Project acronym: EDGeS

Project full title: Enabling Desktop Grids for e-Science

Grant agreement no.:RI 211727

EGEE application porting to the EDGeS platform

Characteristics of portable EGEE applications

Author: Gabor Szmetanko

Version: 1.0

Date: 10/07/2008

Table of Contents

1 Introduction

Grid user communities require large computing power not available or accessible in current scientific e-Infrastructures. In order to support the specific needs of these scientific communities EDGeS interconnects the largest European Service Grid infrastructure (EGEE) with existing Desktop Grid (DG) systems. Making a bridge between these two types of Grid systems will enable the users to transparently execute applications on any arbitrary platform involved in the new infrastructure. Running applications on this joint platform, however, requires the applications to meet certain requirements.

This document describes the characteristics of a typical EDGeS application and the requirements an application must meet in order to be able to run on this joint platform. The intended audience of this document are those EGEE users (application developers) who have an EGEE application running on EGEE and are not satisfied with the performance of the application. Based on the requirements and guidelines described in this paper EGEE users will be able to determine whether their application is capable of running on the EDGeS joint platform.

The first sections of the document describe the characteristics of both EGEE and Desktop Grid based applications. Based on these properties the next section introduces the concept of EDGeS applications and familiarises the reader with their characteristics and the most important requirements they should meet. The last section emphasises the benefits of using the joint EGEE-DG platform.

2 Service grids

Service Grids provide reliable 24 hour service operated and looked after by professional system administrators. A resource, typically a computing cluster, becomes part of a Service Grid by installing and running a highly complex middleware package. Although these Grid systems provide guaranteed service, the number of resources offered by them is restricted to the magnitude of a few tens of thousands of processors. EGEE is a typical service grid. EGEE and applications running on EGEE have the following characteristics:

- wide range of application models are supported (sequential, parallel including MPI and parameter sweep applications as well as complex workflows);
- shared data storages are available;
- applications can be interactive;
- high-level services can be used (monitoring, metadata management, portals, etc.);
- anyone can submit jobs with a valid certificate (no validation of job executables before submission).

3 Desktop grids

Desktop Grid or Volunteer Grid systems are capable to offer computing resources in the magnitude of millions of PCs. The installation and maintenance of the client side software is extremely simple. The resources are donated by individuals or institutions and their computing power can be utilised for Desktop Grid computations whenever the processors are idle. However, this also means that DG systems come with absolutely no guarantees concerning their quality of service. Characteristics of DG systems and applications can be summarised as follows:

- only master/worker or parameter sweep parallelisation is supported;
- no available shared data storage;
- MPI or internal communication between worker nodes is not allowed;
- nodes can use the results of other nodes, but only through the server;
- DG jobs (work units) are typically long running jobs with small or medium-sized (max. 100 MB per slave) inputs and outputs;
- only trusted applications can run on DGs (the DG administrator has full control over the running applications).

4 Characteristics of EDGeS applications

The main objective of the EDGeS platform is to interconnect EGEE with existing Desktop Grid systems. In order to enable the users to transparently execute applications on any arbitrary platform involved in the new infrastructure applications must fulfil certain requirements. Since applications must be able to run on both EGEE and DG systems, characteristics of EDGeS applications can be derived from the intersection of the features of the two grid systems:

- only master/worker or parameter sweep parallelisation is supported;
- no available shared data storage;
- MPI or internal communication between worker nodes is not allowed;
- nodes can only use the results of other nodes through the server;
- typically long running jobs with small or medium-sized (max. 100 MB per slave) inputs and outputs;
- DG systems only allow trusted applications (validated), therefore EDGeS applications must go through a validation process.

As we can see, the requirements of an EDGeS application are pretty much the same as the requirements of a DG application. If these criteria are met, the application might be an EDGeS application.

4.1 Limitations of EDGeS applications

4.1.1 Application types

- Parametric sweeps:
    - the size of inputs and outputs should be under 100 MB per slave;
    - if they are over 100 MB, multiple data servers and load balancing is necessary;
- Master/worker applications:
    - the same size limits apply;
- Multi-threaded applications (worker applications that start multiple threads to utilise multi-core CPUs);
- Applications using GPUs (CUDA).

4.1.2 Job execution times

In order to achieve better performance applications should adhere to the following guidelines regarding the work unit (job) execution times:

- DG systems generate a small overhead caused by work distribution. Because of this execution time of individual work units should be over 10 minutes, otherwise the overhead caused by the DG system will reduce the performance;
- DG donors (users of the computers) can interrupt the computation any time. Therefore, the execution time of single jobs should be no more than 30-60 minutes (if it is more, application level checkpointing is required to avoid loss of computation caused by user interventions);
- The execution of different jobs should take around the same amount of time (better scheduling, less load on the server).

4.1.3 Programming languages

BOINC DG systems provide APIs primarily in C and C++ (there is also a binding for FORTRAN); therefore, ideal applications are written in one of these languages. If the application has not been written in any of the previous languages, it can be still run on EDGeS by using one of the provided wrapper technologies (GenWrapper or BOINC wrapper). Running applications on XtremWeb DG systems does not require any modification in the original code.

4.1.4 Operating systems

EGEE is restricted to Scientific Linux systems. Desktop Grid systems, on the other hand, can support multiple operating systems and versions. A ported EGEE application will automatically run on DG systems where the Linux platform is supported. It is advisable to provide binaries for other platforms as well to involve more resources into the computation.

4.1.5 Other limitations

DG worker nodes have diverse resources (especially public DGs). The resources are donated by individuals or institutions and their computing power can be only utilised for Desktop Grid computations whenever the processors are idle. DG systems come with absolutely no guarantees concerning their quality of service. Therefore, applications should minimise their resource usage (primarily RAM and HDD) to maximise the number of nodes involved in the computation (the worker nodes have different amount of RAM and HD space).

5 Benefits of using the EDGeS joint platform

Grid user communities require large computing power not available or accessible in current scientific e-Infrastructures. Desktop Grids are easy-to-scale systems and able to collect 1-2 orders of magnitude more compute power than Service Grids. By interconnecting EGEE and DG systems EDGeS users can transparently execute applications on any arbitrary platform involved in the new infrastructure. Applications meeting the requirements of the EDGeS platform could utilise much more resources than in the current EGEE infrastructure.

Porting applications to the combined EGEE-DG infrastructure provides the following main benefits:

– reduced turnover time;
– improved fault-tolerance (redundant computing);
– higher throughput.