EGEE

EARTH SCIENCES: APPLICATIONS AND REQUIREMENTS

FOR THE EGEE GENERIC APPLICATION ADVISORY PANEL AND THE EGEE MIDDLEWARE

Document identifier:	**Applications_requirements-eo-egee-fv2.doc**
Date:	**30/08/2004**
Activity:	**NA4**
Authors:	**Monique Petitdidier, Joachim Biercamp, Cathy Boonne, Luigi Fusco, Julian Linford, Olivier Marti, Geneviève Moguilny, Philippe Renard, Wim Som de Cerff, Dominique Thomas, Jean Pierre Vilotte**
Document status:	**Final Version 2**
Document link:	

Abstract: This document describes the Earth Sciences Community, Applications and Requirements for using the EGEE Grid for Earth Observation, Climate, Solid Earth Physics, Hydrology, and Geophysics applications. Partners: IPSL, CGC, CRS4, DKRZ, ESA, IPGP, KNMI, UNINE, UTV; Lead partner: IPSL

Document Log

Issue	Date	Comment	Author
0	27/04/2004	First draft	Monique Petitdidier, Cathy Boonne
0.1	28/04/2004	Modifications to draft	Wim Som de Cerff
0.2	30/04/2004	Modifications to draft, Climate section added	Monique Petitdidier, Olivier Marti, Marie Alice Foujols
0.3	06/05/2004	Requirements definition EO applications filled	Julian Linford, Wim Som de Cerff
0.4	18/05/2004	Modifications to draft, Climate section changed	Monique Petitdidier, Joachim Biercamp
0.5	20/05/2004	Modifications to the draft, new application in EO, recommendations filled, IPGP Geophysics applications filled	J. Linford, M. Petitdidier, J.P. Vilotte
0.6	26/05/2004	CGG Geophysics part filled, IPGP Solid Earth Physics recommendations filled	D. Thomas, G. Moguilny
0.7	27/05/2004	Modifications to the draft and content, Review, Climate requirements filled	M. Petitdidier, Julian Linford, J. Biercamp
0.8	28/05/2004	Review. Data Access limitations	L Fusco
1	01/06/2004	Final version fv1	All the partners
1.1	28/06/2004	Hydrology section added	P. Renard
1.2	22/07/2004	Hydrology section completed	P. Renard, G. Lecca
2	30/08/2004	Final version fv2	M. Petitdidier

Document Change Record

Issue	Item	Reason for Change
0.1	Document	Format, TOC
0.5	Chapter 4	Sections filled
0.6	Chapter 5	Section filled
0.7	Chapter 3, Chapter 6	Sections filled
2	4.3.4; 4.5.4	Hydrology sections added

CONTENT

1. INTRODUCTION

1.1. Objectives

This technical note aims to provide a general description of the Earth Science applications in four different domains: Earth Observation, Climate, Hydrology and Geophysics, and the respective anticipated requirements. Several documents were written during DataGrid to describe the EO application (R7, R8, R9, R13, R15) and to specify the requirements (R1, R2, R3, R4, R10). Some of them are common to High Energy Physics, Biology and EO application work packages (R5, R6). The experience gained in DataGrid has permitted to have a better idea about common and specific requirements. In this document, the requirements, presented in the DataGrid documents, are updated taking into account the new applications in the four domains and also their priority.

Earth Sciences cover a broad range of applications, present in EGEE, like Earth Observation, Climate, Hydrology, and Geophysics. Earth Observation applications, already in DataGrid, cover the more specific tasks involved in gathering and analysing remote sensing information on the conditions of the earth's surface of land, sea, ice and atmosphere. Climate covers the evolution of the earth climate (Atmosphere, Ocean, Hydrology, Atmospheric and Marine chemistry..) from the past up to now and its prediction. Hydrology covers the study of distribution, uses and conservation of water on the earth. Solid Earth Physics covers the study of the earth itself and all the perturbations. Geophysics is concerned by the same topic but for industrial applications.

According to the user perspective, such as research and industry, the application objectives and the requirements on EGEE will be different.

1.2. Application area

The prime focus of this document will be limited to the applications, described in sections 4 and 5, which are proposed by each domain and the related requirements in section 6.

In ***Earth Observation***, all the applications concern an intensive use of data from space borne instruments. Auxiliary data are needed to validate or complement the results obtained. Several scientific topics are approached.

In ***Climate***, the applications concern an ensemble of multi-model earth system simulations, performed on supercomputer. The goal is a collaborative evaluation, processing, exchange and comparison of geographically distributed climate data.

In ***Hydrology***, the applications concern the modelling of the different processes affecting surface and subsurface waters to improve the management of endangered freshwater resources in coastal areas. These applications require intensive numerical modelling and exchange of large amount of spatio-temporal data.

In ***Solid Earth physics*** the applications concern intensive processing and use of GPS data, community sharing of synthetic seismograms and distributed computing for earth's core dynamo.

In ***Geophysics***, the applications are centred on seismic, and devoted to industrial applications like a seismic processing generic platform and intensive computation for modelling and imaging.

1.3. references

[R1] Earth Observation requirements specification, https://edms.cern.ch/document/332411

[R2] Main use case description (annex to the requirements): https://edms.cern.ch/document/346050

[R3] Earth Observation Applications: Security Use cases
https://edms.cern.ch/document/398262
DataGrid Technical note, 17 June 2003

[R4] Earth Observation Application: Metadata use cases
https://edms.cern.ch/document/414416
DataGrid Technical Note 15 July 2003

[R5] Application Working Group: Joint list of Recommendations, WP8, WP9, WP10
https://edms.cern.ch/file/386183/2/joint-list-recommandations-AWG-v1.3.doc
DataGrid Technical note, 15 May 2003

[R6] Application Working Group: Joint list of use cases, WP8, WP9, WP10
https://edms.cern.ch/file/386184/1/joint-list-usecases-AWG-v2.2.doc
DataGrid, Technical note, 22 September 2003

[R7] Earth Observation Applications Approach to Data and Metadata Deployment on the European DataGrid Testbed, L. Fusco, J. Linford, W. Som de Cerff, C. Boonne, C. Leroy, M. Petitdidier, Global Grid Forum, proceedings of the future of Grid data environments workshop, Berlin, March 9, 2004

[R8] Grid Technology for the analysis of atmospheric ozone from satellite data. S. Casadio, F. Del Frate, S. Godin-Beekmann, M. Petitdidier. Proceedings of Data Systems In Aerospace (Dasia), Prague, Tchequie, 2-6 June, 2003

[R9] Grid-Enabled Gomos, Data validation testbed, GRInGO, F. Martini, J. Van Bemmelen, Technical note, ESRIN/ESA, 14 March 2003.

[R10] Using Virtual Organizations Membership System with EDG's Grid Security and Database Access. Marco Niinimaki, John White, Wim Som de Cerff, Joni Hahkala, Tapio Niemi, Mikko Pitkanen. Grid and Peer-to-Peer Computing Impacts on Large Scale Heterogeneous Distributed Database Systems, GLOBE'2004, Zaragoza, Spain, Proceedings published by IEEE Computer Science Press.

[R11] SatSTORE, Generic Client Interface, Ref : SatSTORE/ICD, Issue: 4.1, 30/12/1999, http://styx.esrin.esa.it/grid/docs/docs/SatStoreICD401.doc

[R12] G.M. Pinna, "The European Agency Earth observation data archive management system", paper ID: 3a010 http://styx.esrin.esa.it/grid/docs/docs/ESA_AMS_HSM_paper.pdf

[R13] Earth Observation applications on European DataGrid, Luigi Fusco, Julian Linford, Pedro Goncalves, Annalisa Terracina, Monique Petitdidier, Cathy Boonne, Christine Leroy, W. Som de Cerff, J. van de Vegte, submitted to Grid Computing Journal, March 2004.

[R14] Neural networks for oil spill detection using ERS-SAR data, F. Del Frate, A. Petrocchi, J. Lichtenegger, G. Calabresi, IEEE Transactions on Geoscience and Remote Sensing, vol. 38, n. 5, pp. 2282-2287, 2000.

[R15] Neural networks for the dimensionality reduction of GOME measurement vector in the estimation of ozone profiles, F. Del Frate, M. Iapaolo, S. Casadio, S. Godin-Beekmann and M. Petitdidier, submitted to Journal of Quantitative Spectroscopy and Radiative Transfer.

[R16] Open GRID services for Envisat and EO applications, L. Fusco, P. Goncalves, F. Brito, M. Fulcoli, J. van Bemmelen, J. Linford, EOGEO workshop, London, UK, 23-25 June 2004

1.4. Document evolution procedure

This document is under the responsibility of IPSL, CGC, DKRZ, ESA, IPGP, KNMI and UNINE. Amendments, comments and suggestions should be sent to the lead partner, IPSL.

1.5. Terminology

Definitions

AMS	Archive Mass Storage
BADC	British Atmospheric Data Centre, UK
CEA	Commissariat à l'Energie Atomique, Saclay, France
CGG	Compagnie Générale de Géophysique, in Massy, France
CRS4	Center for Advanced Studies, Research and Development in Sardinia, Italy
CNRM	Centre National de Recherches Météorologiques, Météo France, Toulouse, France
DKRZ	German Climate Computer Centre, Hamburg, Germany
EGEE	Enabling Grids for E-science in Europe
ENVISAT	ENVIronment SATellite http://envisat.esa.int
EO	Earth Observation
ES	Earth Sciences
ESA	European Space Agency
ESRIN	European Space Research Institute in Frascati, Italy
GMES	Global Monitoring for Environment and Security
GOME	Global Ozone Monitoring Experiment
GOMOS	Global Ozone Monitoring by Occultation of Stars
GPS	Global Positioning System
IDL	Interactive Data Language http://www.rsinc.com/idl
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique
IPCC	International Panel of Climate Change
IPGP	Institut de Physique du Globe de Paris, Paris, France
IPSL	Institut Pierre Simon Laplace, Paris, France
KNMI	Royal Netherlands Meteorological Institute, de Bilt, Netherlands
LIDAR	Light Detection And Ranging
NDSC	Network for the Detection of Stratospheric Change
NNO	Neural Network Ozone
MPI-M	Max Planck Institute for Meteorology, Hamburg, Germany
MPI	Message Passing Interface
PVM	Parallel Virtual Machine
RIVM	National Institute for Public Health and the Environment, Netherlands
RMC	Replica Metadata Catalogue
SCIAMACHY	Scanning Imaging Absorption spectrometer for Atmospheric ChemistrY
SRON	National Institute for Space Research, Netherlands
UNINE	Neuchâtel University, Switzerland
UTV	Università di Tor Vergata, Italy
VO	Virtual Organisation
VOMS	Virtual Organisation Membership Service
WP	Work Package

Glossary

Ancillary data	Data from the space craft (e.g. attitude, orbit)
Auxiliary data	Data from other source than the instrument or space craft (e.g. wind data from ECMWF)
Catalogue	Container of product descriptors and product collections
Metadata	Additional information about EO product data and geolocation quality information, quick browse image, user help, keywords, processing parameters, e.g. related to image calibration, algorithm information, e.g. programs, interactive services, documentation
On-line	Data stored on a network accessible disk
Near-line	Data stored in MSS accessible over a network
Off-line	Data that are available only after human operator intervention
Near-real-time	Data required within a defined period after the data was collected
Close-real-time	Data required before future granules arrive and cause backlog problems
Level 1 data	Time referenced and geolocated radiances data
Level 2 data	Derived geophysical variables with the same resolution and geolocation as the Level-1 data.

2. EXECUTIVE SUMMARY

EGEE is a EU-funded project under the Information Society Technologies 6th Framework Programme. EGEE will support common Grid computing needs, integrate the computing infrastructures of the different communities and agree on common access policies. The resulting infrastructure will surpass the capabilities of localised clusters and super computing centres, providing a unique tool for collaborative computer and data intensive e-Science. Two pilot applications areas have been selected to guide the implementation and certify the performance of EGEE: the Particle Physics LHC Grid LCG, where the computing model is based exclusively on a Grid infrastructure to store and analyse Peta bytes of data from experiments at CERN; Biomedical Grids, where several communities are facing equally daunting challenges to cope with the flood of bioinformatics and health care data. Besides these two applications, other communities, partners of the NA4-Generic application, are interested in deploying their applications on the Grid. One of the generic applications is Earth Sciences, which cover five different domains: Earth Observation, Climate, Hydrology, Solid Earth Physics and Geophysics. This document aims to describe to the EGEE Generic Application Advisory Panel the Earth Science community, the applications and the main recommendations for the Grid Middleware.

The objectives and an overview of the applications proposed in each four domains are given. The added values for the community and for EGEE are also expressed. The Earth Sciences requirements and recommendations for EGEE take into account new applications and the requirements that were not fulfilled within the lifetime of the DataGrid project; nevertheless they are absolutely needed in EGEE to deploy ES applications on a large scale.

The document is organized into the following sections:

- **Section 1**. Introduction
- **Section 2**. Executive Summary (this section)
- **Section 3**. Coordination and organisation
- **Section 4**. Earth Sciences Research VO and applications
- **Section 5**. Earth Sciences Industrial VO and applications
- **Section 6**. Earth Sciences Research VO requirements and Recommendation

3. COORDINATION AND ORGANIZATION

It is most convenient to divide the Earth Sciences community and applications into two main groups, as this will allow each group to share the same data policy, the same set of resources and even the same data. The first group is concerned with research, while the second group is dedicated to industrial applications.

Therefore, two VOs are proposed: the Earth Sciences VO, made up of IPSL, ESA, CRS4, DKRZ, KNMI, IPGP, UNINE, and UTV, coordinated by Monique Petitdidier (IPSL, France), and the Geophysics VO, coordinated by Dominique Thomas (CGG, France).

However, bearing in mind the overall aim to explore and exploit commonalities in research as well as industrial applications, the intention is to maintain frequent interactions between the two groups, in order to promote collaboration and sharing of expertise and methodology.

Although we have started to address applications for the whole Earth Sciences community, the evolution of this document will concentrate only on the Earth Sciences Research community VO, leaving to others the responsibility to deal with other VO. This is especially true for the identification and definition of EGEE requirements as described in section 6.

4. EARTH SCIENCES RESEARCH APPLICATIONS

4.1. VO Organization and Coordination

Earth Sciences Research VO partners are IPSL, ESA, CRS4, DKRZ, KNMI, IPGP, UNINE, and UTV. The partners will be coordinated by Monique Petitdidier (IPSL, France).

The applications are divided in four main groups:

- Earth Observation Applications
- Climate Applications
- Solid Earth Physics Applications
- Hydrology Applications

4.2. Grid Awareness, Experience and Expertise

The work done by the EO VO in the DataGrid project represents the state of art of Grid applications in Earth Sciences, and in particular by ESA, KNMI, IPSL, IPGP, UTV, which played a documented role as partners in DataGrid.

All partners have made other specific experiences in other projects, funded by ESA, EC and/or national programmes.

4.3. Added Value for EGEE

4.3.1. EO Applications

The EDG VO partners bring substantial experience gained by deploying applications in EDG. Several of the partners have expertise in many areas of both Globus and EDG middleware (security, job management, replica management, mass storage interfacing, information system usage, etc.). Thus the community is able to make a quick start and will help to disseminate previous Grid expertise to newer members within the VO. Several important grid requirements are introduced by Earth Sciences applications. The ES requirements are well developed and formulated, they include specific application test cases designed to assess and improve performance of specific middleware features. They address several important features including job handling, security, metadata data management and sharing of large data file numbers and volumes.

The access of satellite data from ESA archive by using the GRID and the computing possibilities associated is a pioneering work in the domain. It is important to note the special actions taken by ESA to move in the EO operational environment some mature EDG components [R16], which are ready to evolve according to the EGEE operational demonstration achievements.

4.3.2. Climate Applications

The European Climate community is a well-defined community that has worked together fro many years through EC projects. Once the demonstration will have successfully proven the benefits of collaboration through the grid infrastructure all the Climate partners are expected to deploy the EGEE infrastructure.

The deployment of climate applications on EGEE will be a first initiative to future interconnection with super computer grids, like DEISA.

4.3.3. Solid Earth Applications

The main added value for EGEE is the deployment of the Grid among several communities and institutes or laboratories all over Europe at first. New resources will be implemented. The fact that through given communities, laboratories or institutes are involved is another factor to interest new communities on the Grid.

4.3.4. HYDROLOGY

As many other scientific communities, the Hydrology community is widely distributed. Various EC projects have allowed the creation of a strong European network. However the data and the different models used for investigations are not widely available through the community. This problem is even more drastic when EC researchers are collaborating with South Mediterranean partners. These countries are facing major difficulties related to water management (drought, salinization, etc.) but their local computing resources and data storage systems are inadequate to support complex modelling systems needed to analyse the potential impact of alternative management scenarios.

One of the main added values that the Hydrology community can bring into EGEE is the demonstration of the use of geographically shared resources between Europe and South Mediterranean countries to improve the sustainable management of freshwater in coastal areas.

4.4. Resources available for contributing to EGEE Grid

ESA-ESRIN: UI, CE (15 nodes) and SE (1.4 TB), with priority access by the Earth Science Research VO

IPSL: UI, CE (4 nodes) and SE (500Gb)

KNMI : UI

DKRZ : UI, CE (2 nodes), SE (Several TB, depending on application)

IPGP : UI

UTV : UI

It is expected that, following the initial consolidation phase, all Earth Sciences Research VO partners will make available adequate shared resources.

In agreement with the ESA data policy, access to ESA owned data will be only granted to users, which have approved submission to the ESA ENVISAT CAT-1 research programme.

4.5. Community, Applications and Proposed EGEE Work Programme

4.5.1. EO DATAGRID APPLICATIONS

The EDG VO partners consisted of the European Space Agency (Frascati, Italy), the Netherlands Royal Institute of Meteorology (KNMI, de Bilt, Holland), and Institute Pierre Simon Laplace (IPSL, Paris France). The partners formed a virtual organization of more than fifteen scientists and engineers who collaborated on setting up the EO Grid infrastructure and deploying the application. The partners' main objectives were to (1) demonstrate how Grid can respond to the complexity and the constraints imposed by applications in EO domain and (2) identify the benefits of the technology and how it can improve the work of EO technical and scientific users.

Figure 1. Scheme for GOME profile processing and validation using DataGrid

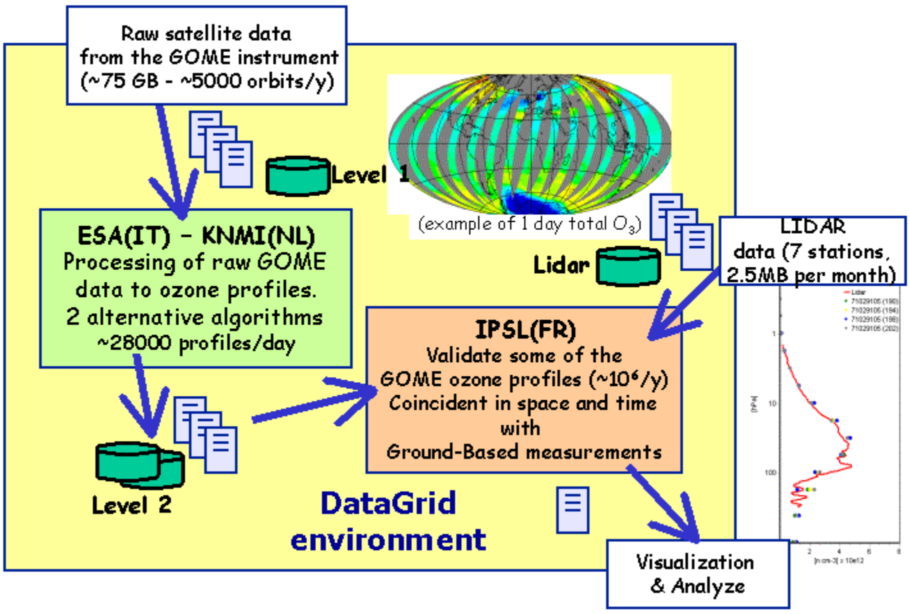

The selected EO Applications use-case involves processing and validating global atmospheric ozone observations made by the GOME instrument flying on board the European ERS satellite, throughout a 8-year mission. The GOME Ozone profiling application was chosen as an ideal candidate for evaluating the DataGrid testbed. The large data volumes and large number of files, the processing-intensive nature of the scientific algorithms and the scattering of datasets, processing resources and participating organizations over an extended geographic area, are all factors where Grid technology can offer improvements over conventional computing solutions. Furthermore, since the application is fairly representative of the product processing, refinement and quality control procedures that routinely take place in the EO applications domain, the problems and solutions encountered can be considered representational of a wide range of Earth Observation applications.

The aim was to reprocess the entire GOME dataset, consisting of 7 years of global ozone observations by satellite, using two different ozone-profiling algorithms, OPERA and NNO. The results of the satellite observations would then be validated against ground-based Lidar observations. The OPERA algorithm, based on Optimal Estimation, is developed by KNMI, while the NNO method, based on Neural Networks, is developed jointly by the University of Tor Vergata and ESA-ESRIN (Italy). The LIDAR observations are extracted from the Network for the Detection of Stratospheric Change (NDSC) database by IPSL, who also developed the validation algorithm.

A use case was developed to describe the complete, end-to-end processing and validation chain (Figure 1). The use case involves four distinct datasets shared among the three EO institutes: Level1 raw data, Level2-OPERA and Level2-NNO products and LIDAR data. Each year of observations requires approximately 267 GB of data contained in several million files (Table 1). The OPERA and

NNO algorithms used different approaches to storing the Level2 profiles. OPERA splits each orbit into several thousand Level2 pixel files, while NNO keeps the Level2 profiles together in a single orbit file.

Dataset	Number of files handled (per year)	File Size
Level 1	4,724	15 MB
Level2 (NNO)	4,724	19.5 MB
Level2 (OPERA)	9,448,000	12 kB
Lidar	12	2.5 MB

Table 1 GOME application data volumes

The discrete steps of the processing chain are shown in Figure 2. This requires transferring the Level1 data from EO archives to DataGrid Storage, running the two different processing algorithms to produce the Level2 data and then running the validation procedure to verify the results. Each institute in the collaboration carried out a particular role: ESA-ESRIN was responsible for transferring the GOME Level1 data to the Grid testbed, both KNMI and ESA-ESRIN were responsible for processing Level2 products, while IPSL was responsible for extracting the Lidar data from the NDSC database, transferring them to the Grid testbed and carrying out the validation.

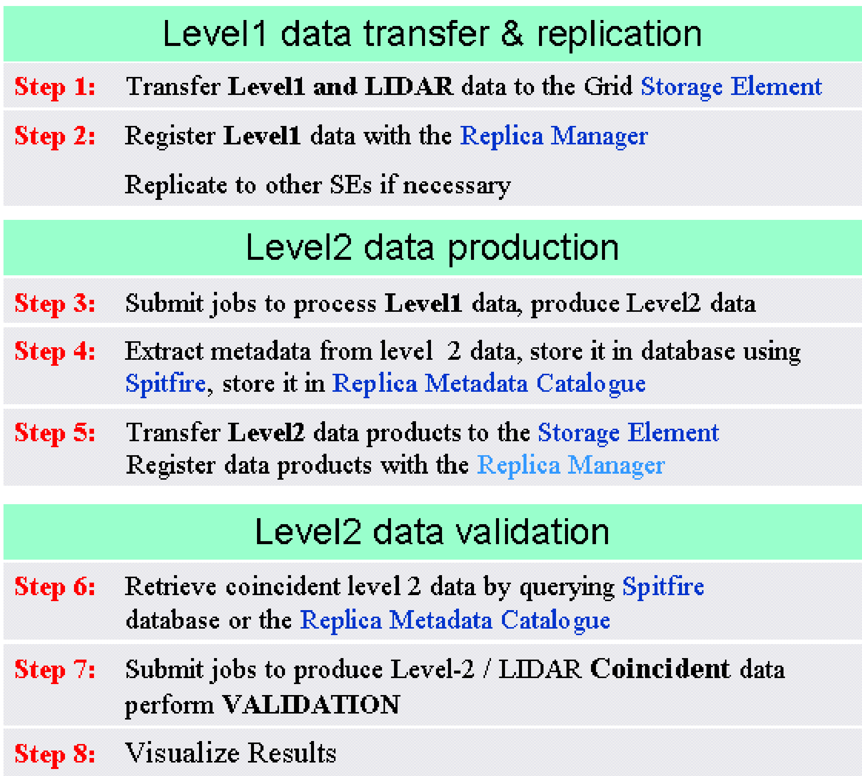

Figure 2. Discrete steps of the end-to-end GOME data processing and validation chain

4.5.1.1. EO Achievements in DataGrid

Handling of large number and large volume of files

The DataGrid EO use case concerns satellite data (Gome, Gomos) and ground based measurements. The characteristics of the satellite data are the large number of files; their volume being relatively small (Table 1).

Approximately 35,000 files were registered in the EO replica catalogue. No known limit was reached. However, 'soft' limits were encountered, due to the time taken to copy and register the files and also the limited space available on the testbed shared storage elements.

Processing with complex algorithms, parallel jobs

Complex algorithms:

In the DataGrid EO use case, Ozone profiles were retrieved by using two different algorithms. The Neural Network approach (NNO) was developed using IDL (Interactive Data Language COTS), a runtime version was installed in many nodes.

The inversion algorithm, OPERA, has used different databases and model. The application represents a very demanding consumer of computing power.

Languages available on the Grid:

The languages used in DataGrid were Fortran 77, IDL and C. In order to port new applications, the following frequently used languages are needed: IDL, Matlab or clones, Scilab or Octave, Fortran 77, Fortran 90, and C.

Parallel Processing

Some Geophysical applications developed at IPGP require a large amount of memory and use MPI (Message Passing Interface) to distribute the load across several processors that communicate with each other; this allows using the shared memory of all the processors. Several different tests were carried out unsuccessfully. The requirements for parallel applications will be described in section 4.5.3.

Creation of and secure access to metadata catalogues and data [R3, R4]

Metadata catalogues are needed by almost all ES applications. Three scenarios were tested in DataGrid by EO Work Package.

1) For some data collections the catalogue of metadata exists and is associated with the archive behind a firewall, like the MUIS catalogue (ESA). In this case a secure interface was developed to access and query it.
2) The metadata catalogue is portable or created on the Grid using secure database access, provided by Spitfire, which permits to securely update and query databases. This EDG module is transparent for the user and provides a secure access to database (e.g. MySQL). It was used for Ozone profile files (~60,000 tuples) and Lidar (<300 tuples) metadata. It is efficient to insert and query tuples.
3) The last possibility is to use the replica metadata catalogue, part of the Grid Middleware. There will be one per VO. Some attributes are defined by the users and the replica metadata catalogue is filled with each metadata tuple. In order to test that new functionality, 32,000 tuples were registered with 26 attributes, describing the Level1, Level2 satellite and Lidar data. Its performances (transaction delay, search, etc.) have to be improved.

In DataGrid, the role (read, write, update) of each VO partner was done by VOMS, which is compatible with Spitfire and the Replica Metadata Catalogue.

Access restriction through VOMS, which is an important requirement, was not implemented in the Replica Metadata Catalogue. This means all VO members can currently modify the data and metadata.

Interfacing Grid Middleware and existing operational tools and infrastructure

In the DataGrid project, an interface with MUIS, the ESA proprietary product catalogue, was elaborated by ESA [R13] and also a secure interface with the Archive Mass Storage (Figure 3). The AMS is described in [R11, R12].

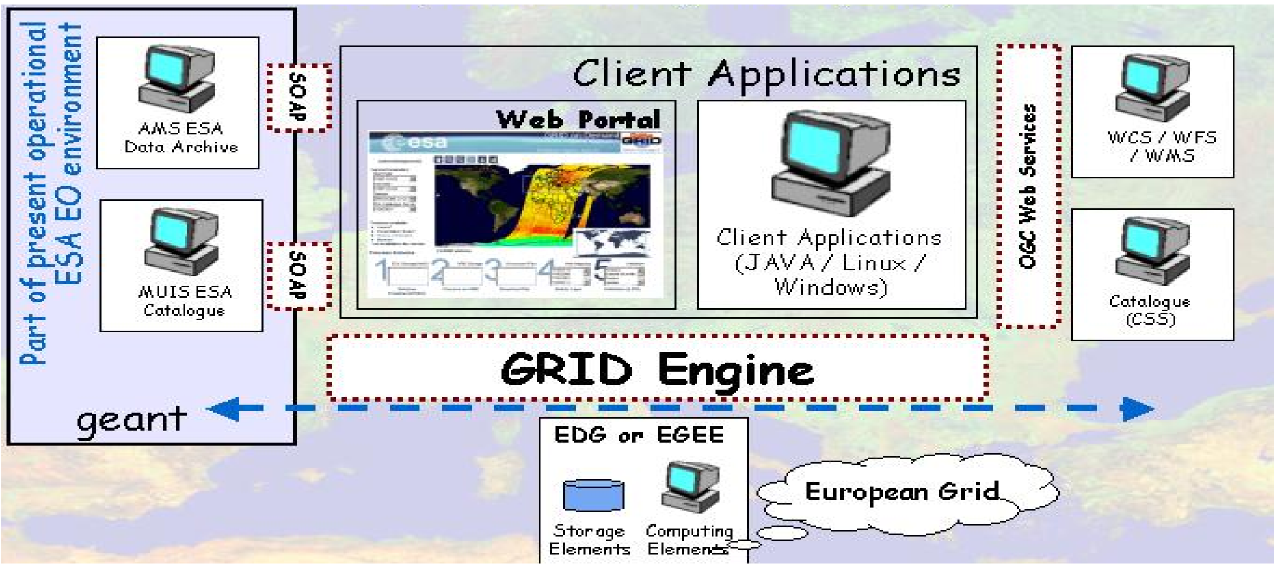

Figure 3: Schema of the interface between the Grid middleware and ESA operational archive and catalogue including also web services like the geographical information service (GIS)

4.5.1.2. Earth-Observation applications in EGEE

The work done during DataGrid, processing and validation of ozone profiles, was presented in different ozone workshops. As a consequence, several new partners have expressed interest to participate in the EO Grid community, e.g. SRON and RIVM.

In the Space community (scientific and operational), ESA has played an important role to facilitate the access to satellite data via the Grid, to convince the international and operational community of the interest of Grid, and to develop tools which facilitate the use of satellite data via the Grid. The role of the institutes has been to use the satellite data and port applications using them.

Within EGEE, the Earth Observation applications community aims to pursue several different goals:

Validation of Ozone profile retrieved from satellite ERS2, Envisat [R9, R13, R15, R16]

The approach is similar to the one described in section 4.5.1. It consists of extending the GOME Grid experience to other datasets like Gomos (tests have started- R9) and Sciamachy. It also involves processing GOME data with new algorithms and validation of the whole set of ozone satellite data. Figure 2 summarizes the procedure. The reference data will be the Ozone profiles obtained by Lidar and/or radiosounding.

Ozone studies and generation of level 3 data on the Grid

The interest of the Grid is to store different ozone databases, and then to conduct different research activities and obtain new scientific results.

There are a large variety of applications with different processing and storage needs. Some of them have been under tests like the destruction of Ozone in the polar vortex, mini holes, assimilation into models, climatology, etc.

Training of neural network algorithms for oil spill detection [R14]

Neural networks have proven to be useful tools for the design of automatic procedures seeking for oil spills in SAR images. In particular the nets are effective in the discrimination process between real oil spills and look-alikes. The accuracy of the neural performance basically depends on the training process. The use of the Grid would be able to speed up such a process and to make it more reliable as much larger databases can be considered.

Other satellite data applications [R13, R15, R16]

ESA aims to provide tools to satellite data users. Some new Grid projects have been funded by ESA and depends on the interest of the relevant communities. ESA has developed portals for demonstration of grid based MERIS (Envisat satellite instrument) mosaic processing and GOME (ERS satellite instrument) profile processing and validation. In addition, new ENVISAT satellite sensor applications, such as sea surface temperature measurement using ATSR, MERIS Level3 product processing (chlorophyll, suspend matter and water vapour) are being ported to the grid. Additional specific applications could be defined for accessing other data types. These applications have to be approved by ESA, following the ENVISAT CAT-1 research programme rules.

4.5.1.3. EO Application Requirements

Most of the EO requirements have been expressed in the different documents delivered during the DataGrid Project. Our major requirements are:

- Improvement on metadata handling (RMC, Spitfire)
- Improvement on security (restricted access, groups and roles within a VO)
- Improvement on scalability (e.g. number of files, file sizes).
- Improvement of the performance of most of the functionalities.
- Support for parallel programs (e.g. MPI, PVM)

4.5.1.4. Added Value for the EO community

Our EDG experience shows that Grid provides not only computing tools but also promotes cross collaboration and communication among diverse partners in the community. The value of the Grid has been demonstrated for carrying out extensive satellite data validation, i.e. involving different geographically distributed locations and covering several years of global observations. The Grid will be important in future to support data exchanging and sharing in the emerging large-scale European projects (GMES).

This community, whose number of members is increasing, provides also computing and storage resources for the Grid. The EO community aims to promote dissemination and use of Grid technology by developing and demonstrating innovative solutions. The Grid-on-demand portal, developed by ESA, aims to establish direct access to the ESA operational archive, combined with 'on the fly' product processing, as an alternative to the 'traditional' way of requesting the data, receiving them stored on tapes or CD-ROM, which have then to be loaded and processed using own facilities.

4.5.2. CLIMATE APPLICATIONS

Understanding and prediction of the climate system is one of the Grand Challenge Problems. For many years, European efforts have been ongoing to improve the prediction of climate change and its impact on the earth system and especially on the living conditions of mankind. Today, it is recognized that global change is much broader than climate change and that the problem of the earth system's evolution must be regarded as much more complex as it had been realized in the past. The problem has also become profoundly interdisciplinary, with the need to develop research teams that belong to different institutions, often located in different countries. Within the European FP5 project PRISM an integrated modelling framework is being created for coordinating and executing multi-institutional climate simulations. A set of portable climate community models (the main components being atmosphere, ocean, land, biosphere, cryosphere, bio-geo-chemistry and atmospheric chemistry) as well as associated diagnostic software is being developed under standardised coding conventions and will be made available to all European scientists.

Most climate studies today are based on results produced by an ensemble of extensive multi-model earth system simulations, e.g. within ENSEMBLES, a new project funded by the EU within FP6, gathering 72 partners. The model simulations are performed on high end supercomputers and produce very large data sets which are generally stored in hierarchical mass storage archives at the site where they had been created. Data volumes produced by typical climate simulations today amount to tens to hundreds of Terabytes and will increase by an order of magnitude within the next few years. There are different kind of users, interested by the outputs, in the first place the scientists working in this field, and then economist researchers interested to evaluate the impact of a climate change on the economy. The last class concerns the general public. The first objective to attain in EGEE concerns mainly the first class of users. They perform similar statistical analysis on all model outputs.

In order to estimate the number and size of files used for one analysis two examples are given. A typical climate simulation is presently 200 years long, shared in 2400 monthly jobs. Data are stored in about 25000 files of 50 MB each. A typical analysis compares around 10 simulations (i.e. 250000 files or 12.5 TB). Within two years, the model resolution will increase the size of files by a factor 10. The number and the length of simulations will increase the number of files by a factor10 to 100. Another example is a typical high-resolution ocean simulation, which is presently 50 years long. Data are stored in about 25000 files of 1 GB each.

4.5.2.1. Climate applications in EGEE

A data grid infrastructure will be implemented which allows for efficient and collaborative evaluation, processing, exchange and comparison of geographically distributed climate data.

For an ensemble of climate simulations, stored on different mass storage facilities, data sets have to be identified, accessed and processed in a coherent and transparent way. An example would be detection and retrieval of the sea-surface temperature average for a given geographical region and time period based on semantic description of the data and the analysis and visualisation of these data using standard and/or customized processing tools.

Two typical studies using the climate grid illustrate possible use of EGEE:

- For an ensemble of climate simulations, stored on different mass storage facilities, get the sea-surface temperature average on a geographical region for all the simulations length (1D fields – time) and perform analysis (Fourier transform, etc..) using either Python, Fortran 90, IDL, Ferret, ….
- For an ensemble of climate simulations, stored on different mass storage facilities, extract the sea-surface temperature on the North Atlantic region for all the simulations length (3D fields – long/lat/time) and perform analysis (EOF, SVD, etc…).
- Same analysis as above, but for an ensemble of high-resolution oceanic simulations. This study implies that all data are described by a catalogue with a unique entry.

In the medium range we plan the unified and coherent analysis and comparison of distributed simulation data as for example the simulation data currently produced for the forth assessment report of the international panel of climate change (IPCC) at different sites worldwide. The various steps of the necessary processing chain will first be implemented and tested between M&D/MPI-M (running the WDC for Climate) and IPSL on the client side, and IDRIS, CCRT/CEA and DKRZ as providers of computing and data services. In a next step the other European laboratories linked via European projects (PRISM, ENES, GEMS), e.g. the Royal Netherlands Meteorological Institute (KNMI), the British Atmospheric Data Centre (BADC) and CNRM/Météo-France will be included.

The final goal will be the provision of an unified climate data processing grid which enables earth system scientists to set up complex workflows (model simulations, data-storage, data-processing, data-analysis, visualisation) in a transparent and efficient way. This environment should rely on existing Grid technologies (Globus, DataGrid/LCG/EGEE) and integrate existing tools and frameworks used by the ESM community- as for example the PRISM system and UNICORE components – wherever appropriate and possible.

In the future, a computing grid with coarse-grid earth system models could be implemented on the Grid.

4.5.2.2. Climate Application Requirements

The deployment of climate applications on the Grid implies the following issues to be addressed by the climate community

- Unified and efficient access to data and information currently described by different meta data conventions and stored in different mass storage archives. Definition of meta-data standards with respect to earth system modelling
- Definition of semantic data models and ontologies to automatically check inconsistencies between e.g. software modules, user preferences and data formats and to support users setting up consistent and coherent workflows.

Basic requirements in the Grid infrastructure are

- Handling of large data sets stored in different places and formats: This may include data sets combining hundred-thousands of files or other data objects for one analysis
- A secure but comfortable access to the mass storage archives (single-sign-on, data security, resource sharing)
- Resource broker services for transparent access to data and processing tools
- Concepts for load balancing, quality of service and assurance of data quality
- Definition of the role and permission for each partner of a VO to access metadata and data
- Integration of different scripting languages and COTS.

4.5.2.3. Added value for the Climate community

The Grid will facilitate the collaboration and communication among the community. The Grid will facilitate cooperative studies implying intercomparisons of Earth System Model outputs (e.g. FPS Motif, FP6 Ensemble…). It will allow dissemination of climate-change simulations towards scientists studying the impacts of climate change in agriculture, economy, different region…

4.5.3. SOLID EARTH PHYSICS

Research in imaging the Earth interior and in understanding the Earth dynamical processes is increasingly reliant on large-scale computational models. In the past ten years the computational methodologies used in the fields of seismology, of earthquake fault dynamics, of core dynamos and of exploration geophysics have improved dramatically, due to the massive use of numerical methods, new observational technologies and increasing density of observations.

Today, modern computational techniques, in combination with parallel computer architectures, allow the simulation of complex, three-dimensional problems. Applications range from fundamental research (e.g. imaging the Earth's interior structure, understanding Earth's core characterizing earthquake rupture processes), to more applied or societal research (e.g. physical description of hydrocarbon reservoirs, earthquake risk assessment, etc.).

At the same time, both the resolution and volume of the surface and satellite observation data, available to the research community, through coordinated international large data-bases, continues to increase. This has led to the recognition that the problems as more complex and multi-scaled than before, and the need to develop sophisticated inversion and assimilation tools. Moreover, as size and complexity of the observational data-base increase, the treatment and the extraction of information also need sophisticated data-mining and visualization tools, along with associated increase in computational resources.

An original, interdisciplinary approach is now needed, combining data-bases and computational resources with research teams that belong to different European institutions.

Within the European FP6 project SPICE, several institutions with specialisations in physical, mathematical and computational aspects of wave propagation and earthquake dynamics, came together with the aim of developing new tools for simulation, inversion, data assimilation and visualisation.

The availability of new grid technologies and infrastructures such as EGEE need to be explored, particularly with respect to their capability to improve efficiency and collaboration. This is needed in particular for promoting the evaluation, processing, exchange and comparison of experimental and synthetic geographically distributed data. In addition, it is necessary to provide the community with common and portable tools. This suggests that, by using the Grid, the simulation, inversion and assimilation processes in geodynamics will experience a quantum jump in resolution and accuracy in the coming years.

4.5.3.1. Solid Earth Physics Applications in EGEE

Processing GPS data

Today, high precision geodetic measurements with GPS are commonly used in geophysics for crustal dynamics studies (plate motion, earthquake transient deformation, etc.). Since ten years ago, the volume of data is increasing exponentially, from both permanent (global or regional) and non permanent (experiments) instrument arrays.

The test application selected for EGEE involves processing and validating GPS observations (positions of ground stations and satellite orbits) made by various FP5 and FP6 European programs, e.g. gulf of Corinthe, and GPS networks, e.g. the French RENAG in the Alpes, together with the IGS worldwide permanent array (400 stations).

This application was chosen as a good candidate for evaluating the EGEE infrastructure due to the large data volumes, large number of files, the processing intensive algorithms (based on GAMIT, a globally adopted tool), and the scattered distribution and heterogeneity of the datasets throughout Europe. Due to the increasing number of GPS experimental campaigns, there is the need for developing a more unified and transparent set of processing tools, which could be achieved by using the Grid.

The test application will be a first step toward developing such tools. Furthermore, the application shares common features with the GOME Ozone profiling application and will therefore be able to exploit the EO experiences and methodologies developed within EDG, and also its future developments within EGEE.

GPS data processing on the Grid will use the GAMIT analysis package developed at MIT and Scripps, which is commonly used in the community. This package allows the estimation of three-dimensional relative positions of instrument ground stations and satellite orbits. The processing is organized into sessions, defined as spans, during which a group of stations simultaneously tracks the phases of two or more satellites. In order to accommodate combination with continuously operating stations, sessions are today usually organized into 24-hr spans covering a single UTC day.

Positions of the ground stations are typically estimated every day, while the extraction of the tropospheric information (tropospheric stretching and water content along the ray) is done every 2 hours. The typical size of a 24-hours station file is 2 MB; several hundred stations can be processed each time. The analysis requires a preliminary set of station coordinates (L-file), the broadcast ephemeris (E-file or RINEX navigation file) for the satellites observed, an ensemble of phase and pseudo-range observations, and auxiliary information available from the log sheets (tracking scenarios, antenna heights and meteorological data). Orbital information of the satellites are retrieved from official institutions or space agency.

The GAMIT software is designed to run under any UNIX operating system and versions have been implemented for LINUX work stations. Aside from the GAMIT package, simple tools will be added in order to extract during the data processing other information like the ionospheric TEC (Total Electronic Content) along the ray.

One typical study will involve the processing of the GPS data of the RENAG national network funded by the CNRS/INSU and will involve a collaboration with MEDIAS France which is charge of the MySQL data-base. An other study will involve the processing of the GPS data of the French Volcanoes observatory, under the responsibility of IPGP and the CNRS/INSU.

The final goal will be the provision of a unified GPS data processing which enables earth scientists to set up complex workflows (data-storage, strata-processing, data-analysis, visualisation) in a efficient way. This environment will rely on the EGEE Grid technology and infrastructure and integrate existing tools used in the GPS community, as for example the GAMIT package.

Synthetic seismograms

The theory and applications of seismic (elastic and acoustic) wave propagation and earthquake source mechanics are entering a new era in fields such as seismology and exploration geophysics. Today, seismic networks contains an increasing number of broadband instruments, accelerometers, GPS continuous stations and borehole instruments : for example, the Global Seismographic Network of the Federation of Digital Seismic Network (FDSN) contains more than 120 stations, and in the future several US and European arrays may blanket continuous parts of the word with hundreds of broadband instruments. The challenge in seismology is therefore to calculate highly accurate and very broadband synthetics in order to match observations and make use of the advances in instrumentation. Computational models, through numerical simulation and assimilation, promises new breakthroughs in areas like global seismology, earthquake fault dynamics and hazard mitigations as well as exploration geophysics.

Large simulations of earthquakes and global wave propagation in complex 3D earth structures are very demanding in terms of computational resources and rely on sophisticated numerical algorithms. For these simulations a key issue is the accessibility of an ensemble of simulation results (e.g. synthetic ground motions and seismograms), stored on different mass storage facilities, to a wide distributed scientific community for efficient and collaborative evaluation, processing, exchange and comparison.

An example would be the location and the retrieval of synthetic seismograms or ground motions for receivers within a given geographical region, based on a semantic description of the data, and the analysis and visualization of these data using standard processing tools. Such seismic test application shares common features with the Climate application and should benefit from collaborations. Furthermore, it shares also common features, but also some major differences, with the common way of processing and analysing actual observational records stored in the seismological international data-bases, such as those in the European data centres like Orpheus or Geoscope. Major differences are here that synthetic observations must make reference to the earth model used in the simulation and have no absolute time.

Two typical examples illustrate the possible use of EGEE:

- For an ensemble of simulations of a given actual earthquake, for a same fault geometry and regional structure, stored on different grid storage elements, retrieve synthetic seismograms at a given set of receivers at the surface or on the fault surface (3x1D fields in time) based on geographical criteria and perform wave-form and Fourier analysis, as well as comparison between the different numerical models and the actual data.
- For an ensemble of simulations of global wave propagation for a given earthquake and a given 3D earth model, retrieve the synthetic seismograms at a given set of receivers at the surface (3x1D fields in time) based on geographical criteria and perform wave-form analysis.

The numerical simulation tools are quite demanding in terms of memory and CPU and involve heavy parallel computations using F90 and MPI. In a first step, the simulation data will be produced either on the computational resources at IPGP or at the French national centres (IDRIS or CINES) with the associated metadata.

The metadata will be extracted and stored in a metadata catalogue on the Grid, and the synthetic data transferred and replicated on Grid storage elements and registered. In a second step, the simulations, on low resolution models, may be submitted on the EGEE infrastructure. These synthetic data will be retrieved, processed and visualized through grid submitted applications after retrieving information using the metadata catalogue on the grid.

The synthetic data should be stored in the commonly used seismological standard (MiniSEED) and the associated Metadata described in dataless and SEED format or XML-SEED. This will involve a close collaboration with the ORPHEUS and Geoscope data centres as well as with the SPICE European project.

One goal will be to demonstrate on the grid a unified processing and retrieval of synthetic data based on seismological computational models which allow earth scientists to share and compare their results in a transparent and efficient way. The next challenge would be to allow simultaneous processing of actual and synthetic observations. Some contact have been taken with the Orpheus data centre that gathers most of the European seismological data and which is located in Holland.

Earth Core Dynamo

The past decade has witnessed a major breakthrough in understanding Earth's core and the generation of the Earth's magnetic field (e.g. dynamo models). At least, 3D numerical dynamo models have succeeded in reproducing the most basic features of a self-excited geomagnetic field, such as geocentric axial dipolar average magnetic fields, magnetic secular variation and occasional polarity reversals, despite the fact that the models are still very far from realistic with regards to the actual physical parameters. The physical regimes of the core dynamos have still to be explored. Computational investigations of numerous problems have to be done in order to make further progress toward the resolution of the problem in the relevant regimes and to contribute fundamentally to the intellectual foundations in this area. Among others, investigations of turbulent flows in the rapid rotation limit, of RMHD turbulence and boundary layer instabilities are still to be explore. Such physical exploration requires a very large number of numerical simulations and the emergence of the grid technologies may help to make progress in the next years.

The test application will be to explore the possibilities of using the EGEE infrastructure and technology for a large number of computational simulations in order to speed-up the physical exploration of physical problems related to the Earth's core dynamo. The computational model is based on the software CONVECT developed at IPGP and that solve 3D convective MHD in a sphere in rapid rotation. The code is currently running on parallel vector and scalar architectures both at the CNRS computer national centre (IDRIS) and at IPGP. Based on MPI and F90, this code does not require a large number of processors but is quite demanding in terms of memory and CPU time, For each run the amount of input and output data remains small.

The first objective here will be to test the capability of mobilizing distributed computational resources within the EGEE grid infrastructure and environment, through the use of computational elements compatible with MPI applications and with a good restitution time. Up to now tests on the DATAGRID project were quite disappointing but within the EGEE program, the situation should greatly improved.

4.5.3.2. Solid Earth Physics Application Requirements

Some requirements are similar to the previous applications, some are different:

- MPI with Mirinet, SCI or other fast network and Fortran 90: The simulation proposed requires the possibility to submit a script and not only the executable. Modeling and imaging are also intrinsically parallel applications using large memory, medium to large exchanges between computing nodes with MPI, local scratch disc space and are heavy cpu intensive. Another requirement, not encountered in DataGrid, is the consistency between the compiler and libraries available on the UI and the ones needed by the cluster at the working node where the job is executed.
- Definition of the role and permission for each partner of a VO to access metadata and data
- Operational databases: update, mirroring, integrity, secure access and metadata handling.
- Secure access to external database such as the GPS and seismological databases

4.5.3.3. Added value for the Solid Earth Physics community

Over the past ten years, computational models have been shown to contribute fundamentally to the intellectual foundations, in areas such as global seismology, earthquake dynamics, earthquake hazard mitigation, Earth core dynamos, and in particular, fundamental for the effective use (inversion and data assimilation) of the continuous advances in instrumentation and observation

The computational models become more and more sophisticated and demanding in terms of computational resources, the Grid should facilitate cooperative studies implying intercomparison of large scale computational models by providing access to the outputs of these models, for processing, evaluation and visualization, to a large and geographically distributed earth science community. The Grid may therefore help the seismological community to achieve such an objective and to improve its organization at the European level. Furthermore as surface and satellite observations are increasing through permanent arrays or short time experiments, the Grid should also help in providing unified tools for retrieving and processing large volume of distributed data in the most transparent and easy way. Such a need is clearly at the heart of many GPS and seismological projects, and the Grid may foster an unified European approach to this problem.

The rapid advances in computer and information technologies are driving individual initiatives out of business and the community must change its mode of operation by organizing its efforts through interdisciplinary initiatives. The selected applications may help in providing some demonstrative operational tools that can be used by a wide range of domains in geophysics as well as in fostering more collaborative work with the Earth Observation and Climate communities.

4.5.4. HYDROLOGY

Freshwater is not only a vital resource for drinking purposes, it is also and mainly required for food production and agriculture. However, in many regions of the world freshwater resources are seriously endangered mainly by over-exploitation and contamination. As the world population increases rapidly, the stress imposed on the water cycle gets always stronger. On the one hand, industrialisation and intensive food production leads to a continuous degradation of the water quality. On the other hand, natural hazards like floods and droughts may have large impacts both in the domain of water supply and land management. .To minimize all these well known impacts on the environment ensuring a renewable use of natural resources, a deeper knowledge of the global water cycle is a prerequisite. Engineering solutions can then be designed with the help of numerical modelling tools allowing to investigate the potential impacts of several management scenarios. As for any earth science disciplines, one of the most important issue in the framework of hydrology is to account for the inherent uncertainty of natural systems (e.g., water stresses, porous media properties, etc.) through the use of statistical models coupled with physical and economical models. The planning and the management of real-world water systems requires to incorporate the uncertainty reflected in the model results in a decision-making process. This process is a complex collaborative task and requires multidisciplinary research activities (geostatistical data analysis, numerical modelling, environmental engineering). It is therefore very clear that Grid technology can improve very significantly the research and the management of water. On the one side it will allow researchers to share in a transparent way data, and even real time observations. On the other side, it will allow them to cooperate in running a series of numerical tools to analyse data, calibrating models, and generating scenario simulations in a stastitical framework.

From a computer science point of view, the hydrology community uses a vast toolbox of different codes (commercial, in-house developed, and open source). It is frequent that a user has to couple, or to chain different codes to conduct his own analysis. For example, a sensitivity study may require to run stochastic simulations to set up the input dataset for a deterministic physical model first and then to analyse statistically its outputs. Basically each of these individual tasks would have to run on a single computing node, but if thousands of simulations have to be conducted they can run concurrently on different nodes.

As a first step for a Grid technology evaluation, we propose to test a case study in Tunisia within the context of the European INCO project **SWIMED: Sustainable Water management In MEDiterranean coastal aquifers: recharge assessment and modelling issues**. In this case, the data belong to the Tunisian partner, one important computer code (CODESA-3D) is developed by the CRS4 in Italy, and a part of the analysis will be conducted in Switzerland. Using the Grid technology, our aim is to share the data through this first kernel of scientists, to analyse them using numerical tools and to simulate probabilistic management scenarios in a way that the Tunisian partner can have a direct access to the results.

4.5.4.1. Hydrology Applications in EGEE

The hydrology community has a series of different objectives within EGEE.

Data sharing

The first important step is to share the database related to the Tunisian case study through the Grid. This data base is stored within a GIS system. (ARCGIS software) The Grid will allow all the partners of the SWIMED project to access the database online, to select different data layers (geology, topography, landuse, piezometry, hydrography, salinization maps) and to visualize it at the appropriate scale This case study will allow us to implement and test security procedures so that only authorized users, with different roles, may access, modify and visualize the data.

Geostatistical analysis

An important source of uncertainty when modelling sea-water intrusion is the imperfect knowledge of the three dimensional spatial distribution of the hydrodynamic parameters of the aquifer system. Geostatistics allows to build a statistical model of the spatial variability of these properties based on punctual observations and indirect measurements, such as geophysical data. The statistical model is then used to interpolate the available point measurements and to provide a set of equiprobable spatial distributions. These distributions represent an important part of the input dataset of the groundwater flow and transport model. The statistical inference based on the available data is mainly an interactive task in which the data and their different statistical properties (first moments, pdfs, covariances, variograms,...) are explored in order to identify the best statistical model. Three different codes are used for this task. ISATIS is a commercial product, available on LINUX platforms, very efficient for the exploratory analysis of the data. The others two codes are HYDROGEN and MATLAB. HYDROGEN is an open source fortran code that can be easily compiled on many platforms.

Groundwater flow simulation

In the Tunisian case study, the most important problem is the intrusion of saline sea water in the coastal aquifer exploitated for irrigation. The salinization process is hardly reversible and can be mitigated with a rationalization of pumping and an artificial recharge. Enacting proper environmental management requires to accurately model the sea water intrusion mechanism. A finite element code named CODESA-3D, developed by the CRS4, allows to model groundwater flow both in the saturated and in the unsaturated zone of the aquifer taking into account density effects related to sea water intrusion. This Fortran 90 code runs on different platforms including LINUX. The aim is to port it on the GRID in order to be used by the Italian, Tunisian and Swiss partners during the first stage.

An alternative code, used by UNINE, is named FEFLOW. It is a commercial software allowing also to simulate the processes of sea-water intrusions in coastal aquifers. The issue that has to be solved in this case is the licensing problem. Having this code available for a certain number of partners over the grid will allow intercomparisons of results between FEFLOW and CODESA-3D.

The groundwater simulation process is much more automatic and repetitive than the geostatistical analysis. After the output of the geostatistical analysis has been stored on the Grid by a user, the groundwater models is run as many times as necessary (from 100 to 100 of thousand times) in a Monte Carlo simulation. Every run generates, at the prescribed time, water pressures and concentrations at each of the nodes (up to one million) of the aquifer grid. The results of the Monte Carlo simulation are then analysed to obtain probabilistic maps of the aquifer contamination. Finally, the maps corresponding to different management scenarios are evaluated by experts for the selection of optimal solutions.

It is then clear that this part of the application is very highly demanding in terms of CPU, memory and mass storage requirements.

4.5.4.2. Hydrology Requirements

The hydrology requirements are the following:

- Secure but comfortable handling of the GIS database and the model input and output datasets among the community.
- Interface for desktop access to remote high performance computing resources via a configuration manager. Partners are able to remotely build, configure, interactive and batch run, control and visualize complex groundwater simulations.
- Fortran compilers (F90 and F77) and UNIX shells/scripts to control job execution.
- One of the key question is to solve the licensing problem for commercial codes.

4.5.4.3. Added value for Hydrology community

The Grid added value for the Hydrology community is expected to be high. It will facilitate and strengthen the cooperation inside the scientific community as well as with end-users (water managers and technicians) by providing a unified platform for data sharing and modelling. Since historically, the hydrology and especially the hydrogeology has always been one of the less funded disciplines in earth sciences, it suffers from a huge lack in computing resources and personnel. By intensifying the cooperation between researchers, the Grid offers an unique chance to this community to overcome these historical difficulties and to exploit the potential of conducting innovative multidisciplinary applications for small but specialized research groups.

5. EARTH SCIENCES INDUSTRIAL APPLICATIONS

This section should be maintained by the Earth Science Industrial VO.

5.1. VO Organization and Coordination

The coordinator of the Geophysics VO is Dominique Thomas (CGG, France).

Initially a group of the general Earth Sciences VO, this VO is expected to expand quickly. Several contacts are on going. The main target is to make available structure, organization, IT resources hardware and software to the European Geophysics Research Community. CGG is acting as a catalyst and will provide its seismic processing toolkit to the Research community.

Most of the requirements are shared with the others groups of ES and specially with Earth Observation applications and justify the pragmatic way to introduce this VO by being part of ES.

The modern seismic data processing requires more and more huge amount of computing power. Research community hardly keeps pace with this evolution resulting in difficulties for small or medium research centres to deliver their innovative algorithms.

Grid Computing is an opportunity to foster sharing of computer resources and to give the access to large computing power for limited period of time at an affordable cost.

Capability to solve new scientific complex problems and to validate innovative algorithms on real size problem is also a way to attract and keep the brightest researchers for the benefit of both academic and industrial R&D geosciences community.

Initial Action Plan

- Create a small European partnership to develop and demonstrate the added value of Grid Computing for research in seismic data processing, with the help of EGEE support team.CGG intends to gather new partners in the coming months.
- Present the results of the experiment, promote collaboration between R&D from both academic and industrial actors. Find a real, neutral and well targeted organization to support and extend the "Virtual Organization" for seismic and other areas of geosciences. EAGE (European Association of Geosciences Engineers) is approached and will give advices.

5.2. Grid Awareness, Experience and Expertise

CGG has an initial experience in Grid and is an official partner of EGEE SA1. It belongs also to the industry forum.

5.3. Added Value for EGEE

Geophysics is a key technology for earth sciences, it shares and complements requirements of actual EGEE applications. "Geophysics" community is large (thousands of researchers across Europe) but very scattered. EGEE will benefit in enabling such a community to collaborate and progress.

One of the applications is an industrial application which will be a reference for other industries and will support credibility of EGEE infrastructure.

5.4. Resources available for contributing to EGEE Grid

CGG : CE (100 nodes), SE (100GB)

5.5. Community, Applications and Proposed EGEE Work Programme

5.5.1. EXPLORATION GEOPHYSICS

The goal of geosciences is to understand the earth sub-surface and the main tool to progress in this knowledge is the seismic surveys which deliver a "scan" of the sub-surface, at the scale of km^3, in the form of signals coming from acoustic wave propagation. If the first economical usage is for the Oil & Gas industry, other natural resources management (water, geothermics, ore), earthquakes analysis or nuclear proliferation surveillance are concerned.

The complexity of seismic data processing at the modelling and imaging steps comes first by the complexity of *algorithms* (mainly 3D wave propagation in complex, heterogeneous - and unknown- earth model) and second by the capability to *operate* the applications due to the size of the problem. And this size will be increased by a factor 10 in the coming 3 years to fully exploit the existing geophysical algorithms. Let's take a large industrial survey at the imaging step:

- 5 Terabytes of input data
- 5 iterations * 1000 jobs (submitted by a user) * 2 parallel tasks per job on 64 cpus* 10 chuncks (domain split due to physical constraints of hardware: memory/disc/or cpu performance)

These 6,4 million elementary tasks are linked, results combined, etc. and some of them should be restarted in case of a failure in the IT system. So the requirements to simplify (resource discovery and job adaptation), to automate the execution of workflows and better fault tolerance are critical.

The evolution of seismic processing in the industry was extremely fast in the last 3 years due to a move towards 3D/4D algorithms thanks to very large cluster computing. However Geophysical Research teams are spread across several small/medium centres and do not have this capacity so there is a strong requirement to share IT resources and software to keep pace with this industrial evolution.

However in the domain of Research and Development, the size of the problem is fortunately a small part of what is described above: several gigabytes of data and few hundreds of jobs, but is still large enough to be out of the capacity of most of the academic research teams. This is where the Grid technology could have a tremendous impact by sharing computing resources and software platform.

5.5.1.1. Geophysics Applications in EGEE

Seismic processing generic platform

Geocluster is the standard processing software used by CGG for its business and also by several Research Institutes across the world, in both industrial R&D and academic environments such as Lausanne UNIL (CH), University Watt (UK), Khantymansiysk Yugra (Russia), Trondheim University (Norway), and several in France (CNRS/INSU, IFP, Ifremer, etc.). Geocluster covers all the steps from data-preconditioning, first arrival picking, refraction statics, signal processing, velocity analysis, dip move out and migration, and include the state of the art geophysical algorithms for seismic processing. Researchers have the possibility to develop their own modules (Fortran 90) and to use them in any workflow.

This application is a batch, mono cpu process but can also trigger parallel algorithms as standalones. Main issues concerning the grid version are currently on the management of the environment: a “job” is a set of input and output seismic data, binary executables, parameter files, scripts and log information. All these components may go through various grid nodes (e.g. seismic data from a database node, software from a “service provider” node, parameter/results from/to user node and the computing on standard working nodes).

Modelling and Imaging in Geosciences

The Imaging application creates a depth image of the sub-surface from pre-processed seismic data (filtered, noise reduction, multiple refractions removed,..) and from a velocity model. It is currently a distributed application, running on cluster of PC Linux.

The Inversion application converts seismic data into acoustic impedance and then rock properties, by inverse modelling. It starts from an initial model and iteratively updates it on a global basis. The updating results from volumetric impedance perturbations and strata interface deformations, which are accepted or rejected according to the Metropolis algorithm implemented within a Simulated Annealing.

The simulated annealing is an approach used in several other industrial domains to solve non-linear problems.

When a problem can be approximated and expressed as a linear system, then the main problem is often to inverse a sparse matrix. In the seismic domain, e.g. in tomography inversion, the matrix can be so huge (millions* millions) that numerical instabilities coming from cpu computing, or hardware failures prevent to reach a solution. Grid computing, by decreasing computing cost, will open completely new, innovative implementations, for example using partial redundant computations.

5.5.1.2. Geophysics Applications Requirements

Some requirements are similar to the previous applications, some are different:

- MPI with Mirinet, SCI or other fast network and Fortran 90: The simulation proposed requires the possibility to submit a script and not only the executable. Modelling and imaging are also intrinsically parallel applications using large memory, medium to large exchanges between computing nodes with MPI, local scratch disc space and are heavy cpu intensive. Another requirement, not encountered in DataGrid, is the consistency between the compiler and libraries available on the UI and the ones needed by the cluster at the working node where the job is executed.
- Definition of the role and permission for each partner of a VO to access metadata and data
- Operational databases: update, mirroring, integrity, secure access and metadata handling.
- Secure access to external database such as the GPS and seismological databases

The metadata and data policies depend on their status (proprietary or public) and their use for research or business. Then, if company and research institute share the same VO access restriction needs absolutely to be implemented. Some restriction may also occur to share computing and/or storage elements.

5.5.1.3. Added Value for the Geophysics community

The Geophysics community is large (thousands of researchers across Europe) but very scattered. Delivering a grid version of the generic seismic processing platform, will have a tremendous impact on the capacity of several Research Institutes to innovate, validate and use their algorithms on real size seismic datasets and more important on 3D models or 4D models (“time-lapse”: successive 3D models in time). This will close the gap between small but innovative research teams and applications in the industrial domain. Industrial, robust software will be available for standard processing focusing the researcher to its real added value on various projects like those mentioned in the above ES projects. The key success factors are clearly the sharing of computing resources and software. Additional benefit is on the contrary to isolate the complexity of software management (versioning, licenses management, installation, support, etc.) and to hide this complexity to the users. At last, this initial VO will be the core of a larger Geophysical community in Europe as soon as the capacity to process geophysical data (seismic, models, simulation, inversion, imaging) on the grid is successfully demonstrated.

6. EARTH SCIENCES RESEARCH APPLICATION REQUIREMENTS AND RECOMMENDATIONS

The basis for describing the applications requirements is the joint list of use cases of the AWG (the DataGrid Applications Working Group), which combines the common requirements of HEP, Biosciences and Earth Observations applications. All requirements and use cases described there are valid for the Earth Science community. This section should be considered in addition to the AWG joint list of use cases [R6] and represents the main Earth Science recommendations.

The following section presents a discussion of the main ES recommendations arising from ES Grid application deployment experiences so far. This is followed by a set of tables summarizing the recommendations discussed.

6.1. Discussion of ES research requirements

6.1.1. VO MANAGEMENT

A VO database contains information about available VOs, groups within the VOs and the users belonging to the VOs.

In addition to defining one or more groups within a VO, it must be possible to define user roles. The VO database must allow the definition of named roles in terms of available VO groups.

It must be possible to assign to a user one or more groups and/or one or more roles.

A VO must have at least one group.

A member of a VO must belong to at least one group or have at least one role.

The VO database must be operated via an administrator interface available over the Internet using a simple web pages interface. The administrator interface must authenticate administration sessions using standard grid security mechanisms (i.e. via the user's proxy certificate).

The VO administrator interface must allow the definition of system administrators and VO administrators. Only system administrators can create or destroy VOs. One or more VO administrators must be assigned for each VO that is created. When adding a new VO to the database the system administrator must specify the details of the VO administrator(s). The administrator interface must ensure a backup (i.e. deputy) person is specified for each system and VO administrator.

A VO database service must provide information about VOs, groups, roles and members to clients (i.e. local grid security services) that request it. The service must be remotely accessible via standard web protocols and must use standard grid security mechanisms to carry out authentication and authorization of service requests. In order to ensure timely propagation of VO changes to the sites, clients which use the database information service must be able to register with the service to receive automatic notifications about VO database updates.

6.1.2. SECURITY

A clear and straightforward procedure must be established for a user to get a valid key pair signed by a CA. A well-established CA infrastructure with clear and verifiable procedures is vital for security.

A single, uniform, standard security model must be developed, applied and used by all grid components that need to implement security.

Security must be seamlessly integrated into all components that need it as a fundamental design consideration right from the start.

Individual grid components should not design or develop individual security solutions but should make use of the standard security components provided by EGEE.

The security solution must be fully compatible, conceptually and functionally consistent with the proxy certificate model used by GSI.

It must be possible, having fulfilled any prevailing security considerations, for the user to belong to any number and combination of VO(s)/group(s)/role(s).

A grid user should only need to possess a single certificate, independently of the VO(s)/group(s)/role(s) he/she is associated with.

The user certificate should only be used for the purposes of identifying the user and as a consequence there should be no specific information about VO(s), group(s) and role(s) of the user associated with the certificate.

It must be possible for the user to carry out the grid single sign on (proxy generation) procedure directly from the users normal machine, whether a desktop, workstation or laptop, whether running under Linux or Windows, without requiring the user to log on to an intermediate machine, and the user must never be required to store his/her private key on a remote machine (e.g. grid User Interface machine, etc.).

Any VO(s)/group(s)/role(s) authorization information, as may be required for a particular task in hand, should be specified by the user at the time the user performs grid "login" or single sign on, during the process of obtaining a valid proxy certificate.

During the single sign on process, the user must specify the required VO/group/role attributes to be associated to the proxy certificate that is created. A process running on the user's machine should contact each of the relevant VO database information services to gather the required authorization information as requested by the user. It must be possible to attach any number and combination of VO/group/role attributes to the proxy certificate, as long as any specific security constraints that may apply (e.g. conflicts between requested roles) can be resolved.

From the users point of view, the security mechanisms should appear as non intrusive, transparent and functionally and procedurally the same, regardless of the type of grid resources being accessed (e.g. resource broker, application executables, application environment, replica catalogue, replicated data, metadata, information system, storage resources).

A standard, uniform security access procedure should be applied whenever a proxy delegated process attempts to carry out an operation on a grid resource, regardless of the type of process and the type of resource. First, the authorized identity (or identities) delegated to the remote process is established by recovering the authorization attributes embedded in the proxy certificate. Next, the local permissions associated with the resource are checked. The local permissions consist of a table of recognized identities (e.g. VO/group/role) and the corresponding permitted operations (read/write/execute). The operation will only proceed if a positive match is obtained between the operation(s) requested and permissions granted to one of the authorized identities, as delegated via proxy certificate attributes to the requesting process.

Whenever the user process running on the grid creates any persistent resources or information, such as data on a storage element, or replica metadata information in a replica catalogue, etc. the user's default or currently set security attributes (e.g. read/write/execute permissions), or as delegated via proxy to the requesting process, shall be attached to the entity that is created. All subsequent accesses to that entity will be checked by the grid security middleware, using the attached permissions. It must be possible for the user to subsequently modify the assigned permissions later on, at any time.

6.1.3. DATA MANAGEMENT

The grid middleware must provide applications with both command-line commands and library embedded APIs to carry out all required data management functions. Where appropriate, (e.g. replica manager commands, information system commands), the same middleware must be available and must function the same way, whether running on a worked node or on a User Interface (i.e. job submission) machine.

All data management operations must follow standard security procedures as defined in the section on security requirements above. For instance, creating replicas, file transfers, inserting, updating and retrieving replica catalogue information and metadata, must include proper security checks on the proxy-delegated permissions.

Users of replicated data must be able to have 100% confidence in the validity and integrity of the replicated data (it must not be necessary for the application to expressly check the integrity and validity of replicated data every time it is used).

Data replication must be available as an atomic operation, consisting of data transfer and catalogue metadata registration. Data transfer must not be started until the middleware has correctly verified that the proceeding registration can be successfully performed (i.e. the replica catalogue service is running and the user has the required permissions, etc.). Integrity of replicated data must be automatically verified by the data transfer procedure and subsequent access to the replicated data must be preceded by an automatic data integrity check, which can only be skipped by an explicit override option. Replica registration must not be attempted until data transfer has been successfully completed and data integrity has been verified.

The replica and data management systems must not expose any internal replica details to end users and applications. For instance, any internal ids, such as globally unique IDs, that are assigned to replicated data and used internally by the middleware to manage files, must be hidden from the user or application, unless the details are specifically requested. The application must be able to refer to replicated data at all times by the normal application/user defined logical filenames; it must not be mandatory to refer to data (whether logical file names or physical file names) using a different id that was assigned automatically by the middleware, which is probably going to be meaningless and devoid of any semantic application information. From the application/user point of view, details such as the globally unique ID are considered internal to the middleware and should be handled transparently.

Replica catalogue performance must be constant and must not be susceptible to degradation, either due to the number of entries in the catalogue or the request loading on the replica cataloguing service. Quality of service, in terms of performance and overheads (i.e. wall clock time consumed) associated with the insertion (registration), search and retrieval of replica information in the replica catalogue, must be quantified as absolute minimum, mean and maximum, guaranteed values. These QoS values must be published by the service in the grid information system and easily retrieved by the end user/application.

Both the replica catalogue, as well as the storage resource services, must be capable of handling (storing, searching, updating, retrieving) file numbers of the order of millions of replicas.

It must be possible to operate any number of independent replica and storage resource services per VO, for instance, each VO group may operate, if necessary, one or more individual replica catalogues and storage services.

An application must be able to create, populate, utilize and destroy replica catalogues on the fly, as required by the application, without the need for human administrator intervention, having performed the appropriate security and resource utilization quota checks.

Applications must be informed in advance when resource utilization (e.g. number of files, storage space) approaches the allowed limits imposed by the system. Actually reaching imposed limits must be handled in a controlled, predictable and organized way; the user/application must be properly informed. Both replica catalogues and storage resources must provide ways to inform users/applications about levels of utilization (e.g. percentage available, used).

The replica catalogue must offer mechanisms to associate user/application defined metadata values with individual replicated data files. It must be possible to define key/value tuples. All the standard, most frequently used data types must be supported, e.g. date, time, polygon, integer, alphanumeric, float. When inserting/updating replica information it must be possible to assign both the associated set of metadata keys, as well as the logical filename, by a single operation (i.e. it must not be necessary to call the API or issue the command line for each associated metadata key/value pair). The application must be able to refer to replicated file(s) either by using assigned logical names or using metadata key/value tuples. Given a logical file name, it must be possible to recover the full list of associated key/value pairs in a single operation on the API or command line. Given a metadata search consisting of a set of target key/value pairs (e.g. as a SQL query), it must be possible to recover the complete list of logical file names that satisfy the query. It must be possible to specify metadata key values to be searched on as conditional expressions.

It should be possible for users to change ownership and/or access rights to files on the Grid, comparable to 'chmod' and 'chown' commands on Unix systems.

6.1.4. WORKLOAD MANAGEMENT

The main function required of the workload management system is automation of resource discovery, job submission (including staging of executables and auxiliary data from the submitting machine to the execution platform), monitoring of job status throughout the submission, and automated retrieval of job result files created on the execution platform, as specified by the job description. The possibility to cancel a submitted job and the possibility to communicate with the job via standard input and output channels are also necessary.

The workload management system includes a resource broker service that carries out automated matching of job requirements, as specified by the user/application, against available resources published in the grid information system. The resource broker uses JDL language to specify job requirements, including replicated data.

Often a job executing on the worker node needs to retrieve certain job parameter information that was specified by the application when composing the JDL script. It must be possible for jobs running on the worker nodes to retrieve parameters that were specified in the JDL, for instance, the full list of logical file names specified using the inputData attribute. This avoids the application job pre scheduler environment having to specify job parameter information twice, e.g. once in the JDL and then again for use by the programme to be executed on the worker node environment.

Since the grid system is by nature highly dynamic, and users usually take some time and experience to understand how to make best use out of the middleware, it will often not be possible to successfully carry out user or application requests, for a wide variety of reasons. Some failures may be due to error conditions and others may be may be due to conflicts in the formulation of JDL scripts, while still others may be due to service outages, malfunctions or misconfigurations. Whichever the case, components of the workload management system (indeed, the same applies most of the grid middleware components that have application interfaces) must provide clear, accurate and comprehensive information about error conditions, or any other kind of failure (for example, reasons for failing to find any matching resources). The error or failure status information returned by the middleware must be detailed and accurate enough to allow the user (or the application) to identify the reason for the failure and to take the appropriate corrective action.

The workload management system should not be a single point of failure. It should be a system containing redundancy to cope with system failures, transparent to the user.

The performance of the workload management system (i.e. response time) should be guaranteed within certain known QoS limits. The service should include mechanisms to protect from response time degradation under conditions of heavy load.

6.1.5. INFORMATION SYSTEM

The information system should provide a reliable and functional service for applications. This implies the information system not to be a single point of failure, i.e. it should be able to cope with system failures, transparent to the user.

In order to use the information system effectively, applications must first be able to obtain the schema used by the information system, so that they can use it to obtain, and make use of detailed information about available grid resources.

Next, applications need to obtain the addresses (i.e. URLs) of the available information system services. This information (including schema information) could be published in a top-level grid directory or grid information web pages. This same top level directory could also provide a running log about overall grid status and serviceability, known problems, outages, scheduled maintenance, useful sources of information, etc. This would be useful for users to locate the required information.

Details published in the grid information system should be protected by security controls just like any other grid resource. Application (or other client) requests to obtain details from the grid information system should first be checked to ensure the requesting entity is authorized to receive the requested details. For instance, details of grid services dedicated for use only by a particular VO should only be available to members of the same VO, and the grid system administrators VO. If a delegated process running on the grid requests information about a given class of resources (e.g. storage resources or replica services), it should only receive details about those resources that are available to the VO(s) which the process belongs to (in DataGrid a request for available Computing Elements returns information about all resources, including those that cannot be used by the VO the application belongs to).

The performance of the information system (i.e. request response time) should be guaranteed within certain known QoS limits. The service should include mechanisms to protect from response time degradation under conditions of heavy load, e.g. due to numerous, frequent updates to the information system by information producers at the same time as numerous, frequent requests by information consumers.

All available grid services should be published and discoverable in the grid information system, e.g. including resource brokers, replica catalogues, replica optimizers, metadata catalogues.

The information system should be able to detect wrongly configured sites. The information system should flag these systems as ‘unusable’, so the RB will not use them.

6.1.6. FABRIC MANAGEMENT

The procedure for a new site joining the grid must be clearly defined, well documented and published so that it is easily accessible to potential new sites. The grid site installation and configuration procedure must be simple, automatic and straightforward. As far as possible, the installation procedure must follow standard software installation procedures and conventions. There must be no assumption or requirement of any previous grid system knowledge or experience by the local site system administrators. Fast turnaround installation support by email and telephone must be available for local site system administrators.

It should not be necessary to wait until a site is properly certified before it can be registered in the production grid information system; a site should be automatically registered as part of the installation and configuration procedure. However, site status published in the grid information system should be properly set, according to the actual site status (i.e. whether fully certified, or undergoing installation and acceptance tests, or offline, down for maintenance, etc). The Resource Broker, or any other potential client services, should specifically check the published site status before using any resource.

The site installation and configuration procedure should include automatic launching of self-checking procedures. Once the automatic procedures are passed, the local system administrators should run specific acceptance testing and certification procedures to confirm the site integrity. The site (or resource) status published in the grid information system must be updated only after the certification has been successful.

6.1.7. MASS STORAGE MANAGEMENT

A Grid aware Mass Storage management system, or Storage Resource, should provide uniform user and local administrator interfaces for different types of mass storage system implementations. The Storage Resource should provide functionality that allows the local system administrator to control the way that the device is shared by grid users, i.e. chiefly who has access to which areas of the storage system and how the available space is used. It should present the grid user or application with a standard grid storage resource interface that abstracts the details of the underlying system.

The basic functionality required by applications is the ability to store, retrieve and delete data files, whether replicas or master copies. In addition the application must be able to find out the amount of storage space available, and to reserve space for a determinate period of time.

Standard, fine grained grid security control and checking functionality, based on VO/group/role identities (as described in the section on security requirements) must be an integral, automatic part of the storage resource service, its operation should be transparent and non intrusive to the user or application use of the service.

The storage service should ensure a high degree of reliability and should perform automatic checks to ensure integrity of data at all times.

The storage resource may be either a stand-alone disk system or a front-end interface to an archive system, or both. In the case of a disk system and a front-end interface to an archive system, the storage resource should provide functionality to automatically manage the disk system as a cache for the archive system. This involves automatically managing which files are allowed to stay in the cache and which files should be moved to the archive storage. A file pinning mechanism must be provided in case the application has to ensure files are not moved while they are being used.

The storage resource service must continually monitor free space available and should never allow the system to become completely full and unusable as a result. The service should warn clients and system administrators well in advance if free space falls below a specified safe threshold, so that advance remedial action can be taken well before the problem becomes critical enough to impact on essential operational services.

As files are continually added to a storage resource, it must be possible to ensure a minimum amount of free space is available in the cache, by automatically moving unused files to the archive; this should be done on a least recently accessed basis. The storage resource middleware should record the times of file accesses; the least recently accessed files should be the first candidates to be moved to the archive.

When a request arrives to access a file that has been moved to the archive, the storage resource middleware should automatically handle the retrieval of the file back into the cache. Files that are created on the storage resource should be assigned a default pinning time to live, that can be specified by the application. After this time has elapsed they can be moved to a secondary storage area.

Since the replica catalogue will keep track of files stored on different storage resources, the storage resource middleware and the replica management middleware must be must be closely integrated. In particular, when asked to return a reference to a physical file location, the replica catalogue should first check the availability (i.e. serviceability) of the referenced storage resource, in order to avoid returning references to physical files that cannot be accessed because the associated store resource service is down (e.g. for maintenance or other hardware fault). If a replica catalogue entry contains a reference to a physical file on a storage element that cannot be accessed (i.e. because the storage resource service in not working, for whatever reason), the replica manager should consider the entry to be invalid and ignore it (it should not return the invalid physical reference to the application to handle the error).

The storage resource must be capable of handling millions of files per VO/group. It must be capable of handling file sizes anywhere between e.g. 10kb and 100Gb.

A single storage resource may operate one or more storage resource services, possibly on a per VO basis. In order to avoid degradation of performance, the number of services running at any time should be dimensioned dynamically according to levels of utilization (i.e. number of requests).

When requested by the application, the storage resource service should provide QoS estimates on transfer times to, or from one or more host locations specified by the application. The service may in turn make use of a data transfer optimization service in order to calculate the requested QoS estimates.

6.1.8. System Integration

Since grid is all about intercollaboration and cooperation, grid components typically work together with other components in order to carry out the required operations. Therefore it is essential that developers take special care to clearly specify the interfaces between components that work side by side, e.g. resource broker and replica catalogue, replica catalogue and storage element, information provider and information user, etc. Often there is a mismatch between interfaces that needs to be clarified; usually this arises when it is not immediately clear which components should be responsible for certain functionality lying close to the borderlines of certain component interfaces.

Standard mechanisms and approach to fault tolerance, error handling and recovery should be pervasive throughout all middleware components. Services that run as background processes are often subject to failure, due to unexpected process termination without warning, entering an endless wait or sleep state, or getting stuck in a loop. To detect such occurrences, essential services should be monitored and restarted automatically when they fail.

A standard set of grid error conditions and associated error codes, comparable to standard HTTP response header codes, e.g. 404 page not found, could be very useful to characterize and classify known, frequent error conditions and failures, whether due to system or user failures (or both).

Middleware components that represent potential single points of failure should be avoided, e.g. by providing redundant backup services which clients can use as an alternative.

There should be recommendations for middleware developers to follow a consistent model/format for APIs and command line commands and options.

6.1.9. Application deployment

Standard procedures must be available for packaging and deploying application environments on all grid sites where a given VO or VO sub group is authorized to use processing resources. This procedure must be well documented and published so that application integrators can easily identify, obtain and use the procedure to deploy their application environments.

To allow for upgrades, it must be possible for applications to install several versions of the same environment, capable of being used side by side, simultaneously.

It is essential to have a well defined and simple mechanism to allow the deployed application to specify the unique id signature which identifies the application environment, that gets published, i.e. as part if the CE information, in the grid information system by the local information provider, which is the same runtime environment identifier value that will be specified by the application or user when composing the JDL script.

6.1.10. Documentation and user support

Middleware components created by different sources making up the overall grid middleware should conform to use a well-defined documentation model and format. The documentation should have clear version management.

A standard documentation hierarchy suite should be specified and adhered to by all grid middleware developers and service operators, e.g. installation guide, administration guide, user guide, troubleshooting guide, etc.

A consistent documentation model would assist system administrators and users to locate useful information, independent of the type of middleware component.

6.1.11. Miscellaneous

An important point for the application is the availability on the grid of different languages and COTS. In order to port new applications, the following languages generally used in ES are needed IDL, Matlab or clones Scilab or Octave, Fortran 77, Fortran 90, and C. The climate community uses NetCDF, BLAS/LAPACK, Python, Ferret…

Other difficulty, that occurs, is related to the version of software available on the editor server and the one required by EGEE. As an example, the version of Netscape required by EGEE is no more available on the Netscape server.

6.2. Summary of Recommendations for EGEE development

The requirements described in the previous sections are summarized here, using a similar format to the one used for the recommendations of the EDG Application Working Group. The priority is assigned from 1 (low) to 5 (high). Even if the average is low some applications (priority 5) cannot run if the functionality is not there. Geophysics indicates that their requirements are similar to EO Observations ones. Hydrology has similar requirements to Solid Earth Physics ones These recommendations should be considered in addition to the AWG recommendations and requirements presented in the two documents [R5][R6].

6.2.1. VO MANAGEMENT

VO management	Priority			
	EO	Climate	Solid Earth	Average
Definition of groups within a VO	5	3	5	4.3
Definition of roles within a VO	2	4	3	3
Assign to a user one or more groups and/or one or more roles	5	5	3	4.3
Web based VO database administrator interface	3	2	3	2.7
Administrator interface login via the user's certificate loaded in browser	3	2	3	2.7

6.2.2. SECURITY

Security	Priority			
	EO	Climate	Solid Earth	Average
Clear and straightforward procedure for obtaining a certificate and joining the Grid	4	4	5	4.3
Unified, common security protocol, standard for all services and resources	4	4	5	4.3
Standard security protocols seamlessly integrated into all components as fundamental design feature	4	4	5	4.3
From user perspective, security mechanisms should be non intrusive, transparent and should present similar functionality, independent of the type of grid resources being accessed	4	4	5	4.3
User may belong to any number and combination of VO(s)/group(s)/role(s)	4	4	4	4
User may possess a single certificate valid for all associated VO(s)/group(s)/role(s)	4	3	4	3.7
Private key stored only on users PC; proxy generation directly from the users PC	3	1	3	2.3

User able to specify the required VO/group/role attributes to be associated to the proxy certificate during single sign on	5	4	4	4.3

6.2.3. DATA MANAGEMENT

Data management	Priority			
	EO	Climate	Solid Earth	Average
Fast performance of replica catalogue operations, resistance to degradation under heavy loads, whether due to increased data or user numbers	5	5	4	4.7
QoS response times published by each replica manager service in the grid information system, retrievable by the end user/application	3	3	3	3
All functions, e.g. creating replicas, file transfers, inserting, updating and retrieving replica catalogue information and metadata, must include proper security checks on the proxy-delegated permissions	5	3	5	4.3
Mechanisms for establishing and operating QoS levels to guarantee validity and integrity of replicated data	3	2	3	2.7
Data replication (i.e. data transfer and catalogue metadata registration) must be performed as an atomic operation	4	4	3	3.7
Correct management of disk space quota checking integrated in the data transfer operations	5	5	5	5
Advance, pre-emptive warnings when available space falls below threshold levels	4	4	4	4
Automatic data integrity checks embedded in replica manager operations	4	4	4	4
Capability to handle (storing, searching, updating, retrieving) file numbers of the order of millions of files	4	3	3	3.3
An application must be able to create, populate, utilize and destroy replica catalogues on the fly, as required by the application, without the need for human administrator intervention, having performed the appropriate security and resource utilization quota checks	2	2	3	2.3
RC integrated support for user or application defined metadata associated to logical filenames, supporting wide range of available data types (e.g. date, time, polygon, integer, alphanumeric, float)	4	5	4	4.3
Register both the logical filename and the associated set of metadata keys (list of key=value tuples) in a single operation	3	4	3	3.3
Retrieve full set of metadata tuples associated to a LFN in a single operation	4	4	5	4.3
Capability to store millions of files per VO	4	3	3	3.3
Mirroring of metadata and data bases	3	3	3	3

6.2.4. Workload management

Workload management	Priority			
	EO	Climate	Solid Earth	Average
Method for jobs running on the worker nodes to retrieve attributes and values that were specified in the JDL	3	3	2	2.7
Improved error messages	4	5	4	4.3
Provision and enforcement of job submission quotas to prevent overloading the RB	4	3	4	3.7
Monitoring of RB state and job submission timings with results dynamically published in GIS, to enable user/application to choose RB service based on current performance and loading	3	4	4	3.7
Mechanisms to protect from response time degradation under conditions of heavy loading	4	3	3	3.3

6.2.5. Information system

Information system	Priority			
	EO	Climate	Solid Earth	Average
Improved reliability and fault tolerance, eliminate vulnerabilities due to single points of failure	4	4	4	4
Schema to include grid descriptor info to allow applications to determine the type of Grid environment e.g. production grid, development grid	3	2	4	3
Mechanisms to enforce security restrictions on published information, e.g. info about VO resources only available to members of the VO	3	2	3	2.7
Guaranteed response times within published QoS figures	2	2	2	2
Schema to include info on all available services, e.g. resource brokers, replica catalogues, replica optimisers, metadata catalogues, etc.	3	3	5	3.7
To allow clients to determine whether a service is usable or not, the schema should support inclusion of current serviceability status for all published services	4	4	4	4

6.2.6. Fabric management

Fabric management	Priority			
	EO	Climate	Solid Earth	Average
Fast, easy site installation and verification procedures for new sites joining the Grid	4	4	4	4
Mechanisms for scheduling maintenance and dealing with outages, forewarning users about service problems, downtime or unavailability	3	4	4	3.7

6.2.7. Mass storage management

Mass storage management	**Priority**			
	EO	**Climate**	**Solid Earth**	**Average**
Mechanisms for administrators to perform storage space allocation and management among VOs and users contending for the same storage resource	4	4	4	4
QoS mechanisms allowing applications to determine service level conditions related to: operation (commands) response times (e.g. for storing, retrieving, or deleting), space management (total, used, free and reserve space available, procedure to request more space), data integrity and privacy levels and guarantees	3	3	3	3
Advance warning when available space falls below reserve levels	4	4	4	4
Automated management of disk cache by file swapping to back end archive storage device	1	5	3	3

6.2.8. System Integration

System Integration	**Priority**			
	EO	**Climate**	**Solid Earth**	**Average**
The grid middleware must provide applications with both command-line commands and library embedded APIs	4	4	5	4.3
Uniform configuration across the grid (i.e. available software and installed versions) for each type of resource or service	3	4	4	3.7
Standard mechanisms and approach to fault tolerance, error handling and recovery supplied to all middleware components.	4	4	4	4
Common standard (guidelines) for command line formats and APIs	4	4	5	4.3
Standardization of error codes, error messages and error handling procedures	4	4	5	4.3

6.2.9. Application deployment

Application deployment	**Priority**			
	EO	**Climate**	**Solid Earth**	**Average**
Standard procedures for packaging and deploying application environments on CEs	3	5	4	4

Ability to deploy and utilize different versions of the same application environment e.g. IDL v5.4, IDL v5.6	3	3	4	3.3
Fortran 90	4	5	5	4.7
Fortran 77	4	1	2	2.3
Matlab or clones: Scilab, Octave	3	1	2	2
MPI	2	3	4	3
BLAS/LAPACK or clones	1	5	0	2
netCDF	2	5	2	3

6.2.10. DOCUMENTATION AND USER SUPPORT

Documentation and user support	**Priority**			
	EO	**Climate**	**Solid Earth**	**Average**
Common standard for documentation, consistent documentation model, supported tools and formats across all middleware packages	4	4	5	4.3

6.2.11. MISCELLANEOUS

Miscellaneous	**Priority**			
	EO	**Climate**	**Solid Earth**	**Average**