Newsletter n°2 October 2007

The BEgrid Tribune

Contact

BEgrid

BELNET

Wetenschapsstraat, 4
Rue de la Science, 4
1000 BRUSSELS

Tél: +32 2 790 33 94

e-mail: begrid@belnet.be

Editorial

Dear Reader,

Here is the second issue of the BEgrid Tribune. We have chosen to dedicate a large part of this newsletter to a description of the BEgrid infrastructure and services. There is a general overview of BEgrid and the services offered. Then each BEgrid participant tells about their resource contribution, their users, applications running on the grid, …

Of course we did not want to forget or ignore other topics so we have a contribution about grid research, about two interesting projects in BEgrid and a look back on the grid event of June.

Again, do distribute the BEgrid Tribune to your colleagues, friends, bosses, … so that as many researchers as possible can learn about grid computing and the advantages it can represent for their research. If you have any questions about BEgrid then do not hesitate to contact us.

Rosette Vandenbroucke

BEgrid: the Belgian computing grid infrastructure for research

Rosette Vandenbroucke
BEgrid Coordinator

History

The BEgrid initiative started early 2003 with the goal to stimulate knowledge, use and research in grid computing in Belgium and to bring Belgium on the European map of grid computing. The initiative attracted quite some interest in the research world and there were more then 100 attendees at the announcement meeting. Soon interested institutions met to discuss and to set up a small pilot infrastructure. Different from other grid initiatives in Europe, the BEgrid initiative had no money to build a new infrastructure and from the beginning depended on the contribution of the participants. The Flemish Government was very interested in the initiative and contributed via a project providing funding for equipment. The institutions that participated from the start were the Katholieke Universiteit Leuven, the Universiteit Antwerpen, the Universiteit Gent, the Université Libre de Bruxelles and the Vrije Universiteit Brussel. They were joined later by the Vlaams Instituut voor de Zee.

Middleware

The grid middleware chosen for BEgrid was the production middleware of the main European grid infrastructure DataGrid and later EGEE (Enabling Grid for EsciencE). The current EGEE production middleware is gLite 3.0 and implemented on the whole of BEgrid.

Current Status

BEgrid includes for the moment 650 CPUs worker nodes together. Thirty Tbytes of storage are publicly available, however more can be attributed if necessary. There are more than 500 valid certificates issued.

BEgrid Services

BEgrid offers several services for the grid community including the issue of grid certificates, the operation of the main grid services, the Quattor installation service and the organisation of grid courses. Grid certificates are needed by users and servers for authentication and authorization. Those grid certificates are valid throughout the world on grid infrastructures of all organisations and projects that are member of IGTF (International Grid Trust Federation). The main grid services offered are a VOMS (Virtual Organisation Management Server), a User Interface to give users access to BEgrid, the operation of the BDDI, a database of all resources in BEgrid. The Quattor installation service enables BEgrid participants to install their grid cluster with a minimal effort. Finally a series of grid courses are scheduled: Hands on BEgrid, BEgrid for Computer Scientists and a grid installation course.

Virtual Organisations

Virtual Organisations (VOs) are an important concept in grid computing. A VO is a group of users that share a certain part of the grid infrastructure. VOs can be defined locally on a grid cluster, on the whole grid or on part of the grid. BEgrid supports four VOs: BEtest (for novice users), BEapps (for production), BEcms (for researchers participating in the CMS experiment) and NCF (the Dutch VO).

User access to BEgrid

Belgian researchers can request to use the BEgrid infrastructure. If their home institution is already a participant in BEgrid then they can get in a short timeframe a grid certificate, a membership to a Virtual Organisation and an account on a User Interface. Access for users of institutes that are not yet participants in BEgrid can also be organized.

Connection to other infrastructures

BEgrid is not an isolated infrastructure but is already connected to the EGEE (Enabling Grids for EsciencE) infrastructure and to NLgrid. Belgian jobs that are submitted in the VO BEapps can be executed in Belgium as well as in the Netherlands. The corresponding Dutch VO is NCF.

BEgrid projects

BEgrid is also trying to incorporate existing computing resources. A project to migrate the PCs of the student PC rooms to grid worker nodes during nights, weekends and holidays is successfully terminated and its results are going now in production. A project to turn a common user PC into a grid worker node with a minimal effort has started. At the high end of the grid a project with HP, Intel and Voltaire is signed to incorporate a small HPC machine as a resource in the grid.

Future of BEgrid

More and more interest is shown in BEgrid and before the end of the year at least three new major institutions will have joined. Researchers slowly find their way to BEgrid and are often enthusiast when they see the new possibilities it brings to their research. So BEgrid will grow as well in available resources as in users and applications.
On the European scene BEgrid is considered as the Belgian national grid for research and is involved in the study for the definition of building a sustainable European Grid Infrastructure.

FUNDP steps towards BEgrid

Badr Jabari, Jean-Pol Vigneron

In order to enhance the computing power and storage capacity in the University of Namur (more precisely, Solid-State Physics Laboratory (LPS) and Faculty of Computer Science), we have chosen to rely on the increasing importance and potential offered by the grid computing. In addition, we wanted to have a profitable open cluster; this is why we preferred to join the Belgian Grid for Research: BEgrid to test and integrate this new platform.

The current Interuniversity Scientific Computing Facility (iSCF) resources contains: 80 Sun Fire V60x, 19 Sun Fire V20z and 2 Sun Fire V880 as shown in the following figure.

The cluster has in total 101 nodes with 214 microprocessors, 1103.4 GFlops, 480 GB of memory and

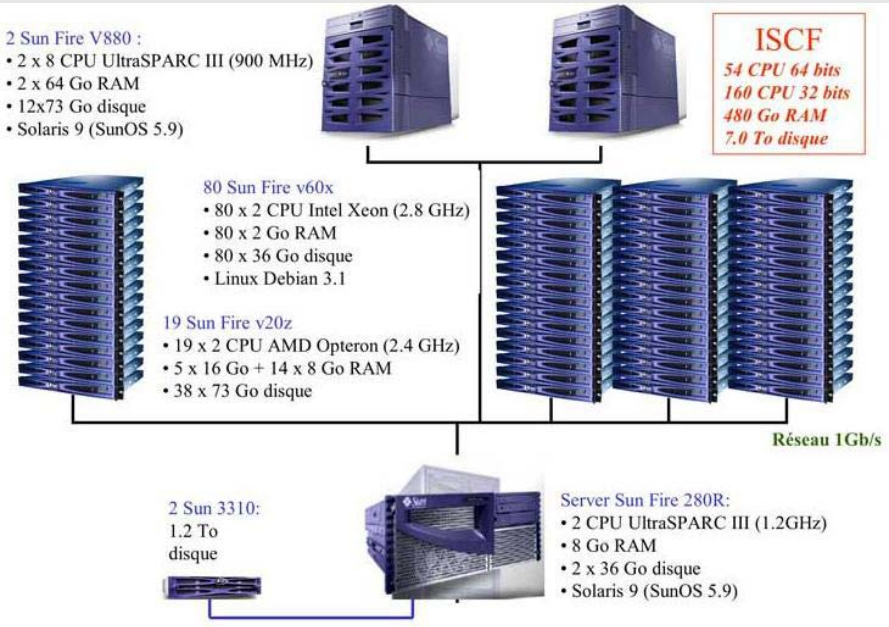

can store up to 5.7 TB of data. There are mainly three main system architectures: UltraSPARC® III @ 900 MHz, Intel® Xeon™ @ 2.8 GHz and AMD Opteron™ 250 @ 2.4 GHz. The nodes operate using two operating systems which are SunOS or GNU/Linux. They are interconnected by a 1Gbps Ethernet switched networks. We set out a Sun Grid Engine Middleware to manage all these resources.

Our system is in use by different researchers and scientists for different applications such compiling programs (using mainly GCC, Intel Fortran, Intel C++ compiler…), running intensive scientific calculations (using Abinit, Gaussian, Turbomole…).

We have chosen to join BEgrid progressively to simplify our testing and evaluation of the future platform and familiarize our users with this new working environment. We have already started by connecting 7 machines (1 Computing Element and 6 Working Nodes). Finally, we want to mention that our cluster is a third level one.

BEgrid node at the Hogeschool Antwerpen

Filip Van der Schueren

The Hogeschool Antwerpen (HA) operates a very small test site on BEgrid. It is used to test installation procedures and to allow some research into grid enabling of applications.

An example of a research project is the gridification on BEgrid of the rendering process for the design of 3D animations. Through this we investigated the feasibility of business applications to run on grid infrastructures. It is the outgrow of a Master thesis, realized at Hogeschool Antwerpen, department. Elektronica-ICT by S.Dirix and K.W. Cheung.

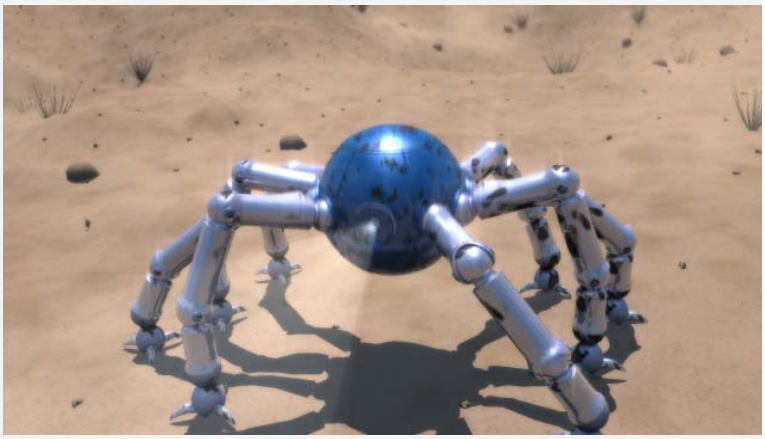

Since quite a time 3D computer graphics are often used to add extra features and special effects to movies. When talking about the creation of 3D computer graphics we usually divide the process into three basic steps. The first step is the modelling where we create or describe objects by using some kind of 3D modelling tool like 3D Studio Max, or Blender. The next step is placing our objects within a scene. In this scene we network dedicated to rendering. Render farms have a few downsides however, they are costly to hire or to install and they have limited scalability. Here we make the step from a cluster to grid computing.

In our project we described a solution to 3D animation rendering on grid-based computer networks. We created a web-based user interface where a user can upload scenes to BEgrid, start the rendering of the scene, monitor running jobs and the possibility to preview the finished frames. Next to this, a plug-in for Blender was created, which allowed the user to do on-demand rendering while he was modelling. During tests, promising results were accomplished in speedup time as can be seen in the graph. In continuing research more attention will be paid to the optimal job submission, the monitoring of jobs in which failures get well intercepted. Furthermore, a final version web-based user interface will be built.

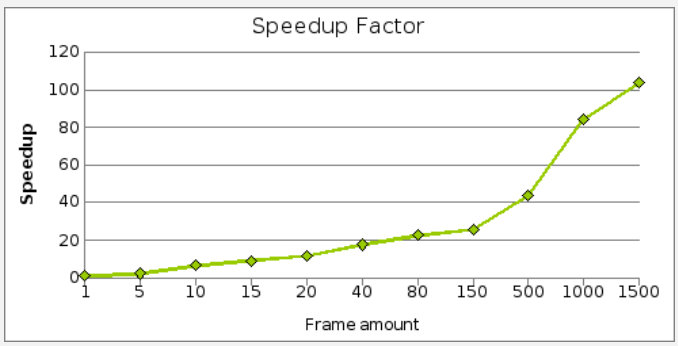

The BEgrid node at the KU Leuven

Wim Obbels, Anne-Marie De Meyer

In June 2004, the KU Leuven installed a first test cluster with grid middleware on 20 old PC's. It had all grid services, and one of these, the central VOMS server that manages user authorizations on BEgrid, remained in service until a few months ago. But, for real computing power, in September 2004 a 'real' cluster, with new high-performance hardware was installed. Still, it had initially only 12 work nodes (24 CPU's). Since then, it has been expanded two times, and now has 264 cores. The Storage Element has now several Terabytes of disk space available, up from the 12GB IDE disk in the initial setup ... ! Since beginning of 2006, the KU Leuven BEgrid site is registered in the pan European EGEE grid infrastructure.

The first cluster had EDG middleware running on the -already at that time- very old Redhat 7.3 Linux distribution. Installation and configuration was done through the LCFG software, later replaced by YAIM, and now we are using the central Begrid Quattor repository for the installation and configuration of the cluster. This allows us to more easily keep up to date with new versions of the KU Leuven, so that users of that cluster can also submit jobs onto the Grid cluster, and the other BEgrid sites.

The past three years the Grid cluster has been used for both scientific calculations and investigations in Grid technology. So was the Grid cluster used for several Master theses, including some by Erasmus students end also for large scale simulations in eg Hydraulics.

The BEgrid node at the University of Antwerp

The BEgrid node at the University of Antwerp has at present a cluster of 77 nodes with HyperThreaded PIV CPU's, with 1GB of memory and 80GB local disk on-line. Soon, 8 new machines will be installed, with 2 quadruple CPU's, 16GB of memory and 250 GB of local disk space each, bringing the total to 205 cores. A small storage system of 1 TB mirrored disk space will also be installed next month.

Currently the cluster is operating with the LCG middleware of the EGEE project. Very soon gLite middleware will be used instead. Connection to EGEE is pending.

One of the key applications that use the UA Begrid node has been developed by F. Arickx of the UA, in cooperation with scientists of the Bogolyubov Institute of Theoretical Physics in Kiev, Oekraine. It deals with quantum scattering of light nuclei and nuclear fragments and calculates phaseshifts and cross sections based on ab initio microscopic quantum calculations. This means that each nucleon in the fragments is explicitly taken into account in the calculation, including the effects of the Pauli exclusion principle throughout the scattering process. The method allows the calculation of resonance properties as well as resonance wave functions making an in-depth analysis and interpretation of the results possible.

Another application that also ran extensively on the UA BEgrid node, deals with workload characterization on modern processor architectures. It was run by the ELIS group of the Ugent Characterizing computation extensive computer applications involves figuring out how the microprocessor in the system is being stressed. Commonly, this is done using hardware performance counters, by measuring metrics such as cache miss rates, branch prediction miss rates, etc. The problem with these performance metrics is that they are very dependent on the system they are

measured on, and thus hide true inherent application behaviour.

In the ELIS research group a way to characterize programs independent of the micro architecture of the system, thus allowing reasoning about inherent program behaviour across systems (for example Intel Pentium 4, Intel Core 2, Intel Xeon, ...) has been developed. This micro architecture-independent workload characterization is done using an instrumentation framework by Intel, called Pin, which allows analyzing the runtime behaviour of programs. Due to the number of programs used in that research, and the significant slowdown in total execution time caused by instrumentation, the services provided by BEgrid are a valuable tool. Because the Pin framework is provided by Intel, running it on systems supported by an Intel microprocessor was required until recently.

The BEgrid node at Ugent

Stijn De Smet

The initial installation of the Grid located at UGent started in December 2003 - January 2004. The grid consisted of 6 nodes, of which only 2 were worker nodes doing computing jobs. The other nodes were running the services required for the – at that time – EDG middleware. During the following years, three expansions with additional equipment added a total of 220 CPU's to the 2 computing nodes.

Also, the server infrastructure was expanded with 2 new servers for storage and additional services. We have another expansion planned, which will add about 100 CPU's to the current setup.

The middleware running on the grid changed over these three years. The EDG middleware running on RedHat 7.3 was replaced by the LCG middleware. With this change came the possibility to use Scientific Linux 3, and we were one of the first sites in the European grid to use this setup together with the installation management system Quattor.

By the final upgrade in 2006, we moved to the newer gLite middleware.

The Quattor system management was also replaced, and we are using a combination of kickstart, gLite's yaim (Yet Another Installation Method) and shell scripts.

During the last upgrade in 2006, we also started using Linux Vserver to replace the large 4 servers running all required gLite services together, with 10 separate vservers running on 3 servers. This allowed us to use easier yaim configs, and made it possible to upgrade one service without touching another. The vservers have proven to be very useful in the last year, when we were able to move heavily loaded vservers to different server hardware without interfering the running jobs. This way we can also run a more current Linux on the servers' hardware, and gaining from performance improvements, while keeping the distribution used for the gLite services at SL3.

The Ugent grid site is heavily used by Ugent researchers. Over the months September, October, November 2006 our site ran more than 20.000 jobs giving an efficiency of 80% over this period. During a week in September this year, our site ran 871 jobs with an efficiency of 95%. These jobs were CPU simulations, 2d and 3d parallelized electromagnetic crystal simulations and java-based network simulations and analysis of embryonic cell lineage data.

ULB-VUB Contribution

Shkelzen Rugovac , Stijn De Weirdt

Several departments of the VUB and ULB have chosen to put their computing resources into one large cluster in order to reduce the equipment overhead to a minimum and to solve, partially, the problem of a shortage of manpower to manage the grid cluster. This equipment, mainly coming from the High Energy Physics departments of both universities and from a pilot grid project of Vlaanderen, is now mainly managed by members of the IIHE (Interuniversity Institute for High Energies).

For the moment the ULB-VUB grid cluster includes about 200 cores and 50 Tbytes of storage and runs gLite 3.0.1 as the middleware. The cores are a mix of AMD and Intel type of processors. The installation of the middleware is done via the Central BEgrid Quattor Server.

A test bed consisting of 8 machines is set up to test upgrades and new features of the grid middleware.

ULB-VUB is heavily involved in High Energy Physics and is participating in CMS (Compact Muon Solenoid) one of the four experiments of LHC (Large Hadron Collider). The researchers active in this domain need grid computing to being able to contribute in the work of such large international collaborations. They not only use BEgrid resources but also EGEE resources that are attributed to the Virtual Organization "CMS". A virtual organization BEcms .exists on BEgrid and most participants have contributed resources to this VO. Their applications run in production on the grid.

Shkelzen Rugovac and Stijn De Weirdt in front of part of the ULB-VUB gridcluster

Other groups and individual researchers use BEgrid more on a project basis, resulting in peak moments of use alternating with low use of even no use. Once users have managed to run their application on the grid then they are mostly happy with the benefits brought by BEgrid.

The BEgrid node at VLIZ.

Francisco Hernandes

The grid node at VLIZ has 1.5 Tbyte of storage and a small number of processors. The node has been available to the scientific community for some years now. But though the node is on a 850 Mb backbone to Belnet the lack of processing power near the storage was a problem. It was noticed that a few processing nodes near the storage should be considered.

Also the plans to install backup tape libraries to archive the data on the grid were re-evaluated, and considered not required at this stage of the project.

So we decided to buy and install 8 double Quad processor machines to be used as worker nodes. Totaling 64 cores this small grid cluster will be more practical in use.

The installation also involves an upgrade of the VLIZ node to gLite, and should be finished somewhere in October.

Meanwhile at VLIZ we further developed the first version of our Marine Data Archive, and hope to have the link to the grid ready by the end of the year. The Marine Data Archive will allow to collect, annotate, archive and share data files by and for marine scientists working in more than 30 marine research centers in Belgium.

We would like to further develop the MDA as an interface to BEgrid.

Our collaboration with the Ukrainian Marine Hydrofysical Institute to optimize and run their models on our grid node should give results at the start of next year. The MHI runs a nested hydrodynamic model for the Black sea. A grid version of this model will allow them to run more simulations.

The models will also be made available for use in the Virtual Lab of the IODE, the oceanographic data and information exchange program of IOC.

Future BEgrid Participants

VITO is exploring the opportunities of grid-computing

Stijn Jassen, Erwin Goor, Jo Vliegen, Patrick Claes and Geert Borstlap
VITO - Flemish Institute for Technological Research

Recently, researchers at VITO started to explore the opportunities of grid computing as a solution for new challenges in their ongoing activities.

Air quality modelling

The *Air Quality Modelling* group of VITO has a huge experience in the development of computer models to determine the air quality on local, urban and regional scale. Apart from the development of models and the investigation into new improved model techniques, model applications are also implemented extensively in policy support and environmental impact studies in Flanders, Europe and abroad.

One of the spearheads of this research unit is the AURORA grid model. AURORA allows to simulate meteorological and air quality fields on an urban scale with a resolution of a few hundreds of meters. The model includes an advanced chemical module with tens of compounds. The numerical treatment

of the dispersion in the atmosphere of all chemical species and their interactions needs extensive computing resources. On a powerful Linux-pc, simulation times range from days to months depending on the configuration of the model setup. Currently, longer and higher-resolution simulations are requested by end-users and EU-legislation. As such grid computing presents itself as an interesting platform to cope with the new challenges of these model applications.

A second CPU-intensive and promising atmospheric modelling tool at VITO is the CFD- model ENVI-Met. ENVI-met is a micro scale model specifically designed for the simulation of street canyons and urban micro environments with a resolution of a few meters. The model is extended with a detailed vegetation module for the description of the impact of vegetation on the state of the lower atmosphere and the air quality. This 3D grid model is extremely time consuming and grid computing is also offering here new perspectives for upcoming challenges.

Earth Observation applications

The remote sensing centre of expertise (named TAP) within VITO conducts research, develops and operates multiple earth observation (EO) applications. Typically for these applications, the amount of data needed is huge and the data is stored in different institutions at different geographical locations.

It might take a very long time while consuming a lot of bandwidth to copy the needed data to the user's own local computer in order to run the user's applications. This situation also puts high demands on the user's storage capacity. As EO continues, satellite image datasets only grow larger, making it more and more impractical for users to work with.

With the introduction of grid technology this can be solved by bringing the processing algorithms to the data instead of having to move large amounts of data to the user's processing site. An EO grid infrastructure offers processing power, data storage and direct access to EO input images. The users could have the possibility to use existing algorithms made available to them by the EO provider, or they can run their own algorithms, or they can even use a combination of both. In addition, by coupling different grids from different EO providers, we bring all data and services together in one virtual world. Hence it's possible to build automated chains of services and data from different providers.

Another advantage is that processing of EO data can be done in near-real-time, since data processing can start as soon as the data becomes available.

The von Karman Institute Foray into Grid Computing

Raimondo Giammanco

The von Karman Institute, as from September 2007, is in the process of joining the Belgian Grid. The current HPC resources of the Institute consist of two separate clusters, a Beowulf class one, sporting 120 CPUs, and a rack mounted one powered by Infiniband connections, used for high end computations with its 44 CPUs.

For our first trials on the Grid we will use part of the Beowulf Cluster, currently attested at 80 single core PIV with 2GB of memory and 20 Intel Core Duo 2 E6600 with 4GB of memory running as native 64bit CPUs. Part of the PIV will be dedicated to testing the Grid before committing more resources. The cluster is scheduled to receive a significant upgrade during the present academic year, and this will allow to further devote resources to the Grid once all the testing is done.

Due to the peculiar nature of the Institute, the class of jobs that are currently submitted to our cluster,

and that will be later on submitted from the grid, vary from heavy duty data manipulation for post processing gigabytes of data extracted from PIV experimental images, to massively parallel numerical simulation of inductively coupled plasma flows and atmospheric re-entry vehicles simulations running on hundreds of nodes. Between these two extremes different jobs types exist, ranging from optimization procedures for turbine blades designs and parallel Large Eddy simulation codes running on dozens of CPUs.

The heavy users of parallel computational resources are noticeably PhD Candidates in the Aerospace Department and the students that they supervise during the Diploma Course at the Institute. More oriented towards arrays of serial jobs are Turbomachinery researchers and their students.

Many students arriving at the von Karman Institute are here exposed for the first time to HPC facilities, and thoroughly educated to their use and consumption: seminars are prepared, tutorial on parallel programming, shell programming and job submission held, cluster usage demystified.

For many it will be as well the first contact with grid computing, and in this way the Institute will continue in its now half centenary tradition in educating in research through research, introducing young scientist to the vast possibilities that grid computing has to offer to the scientific community at large.

Grid Research

Research at the University of Liege on the integration of P2P file sharing with P2P Grid computing

Cyril Briquet, Xavier Dalem, Sebastien Jodogne and Pierre-Arnoul de Marneffe
EECS Department, University of Liège

Université de Liège

An operational P2P Grid middleware (Lightweight Bartering Grid) has recently been developed at the University of Liege in the context of a PhD thesis (2003-2007) and a thesis of a last year student in Computer Science (2006-2007).

The major contribution of the research presented in this article is the integration of P2P file sharing technology with P2P Grid computing technology.

The Lightweight Bartering Grid enables Peers to exchange computational time and to distribute the computing of Java applications. It targets applications structured as sets of independent computational Tasks (so-called Bags of Tasks) often encountered in GIS, computer vision, data mining, bioinformatics ... To process Tasks, supplier Peers (i.e. Peers which accept to compute Tasks on your behalf) must first download input data files, which may considerably slow down response time. Our basic hypothesis is that P2P file sharing is a promising technology to remove such performance bottlenecks.

While Peer-to-Peer Grid computing has been gaining considerable importance in the Grid community, many state-of-the-art implementations still depend on a centralized data transfer architecture not well adapted to P2P Grids, or require hard-to-obtain information. In the last few years, though, a few P2P-based data transfer protocols have emerged and have been considered for Grid data transfer architectures. BitTorrent is one of them, and is most efficient when multiple clients download a large (from a few dozens megabytes to gigabytes) file. The power of BitTorrent technology comes from the cooperation of downloaders which utilize the so-called orthogonal bandwidth between them rather than connecting directly to the server. The time for a group of Peers

to download a given file is roughly independent of their number, and remains close to the cost of transferring the file only once.

Input data files are often identical for many Tasks of a Bag of Tasks. In this case, using the BitTorrent P2P file sharing protocol to transfer identical files is very efficient, but requires that Tasks with identical files are scheduled (relatively) simultaneously. To benefit from the excellent BitTorrent performance, we use a scheduling policy called Temporal Tasks Grouping, that schedules as simultaneously as possible Tasks depending on identical data files. This way, we are able to create so-called "flash crowds", i.e. forcing many Grid nodes to download the same file concurrently, in a controlled way and thus recreate in a timely fashion the conditions in which BitTorrent is most efficient.

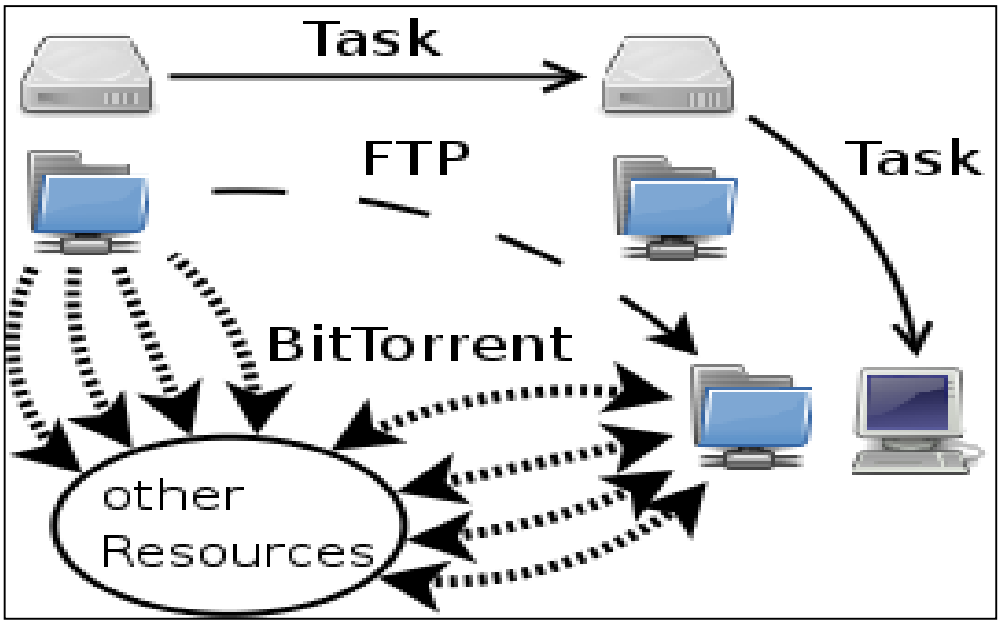

Some Tasks cannot be scheduled concurrently with other requiring the same input data files (e.g. because there are not enough resources simultaneously available). It also happens that some input data files may be required by different Bags of Tasks spread over time. To further increase the performance of data transfers in this case, we also include a caching capability on Grid nodes. This yields two interesting results:

(1) cached data files remain shared using BitTorrent even after they have been processed, which speeds up future downloads by other Grid nodes;

(2) some downloads in the near future may be avoided if a data-aware Task scheduling policy is used.

The following figure illustrates data paths from a Peer A (left), a Peer B (right) and a computational resource (bottom right). B supplies its resource to A to compute one Task. BitTorrent transfers involve A, the supplied resource and possibly many other resources that either download the same file concurrently or already store it into their data cache. FTP is used to avoid BitTorrent overhead when files are too small or too diverse.

Experiments have showed that the combination of BitTorrent, Temporal Tasks Grouping, data caching and data-aware Task scheduling delivers excellent performance. In the case of a Bag of Tasks with sets of 4 Tasks sharing an identical input data file, BitTorrent is up to twice as fast as FTP. In similar conditions, the presence of data caching support speeds up BoT response time by a factor of 2 to 3, for both FTP and BitTorrent.

Our operational implementation is extremely scalable and easily deployable because P2P technology is built into the system both at the computing layer and at the data transfer layer. It is developed 100% in J2SE 5.0 and based only on Open Source technology (Azureus, Apache FTP server, edtFTPj FTP client).

Further reading:

C. Briquet, X. Dalem, S. Jodogne and P.A. de Marneffe.
Scheduling Data-Intensive Bags of Tasks in P2P Grids with BitTorrent-enabled Data Distribution.
In Proc. UPGRADE-CN'07, HPDC Workshops, Monterey Bay, CA, USA, 2007.

BEgrid Projects

Migration of the PCs in the student PC rooms to grid worker nodes during off-time

Eric Robette

Student PC rooms are often not used during nights, weekends and holiday periods hence all the computing power of such infrastructures is wasted. This was also the case at the VUB.

Eric Robette, a ULB student, took up the challenge to make use of this dormant computing power. For his Master thesis he developed the whole set-up needed to realize this idea.

The integration is realized by automatically switching between the normal Windows environment and a Linux grid environment. The main idea is that, automatically, the gLite client software is installed on a separate partition on the PCs of the student PC rooms. A server controls when the PCs will be available for becoming grid worker nodes and gives the start to switch from the normal operating environment to the grid environment and vice versa.

The Student PC room managers can interactively enter the times of availability for the grid environment. As those machines will not be in the grid on a 24/7 scheme, they have only short execution queues defined. Jobs that are running on those PCs when an environment change happens are aborted and rescheduled on other available worker nodes on the grid.
This project will enter in a pre- production phase in October, meaning that the set-up will be used by one PC room. After a positive evaluation more PC rooms will be integrated.

A guide is available to help other groups that want to do a similar set-up.

A BEgrid, HP, Intel and Voltaire project to bring a HPC resource into BEgrid

Alfredo Bonafede, HP and Rosette Vandenbroucke, BEgrid

BELNET, HP, Intel and Voltaire are signing a contract to bring a HPC blade cluster with Intel Quadcores and Infiniband from Voltaire into BEgrid. The equipment is being ordered and will arrive before the end of the year. The goal of the project is to show that a HPC cluster can be a - very valuable - resource in the grid. Users will prove this concept. The BEgrid Tribune will report more about this project in its next issues.

BEgrid Participants

BELNET- Facultés Universitaires Notre Dame de la Paix Namur - Hogeschool Antwerpen - Katholieke Universiteit Leuven - Universiteit Antwerpen - Universiteit Gent - Université Libre de Bruxelles - Vlaams Instituut voor de Zee - Vrije Universiteit Brussel

International Symposium on Grids for Science and Business

12 June 2007 - Short Report

The grid event "**International Symposium on Grids for Science and Business"** was organised by Ugent, ULB, VUB, BEgrid with the support of IBBT. IBBT took care of the administrative matters: reserving the location, the catering, the corresponding website including the registrations. Ugent, ULB, VUB and BEgrid were responsible for the program and the invitation of the speakers.

A very reasonable contribution was asked from the participants: 50 € for members of universities, high schools and research organisations and 100 € for industry. 92 persons registered for the day. Only 5 did not show up.

The day went well, all speakers showed up and the presentations were all at a very good level. There were no evaluation forms but participants expressed their appreciation during the cocktail and also by e-mail received the day after.

Our BEgrid leaflets were in high demand and Antal and me were quite busy during the breaks with answering BEgrid questions. Some interesting new contacts were made (FUNDP, Alcatel-Lucent, …).

The presentations and the list of participants are available at the website:
http://events.ibbt.be/grid2007

Upcoming BEgrid events

Courses:

BEgrid for Computer Scientists
23 November 2007
Location: BELNET

BEgrid cluster installation
3 December 2007
Location: BELNET

For more information see http://www.begrid.be/