

# UNICORE: Beyond Web-based Job-Submission

Mathilde Romberg  
Research Center Jülich  
Central Institute for Applied Mathematics  
D-52425 Jülich, Germany  
m.romberg@fz-juelich.de

## Abstract

UNICORE (Uniform Interface to Computer Resources) is a software infrastructure to support a uniform, secure Web-based access to distributed resources. The talk will give an overview of the architecture, its security features, the user functions, and the current implementation status.

## 1 The UNICORE Plus Project

In 1997 the German Ministry for Education and Research (BMBF) set up a two years project to develop a software infrastructure for seamless access to distributed supercomputer resources. Motivation for the UNICORE (UNiform Interface to COmputing REsources) project<sup>1</sup> has been that users who have to solve large problems in computational science usually need resources on a variety of systems at different locations. These users are faced with different site policies and practices (security, user identification, data management, ...), different system architectures, and system software. For the efficient use of the resources the users need to learn about these differences. UNICORE is going to overcome the additional effort. The project goal is to leave batch users of the German supercomputer centers with a seamless, secure, and intuitive Web-based interface for job preparation, submission, and monitoring. The developed UNICORE prototype runs in test-mode at the centers in Berlin (ZIB), Jülich (FZJ), Karlsruhe (RUKA), München (LRZ), Offenbach (DWD), and Stuttgart (RUS) for the target systems T3E, SP2, and VPP. It provides the user with a single sign-on environment via X.509 certificates. Through the graphical interface the user can prepare multi-step UNICORE jobs consisting of script jobs for shell-scripts of existing applications, Compile-Link-Run-jobs for new applications, and file transfer jobs. Jobs to be run at the same target system are grouped together in job-groups. The user specifies general information like the target system per job group. Resource requirements for CPU time, number of CPUs, amount of main memory, and disk space are specified per job. The user input is mapped by UNICORE to the target system specific commands and options. The job monitor allows for checking the job status and accessing the standard output and error files.

A follow-on project, UNICORE Plus<sup>2</sup>, started in January 2000 and is planned for a duration of three years. Based on the work done in the UNICORE project UNICORE Plus is going to develop a grid (see [5]) infrastructure for seamless and secure access to supercomputer resources. This includes a sophisticated graphical user interface and security mechanisms as well as server components for the translation of the uniform job specification into batch jobs for the target system, for scheduling etc. The project goals are to develop an improved, robust prototype which allows for multi-part, multi-site jobs with elaborate job flow functions and which supports application specific interfaces, efficient data transfer, resource modeling, and meta-computing at application level. As within UNICORE it is the design principle to keep site autonomy. Project partners are the German Weather Service (DWD), Research Center Jülich, Computer Center of the University

---

<sup>1</sup>See <http://www.fz-juelich.de/unicore>; UNICORE has been funded by BMBF, grant period 1.8.1997 - 31.12.1999, grant id 01 IR 703

<sup>2</sup>See <http://www.fz-juelich.de/unicoreplus>; UNICORE Plus is a BMBF funded project, grant id 01 IR 001

of Stuttgart (RUS), Pallas GmbH, Brühl, Leibniz Computer Center, Munich (LRZ), Computer Center of the University Karlsruhe (RUKA), Paderborn Center for Parallel Computing (PC2), Konrad Zuse Center, Berlin (ZIB), and Center for High Performance Computing at TU Dresden (ZHR). The project is structured in eight sub-projects each dealing with one or more aspects, including software development and project-management. Topics of the sub-projects are test, quality management, administration, and public key infrastructure (PKI), resource modeling, application specific interfaces, data management, job control flow, and meta-computing.

While in the UNICORE project the hardware vendors<sup>3</sup> were non-funded project partners they now are members of the recently founded friendly society UNICORE Forum e.V. to further support the UNICORE development.

## 2 The UNICORE Architecture

The architecture in the UNICORE Plus project is based on the three tier architecture developed in the predecessor project (see [1], [6], <http://www.fz-juelich.de/unicore>). It consists of user tier, UNICORE server tier, and target system tier as shown in Figure 1. User and server tier are Java applications while the system tier is realized in perl.

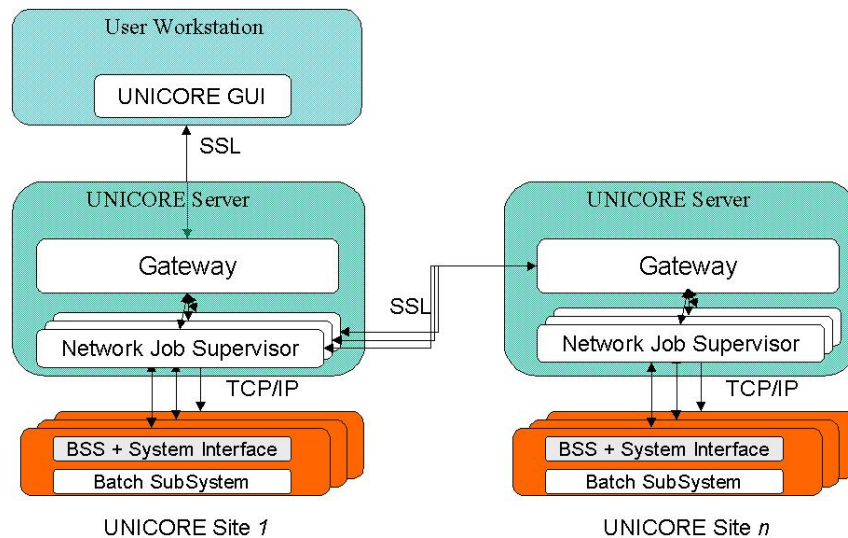


Figure 1: Architecture Overview

The user tier consists of the graphical user interface. It offers the functions to prepare and control UNICORE jobs (for details see Chapter 4) and to set up and maintain the user's security environment. As described in Chapter 3 in more detail, the security architecture is based on the Secure Socket Layer (SSL) protocol. SSL is used for the communication between the components which talk to each other over public networks. From the user input the user interface generates an Abstract Job Object (AJO) which is sent via SSL to the Gateway. The AJO is a key component in the architecture. It comprises the UNICORE protocol between user interface and the Network Job Supervisor together with the abstract job specification generated from the user input. The AJO is realized as a Java class library. [6] and <http://www.fz-juelich.de/unicore/DARR.ps> explain the AJO in more detail.

The Gateway is the first part of the UNICORE Server tier. It takes care of the user authentication, secure communication between client and server, and provides the client with information about the available resources at the site. It also talks to the Network Job Supervisor (NJS) servers at its site to send jobs and data, status requests and control commands for further processing as well as receive data to make it available to the user. Each NJS controls one target system or one cluster of systems which are governed by a single resource management system. It fulfills the following tasks:

<sup>3</sup>Hitachi, HP, IBM, NEC, SGI/Cray, Siemens/Fujitsu, Sun

- analyze the AJO (representing the UNICORE Job) for local and remote jobs;
- map the UNICORE user id to the local userid for the target system;
- translate the local jobs contained in the AJO into real batch jobs for the target system;
- send job-groups to be executed at other UNICORE sites to the corresponding gateway;
- provide local resource information to the gateway;
- take care of necessary file transfer;
- provide job status information and job output.

The underlying recursive job model is such that a UNICORE job consists of

- job groups, which build a frame with general information for this part of the job like the target system; they contain them-self jobs groups, jobs, and dependencies,
- jobs or tasks, which are going to be translated into batch jobs for the target,
- dependencies between the elements to reflect the necessary synchronization; they build a directed acyclic dependency graph.

Each UNICORE job is assigned a UNICORE Job directory which is the unix working space for the job. It is a temporary directory existing only during the lifetime of the job at the site. All data needed for the job execution have to be imported from permanent file space to the job directory. All data which is needed after the job has finished has to be saved to permanent file space. The user has to specify the data to be imported and those to be exported. The transfers are done by UNICORE transparently to the user.

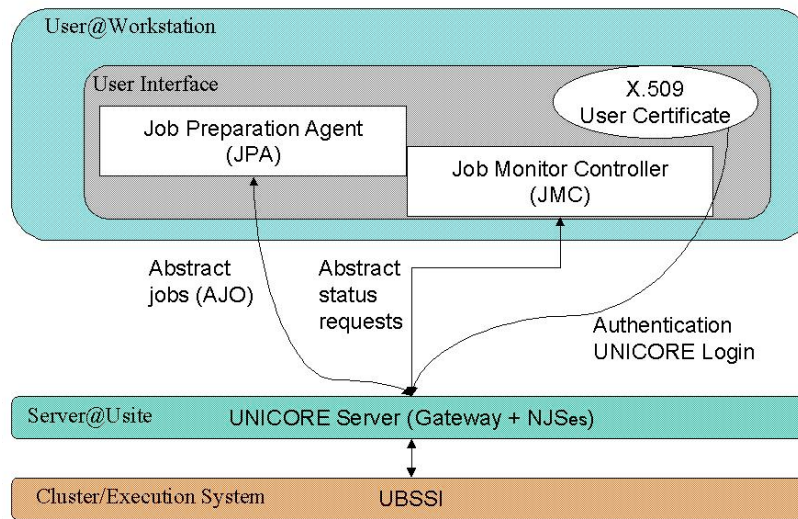


Figure 2: The User Tier

### 3 The Security Architecture

This chapter takes a closer look at the architecture described above explaining the three tiers in more detail with respect to the security features. The security is based on the Secure Socket Layer (SSL) protocol and the X.509V3 type certificates (see [4]). SSL uses public key cryptography for establishing connections between client and server. Therefore each component has to have a public-secret key pair with the secret key kept secret and the public part being known by the others. The keys have to be certified by a certification authority (CA) so that the components can be sure that they communicate with the user (or program) he or she claims to be. By default certificates signed by unknown signers are not accepted.

The UNICORE user's X509 certificate is his or her UNICORE user identification. It is maintained by the user interface application in a crypted data base. The user interface also needs to know about the Certification Authority which signs the user and the gateway certificates. With this information a secure connection with mutual authentication of both components can be established. Figure 2 shows the user tier hosting the main GUI components Job Preparation Agent (JPA) and Job Monitor Controller (JMC) as well as the user's certificate. The user is authenticated by the Gateway when presenting his or her certificate. The certificate is also part of the AJO for authorization at different sites and the AJO parts as well as the whole AJO are signed with the user's certificate. This is necessary to make sure that no one can tamper with the AJO on its way without being noticed.

UNICORE uses the UNICORE Protocol Layer (UPL) to send the AJO to the Gateway which just hands it over to the NJS controlling the specified target system. In case a UNICORE site uses additional authentication methods like DCE (Distributed Computing Environment) or SecureID cards, the Gateway allows for site specific additional authentication as is depicted in Figure 3. Gateway and NJS communicate via sockets. The NJS first checks whether the AJO is an correct AJO from the user who signed it. It unpacks the AJO with the user's public key and analyses it for parts to be send to other target systems. These are send either to another local NJS or to a gateway at another UNICORE site at the point in time the dependency graph defines. Each NJS has its own X.509 certificate which it uses to communicate to a Gateway at another site via SSL. The Gateway then knows, that the AJO is part of another AJO and that it has to use the user's public key to unpack the AJO instead of the one from the sending NJS.

For tasks to be processed locally a NJS does the userid mapping which means that it maps the user's certificate to his or her user identification at the local system. The NJS then does the translation into batch jobs for the target system and sends them according to the dependency graph to the UNICORE batch subsystem interface (UBSSI) on the target.

The UNICORE server can operate in a firewall protected environment. The server can either be completely behind the site's firewall or the firewall is located between Gateway and NJS.

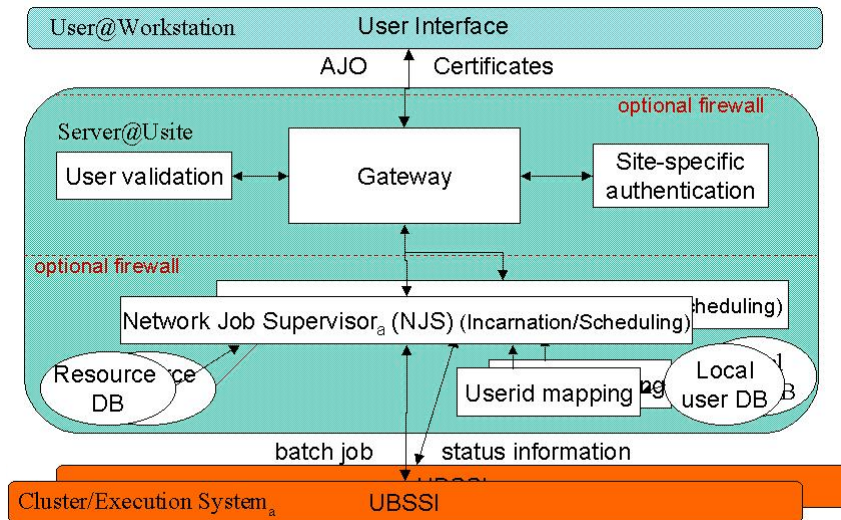


Figure 3: The UNICORE Server Tier

The target system may be a cluster of systems which are governed by one cluster resource management system or the target is a stand-alone system. Each of the systems is governed by a UNICORE Batch Subsystem Interface (UBSSI). The UBSSI talks to the local batch system to maintain the job directory, submit jobs on behalf of the user, and to control the job status.

A key component for the security architecture in UNICORE is the Public Key Infrastructure (PKI). It is needed to generate the certificates for users, Gateways, software developers, and NJSes. The UNICORE PKI policy is based on



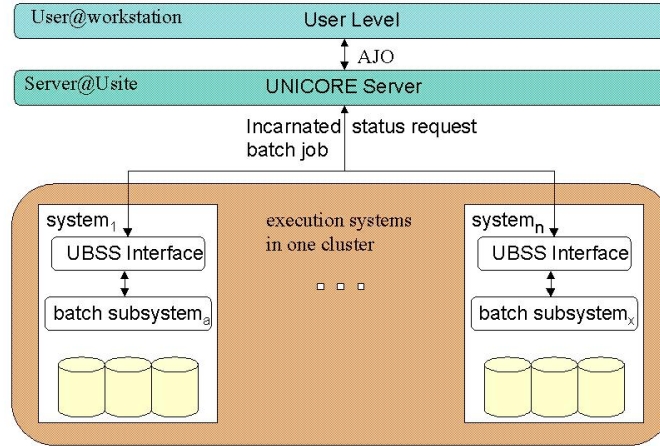


Figure 4: The Target System Tier

the regulations defined by the DFN-PCA (German Research Network - Policy Certification Authority<sup>4</sup>). The UNICORE Certification Authority (CA) is located at LRZ, all partner centers run a Registration Authority (RA).

## 4 The User Interface

Within the UNICORE environment the user has a comfortable way to use distributed computing resources without caring for site or system specifics. It is all done in a seamless way. As described above the specifications of resource requirements and command options etc. are done independently from the selected target system. They are translated by the NJS into the real batch job for the target.

A key point is that a user has a single UNICORE userid, his or her UNICORE X.509 certificate, to get access to the resources at the various UNICORE sites. The graphical user interface offers at top level the functions to maintain the security settings, to prepare UNICORE jobs, and to monitor them.

A UNICORE job can be build from multiple parts which can be executed asynchronously or sequentially dependent<sup>5</sup> on different systems at different UNICORE sites. Currently the following elements are offered by the Job Preparation Agent (JPA):

- Script task, to submit existing job scripts via UNICORE;
- Compile-Link-Run task, to prepare new applications;
- Job groups, to build subjobs for other target systems.

A new UNICORE job is created by selecting *New* from the JPA's *File* menu. This function establishes a frame with general information for the job, as there are

- the job name;
- the target system, which can be selected from the list of all systems available to UNICORE;
- the account id, it is currently not used but will be when it becomes necessary for a UNICORE user to have different accounts;
- the user's e-mail address, where the system should send messages to.

Figure 5 shows the input panel and an example list of known target systems the user can select the *Destination* from. User defaults for these values can be set via *Preferences*.

<sup>4</sup>see <http://www.cert.dfn.de/dfnpca>

<sup>5</sup>This is the current implementation status; during the UNICORE Plus project more kinds of dependencies will be included, i.e. for parallel execution.




Figure 5: General Information in a Job Group

When the frame is set up a tree icon is placed in the upper left corner of the GUI window with the specified job name attached to it. The right hand side of the window is initially left blank. The JPA's *Job* menu offers a choice of tasks a user may add to the current job level: *Script*, *Compile-Link-Run*, *Transfer*, and *Job group*. Selecting one of these a new icon is added to the job tree on the left and an icon is placed on the right side. The added task can be filled with input when the icon is selected on the left side which opens the appropriate input window on the right. Figure 6 shows an example of the JPA with a job displayed at top-level. The job frame icon with the job name (*gm\_and\_lm*) is highlighted.

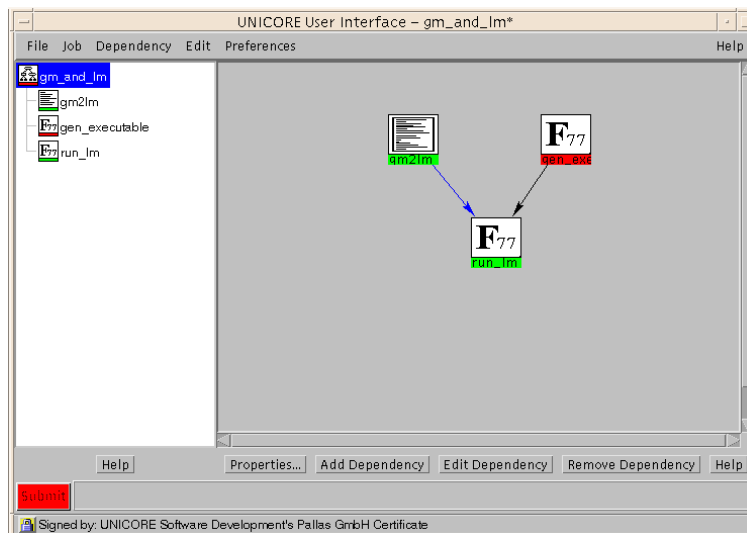


Figure 6: A Job Example

This UNICORE job consists of three different job steps, all of them are to be executed on the same target. *gm2lm* is a script as can be seen from the icon. The other two are compile-link-run tasks. On the right hand side of the window the icons of the selected job-group are shown. They are marked by colors representing their preparation status (red - not yet ready for submission; green - ready for submission). In case all steps of a job are marked *ready for submission* the *Submit* button also turns color and is activated for usage. The icons on the right can be moved by the user to an appropriate location and they can be connected via dependencies to reflect the necessary synchronization between the job steps. There are two dependencies already defined via the *Add Dependency* button at the bottom. They result in an independent execution of *gm2lm* and *gen\_executable* while *run\_lm* can be scheduled only after both predecessors have finished successfully. Within the dependency definition (see Figure 7) the user can specify the data set names of the files created by the predecessor and which are needed by the successor job step.

Another example shows the input panel for a compile step of a compile-link-run task (see Figure 8). It is divided into three major areas: general information, Input, and Output. Besides job step name the resources for the job have to be

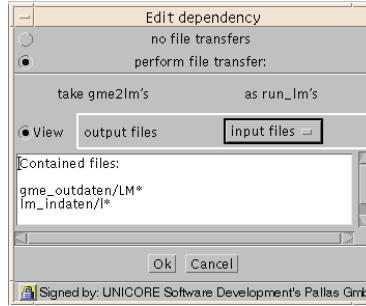


Figure 7: Dependency Definition

specified. Figure 9a shows the current resources input panel. The number of processors or CPUs, the connect or CPU time, the amount of memory, and job disk space requirements can be specified. From the resource information received from the NJS systems the JPA knows about the maximum values for the selected target system and will not allow a user to specify a value higher than that. This supports the user in creating jobs suitable for the target.

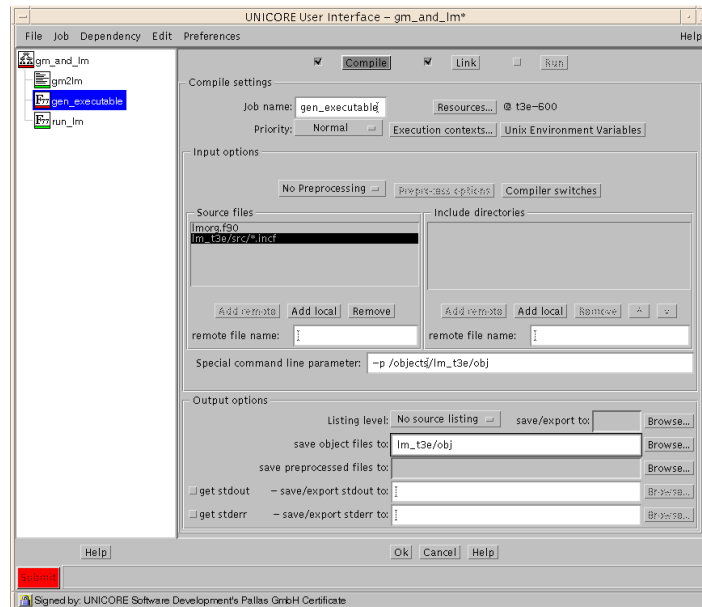
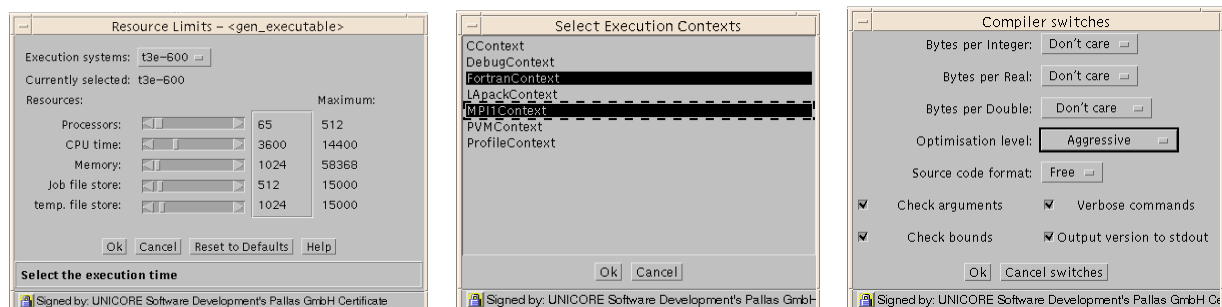


Figure 8: GUI for COMPILE Task

With the *Execution Contexts* (see Figure 9b) a user specifies the needed compiler and special libraries. Again the JPA knows from the resource pages which libraries and application packages are available. Also the *Compiler Switches* are a good example for seamless job preparation. The options specified in Figure 9c are going to be translated by the NJS in what the compiler at the target system understands. As an example, there are three *Optimization levels* offered by the interface: *No Optimization*, *Safe Optimization*, and *Aggressive*. The optimization level *Aggressive* would be translated for a Cray T3E into the Fortran 90 option `-O3,unroll2` while for a NEC SX-4 into `-Wf,-dir par -P auto -C hopt -pi -Wf,-pvctl,fullmsg, vwork=stack,noassume,loopcht=10000000`.

The Job Monitor Controller (JMC) offers functions to display the job status, to list and save job output, and to kill jobs. Figure 10 shows the JMC interface with a list of jobs submitted by the user to UNICORE. The status display lists the jobs sorted by target sites. The job icons on the left can be expanded to show the complete job tree. As within the JPA selecting an icon on the left results in displaying the dependency graph of the job at the selected level, or for script and



(a) Resource Specification

(b) Execution Contexts

(c) Compiler Switches

Figure 9: Compile Task Selection Panels

compile-link-run jobs standard output and error files. The icons are marked with colors to represent the current status within UNICORE: blue - queued, yellow - running, red - abnormally ended, green - finished successfully. The user has access to the job status and the output files until he or she explicitly removes them from the status display.

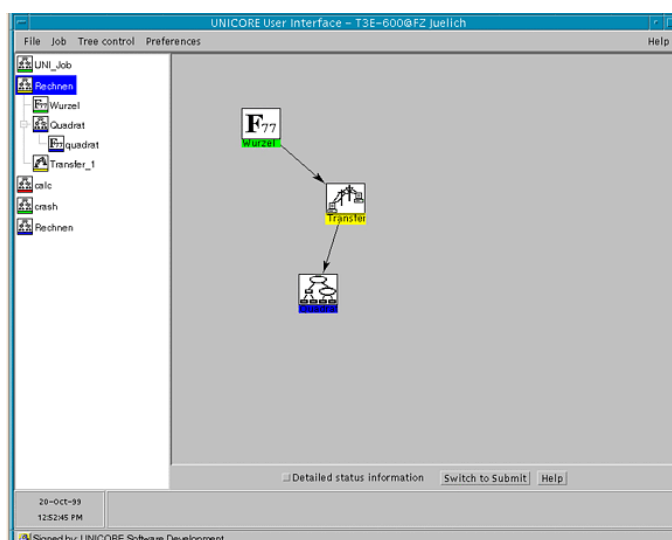


Figure 10: Job Monitor Controller

## 5 Outlook

The UNICORE system as it is today already provides seamless and secure access to remote resources, supports data transfer between the connected sites, and has established a public key infrastructure. The software written in Java allows for portability, so it can run on a variety of client and server platforms. The data base for the translation of the abstract job specification into a real batch job for the target system is easily adaptable to new platforms. Currently CRAY T3E, NEC SX-4, Fujitsu VPP 700, and IBM SP2 are integrated and the batch subsystems NQS and LoadLeveler.

Within the UNICORE Plus project the UNICORE system is going to be extended and put into production at the partner sites. In Chapter 1 the areas of development have already been mentioned. Within the short term application specific interfaces for CPMD (Car-Parrinello Molecular Dynamics), MSC-NASTRAN, FLUENT, and STAR-CD are developed.



These interfaces will guide users through the configuration step for their application and use the UNICORE mechanisms to seamlessly submit and control the jobs. The interfaces are going to be integrated into the JPA so that e.g. a NASTRAN job can be one job step in a larger UNICORE job. For the longer term it is planned to build a generic interface to easily integrate other applications.

The current resource model includes a small set of static resource information. The project will extend the set and will include dynamic information as well so that, for example, a user gets hold of current system load as a decision criterion for target system selection. The new resource model uses XML.

One of the most important issues in the UNICORE Plus project is the data management: How to efficiently transfer data between UNICORE sites, especially if large amounts of data need to be staged to a site for job processing. In addition, access to data archives will be integrated. It is also planned to add a file transfer feature for data transfer independent from UNICORE jobs but using the UNICORE security mechanisms.

In the area of meta-computing on application level three topics are under research: Scheduling of mpp-applications to be run in parallel at different sites (most of the batch subsystems do not support features like advance reservation), techniques for distributed execution of applications including a fault tolerant, batch oriented startup procedure, and mechanisms for visualization of the performance of such meta-computer applications.

All these efforts together with strong testing and quality management will lead to a user, site, and system friendly infrastructure for grid computing.

## References

- [1] ALMOND, J., AND ROMBERG, M. The uncore project: Uniform access to supercomputing over the web. *Proceedings of 40th Cray User Group Meeting, Stuttgart, Germany* (1998).
- [2] ALMOND, J., AND SNELLING, D. Unicore: Uniform access to supercomputing as an element of electronic commerce. *Future Generation Computer Systems* 613 (1999), 1–10.
- [3] DURBIN, D., MCGREGOR, R., OWLETT, J., AND YEOMANS, A. *Java Network Security*. Prentice Hall, 1998.
- [4] FEGHHI, J., FEGHHI, J., AND WILLIAMS, P. *Digital Certificates - Applied Internet Security*. Addison-Wesley, 1998.
- [5] FOSTER, I., AND KESSELMAN, C., Eds. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufman Publishers, 1998.
- [6] ROMBERG, M. The uncore architecture: Seamless access to distributed resources. *Proceedings of the eighth IEEE International Symposium on High Performance Distributed Computing, HPDC-8, August 3-6, Redondo Beach, CA* (1999), 287–293.