnerships for Advanced Computational Infrastructure (PACI).

TeraGrid envisions the following projects to use their grid computing resources:

- The MIMD Lattice Computation (MILC) collaboration
- NAMD - simulation of large biomolecular systems

[27]

Goals and Objectives TeraGrid does not intend to expand their computational grid beyond a limited amount of sites; rather other projects are expected to interoperate with TeraGrid based on standard SLAs.

Project Status The system has been on-line since 23 January 2004.

Application Environment and Tools TeraGrid utilizes the middleware of the NSF Middleware Initiative's (NMI), which to a high degree builds on GT. Teragrid has support for MPI, BLAS and VTK.

Job Submission and Scheduling TeraGrid uses GT's GRAM for job submission and scheduling.

Security GSI is used for authentication.

Resource Management TeraGrid uses Condor for job queueing, sheduling and prioritizing, and for resource monitoring.

Data Management TeraGrid employs SRM for storage allocation, Globus' Global Access to Secondary Storage (GASS) for simplification of data access and GridFTP for data transfer.

Monitoring and Performance GT's Monitoring and Discovery System (MDS) presents information on clusters, jobs and users on the TeraGrid.

Fabric The TeraGrid has 20 teraflops computing power and 1 petabyte storage, distributed over nine sites, by 2004. The components of TeraGrid are connected through a 40 gigabit per second network. [27]

The TeraGrid computing infrastructure has been established through three programs within the NSF initiative, TeraScale:

- A 6 TFLOPS computational resource at the Pittsburgh Supercomputer Center, the TeraScale Computing System (TCS-1),
- a 15 TFLOPS in the process, through the Distributed Terascale Facility (DTF), which is a computational grid composed of resources from Argonne National Laboratory, Caltech, the National Center for Supercomputing Applications (NCSA), and the San Diego Supercomputer Center (SDSC),
- a proposal to combine resources from TCS-1 and DTF into the Extensible TeraScale Facility (ETF).

7 Projects Beyond Prototyping

The following projects are involved in different areas of distributed computing and provide a production quality environment for scientific or other computational applications.

7.1 Academic

The following projects employ grid computing technology for scientific computation and Grid research projects.

7.1.1 Enabling Grids for E-science in Europe

The Enabling Grids for E-science in Europe (EGEE) project is a European Grid project that aims to provide computing resources to European academia and industry. Working areas include the implementation of a European grid infrastructure, development and maintenance of grid middleware and training and support of grid users. Many of its activities are based on experiences from the European DataGrid (EDG) project.

The project was launched 1 April 2004.

URL http://public.eu-egee.org/

Clients and Application Areas EGEE wishes to involve industry and academia both as partners, users and providers. The main applications for EGEE are the LHC experiments. EGEE has chosen as pilot projects, LCG and Biomedical Grids. [13]

Goals and Objectives By the end of the second year of the project aims to have 8000 CPUs and 3000 active users from at least five different disciplines.

The project intends to, in addition to demonstrating quantitatively the benefits of grid computing, to achieve qualitative improvement by introducing new functionality to grid computing platforms. [13] The main working areas of the project are the operation of a production grid computing infrastructure and services and supporting the end users.

The LCG and EGEE are working towards the common goal of being able to let jobs, that today can run on LCG, run on the grid of EGEE. In order to coordinate and facilitate the merging of their respective computing platforms, LCG has organized ARDA (A Realisation of Distributed Analysis for LHC), whose mission is to advice on the work on the projects' respective middleware, possibly converging towards a common middleware. ARDA has ten members with software developing backgrounds; among these are four involved in LCG and another four involved in EGEE.

Project Status Basic Grid services and middleware is deployed. [13] Data Challenges 2004 have revealed an overall Grid efficiency of 90-95% on the LCG-2 Grid. [4]

Security The security team of EGEE will lead the EU-GridPMA project. EU-GridPMA has established a body of trust between EU and national grid computing projects and international partners. The involved partners are commited to a common policy which they agree to operate according to or better than.

Access control and accounting will be managed by VO authorization mechanisms, while allowing transparent authorization. [7]

Community and Work Force The project has 600 man years allocated over two years. 48% is allocated to activities related to Grid services, 24% to the re-engineering of middleware and 28% for dissemination and training. [13]

Application Environment and Tools The middleware will integrate middleware from the VDT, the EDG and the AliEN project. [13]

Fabric The infrastructure of the EGEE computation grid will be built on the EU Research Network GEANT and national research and education networks across europe. The amount of CPUs will grow from 3000 CPUs at the beginning of the project to over 8000 by the end of the second year. [13], [11]

7.1.2 Berkeley Open Infrastructure for Network Computing

The BOINC project is developing distributed computing technology for employment on regular computers connected to Internet.

URL http://boinc.berkeley.edu/

Clients and Application Areas The BOINC project envisions deployments in many scientific fields such as astronomy, biochemistry, physics simulation, biengineering, weather modeling, but BOINC places no restrictions on the application of the software. The following projects have interest in the BOINC software:

- SETI@Home - the project is currently using BOINC in searching for extraterrestrial life. The computation power is used for analyzing signals from space.
- Predictor@home - the project is currently using BOINC for predicting protein structures.
- ClimatePrediction.net - the project is currently testing BOINC for climate prediction.
- Einstein@Home - will use BOINC for computation in the LIGO experiment.
- LHC@Home - will use BOINC for experiments related to the LHC.

Project Status The BOINC architecture is currently deployed in the SETI@Home distributed computing project as a production computing environment. SETI@Home currently has more than 25000 users through BOINC. [3] Other projects are using BOINC in a development environment.

Data Management Any relational database may be used for the storage of data.

Licensing The license of the BOINC server and client software will in January 2005 become GPL. The BOINC API license will become either GPL or LGPL. Currently, the BOINC software is licensed after a modified Mozilla Public License scheme, which prohibits commercial use until the transition to GPL.

Job Submission and Scheduling The submission of jobs to a BOINC network is performed either through the BOINC API or by building a wrapper for the program to run.

Monitoring and Performance Database tools and web tools exist for monitoring the BOINC infrastructure.

Application Environment and Tools BOINC has a GUI for Windows and a CLI for Linux, Windows and Mac OS. Jobs that run on BOINC, interface through the BOINC API; it has functionality for graphics and message passing, among others.

Fabric BOINC may be deployed on any set of interconnected computers running Linux, Windows or Mac OS.

Problem Characteristics The problems that currently run on BOINC infrastructure amount to 0.1 to 100 TFLOPS.

7.1.3 NorduGrid

The NorduGrid project started as a Nordic grid computing collaboration between the Scandinavian countries, Norway, Finland, Sweden and Denmark, to be employed on NorduNet, hence the name. NorduGrid has long operated in a production environment, and, despite the name, is open for collaboration with parties outside of Northern Europe.

URL http://www.nordugrid.org/

Clients and Application Areas NorduGrid is open for participation from all countries and consists of academic research sites from all over the world. The research areas are mainly within HEP, but the grid is not restricted to these applications. The core users of NorduGrid are ATLAS users. [5]

Goals The goals of the NorduGrid project include implementing a easy to use and maintain, functional grid computing system, while adhering to de facto standards to a feasible extent.

Project Status The NorduGrid has been operational in a production environment since July 2002. The grid is continuosly evolving through middleware upgrades and dynamical resource changes. [5]

Licensing The software produced and used by NorduGrid is GPL licensed.

Application Environment and Tools The middleware of NorduGrid is mainly based on Globus grid software, where services of GT are replaced by services that are produced by the NorduGrid project, in which only the framework and libraries of GT are used. NorduGrid's self-produced alternative service solutions consist of Grid-manager, Jobplugin, Gridftp server, Userinterface &

Broker, Information model and XRSL, among others. The project has produced the Advanced Resource Connector (ARC) middleware, which is built upon standard open source solutions such as the OpenLDAP, OpenSSL, SASL and GT2 libraries.

Security NorduGrid uses GSI for security.

Job Submission and Scheduling NorduGrid's Grid Manager replaces GT's GRAM. For job descriptions, NorduGrid have developed xRSL, which is an extension of Globus' RSL. Each user has her own job broker for submitting jobs.

Data Management ARC introduces the concept of Storage Elements (SE). SEs are GridFTP-enabled. Clients use GSI enabled FTP transfer. All data transfers use GridFTP.

For data storage, NorduGrid uses The NorduGrid Replica Catalog, which is a slightly patched Globus Replica Catalog.

Monitoring and Performance The NorduGrid information system uses MDS 2.2 to present information on clusters, jobs and users. [5]

Fabric The computing resources connected to NorduGrid are mostly clusters from small test clusters to university production-class clusters. As well as two world-class Top500 clusters, Monolith and Seth, NorduGrid uses ordinary idle workstations at institutions as computing resources. [5] Currently, 40 to 42 sites amounting to 2100 to 2900 CPUs are contributing to the NorduGrid. [1]

7.2 Commercial

The following commercial projects produce proprietary products and are mostly driven by industry interests.

7.2.1 Liberty Alliance

The Liberty Alliance defines Liberty, which specifies a framework for a "single sign-on" system, mainly for web-based services, similar and competitive to Microsoft's Passport, only with a decentralized architecture. The organization provides the technology, expertise and certifications for employing such an infrastructure.

URL http://www.projectliberty.org/

Clients and Application Areas The project has more than 150 members from various industries from many countries and is open for all.

7.2.2 Xgrid

Xgrid is a self-configurating and easy to use and set up technology to facilitate heterogenous distributed computing over platforms such as Mac OSX and Linux.

URL http://www.apple.com/acg/xgrid/

Clients and Application Areas Some scientists are beginning to see a potential in deploying a Xgrid infrastructure. A ”community supercomputer” project, Wolfgrid, is running at NCSU, which lets Mac OSX users join an any-purpose Xgrid computing infrastructure. Another project, Xgrid@Stanford, is running at Stanford University. [2]

Goals and Objectives The project aims to provide an easy to use and fast to set up solution for grid and cluster computing. The software is to be built into the next version of Mac OSX, Tiger, and so to include computers running this OS as a potential part of a distributed or parallel computing platform.

Application Environment and Tools Xgrid has both a GUI and a CLI and supports MPI.

Fabric Xgrid uses interconnected computers running Linux or Mac OSX to form a cluster. A feature of Xgrid, Rendezvous, facilitates the set-up of the cluster through automatic discovery of available resources. [18]

8 Conclusion

Table 1 shows a summary of the selection of projects. The summary gives indicators of the projects' sizes, through their partners, funding and fabric attributes, and capabilities, through their job efficiencies. With help from these attributes, a subjective estimate is made of the likely impact of the projects.

Name	URL	Lead Partner, Country	Size Indicators	Capability Indicators	Subjective Evaluation
GRACE	http://www.grace-ist.org/	Telecom Italia, Italy	5 partners	Depends on underlying Grid - currently LCG-2	The success of GRACE depends on the usefulness and usability of the product.
Grid2003	http://www.ivdgl.org/grid2003/	US CMS and USATLAS	5 partners, 22 institutions, 27 sites, 2700 CPUs	90% efficiency in well run sites	High impact in US HEP
TeraGrid	http://www.teragrid.org/	NSF, USA	USD 98 mill. in funding over three years, 10 mill. in 2003; 14 partners, of which 10 are institutions, 9 sites; 20 TFLOPS	Most sites are unstable; a few sites have > 90% efficiency on test units.	Medium to high impact in general grid computing
EGEE	http://public.eu-egee.org/	CERN, EU	€32 mill. in funding from EU, similar amount in other funding; 71 partners; between 3000 and 8000 CPUs	90-95% efficiency on LCG-2	High impact in HEP and the general development of grid computing
NorduGrid	http://www.nordugrid.org/	Nordic universities	40 to 42 sites amounting to 2100 to 2900 CPUs	Very high efficiency	Underrated and wrongfully ignored
BOINC	http://boinc.berkeley.edu/	SETI@home, USA	Limited by the popularity of the employing project. SETI@home has more than 25000 users.	High job completion rate through redundancy, low utilization.	Increases the availability of distributed computing; moderate to high impact in distributed computing in general.
Xgrid	http://www.apple.com/acg/xgrid/	Apple, USA	High potential - essentially all Mac OSX Tiger users		Increases the availability of distributed computing; moderate to high impact in distributed computing in general.

Table 1: A summary of the projects.

A Projects for Further Review

The following projects are set aside for further review.

A.1 DEISA

URL http://www.deisa.org/

A.2 Butterfly.net

The Butterfly.net project employs grid technologies to operate online games.

URL http://www.butterfly.net/index.html

Project Status The project is currently adapting to running games.

Application Environment and Tools The Grid architecture uses GT2.2.

A.3 Platform Computing

URL http://www.platform.com/

A.4 Corporate Ontology Grid

The Corporate Ontology Grid (COG) is a commercial project that aims to apply grid technology to integrate corporate information via ontological modeling.

URL http://www.cogproject.org/

A.5 Avaki

URL http://www.avaki.com/

A.6 Shibboleth

URL http://shibboleth.internet2.edu/

A.7 The GridPP Project

The GridPP project is building a computing grid for particle physics, in a collaboration with particle physicists and computer scientists from the UK and CERN.

URL http://www.gridpp.ac.uk/

Clients and Application Areas The intended application for the GridPP grid is particle physics.

Project Status The project is currently developing a prototype grid.

A.8 GridBus

The GridBus project is engaged in developing open-source specifications, architecture and a reference Grid toolkit implementation of service-oriented grid and utility computing technologies for scientific and commercial applications.

URL http://www.gridbus.org/

A.9 Biomedical Informatics Research Network

The Biomedical Informatics Research Network (BIRN) aims to bring a distributed computing infrastructure to biomedical science in the USA, in order to stimulate the collaborations between the involved parties.

URL http://www.nbirn.net/

A.10 Earth System Grid

The Earth System Grid (ESG) aims to integrate supercomputers with large-scale data and analysis servers located at different american labs and research centers to create a distributed computing environment for climate research in the USA. [17]

URL https://www.earthsystemgrid.org/

Application Environment and Tools The ESG project utilizes GT.

A.11 National Grid

The National Grid (NG) project aims to build a secure and efficient computational infrastructure for use in scientifical engineering and biomedical research and development. [24]

URL http://www.ngp.org.sg/

Project Status The first phase of the project is named "National Grid Pilot Project" (NGPP). The project has entered NGPP, which has included the implementation of a physical network and courses in grid technology.

A.12 ChinaGrid

URL http://www.chinagrid.com/

A.13 GridLab

The GridLab project is a European grid technology development project that develops services and tools for grid computing that will provide capabilities such as dynamic resource brokering, monitoring, data management, security, information and adaptive services. [22]

URL http://www.gridlab.org

Project Status The project is developing the Grid Application Toolkit (GAT). A prototype GAT is ready.

Licensing GridLab software is licensed after the "GridLab Open Source License" (GL License). This license allows for use of the software for any purpose by anyone, with the only obligation beeing that redistributions inherit the license.[6]

Application Environment and Tools GridLab is developing the Grid Application Toolkit which contains a set of tools and GAT API. The goal of the GAT API is for it to have interfaces to at least the fundamental services of the GridLab architecture and to organize and simplify the access to these services through the API. The API is planned to include interfaces to basic grid services of the GAT such as job submission and management, resource management, data storage, monitoring and security. The GAT consists of GAS (Grid Authorization Service) (see "Security"), GridSphere - a grid portal development framework, GRMS (Grid Resource Management and Brokering Service) (see "Job Submission and Scheduling"), Data Access and Management (Grid Services for data management and access), iGrid (GridLab Information Services), Delphoi (Grid Network Monitoring & Performance Prediction Service), Mercury (Grid Monitoring infrastructure), Visualization (Grid Visualization Services), Mobile Services (Grid Services supporting wireless technologies). [22]

A.14 DataTAG

DataTAG was a research and technological development grid project that focuses on transatlantic data transfer. It employs a collaborative grid testbed between Europe and USA.

Clients and Application Areas The grid testbed will focus on advanced networking issues.

URL http://datatag.web.cern.ch/datatag/index.html

A.15 Particle Physics Data Grid

The Particle Physics Data Grid (PPDG) project is a field specific grid project located in the USA.

Clients and Application Areas PPDG is a collaborative project among physicists and is to provide distributed computing for high energy and particle physics.

URL http://www.ppdg.net/

Project Status The project is currently running as a pilot project. The funding partner, SciDAC, has approved a two year extension for until 2006, to the project. [26]

A.16 NEESgrid

NEESgrid is an American field specific grid infrastructure being developed by Network for Earthquake Engineering Simulation (NEES) for earthquake engineering simulations. NEESgrid's system architecture aims to be able to link up to twenty earthquake laboratories into one collaboratory. This collaboratory's purpose is to allow remote operation and observation of experimental equipment, allow remote and shared viewing of real-time data, provide sharing of output from numerical simulations and joint document editing, viewing of data visualizations and operation and monitoring of physical and numerical simulations. [29]

URL http://www.neesgrid.org/index.php

Project Status NEES grid has, through experiments proved the effectiveness of their architecture, which is still in development.

Application Environment and Tools The NEESgrid software is to integrate significant elements of the NSF Middleware Initiative (NMI) software system. [29]

Fabric The fabric of the NEESgrid is intended to consist of the involved earthquake research sites' experimental facilities, computer resources, campus and wide area networks, and a central monitoring facility. [29]

A.17 CrossGrid

The CrossGrid project aims to extend a grid environment across European countries and to new application areas. [19]

URL http://www.eu-crossgrid.org/

Clients and Application Areas CrossGrid aims to extend a grid environment to include applications such as medicine, environmental protection, flood prediction, and physics analysis.

Application Environment and Tools The Crossgrid project plans to build a software Grid toolkit, which will include tools for scheduling and monitoring resources. [19]

A.18 LHC Computing Grid

The LHC Computing Grid (LCG) project is building and maintaining a grid infrastructure for the high energy physics community in Europe, USA and Asia. The main purpose of the LCG is to handle the massive amounts of data produced from the LHC (Large Hadron Collider) experiments at CERN.

URL http://lcg.web.cern.ch/LCG/

Clients and Application Areas The applications for the LCG are, among others, high energy physics, biotechnology and other applications that EGEE brings in. The main application is the gathering of data from the LHC experiments ATLAS, CMS, Alice and LHCb.

Problem Characteristics The LCG project has estimated that an annual amount of about 10 petabytes of data, gathered from the LHC experiments, need to be stored in the LCG. For the analysis of the data, the LCG project has estimated that the LCG will need the processing power of in the order of 100 000 CPUs.

Fabric The amount of computers the centers participating in the LCG project have to manage is so massive that manual maintenance of the installed software of these computers is too labour intensive. Also, in such an amount of components, the failure of one component should automatically be overridden and not affect the overall operability of the system. The LCG project has designed a fabric management software which is to automate some of these tasks.

B Partners and Funding

The following shows from where projects receive funding and with whom they cooperate.

B.1 Grid2003

The Grid2003 facility is operated jointly by US projects,

- iVDGL,
- GriPhyN,
- PPDG,

and is principally led by the US participants in the LHC experiments ATLAS and CMS, US CMS and USATLAS. 22 institutions participate in this project.

The project is supported by the National Science Foundation and the Department of Energy. [23] [14]

B.2 BIRN

The BIRN project is a National Institutes of Health initiative.

B.3 TeraGrid

NSF initially funded USD 53 million to the four sites,

- National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign,
- San Diego Supercomputer Center (SDSC) at the University of California, San Diego,
- Argonne National Laboratory in Argonne, IL,
- Center for Advanced Computing Research (CACR) at the California Institute of Technology in Pasadena.

In October 2002, NSF supplemented USD 35 million in funding, and Pittsburgh Supercomputing Center (PSC) at Carnegie Mellon University and the University of Pittsburgh joined the TeraGrid project.

In September 2003, NSF awarded USD 10 million, and

- Oak Ridge National Laboratory (ORNL), Oak Ridge, TN,
- Purdue University, West Lafayette, IN,
- Indiana University, Bloomington,
- Texas Advanced Computing Center (TACC) at The University of Texas at Austin

joined the project.

Other partners are

- Myricom,
- Sun Microsystems,
- Hewlett-Packard Company,
- Oracle Corporation.

[27]

B.4 Corporate Ontology Grid

The following are involved in the COG project:

- Unicorn
- Centro Ricerche Fiat
- LigicDIS
- Universitt Innsbruck
- IST

B.5 EGEE

EGEE is funded by the European Commission. The EU has funded the project with €32 mill., while partners' funding amount to a similar number. [13] EGEE's partners are: [15], [8]

- TERENA
- European Organization for Particle Physics
- Institut fr Graphische und Parallele Datenverarbeitung der Joh. Kepler Universitat Linz
- Institut fr Informatik der Universitaet Innsbruck
- CESNET, z.s.p.o.
- Budapest University of Technology and Economics
- Eotvos Lorand University Budapest
- KFKI Research Institute for Particle and Nuclear Physics
- Magyar Tudomanyos Akademia Szamiastecnikai es Automatizalasi Kutato Intezet
- Office for National Information and Infrastructure Development
- Akademickie Centrum Komputerowe CYFRONET akademii Gorniczo-Hutniczej im.St. Staszica w Krakowie
- Warsaw University Interdisciplinary Centre for Mathematical and Computational Modelling
- Institute of Biorganic Chemistry PAN, Poznan Supercomputing and Networking Center
- Ustav Informatiky, Slovenska Akademia vied
- Jozef Stefan Institute
- The Provost Fellows and Scholars of the College of the Holy and Undivided Trinity of Queen Elizabeth near Dublin
- Council for the Central Laboratory of the Research Councils
- The University of Edinburgh
- Particle Physics and Astronomy Research Council
- University College of London
- Commissariat l'Energie Atomique, Direction des Sciences de la Matire
- Compagnie Genrale de Grophysique
- Centre National de la Recherche Scientifique
- CS Systme d'Information Communication & Systemes

- Centrale Recherche S.A.
- Deutsches Elektronen Synchrotron
- Deutsches Klimarechenzentrum GmbH
- Fraunhofer-Gesellschaft zur Frderung der Angewandten Forschung e.V.
- Forschungszentrum Karlsruhe GmbH
- Gesellschaft fur Schwerionenforschung GmbH
- DATAMAT S.p.A.
- Istituto Nazionale di Fisica Nucleare
- Trans-European Research and Networking Association
- Vrije Universiteit Brussel
- Faculty of Science University of Copenhagen
- University of Helsinki
- Foundation for Fundamental Research on Matter
- Stichting Academisch Rekencentrum Amsterdam
- Universiteit van Amsterdam
- University of Bergen
- Vetenskapsrådet, The Swedish Research Council
- Institute of High Energy Physics
- Institute of Mathematical Problems of Biology of Russian Academy of Sciences
- Institute of Theoretical and Experimental Physics
- Joint Institute for Nuclear Research
- Keldysh Institute of Applied Mathematics of Russian Academy of Sciences Moscow
- Petersburg Nuclear Physics Institute of Russian Academy of Sciences
- Russian Research Centre "Kurchatov Institute"
- Skobeltsyn Institute of Nuclear Physics of Moscow State University
- Central Lab. for Parallel Processing, Bulgarian Academy of Sciences
- University of Cyprus
- Greek Research and Technology Network
- Tel Aviv University

- National Institute for Research and Development in Informatics
- Laboratorio de Instrumenta o e Fsica Experimental de Part culas
- S.A.X. Centro de Supercomputacin de Galicia
- Consejo Superior de Investigaciones Cientificas
- Institut de F sica dAltes Energies
- Instituto Nacional de Tcnica Aeroespacial
- Universidad Politecnica de Valencia
- University of Chicago
- University of Southern California, Marina del Rey
- The Board of Regents for the University of Wisconsin System
- Royal Institute of Technology - Center for Parallel Computers (PDC)
- Ente per le Nuove Tecnologie, l'Energia e l'Ambiente
- Universit degli Studi della
- Universit degli Studi di Lecce
- Universit degli Studi di Napoli Federico II
- Delivery of Advanced Network Technology to Europe Limited
- Verein zur Foerderung eines Deutschen Forschungsnetzes e.V.
- Consortium GARR

TERENA is the lead parter in dissemination for EGEE. [8]

The UK National e-Science Centre leads the NA3 User Training and Induction work package, which involves 22 of EGEE's partners. [9]

EGEE will participate in the eInfrastructure Reflection Group (eIRG). [10]

B.6 GRACE

GRACE is funded under the IST 2002 Framework 5, CPA9: "Grid technologies and their applications".

The project has the following partners:

- Telecom Italia
- CERN
- GL 2006 Europe
- Sheffield Hallam University - School of Computing and Management Sciences
- Stockholm University Library
- Stuttgart University Library

B.7 Particle Physics Data Grid

The following institutions are involved in the PPDG project:

- Stanford Linear Accelerator
- Lawrence Berkeley National Laboratory
- Fermi National Accelerator Laboratory
- Argonne National Laboratory
- Brookhaven National Laboratory
- Thomas Jefferson National Accelerator Facility
- California Institute of Technology
- Univeristy of Wisconsin
- University of California San Diego

B.8 NEESgrid

The following institutions are involved in the NEESgrid project:

- National Center for Supercomputing Applications
- University of Illinois at Urbana-Champaign
- Argonne National Laboratory
- Information Sciences Institute
- University of Southern California
- School of Information, University of Michigan
- Department of Civil and Environmental Engineering, University of California, Berkeley
- ERC, Center for Computational Sciences, Mississippi State University
- Pacific Northwest National Laboratory
- Department of Civil and Environmental Engineering, Stanford University
- Department of Civil Engineering, Washington University in St. Louis

B.9 DataTAG

The following institutions are involved in the DataTAG project: DataTAG-Funded Partners:

- CERN
- INFN, Italian National Institute for Nuclear Research
- INRIA, Institut National de Recherche en Informatique et en Automatique
- PPARC
- University of Amsterdam

U.S. Partners:

- California Institute of Technology
- StarLight
- Argonne National Laboratory
- Northwestern University
- University of Illinois at Chicago
- University of Michigan

Collaborating Networks:

- CANARIE Inc.
- DANTE
- GANT
- ESnet
- Internet2 and Abilene
- MB-NG
- SURFnet

Collaborating Institutes:

- Stanford Linear Accelerator Center
- Fermi National Accelerator Laboratory
- Lawrence Berkeley National Laboratory

B.10 CrossGrid

The CrossGrid project is co-financed by IST.

B.11 The GridPP Project

References

[1] Nordugrid. http://www.nordugrid.org/

[2] wolfgrid. http://packmug.ncsu.edu:16080/wolfgrid/index.html

[3] Boincstats. http://www.boincstats.com/ July 2004.

[4] A.Fanfani. Grid experiences in cms, Jul 2004.

[5] B.Knya. The nordugrid production grid infrastructure, status and plans. 2003.

[6] G. Consortium. Gridlab open source license.

[7] EGEE. Jra3 - security.

[8] EGEE. Na2 dissemination and outreach.

[9] EGEE. Na3 - user training and induction.

[10] EGEE. Na5 - policy and international cooperation.

[11] EGEE. Sa2 - network resource provision.

[12] EGEE. Egee na4: Applications identification and support grace proposal for collaboration. 2004.

[13] E. G. for E-science in Europe. Egee fact sheet.

[14] Grid2003. Grid2003.

[15] http://public.eu egee.org/. Enabling grids for e-science in europe.

[16] http://setiathome.ssl.berkeley.edu/polls.html. Seti@home poll.

[17] https://www.earthsystemgrid.org/. Earth system grid.

[18] http://www.apple.com/acg/xgrid/. Xgrid.

[19] http://www.eu crossgrid.org/. The crossgrid project.

[20] http://www.grace ist.org/. Grace - grid search & categorization engine.

[21] http://www.gridforum.org/. Global grid forum.

[22] http://www.gridlab.org/. Gridlab.

[23] http://www.ivdgl.org/grid2003/. Grid2003.

[24] http://www.ngp.org.sg/. National grid singapore.

[25] http://www.nua.ie/surveys/how_many_online/. How many online?

[26] http://www.ppdg.net/. Particle physics data grid.

[27] http://www.teragrid.org/. Teragrid.

[28] e. a. I. Foster. The grid2003 production grid: Principles and practice. IEEE, 2004.

[29] C. Kesselman, R. Butler, I. Foster, J. Futrelle, D. Marcusiu, S. Gulipalli, L. Pearlman, and C. Severance. Neesgrid system architecture version 1.1. 2003.

[30] H. Mikkonen and T. Nissi. Gsi and liberty alliance framework: Goals, architectures and feasibility study for integration. 2003.

[31] U. of Chicago and T. U. of Southern California. Globus toolkit public license version 2. 2003.