

INTEGRATION OF THE ENANOS EXECUTION FRAMEWORK WITH GRMS FOR GRID PURPOSES *

Ivan Rodero, Francesc Guim, Julita Corbalan and Jesus Labarta
Barcelona Supercomputing Center, Universitat Politècnica de Catalunya (UPC), SPAIN

ivan.rodero@bsc.es

francesc.guim@bsc.es

julita.corbalan@bsc.es

jesus.labarta@bsc.es

Ariel Oleksiak, Krzysztof Kurowski and Jarek Nabrzyski
Poznan Supercomputing and Networking Center

ariel@man.poznan.pl

kikas@man.poznan.pl

naber@man.poznan.pl

Abstract The eNANOS is an execution framework developed in the Barcelona Supercomputing Center. One of its main objectives is to provide a framework to execute multilevel parallel applications with low-level support. It is also able to provide information about the execution behavior of applications in run time. This information can be used by a Grid Resource Broker or metascheduler to improve its scheduling and resource strategies and the execution platform can improve the execution time of applications and resource usage as well. In this paper we discuss the steps that we have to follow to integrate the eNANOS execution environment into the GRMS infrastructure, developed by PSNC. In particular we are interested in the mechanisms to allow the integration of the different components and how to use the information provided by eNANOS to improve the scheduling strategies in the GRMS system.

Keywords: Grid resource management, scheduling strategies, application performance, eNANOS, GRMS.

*This work has been partially supported by the CoreGRID European Network of Excellence (FP6-004265) and by the Ministry of Science and Technology of Spain under contract TIN2004-07739-C02-01.

1. Introduction

The eNANOS is an execution framework developed in the Barcelona Supercomputing Center. One of its main objectives is to provide a framework to execute multilevel parallel applications with low-level support. Furthermore the eNANOS architecture is based on the idea of coordination between the different layers [9]. Currently the eNANOS Execution platforms uses the eNANOS Grid Resource Broker [8] which manages the jobs from the Grid layers in coordination with the local resource environments.

The eNANOS System is also able to provide information about the execution behavior of applications in run time, such as its progress or the obtained performance in a given moment. This information can be used by a Grid Resource Broker or metascheduler to improve its scheduling and resource strategies and the execution platform can improve the execution time of applications and resource usage as well.

The main effort of the PSNC in the Grid resource management is the GRMS resource broker [2]. GRMS is an open source meta-scheduling system for large scale distributed computing infrastructures. Based on the dynamic resource selection, mapping and advanced grid scheduling methodologies, it has been tailored to deal with resource management challenges in Grid environments, e.g. load-balancing among clusters, setting up execution environments before and after job execution, remote job submission and control, files staging, workflow management and more.

In this paper we discuss the steps that we have to follow to integrate the eNANOS execution environment into the GRMS infrastructure. In particular we are interested in the mechanisms to allow the integration of the different components and how to use the information provided by eNANOS to improve the scheduling strategies in the GRMS system. In Section 2 we present the eNANOS approach and its main characteristics and in Section 3 we introduce the GRMS system. In Section 4 we study the possibilities of integration between the eNANOS and GRMS systems, and, finally, in Section 5 we present the conclusions and future lines of work.

2. eNANOS Execution Environment

The eNANOS project aims at developing an execution platform for parallel applications on Grids. The objective of this platform is to provide support to implement and evaluate resource management and scheduling policies on a computational Grid. The eNANOS project is based on the idea of having a good low level support for performing a

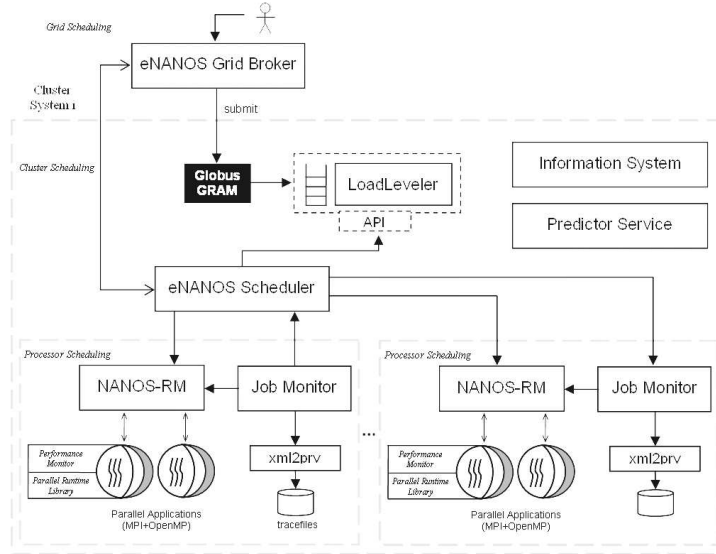


Figure 1. eNANOS Overall Architecture.

good high level scheduling taking into account these ideas: A fine grain control between scheduling levels, dynamic allocation (MPI+OpenMP jobs) to improve system performance, detailed information about current scheduling and performance to improve future scheduling decisions, and, based on this accurate information, efficient scheduling (in terms of slowdown) based on performance prediction.

The scheduling strategies are based in the coordination between the different layers involved in the execution of a job: Grid scheduling, cluster scheduling and processor scheduling. The scheduling decisions are known by the other elements in the system and are taken based on direct information, not based neither on estimations nor just observations.

The idea is to provide well defined API between levels to both forcing scheduling decisions (for instance specific allocations) and getting detailed information (for instance the real performance reached by a job in run time). A general overview of the system architecture is shown in Figure 1.

The Grid jobs are managed by the eNANOS Resource Broker and submitted to the local HPC resources through the Globus infrastructure. The eNANOS Resource Broker is developed on top of Globus Toolkit 3 as a Grid Service and it is compatible with both GT2 and GT3 services. It implements the Resource Discovery, Selection and Monitoring, the Job Submission, and the Job Monitoring. Moreover, the eNANOS Bro-

ker provides a set of Grid service interfaces and a Java API that can be used from command-line clients, applications or portals. The eNANOS broker has been extended to provide more functionality and supporting JSDL 1.0 with a gateway implemented as a GT3 Grid service as well. More detailed information about this broker can be found in [8].

For the cluster management eNANOS uses LoadLeveler as a queuing system and the eNANOS Scheduler as an external scheduler to manage the local jobs. It implements scheduling policies based on FCFS and backfilling (currently in development) guided by the information regarding the behavior of applications and performance obtained in runtime from the processor scheduler. So, the scheduler centralizes the information from the runtime about the allowed multiprogramming level of multiple computational nodes and, also some information related to the applications and hardware (e.g. load of nodes or CPUs used by each job). It also communicates with the eNANOS Broker in order to provide information to the upper level.

The CPU scheduling is performed by the NANOS-RM. Its main goal is to efficiently distribute a set of processors between a set of applications that are under its control. It implements scheduling policies based on dynamic allocation in a two-phase fashion (multilevel). The first phase is between applications and implements a FIFO policy, and the second phase is between processes of a MPI+OpenMP application and implements Equipartition [7] and Dynamic Processor Balancing [3]. The idea is sharing the required information for improving the whole system performance without penalizing the applications performance independently. The NANOS Job Monitor provides the monitoring information about the execution of workloads in XML that can be translated to trace files format to visualize and analyze them [1].

The information system can be seen as a meta-information system that can collect a very large kind of different information, providing a uniform access to it. The Predictor Service provides predictions about the job performance that can be used by both the eNANOS Scheduler and the eNANOS Broker. It implements two kinds of prediction techniques: those that are based on statistical approaches that use estimators; and those that are based on data mining algorithms. Both information and predictor systems are not integrated yet but their services are available.

3. Grid Resource Management System (GRMS)

GRMS is an open source meta-scheduling system for large scale distributed computing infrastructures. Based on the dynamic resource selection, mapping and advanced Grid scheduling methodologies, it has

been tailored to deal with resource management challenges in Grid environments, e.g. selection of the best resources for end-users, setting up execution environments before and after job execution, remote job submission and control, files staging, managing workflows and more. GRMS makes use of low-level Globus services (including the GT4 version). GRMS connects to the core services through a set of Java and C APIs. In particular, GRMS uses GRAM, GridFTP and GRIS/GIIS services. As a persistent service, GRMS provides a set of well-defined GSI-enabled Web Service interfaces for various clients, e.g. applications, command-line clients or portals. Moreover, GRMS is able to take advantage of middleware services, e.g. the Grid Authorization Service, data management systems (e.g. SRB) as well as to interoperate with the infrastructure monitoring tools such as the Mercury Monitoring System. Therefore GRMS is in fact one of the main components of a grid middleware layer that can be organized in many different ways depending on particular infrastructure and applications. The architecture of GRMS together with a set of its internal modules, namely Job Queue, Job Registry, Job Manager, Resource Discovery and a central unit called Broker Module is presented in 2. Worth mentioning is also the Resource Discovery Module, that monitors a status of distributed resources. It uses a flexible hierarchical access to both central (GIIS) and local information services (GRIS). The aim of the Broker Module is to control the whole process of resource and job management. The broker was designed in such a way that it allows us to implement various scheduling and policy plug-ins. One of the plug-ins mentioned later on in this paper, called Reschedule plug-in, is responsible for jobs migration and rescheduling within GRMS.

GRMS uses the multicriteria decision support methods for job scheduling. Based on user preferences it evaluates and selects adequate resources. The multicriteria techniques used in Grid resource management were described in details in [4]. Examples of use of this model for specific Grid resource management problems were presented in [5][6].

4. Integration Issues

To start with the integration of the eNANOS execution environment with the GRMS system we need to identify the common features and what are the functionalities provided by eNANOS that can be useful for the GRMS. Currently, in the complete infrastructure of the eNANOS Execution Framework the Grid resource management is performed by the eNANOS Resource Broker. The idea was to implement a customized environment to execute efficiently multilevel parallel applications, but

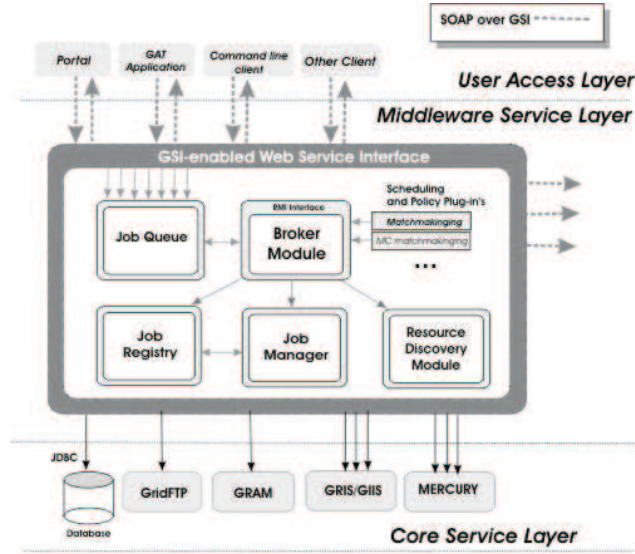


Figure 2. General GRMS architecture

one of the desirable features for one of this kind of systems is to provide generic interfaces. Therefore, the local execution system should be able to be integrated to other Grid Resource Brokers or metaschedulers such as GRMS. Since the eNANOS system supports both traditional jobs and multilevel parallel jobs, the resources supporting eNANOS can be seen as a normal resource or as a specific case for those jobs that require support from the local environment. The idea of integrating this execution environment into a Grid is to improve the behavior of the Grid in general, taking advantage of the complementary functionalities provided by eNANOS to improve the the execution time of applications and the resource usage.

On the other hand, one of the important capabilities of the eNANOS System is the coordination between the different layers involved in a job execution and, in particular, the detailed information that is able to provide to the Grid level regarding the behavior of the applications in run time. The eNANOS offers some libraries and tools to obtain detailed information in run time about the progress of the applications and the performance as well.

As a Grid Resource Broker, GRMS performs the scheduling of jobs and the resource management following some particular strategies. Hence, the GRMS system can consider the information offered by eNANOS regarding progress, performance and fine-grain monitoring to improve

the scheduling strategies implemented into the scheduling engine of the GRMS.

Therefore, we have identified two different possible integration issues: using eNANOS as a new execution platform, and using eNANOS as a new information provider.

4.1 eNANOS as a new execution framework

Even though GRMS has been designed as an independent set of components for resource management processes, the core services used by GRMS are included into the Globus Infrastructure. As the eNANOS system is based on Globus, the integration of both components should be possible without changing the interface to the local resources.

In the eNANOS Execution Framework Globus is used as a middleware and a customized LoadLeveler jobmanager in local resources. Since the GRMS expects to find a Globus-based resource with a particular jobmanager for the queuing or fork system, the resources with eNANOS system installed should be presented as a new resource for the scheduling process. But we have to take into account that the eNANOS-enabled systems have some differences so it is needed to perform a special treatment for this kind of resources.

Since the main approach of the eNANOS system is the low level support for multilevel applications, the GRMS system can perform a special use of the eNANOS resources. In particular the most important consideration regarding this issue is the specification of parallel jobs.

The eNANOS system supports JSDL 1.0 for the specification of jobs. Moreover, in the eNANOS context is being implemented an extension of the JSDL 1.0 for multilevel parallel jobs [10]. In Figure 3 it is shown the XML schema of this extension for the JSDL.

The most important attributes for the local environment are:

- Application type. It is an enumeration specifying the kind of application regarding its programming model (MPI, OpenMP, etc.).
- Number of MPI processes. It is just a positive integer value.
- Number of OpenMP threads. It is just a positive integer value.
- Malleability. It is a boolean that indicates if the application is malleable or not.
- Topology. It is an enumeration type specifying the topology of the application. This value can be taken into account by the LRMS to decide the number of processes or threads that will be spawned for the application (specific, power of 2, etc.).

GRMS uses its own job description language (GJD). It is a XML-based language, which allows specifying a description of the job executable and also the job resource requirements. There are available several parameters in the GRMS Job Description, including the location of files, arguments, file arguments, executable, environment variables, standard input/output/error, checkpointing definition, name of hosts for the job execution, operating system, required LRMS, network parameters, system paths, minimum memory required and so on. Regarding the specification of the parallelism details of the applications, the GJD allows specifying the type of the application with the `<executable>` tag and the number of required processors with the `<cpucount>` tag. It only allows the threads and MPI programming models for parallel applications. The `<executable>` tag contains 'count' and 'type' attributes that denotes the number of executions of the executable and the way the job-manager submits a job respectively. For the 'type' attribute the following values are available: single (only 1 process or thread will be started), multiple (start count processes or threads), mpi (use the appropriate method to start the job compiled with a vendor-provided MPI library, the job is started with count nodes).

There are some attributes required by eNANOS not supported by the GRMS job description language. Thus, it is required to extend the specification of details about parallelism for applications in GRMS. There is basically two different ways: extending the GJD language and using a mechanism to specify more details with the current semantic.

We have decided not extending the language used to describe jobs in the GRMS system. We have used some simple mechanism such as

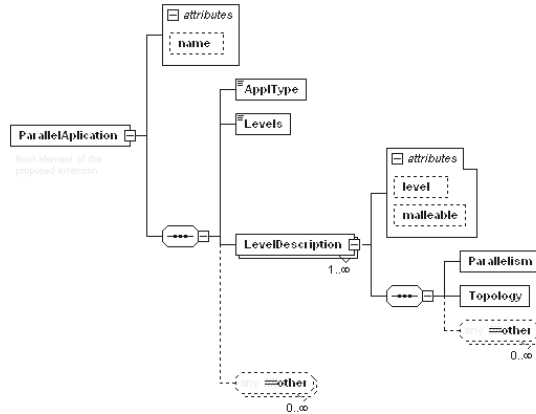


Figure 3. JSDL 1.0 extension for parallel jobs schema.

environment variables. We have chosen existing variable names from OpenMP for a good understanding and some other for other semantics. The complete list of considered environment variables is the following:

- OMP_NUM_THREADS. Indicates the number of OpenMP threads.
- OMP_SCHEDULE. Indicates the scheduling policy to follow by the OpenMP runtime.
- PAR_MALLEABLE. Indicates if the parallel application is malleable or not.
- PAR_TOPOLOGY. Indicates the kind of topology followed by the application.

These environment variables can be included externally to the GRMS system just by the job description file.

There is another important issue to take into account regarding the definition of jobs: the identification of the jobs. Since this is a coordinated infrastructure, the mapping of Job IDs is pretty important. It is very important the mapping of IDs and the description of jobs because there are several layers involved in the architecture. From both the Grid and the local environment it is required to know the mapping of a given job into the other layers (for example to ask for information to a local system). For this reason we need to include into the GRMS system another environment variable, but in this case it has to be assigned inside the GRMS system not by the user (because the Job ID is unknown before its submission). The GRID_ID_ENV variable identifies the Grid job (given by the GRMS system).

4.2 eNANOS as a new information provider

In order to provide more details of job execution in the local environment (where eNANOS gives support for parallel applications) we have found two possible ways of integration: provide an API for eNANOS functionalities that can be used directly from the GRMS system, or using an Information System (such as Mercury) and adding a new plugin to eNANOS collect the information provided by eNANOS (information queried with general mechanisms).

The second solution should be better in terms of extensibility and generic, but for the first steps of integration is more complicated since there is no plugin implemented yet for eNANOS and the GRMS system has to be modified to support the new data in the scheduling engine.

Therefore, in this first step we have implemented an API for the eNANOS Execution Framework that allows the GRMS system to ob-

tain information related to the application behavior in run time. As it is discussed in the previous section, the identification of jobs is very important to map the Grid jobs into the local execution platform when getting data. The API has been implemented in Java and consists of the following basic functionalities:

- **geteNANOSProgressInfo** (JOB_ID, RELATIVE/ABSOLUTE)
- **geteNANOSPerformanceInfo** (JOB_ID, RELATIVE/ABSOLUTE)
- **geteNANOSLocalInfo** (JOB_ID)

The schema of the information returned by the API methods includes several data from the local execution systems. Actually, when a job is in the RUNNING state, the XML includes a set of applications with their respective processes and threads. For each of these processes and threads there is also included information about performance metrics (such as MIPS or MFLOPS). This information is very important to know the behavior of the job since our project is targeted to parallel HPC applications, MPI and MPI+OpenMP as well.

This fine-grain information is provided by the NANOS Resource Manager (NANOS-RM) through its PS interface. The NANOS-RM is the responsible of the CPU scheduling in the local system, so it is able to know the details of both the parallel applications and the CPU resources in real time. It provides for each application the number of processes (e.g. the number of MPI processes), the number of threads per process (e.g. the number of OpenMP threads), the status of these threads and the CPU where the threads are allocated in case of running threads.

The eNANOS execution environment also provides information about the progress and performance of the applications in run time. This information can be queried as an absolute value (for example 1500 MFLOPS) or as a relative one (value from 0 to 100 percent). Currently there are two different ways to generate this information: using a library to instrument the application or just using the library for performance developed on top of the progress indicators infrastructure. In the local system, the progress information is managed by a dedicated daemon which provides an API to access to the indicators. The NANOS Scheduler also queries to the daemon API to collect this kind of information.

GRMS can use various scheduling strategies depending on its configuration and available environment. In most cases it uses multicriteria models for Grid scheduling (e.g.[4][6]). Some of these models allow modifications of schedules in runtime based on dynamic information gathered from an environment. These methods can particularly take advantage

of functionality provided by eNANOS since it allows getting information about performance of applications in runtime.

In particular, in [5] the dynamic rescheduling procedures implemented in GRMS are presented. They use checkpointing and migration mechanisms supported by GRMS. The scheduling model is also based on multicriteria methods. The complete description of the algorithm is available in the quoted paper. In a nutshell, GRMS reschedules automatically 'smaller' jobs if a 'large' job waiting in a GRMS queue cannot be executed due to lack of sufficient amount of resources (e.g. CPUs). By 'small' and 'large' jobs we mean here jobs that require small and large amount of resources respectively. A decision concerning which job should be migrated to which host is done based on multiple criteria.

We used in our experiment the following criteria for the evaluation of destination hosts: available memory, mean load, CPU count, and CPU speed. The following set of criteria was applied to evaluate checkpointing and migration costs: the number of hosts a job can migrate to (to minimize the risk of a failure), the size of a migrating job (memory allocated by this job), the job's current runtime (in order to migrate jobs that are not to finish soon).

Of course, the current runtime is not the best criterion for evaluation of migration costs of a given job. A job progress is much better metric for this. Taking advantage of this eNANOS functionality GRMS can better evaluate candidate jobs for migration and in this way improve efficiency of the whole schedule (for instance mean job completion time).

Another possibility of exploitation of eNANOS functionality for efficient dynamic rescheduling is use of information about job performance. If this performance is below a certain level (and progress is not advanced too much) then a better resource is found and a job is migrated to that resource (of course if checkpointing of this job is possible).

Ideally the information about a level of performance (or even better relative performance) below which a job should be rescheduled could be provided by end-users. They would have to specify it in a certain kind of agreement. Since currently such information cannot be passed in a GRMS job description we will test these issues in the future.

5. Conclusions and Future Work

In this paper we have presented the eNANOS Execution Framework as a target infrastructure to be integrated with Grid services such as a Resource Broker or metascheduler as in the case of GRMS. We have studied how to integrate these two components and which of the required changes are realistic. We also have studied how to describe accurately

the parallel jobs in the Grid layer and how it influences the whole architecture (specially the lower layers). We have solved this problem with a very simple solution based on the use of environment variables. We have implemented an API to allow the GRMS system to obtain specific information from eNANOS. This API has been implemented in Java and provides the functionality to get information about the progress, performance and accurate description of them in a run time. Finally, we have discussed the possibilities for improving the scheduling strategies of the GRMS system using the functionality provided by the eNANOS system.

As future work we should include the implementation of new plugin for the Mercury Information System to allow GRMS to get the information provided by eNANOS with more generic mechanisms. However, some changes in GRMS are needed. We also expect to implement new scheduling strategies and policies into the GRMS system and, finally, evaluate some real workloads on a testbed composed by resources of both PSNC and BSC institutions.

References

- [1] CEPBA Tools Web Site. <http://www.cepba.upc.edu/tools.i.htm>
- [2] Grid Resource Management System (GRMS), <http://www.gridlab.org/grms>
- [3] J. Corbalan, A. Duran, J. Labarta. Dynamic Load Balancing of MPI+OpenMP applications. ICPP04, Montreal, Quebec, Canada. 2004.
- [4] K. Kurowski, J. Nabrzyski, A. Oleksiak, J. Węglarz. Multicriteria Aspects of Grid Resource Management. Grid Resource Management edited by J. Nabrzyski, J. Schopf, and J. Węglarz, Kluwer Academic Publishers, Boston/Dordrecht/London, 2003.
- [5] K. Kurowski, B. Ludwiczak, J. Nabrzyski, A. Oleksiak, J. Pukacki. Improving Grid Level Throughput Using Job Migration and Rescheduling Techniques in GRMS. Scientific Programming. IOS Press. Amsterdam The Netherlands, 12:4, pp.263-273, 2004.
- [6] K. Kurowski, A. Oleksiak, J. Nabrzyski, A. Kwiecień, M. Wojtkiewicz, M. Dyczkowski, F. Guim, J. Corbalan, J. Labarta. Multi-criteria Grid Resource Management using Performance Prediction Techniques. CoreGrid Integration Workshop, Pisa, 2005.
- [7] C. McCan, R. Vaswani, J. Zahorjan. A Dynamic Processor Allocation Policy for Multiprogrammed Shared-Memory Multiprocessors. ACM Transactions on Computer Systems. 1993.
- [8] I. Rodero, J. Corbalan, R.M. Badia, J. Labarta. eNANOS Grid Resource Broker. P.M.A. Sloot et al.(Eds.), LNCS 3470, pp. 111-121, Amsterdam, 2005.
- [9] I. Rodero, F. Guim, J. Corbalan, J. Labarta. eNANOS: Coordinated Scheduling in Grid Environments. Parallel Computing (ParCo), Malaga, Spain, 2005.
- [10] I. Rodero, F. Guim, J. Corbalan, J. Labarta. How the JSDL can Exploit the Parallelism?. 6th IEEE International Symposium on Cluster Computing and the Grid, pp. 275-282, Singapore, 16-19 May 2006.