UNICORE

A Grid Solution for Your Business

Synopsis

Although there is often an abundance of hardware and software such resources seem never available when they are needed to solve a business problem.

Grid computing offers a viable solution to exploit your company's resources and to augment them with offerings from external providers in a seamless fashion when you need them and where you need them.

UNICORE is a Grid system with a proven history in scientific and commercial environments. It provides access to distributed computing and storage systems as well as software and services.

UNICORE is freely available as open source for Unix and Windows. Commercial support can be obtained from T-Systems.

Contents

The IT Flexibility Gap

How can a CIO meet the challenge of providing a timely and cost-effective solution to an emerging business opportunity or to an ad-hoc business problem? The luxury to engage in an extensive analysis, a lengthy procurement of hardware and software, or the development of a custom solution is no longer an option. In fact, it is no longer necessary. There is an abundance of compute power, storage space, and software available either within their own company or from external service providers. The challenge is being able to exploit the available resources in the most flexible way. Grid technology offers the solution and UNICORE is an advanced implementation of a Grid system.

Why Should My Company Care About Grids?

Grids were originally conceived by researchers who endeavored to harness distributed resources – from idle desktops to the most powerful supercomputers – to solve the most challenging compute and data intensive problems that exceeded the capacities of even the largest computer. Furthermore, global access to petabytes of data, usage of sophisticated software, and collaboration with the international community motivated and drove the development of Grid software.

Today, Grids are no longer the playground of leading researchers. Commercial companies, from small start-ups to global players, have embraced Grids and offer a wide spectrum of software, solutions, and services for Grids.

Early adopters were motivated to invest in Grids to accomplish

— Reduction of costs
— Competitive advantage
— New business opportunities

Cost reduction through Grids can be accomplished in more than one way:

Grids allow the utilization of otherwise idle resources within a company by adopting a global view of the resources rather than a focusing on each individual system and balancing the load between them.

Peak demands can be met by dynamically adding external resources offered by providers. This can transform buying decisions; instead of planning for maximum demand and consequently under-using the expensive equipment most of the time, only the normal use needs to be covered freeing capital for other investments.

Many services need not be installed and maintained in house but can be purchased on demand from Grid Service providers.

Competitive advantage will result from any of the following:

Earlier time to market when developing products and services. Grids complete critical steps faster by exploiting all available resources within the company and – if necessary – from external services.

Improved quality of the end-product through more intensive and automated testing – for example parallel simulation runs – is possible within the given time and cost framework.

New products and services can be conceived and prototyped by using resources that would otherwise remain idle.

New business and organizational opportunities exist in many ways:

Manufacturers can dynamically interact with customers and integrate their suppliers on demand by making them a well defined part of their Grid.

Suppliers can become partners to many of their customers in a consistent way while maintaining full confidentiality of their business relations.

Companies can optimize their internal resources (personnel, software, devices) while making them available to all that need legitimate access within and outside the company.

Services that a company uses internally can also be made available to external customers as Grid Services without disclosing confidential data.

Additional services can be purchased on demand without requiring heavy investments in hardware, software, and support. This approach is potentially superior to outsourcing business functions.

None of the above opportunities are completely new. For each aspect singular solutions have been applied successfully. The unified and novel approach that Grids bring to the table is the ability to seamlessly build flexible solutions based on standards, using built-in security, and the growing portfolio of interoperable Grid products and Services.

Is There a Single Approach to Grid Computing?

Since Foster and Kesselman in 1998 first described the Grid concept in their famous *Blueprint for an New Computing Infrastructure*, many approaches have been tried to address specific scenarios. They range from the development of innovative Grid systems to the adaptation of existing products to encompass the new specifications. Most of them build on existing standards for distributed, internet based computing and move to comply with the Grid specific standards that are being developed by organizations like the Open Grid Forum (OGF), the Organization for the Advancement of Structured Information Standards (OASIS), or the World Wide Web Consortium (W3C). Typical examples are:

— The use of Grid concepts to achieve intra-organizational optimization of compute-resource based distributed batch-systems like Platform's LSF, SUN's Grid-Engine or Altair's PBS Pro. The important characteristic of such solutions is that they take place within one security or administrative domain even if they are geographically distributed.
The advantage of this approach is that pre-Grid versions of software can be migrated to a Grid environment relatively painless. However, each of the systems represents an island and interoperability between them remains an objective for the future.

— The development of new applications or re-engineering of existing ones to run them on a Grid infrastructure. Such developments would typically employ the base level toolkits, for example the Globus® Toolkit, which is developed by Argonne National Laboratory and the Globus Alliance. The advantages are that the resulting solution is best suited to execute in well defined virtual organizations once the required standards become pervasive.

— The deployment of Application Grids and Service Grids for existing application scenarios and workflows. Typically, the user of an application or service belongs to a different organization than the provider of the service or the application. Whenever a part or an entire business-process is out-tasked and the business-process is carried out on a service-oriented architecture (SOA), we speak about an application or service Grid. The underlying Grid technology is used to link consumers and service providers in a dynamic way. This is a typical scenario for the use of UNICORE, especially if the processes build upon a service oriented architecture.
The key advantage to this approach is that the application and business processes can remain unchanged and the transition into Grid technology can take place gradually over time. Of cause, an initial investment to deploy a Grid infrastructure and to integrate the applications has to be made.

Showcases of Successful Grid Applications

Improved weather predictions and disaster warnings

Over the last decade, the world has experienced an increasing number of high-impact meteorological events which have disrupted national services, damaged or destroyed infrastructures and resources, and has resulted in a loss of human life. High impact weather is defined by its disastrous impact on society and economy; it includes e.g. tropical cyclones, extreme rainfalls and high wind speeds, etc. It is, therefore, a major scientific challenge of the next decade to reduce and mitigate natural disasters relying on improved weather predictions and by giving early and reliable access to data and information to authorities and services which are in charge of warning and prevention.

Methodological developments and improved databases from weather simulation and observation combined with advanced IT-Technologies are key for this. The vision is of a cooperative of Met Services who combine their various skills and capabilities to ensure all are provided with the best environmental forecast possible, allowing impact scenarios to provide optimum planning information world wide. This vision relies upon supporting technology to enable multiple, geographically separated centers to generate a timely and reliant ensemble product. In the near future virtual computing may allow the easy exchange of data and products in an effective and efficient manner.[1]

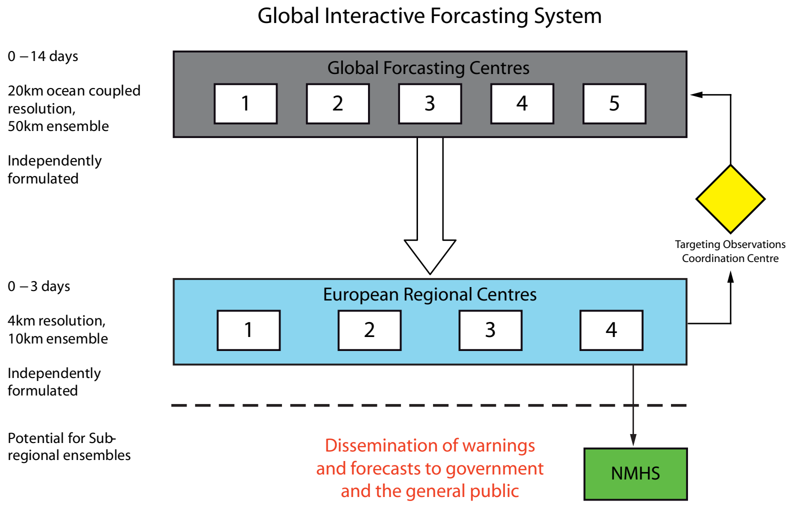

Figure 1: © A Vision for Numerical Weather Prediction in Europe by Dr. Dave Rogers.

---

1 The developments were funded in part by the European Commission in the project GRIP under grant IST-2001-32257.

America's Cup race yacht design

Team Shosholoza is the South African America's Cup Challenge Team for 2007 with T-Systems as the main sponsor. High-tech sailing today is not possible without sophisticated IT-support in the design-phase of the boats as well as in the races.

The real challenge for the IT-support comes from the fact, that the design teams are very dynamic with respect to the people and organizations involved as well as concerning the locality at which the compute power is needed. Furthermore, due to the fierce competition no security compromises are allowed.

Team Shosholoza has access to some of the fastest computers in the world for shape optimizations of the boat based on computational structural mechanics and computational fluid dynamics methods as well as for the calculation of mission scenarios just before the races. This includes online local weather-predictions based on the coarse grain calculations from the weather-services and measurements of the actual conditions.

The required flexibility could only be achieved using UNICORE Grid technology as the main component for the access to computing resources.[2]

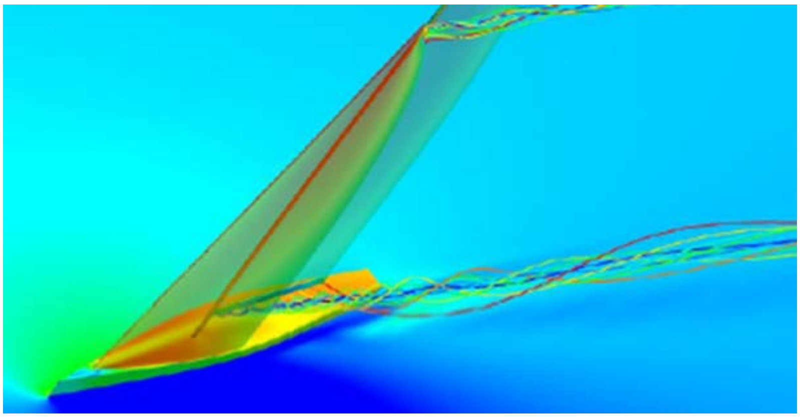

Figure 2: Visualization of simulation results.

2 The developments were funded in part by the European Commission in the project UniGrids under grant IST 004279.

Drug design

Molecular modeling, molecular engineering and drug discovery, as closely related fields, provided the set of real life applications central to the OpenMolGRID project. OpenMolGRID was conceived to exploit the power of Grid Computing to shorten the time to solution for drug discovery, specifically the identification of promising new compounds as potential drug candidates.[3]

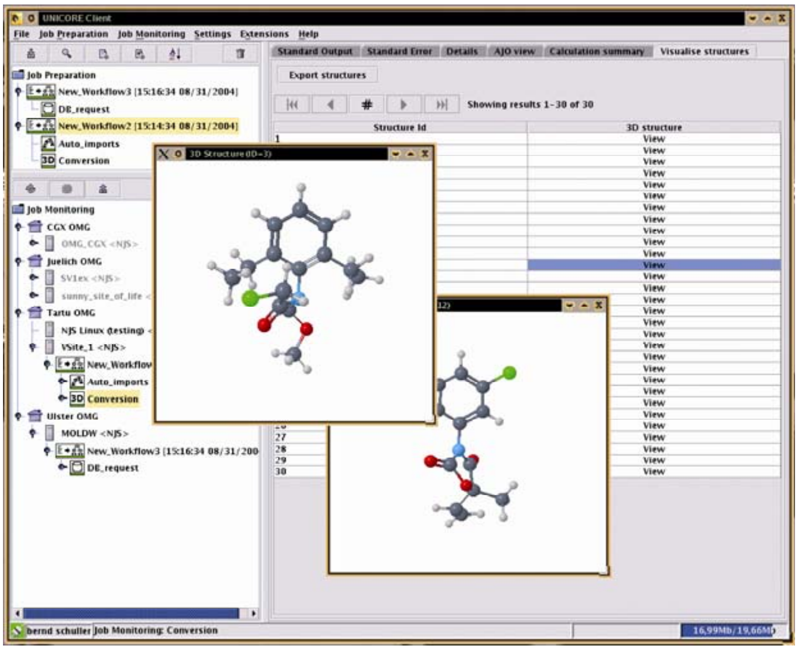

Figure 3: The OpenMolGRID extensions of the UNICORE Client with the molecule structure view.

How can UNICORE Solve Your Business Problem?

The following gives a brief overview of the key features of UNICORE as it is used in production today. The full documentation can be found together with the software on SourceForge.

Key Characteristics of UNICORE

The UNICORE system – *Un*iform *I*nterface to *Co*mputing *Re*sources – was originally conceived in 1997 to enable German supercomputer centers to provide their users with a seamless, secure, and intuitive access to the heterogeneous computing resources at the centers. As a result, the projects UNICORE[4] and UNICORE Plus[5] were funded by BMBF, the German Ministry for Education and Research, with the following objectives:

---

3 The developments were funded in part by the European Commission in the project OpenMolGRID under grant IST-2001-37238.
4 UNICORE was funded in part by the BMBF grant 01 IR 703.
5 UNICORE Plus was funded in part by the BMBF grant 01 IR 001.

— UNICORE was designed to hide the seams resulting from different hardware architectures, vendor specific operating systems, incompatible resource management systems, and different application environments.
— Retaining organizational and administrative autonomy of the participating centers was a key objective of UNICORE. None of the service providers should be forced to change historically grown computer center practices, naming conventions, and security policies to be able to use the full benefits of UNICORE.
— Security was built into the design of UNICORE from the start relying on the X.509 standard. Certificates are used to authenticate servers, software, and users as well as to encrypt the communication over the open internet.
— Finally, UNICORE had to be usable by scientists and engineers without having to study vendor or site-specific documentation. A graphical user interface was developed to assist the users in creating and managing jobs.

In parallel, the term Grid was coined by Foster and Kesselman and it turned out that UNICORE had the characteristics of a vertically integrated Grid system. In addition to the original objectives, UNICORE demonstrated its viability in structural analysis and design, life sciences, or weather forecasting both in industrial and research environments.

As an example, UNICORE provides access to the DEISA[6] (Distributed European Infrastructure for Scientific Applications) production environment which aggregates over 145 TFlops[7] of compute power managed by ten leading European HPC centers. Furthermore, UNICORE is one of the Grid middleware systems used to build up and operate the D-Grid infrastructure within the German e-Science initiative.

UNICORE supports access to resources – computers, data, and applications – for the full spectrum of organizational structure:

— Consistent, standardized utilization of the different architectures and systems within one organization
— Secure, internet based access to geographically distributed resources of a global organization
— Virtual organizations combining the resources of independent companies in support of selected user communities.

Security and Trust

Security and trust are the key criteria that ultimately determine if Grid technology will succeed. UNICORE is architected to support the most stringent security requirements. Certificates according to the X.509 standard provide the basis of UNICORE’s security architecture. Certificates serve as Grid-wide user identifications which are mapped to accounts at the different systems. This approach has three significant advantages especially for virtual organizations:

6 DEISA is funded in part by the European Commission grant IST 508830.
7 145 TFlops is about 50,000 times more than a modern desktop system or notebook with an average performance of about 3 GFlops.

— The participating sites do not have to change their established conventions
— Users need not know the local users names and passwords at the different systems
— The resource providers can always identify the originator of a request and thus can charge the user for the resource consumption reliably.

The user's private key is used to sign and encrypt request, for example job submissions, which allows the server components of UNICORE to detect any tampering with the data while it was in transit over the internet and to map the certificate to the appropriate account at the target system.

Certificates also mutually authenticate peer systems in a UNICORE Grid. The Gateway will accept only requests from a client or a peer system that can provide a trusted certificate. Symmetrically, the client will communicate only with trusted Gateways. Gateways and other servers normally reside behind a firewall. The client software is run from the user's personal system – a desktop that may reside at the office or a notebook connected to the internet from anywhere. The UNICORE servers require that only one port and one service is opened in the firewall per site. The Java code that implements the Gateway is small in size; moreover it has been thoroughly inspected by security experts from leading IT companies like T-Systems as well as leading scientific organizations and no vulnerability has been found so far.

Certificates as a basis for security require the existence of a Public Key Infrastructure (PKI). Essentially, participants in a UNICORE Grid have to agree which Certificates Authorities (CA) will be trusted. However, UNICORE is not limited to exactly one CA, during operation of a UNICORE Grid several CAs can be accepted concurrently. As an example, the DEISA Grid accepts certificates issued by each of the national CAs in the participating countries, that all belong to the EUGrid PMA. A resource provider may accept one CA for one virtual organization and a different one for another VO.

Providers retain full control over the resource they make available as part of a UNICORE Grid. They decide which part of the total resources is available to each virtual organization or to an individual user. The policies depend on the service level agreements between the provider and the user communities.

The strict and straightforward security and trust model of UNICORE has been expanded to support portal architectures by adding Explicit Trust Delegation (ETD). Users can decide if they will trust a portal to execute certain requests on their behalf and providers may decide if they trust certain portals. This architectural extension of UNICORE's security model enables interoperability with other implementations of OGSA[8] compliant systems that rely on delegation.

8 Open Grid Services Architecture.

The UNICORE Client: Bringing Applications and Services to the User

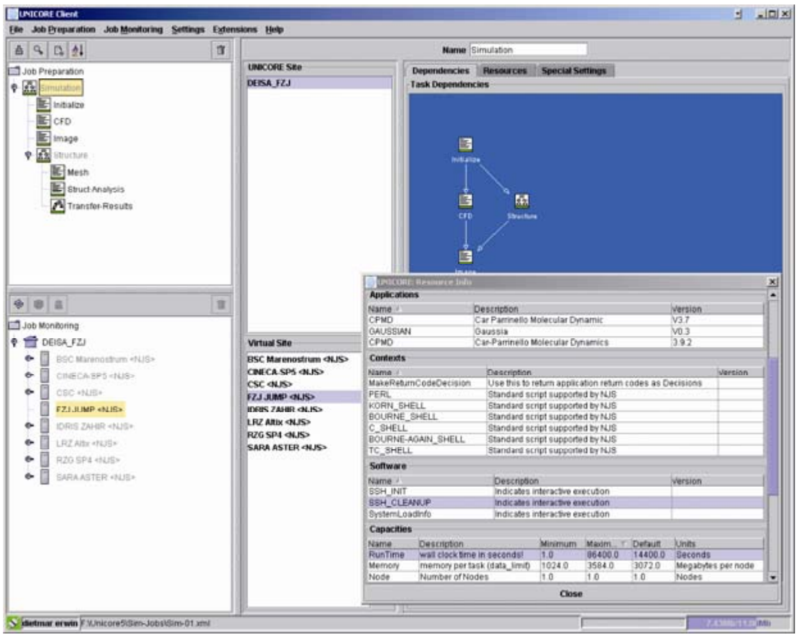

Figure 4: Top-level view of multi-site workflow.

UNICORE provides the users with an intuitive graphical interface to all resources and services available in the Grid. Figure 4 shows the top level view of the UNICORE Client. The pull down menus *File*, *Job Preparation*, etc, group the rich set of functions that UNICORE offers. Depending on the selected function the contents and structure of the panels in the right hand side may change. The Client supports the user in the creation of complex multi-system, multi-site workflows, the management of jobs and data in the Grid, the control over and monitoring of running jobs and results, the definition of the Grid environment and the client look and feel. In detail the functions are:

Job Preparation: A UNICORE job - more precisely a job group - may recursively contain other job groups and/or tasks. Figure 4 shows the structure of a job in a typical folder view. The *Job Preparation* pull down menu allows adding tasks, commands, applications, data management functions, and controlling structures to the job. They can be nested recursively to as many levels as required. For each job group, the user selects a UNICORE site to which the job group will be submitted and the system on which all tasks belonging to the job group will be executed. Similarly, access to information stored in data-bases or archives can be integrated in the workflow. The user may specify temporal dependencies between the entities contained in a job group (see the right hand side of Figure 4). The example shows a workflow consisting of a CFD computation and a structural analysis that are executed on two different systems at different sites. The initial step (*initialize*) is followed by the CFD code that runs at the same site, while the structural analysis (the *structure* job) is processed on a large shared memory machine at a different site. The *structure* job is itself a workflow consisting of a mesh generation, the computations, and a *transfer* task that returns the results. When all tasks have completed successfully, the image task creates the output which may be viewed by the user and stored for later processing.

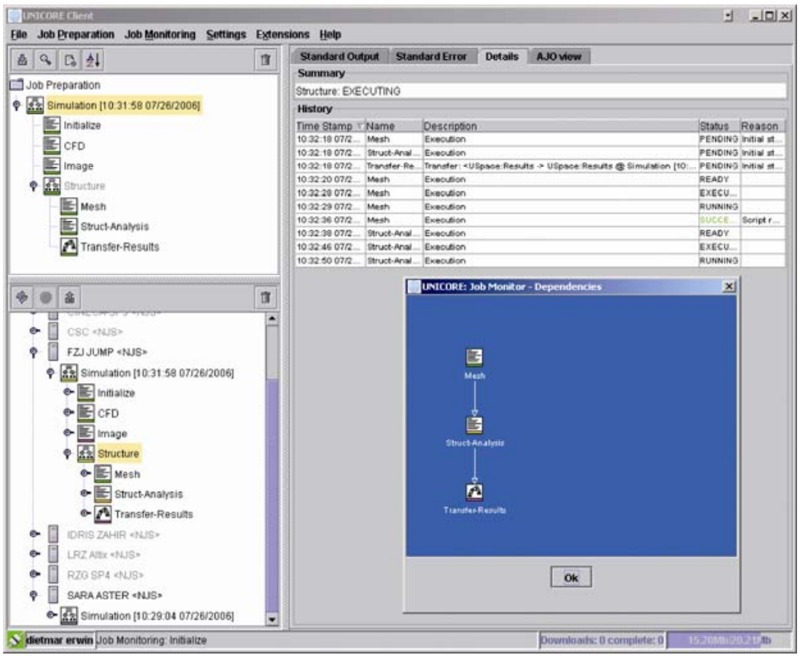

Figure 5: Status information of a UNICORE job

In addition to the dependencies, the user may define control structures such as loops (DO n; REPEAT until), IF-THEN-ELSE constructs, or WAIT for an event. This allows automating complex workflows while interacting with the jobs if necessary. For example, a Web Service, that sends an SMS message to the users mobile phone upon completion can be inserted.

Job Management: Jobs created by the user can be saved to any directory on the client workstation using the *Save* or *Save As* commands in the *File* menu. The jobs are stored both in the UNICORE internal format as Abstract Job Objects (AJO) and as XML files. They can be loaded again using the *Open* command. Sending a job to a different system is accomplished by selecting a new UNICORE Site. The Client automatically checks if the new system can fulfill the resource requirements of the job. If the check is positive this will be indicated by green bars on the icons. If one or more resource requirements can not be met, the color bar turns red and the reason can be seen by clicking on the Info button. Jobs are then save to be executed and can be submitted.

Job Monitoring: The user has full control over jobs and data. A color code presented along with the job icon in the folder view in the job monitoring area shows the overall status of a job: green, red, yellow, blue, magenta, grey, or black to indicate successful completion, failure in execution, queued, or waiting for completion of a predecessor, respectively. The example (Figure 2) shows that *Initialize*, *CFD*, and *Mesh* completed successfully, *Struct-Analysis* is currently executing (status: yellow) whereas *Transfer-Results* and *Image* are waiting to be run upon completion of their predecessors. The status is available at each level of recursion down

to the individual task. It may be refreshed by clicking on the icon or by an automatic timer-driven status update. In addition, detailed log information is available to analyze error conditions. The job output that is written to *stdout* and *stderr* by the execution systems can be reviewed or transferred to the client workstation.
A user may cancel jobs queued for execution anywhere within the UNICORE Grid and terminate executing jobs.

Data Management: UNICORE jobs contain tasks that can be executed at different computing centers. Output created by one task may be used by any of its successors. A temporary UNICORE space, called Uspace for short, is created for each job group. During job creation the user specifies

— data sets which are to be *imported* into the Uspace from the client workstation or any file system or data archive at the UNICORE site to which the user has access,
— data sets to be *exported* from the Uspace to retain them permanently, and
— data sets which are to be transferred to a different Uspace.

The safety of the UNICORE job space and confidentiality of this data is the same as of any other directory of the user on the target resource.
At run time UNICORE performs the necessary data movement without user intervention.
The import, export, and transfer functions are available to the user through the UNICORE Client in the *Job Preparation* pull down menu as explicit tasks. Data sets to be imported or exported can also be specified as part of the job. The data management function can also be imbedded into applications and can be performed fully transparent to the user.

Application Support: The four services described above provide an effective tool to use resources of different computing centers both for capacity or capability computing. Many scientists and engineers use application packages. UNICORE provides extensive support for these users: For applications without a graphical user interface, a custom built plug-in for the UNICORE client can be provided. An API library and rich documentation is available that simplifies the creation of plug-ins. Plug-ins for many scientific and commercial applications are available, like CPMD (Car-Parrinello Molecular Dynamics), Gaussian98, Amber, Gamess, VampirTrace, BLAST, Nastran, Fluent, StarCD, SQL Access, VISIT, or POV-Ray.

Automation of work: One of UNICORE's strength is the ability to define workflows that run unattended in the background on the most suitable systems in the Grid. UNICORE interfaces with any resource management system, making even the most arcane batch systems easy to use. User attention is only required to review the results, and even this step can be automated as was shown in the OpenMolGRID project.

Interactive Access: UNICORE supports an interactive use, e.g. to debug code or to steer appropriately instrumented applications, of Grid resources in two ways:

- The UNICORE Client has an integrated file browser that provides the users with secure access to his or her data stored in remote file systems or data storages within the whole Grid
- In addition UNICORE supports interactive access to computing resources by means of the UNICORE-SSH solution which presents the user with a Unix secure shell validated by the user's certificate.

Support for legacy jobs and easy migration: Even the best solution will not be embraced immediately by users, especially if the existing techniques work. Therefore, UNICORE supports traditional batch processing by allowing users to include their old job scripts as part of a UNICORE job. This approach does not guarantee seamlessness but it simplifies the migration. UNICORE immediately helps users in the following scenario: a user submits a job to a computing resource, periodically checks for its completion, transfers results to a different system, using ftp, and submits a successor job to this system. A simple job group can automate these steps without changing the existing job scripts.

Seamless Integration with Existing and Future Infrastructures

Seamless integration – both on the technical and the administrative level – has always been a key objective of UNICORE. Maintaining the administrative autonomy of the participating sites has been a winning factor for UNICORE. UNICORE adapts to naming conventions by mapping user certificates to local user names. UNICORE does not require changes to existing management policies and accounting procedures.

Since UNICORE is implemented in Java and based on the concept of abstract jobs and incarnation into system and vendor specific jobs, it supports any Linux/Unix based operating system and a wide range of Resource Management Systems, like LoadLeveler, LSF, PBS, CCS or NQS. Support of a new system requires the adaptation of a generalized copy of the Target System Interface module (implemented in Perl, alternatively Java) that executes on the target system. In addition, the UNICORE Incarnation Data Base (IDB) has to be configured to provide a site and system specific mapping. All this is available and used in production in version 5 of UNICORE.

UNICORE version 6, will incorporate the developments of the UniGrids project. Sever Components have been re-factored to support Web Services that can be integrated into the classical UNICORE workflow. All developments follow closely the evolving standards in OGF (JSDL[9], BES[10]) and OASIS (WS-RF[11], WS-N[12])

9 Job Submission Description Language.
10 Basic Execution Services.
11 Web Services Resource Framework.
12 Web Services Notification.

and in some cases provide prototypes since UNICORE developers were part of the teams that defined the standards.

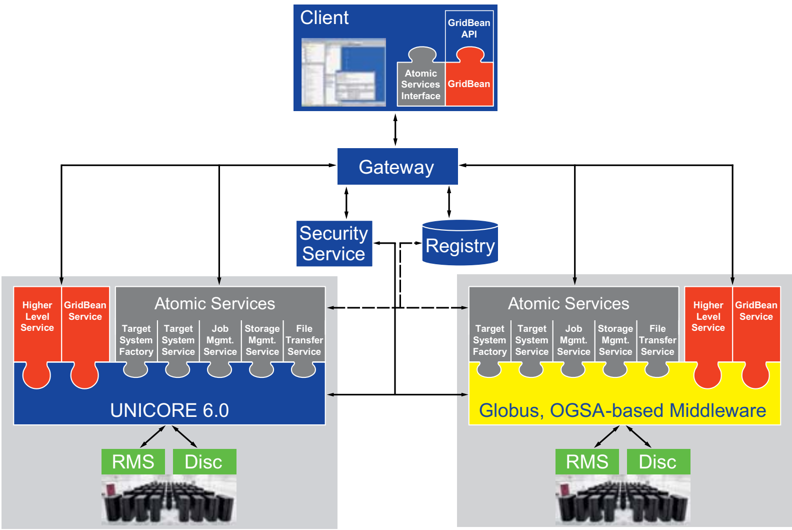

Figure 6: Architecture of UNICORE 6.0.

The new architecture is depicted in Figure 6. It allows direct utilization of Web Services that are developed elsewhere and thus enriches the portfolio of services offered to the user. Over time, high level services that are unique in UNICORE can be converted into standards based Web Services once the required standards become available and mature. Conversely, UNICORE components can be accessed from other OGSA compliant systems.

What is Required to Run UNICORE?

Running the UNICORE software is as easy as using any internet application. Substantial effort has been invested to provide users and administrators with self-installing software. UNICORE clients run on any Windows or Unix desktop. The only prerequisite is a current Java run time library.
The site that provides the UNICORE Grid has to install a small server component on each compute resource that is part of the Grid. The UNICORE Gateway and the UNICORE Network Job Supervisor should be installed on dedicated Unix servers for maximum security. The Servers should reside behind the firewall. Only one port has to be opened in the firewall to give users access to the Gateway. In addition, a Public Key Infrastructure has to be available, either in-house or through a commercial Certificate Authority.

Getting to know UNICORE

To try out UNICORE and explore its capabilities the Unicore Forum maintains a full-function UNICORE Grid at *http://www.unicore.org*. Simply click on *UNICORE Test Grid Version x.y* to follow the three steps to install the client on your Windows or Unix desktop. You will automatically receive an account and the corresponding certificate on the test Grid. Both are valid for 30 days and may be renewed as often as required.

Installation

The latest installation packages can be downloaded from SourceForge at *http://unicore.sourceforge.net*. The core software, i.e. the *Client* and the server components, *Gateway*, *NJS*, and *TSI* can be downloaded individually. The Client installation is similar to the one described above. The proper Gateway addresses have to be configured to point to the Gateway and the certificates have to be obtained from the Certificate Authority your company is using. Your company may elect to provide a preconfigured, customized version of the client.

The system administrator has to download the server components and install them as required by the security policies of the company. Both for performance and security the Gateway and the NJS should be installed on separate dedicated systems. The Target System Interface (TSI) must be installed on each compute resource. To try-out a server installation a q*uickstart bundle* is available that produces a UNICORE Grid in less than 15 minutes.

Optional software such as application plug-ins is also available; they may especially be used as prototypes to create custom extensions to UNICORE.

Administration

Users will expect that UNICORE is operational on a 24x7 schedule. Therefore all UNICORE serve components should be integrated in the company-wide monitoring system that checks the availability of essential servers and software and alerts administrators.

Needless to say your helpdesk and system administration should be familiar at least with the existence of UNICORE and know whom to contact for additional support.

If UNICORE users on the internet are to have access to the Grid the firewall administrators have to open one port that allows communicating securely with the Gateway component on the dedicated Gateway server.

Who Provides Support and Services for UNICORE?

System administrators will have no problem to maintain and administer UNICORE using the information available on the UNICORE SourceForge project pages. As with any software is good practice to establish help desk functions that can be contacted by the end-user in case of problems.

Academic organizations, such as universities or research centers, may contact Forschungszentrum Jülich at *unicore-info@fz-juelich.de* for support to setup and manage a UNICORE Grid. Support will be provided on a best effort basis free of charge.

Commercial support is available from T-Systems on the basis of an individual contract. The options may include:

— Administration of customer-site UNICORE-environments including release-management based on quality gates, management and support-processes.
— Access to compute-resources and applications running in utility data-centers of T-Systems.
— Consulting-Services
— Seamless integration of customer computing-resources with resources of external providers.

For further information please contact *unicore@t-systems-sfr.com*.

Imprint

UNICORE
FORUM

Director of the board
c/o Deutscher Wetterdienst
Prof. Geerd-Rüdiger Hoffmann
Leiter Abt. TI1
Kaiserleistraße 42
63067 Offenbach

Cover-Photo: South Africa's America's Cup Challenge Team Shosholoza uses UNICORE to access some of the world's most powerful computers provided as a service by T-Systems.