RESEARCH ARTICLE

Data issues at the Euro-Mediterranean Centre for Climate Change

S. Fiore · S. Vadacca · A. Negro · G. Aloisio

Abstract Climate Change research is even more becoming a data intensive and oriented scientific activity. Petabytes of climate data, big collections of datasets are continuously produced, delivered, accessed, processed by scientists and researchers at multiple sites at an international level. This work presents the Euro-Mediterranean Centre for Climate Change (CMCC) initiative, discussing data and metadata issues and dealing with both architectural and infrastructural aspects concerning the adopted grid enabled solution. A complete overview of the grid services deployed at the Centre is presented as well as the client side support (CMCC data portal and monitoring dashboard).

Communicated by H.A. Babaie

S. Fiore (✉) · S. Vadacca · A. Negro · G. Aloisio
Euro-Mediterranean Centre for Climate Change,
Lecce, Italy
e-mail: sandro.fiore@cmcc.it

S. Vadacca
e-mail: salvatore.vadacca@cmcc.it

A. Negro
e-mail: alessandro.negro@cmcc.it

G. Aloisio
e-mail: giovanni.aloisio@cmcc.it

S. Fiore · G. Aloisio
University of Salento,
Lecce, Italy

Introduction

Climate change represents an important and critical challenge for several scientists and researchers. Increasingly complex simulation models, management of petabytes of datasets (which are already too massive for current storage devices) are issues that must be faced up to in the related centres. Key elements that must be taken into account are strongly connected both with data and metadata management.

In this paper we introduce the *Euro-Mediterranean Centre for Climate Change* (CMCC) initiative and the adopted data grid solution for the management of climate datasets. Despite the classical approaches, data-grid-enabled solutions (Berman et al. 2003; Foster 2005; Foster et al. 2001) greatly address scalability (users, data, queries, etc.), transparency (access, integration, management, presentation) and efficiency (performance) allowing the management of huge and distributed datasets.

The CMCC represents a fully distributed environment, is comprised of several sites, partners, etc. and it is an harmonic mix of different skills in the field of climate modeling, economy, impact studies and information technology. Taking into consideration the climate data growth rate, it is our considered opinion that a full decentralized schema for the management of data and metadata (addressing data availability, scalability, site autonomy and efficiency) represents the most suitable solution in the proposed environment.

In this paper we present and discuss in detail the data grid management solution adopted at the CMCC. First of all, before presenting the overall architecture designed at the Centre (a view *in the large* related to the involved data and metadata services/components) we provide a complete analysis concerning the main challenges driving our work

(secure, efficient and transparent distributed data management, interoperability, metadata search and discovery, etc.). Then we concentrate our attention and we delve into details of three fundamental pillars: *data management*, *metadata management* and *user support* providing technical motivations behind our choices and additional information about how the data/metadata related issues have been faced and solved at the Centre. Concerning metadata management we present the adopted CMCC metadata schema and the related implementation, the CMCC metadata handling architecture and infrastructure, the distributed metadata search, etc. On the other hand, for the data management part, we deal with data transfer, access, replication and management services and issues. Security is also discussed from several points of view. Moreover, we talk about the available user support presenting the CMCC data portal, the available command line interface and the CMCC monitoring dashboard.

Finally we discuss related works highlighting differences and analogies with the proposed solution and we draw our conclusions in the last section.

The CMCC initiative

In 2005, the Italian government, through the Ministry of the Environment and Protection (MATT), the Ministry of Education, University and Research (MIUR), and the Ministry of Economy and Finance (MEF) started a scientific initiative (namely the *Euro-Mediterranean Centre for Climate Change*, CMCC) aimed at establishing a national research centre devoted to climate change research.

The main partners of this initiative are six Italian research Institutes (the National Institute of Geophysics and Vulcanology, the Fondazione Eni Enrico Mattei, the University of Salento, the Italian Aerospace Research Center, the University of Sannio, and the Consorzio Venezia Ricerche).

As it can be argued this Centre is distributed in nature among several sites at a geographical scale (see Fig. 1) and is comprised of several research divisions which provide support for computing and operations activities, numerical modeling, impact studies (on health, energy, economy, coastal zone, Mediterranean sea, agriculture, etc.), training and dissemination.

This Centre represents the most ambitious initiative undertaken in Italy, within the framework of the National Research Plan, and specifically the National Research Plan on Climate. One of the basic idea behind CMCC is to create a unified environment able to concentrate in the same place numerical models, simulations, big amount of data as well as metadata, post-processing, visualization and analysis tools, etc. exploiting and joining knowledge and skills in the field of climate modeling, impact studies and information technology.

Fig. 1 Euro-Mediterranean centre for climate change

In particular, a data grid solution for such a distributed environment has been chosen as an enabling technology to access, organize, share, analyze, deliver, manage and store huge amount of data (few Petabytes in 2009, tens of Petabytes in some years) produced in the Centre. The CMCC stores and publishes datasets for the study of climate variability and for the validation of simulation models.

Three main sites will strongly contribute to CMCC production activities and are located in Lecce (Tier-0), Bologna and Capua (Tiers-1). Other sites will join after 2009 in the form of Tiers-2. It is worth noting here that at CMCC, Tier-0 refers to the main site hosting the CMCC Supercomputing Centre (about 30Tflops, 1,5 PB of storage), Tiers-1 host some own vector/parallel machines (few Tflops) as well as storage resources (hundreds of TB) and Tiers-2 concerns with peripherical data providers/producers that will contribute with their own resources (tens of TB) to the CMCC infrastructure. Anyway, the production phase will start at the beginning of 2009 and will involve just Tier-0 and Tiers-1. Tiers-2 will join the CMCC data grid infrastructure in 2010.

Right now the CMCC Supercomputing Centre in Lecce (Tier-0) is in the deployment phase. The computational infrastructure is comprised of both IBM p575 servers (parallel scalar machines) and NEC SX-8/SX-9 nodes (parallel vector ones). The overall computational power (in terms of peak performance) of the acquired systems is about 30 Tflops, which relates to 1,100 cores. The high-speed interconnection is INFINIBAND 4X DDR for the scalar cluster and NEC IXS for the vector one. The storage includes 470 TBs in cluster file system (GFS for the NEC cluster and GPFS for the IBM cluster) and a high-performance and high-capacity (1 PB) tape library for the

backup needs. This infrastructure will be soon made available to the CMCC scientists for their research purposes and activities. Moreover it is intendend, in the next future, to extend the infrastructural part with regard to storage capacity and computational power.

CMCC data grid

The CMCC represents a multifaceted and distributed environment for climate change with data and metadata distribution, management and handling issues.

A Data Grid architecture/solution has been chosen in order to (i) ease the management of such distributed data environment for climate change, (ii) deploy a uniform (i.e. from the security point of view) set of data grid services (data grid plane), (iii) provide the proper basis for higher level and distributed activities (i.e. workflow and post-processing ones at the computational plane), (iv) ease the management of several set of users through the virtual organization concept and (v) create a unified, advanced and complete environment for climate scientists.

Main challenges and requirements

The data management architecture design at CMCC needs to take into account many challenges, needs and user/system requirements. We started from them in our analysis and we present them as follows discussing about the proposed solutions:

Management of distributed petabytes of data a centralized approach is easy to manage, but introduces a lot of inherent problems such as: (i) *performance* (a unique/central point to store data and metadata can result in a performance bottleneck for the whole system), (ii) *scalability* (a centralized approach does not scale well when the number of concurrent users, clients, datasets, experiments increases), (iii) *fault tolerance* (when the centralized server is not active the whole system temporarily can not work), (iv) *autonomy* (sometimes datasets produced at several sites can not be moved to a central location, outside a local domain, even if they could be published into the system).

It is worth noting here, that a decentralized approach (both for data and metadata) as the one proposed at the CMCC (see following sections) aims at guaranting performance, scalability, fault tolerance and local autonomy. Moreover replication of data aims at addressing high data availability and fault tolerance (Tiers-1 will act as cache storage devices into our system, even if in a first stage we will not provide data consistency and data coherence management).

Security in a distributed environment like the one proposed in this work, security is extremely complex and multi-faceted. The system is secured by design and we fully adopt Globus Grid Security Infrastructure (GSI) protocol to address every security aspect: (i) authentication between couple of actors in the scene (users, servers, portal, etc.), (ii) communication protection (data encryption and data integrity), (iii) delegation (concerning distributed metadata search and discovery). Concerning authorization requirement and access policies, CMCC users are able to control which principals can access to what resources (datasets, experiments, projects, metadata information, etc.) and under what conditions (policies: read, write, etc.). The visibility of each object can be: private (to a single user, group, site, CMCC) or public. Moreover role membership management is also performed to manage in a scalable and flexible way the CMCC classes of users (i.e. full admin, data provider, metadata contributor, basic internal user, guest or external user, etc.)

Search & discovery in a distributed environment like the one we propose in this work, several metadata sources are in the scene and contribute to index available datasets. It is fundamental to query, combine, retrieve and filter metadata information producing query results in a fast, scalable and efficient way. The proposed P2P and grid solution addresses these important issues through a 2-step search and discovery process (see Section Distributed Search).

Transparency such a distributed environment exploits a wide set of services, protocols, standards, etc. All the underlying details need to be completely transparent and concealed to the end users as well as the access to the system must be pervasive and ubiquitous. For this reason a data grid portal solution is our considered choice for CMCC users and represents the proper container for services, data, metadata, visualization tools, etc.

Interoperability since data sharing is an important issue at geographical and cross-institutional scale, choosing and adopting an interoperable solution (exploiting SOA approach as well as standards for metadata, network protocols and data format) represents a basic requirement for our system. Interoperability (a leading requirement for the CMCC data management system) successfully contributes to really set up an "open" environment.

Global overview

In this Section we present the overall data grid system. In particular we show the involved services/components as well as their distribution, interaction and role. The CMCC data grid architecture has been designed to fulfil and meet

the requirements highlighted in the previous section. In particular the *Grid Metadata Handling System* (GMHS) and the *Grid Data Handling System* (GDHS) represent two basic pillars in the CMCC data management infrastructure. Additional security services/components are discussed as well, since they play an importal role into the proposed system. Let us now delve into details of the GMHS and GDHS.

The key components in the GMHS infrastructure (see Fig. 2) are:

– *GRelC Data Access and Integration Service (GRelC DAIS)*: it performs distributed and efficient metadata search and discovery activities by exploiting P2P/grid protocols (Aloisio et al. 2005; Aloisio et al. 2007). The GRelC DAIS are P2P connected in the CMCC environment, provide support for local (access) and global (integration) metadata management, completely hide the topology of the network and the underlying complexity/heterogeneity of the involved subsystems (RDBMS, XML-DB engines), etc. The GRelC DAIS architecture exploits a super-peer model. The 2-level schema is composed of: (i) normal-peers GRelC DAIS which collects/extracts metadata from the avaaialable data sources (data providers) and send them to the related superpeer; (ii) super-peers GRelC DAIS which are P2P connected with other super-peers and manage routing of queries among the GRelC DAIS backbone. In our system super-peers GRelC DAIS are currently deployed on three sites (Tier-0 and Tiers-1) whereas normal-peers GRelC DAIS will be mainly deployed at Tiers-2. Additional details about the metadata schema, metadata database as well as the 2-step search and
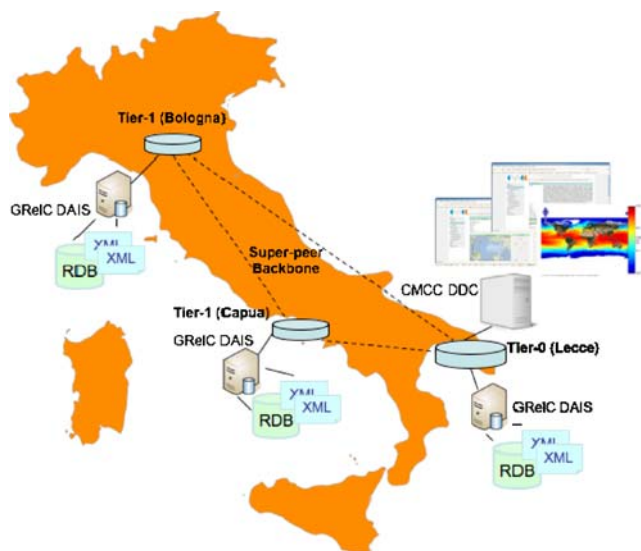

Fig. 2 Grid Metadata Handling System

discovery protocol will be discussed in the related *Metadata Management Section*.

– *CMCC Data Distribution Centre (CMCC DDC)*: it represents the access point to the entire CMCC production activity. It mainly provides several metadata search pages (basic and advanced search, search by experiment, variable, project, etc.), metadata browsing, annotation, validation, etc. as well as data access, transfer, visualization support. Further details are available in the CMCC Data Distribution Centre section.

Concerning the GDHS infrastructure (see Fig. 3) the main components/services are listed below:

– *Storage Resource Manage:* SRM implementations (Shoshani et al. 2004) are now being deployed at Tier-0 and Tiers-1 to ease dynamic space allocation and file management of shared storage components on the CMCC grid. An SRM is able to manage a disk cache (Disk Resource Manager) as will be the case of Tiers-1, a tape archiving system (Tape Resource Manager) or a combination of both, called a Hierarchical Resource Manager (HRM) as will be the case of Tier-0. Several Tier-1 solutions are now under evaluations, in particular possible candidates as SRM are dCache, DPM and StoRM. An HRM evaluation for Tier-0 will start in January 2009 when all the storage devices (tape library, online and near-online storages) will be up and running.

– *Reliable File Service:* gLite File Transfer Service (Stewart and McCance 2006) represents our choice for efficient data delivery (based on GridFTP data transport protocol) and reliable file transfer among Tier0 and Tier1 sites, addressing high data availability requirement.

– *OPeNDAP/THREDDS*: data is efficiently managed and accessed via OPeNDAP services (Gallagher et al. 2006). THREDDS installations will provides OPeN-DAP, HTTP, Web Coverage Service as well as NetCDF-CF Subset service to access data, download it, manage metadata information (Dataset Descriptor Structure, Data Attribute Structure), perform subsetting activities, etc. Moreover last version of OPeNDAP (Hyrax service) will be strongly evaluted (on x86_64 architectures) to access large Netcdf-CF files (some file size limit of 1.5 GB on IA32 architectures arose with former versions of OPeNDAP). Finally, in the case of OPeNDAP, we plan to test and evaluate the grid-enabled version adopted in the ESG project and named OPeNDAP-g to join the OPeNDAP features with the efficient GridFTP protocol support.

– *GridFTP*: it is widely adopted as data transport protocol (Allcock et al. 2001) in almost each one of
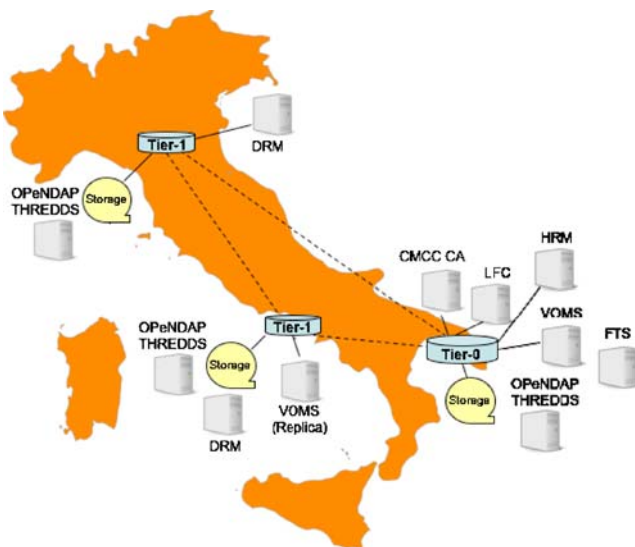
Fig. 3 Grid Data Handling System

the presented data grid services. It addresses security (through Globus GSI) and efficiency (by means of parallel streams and striping capabilities).

– *LHC File Catalog (LFC)*: The CMCC data grid environment must be able to provide the user with a global view of the entire grid as a single logical storage device. This component provides grid file system capabilities and starting from the root folder */cmcc/*, several subfolders will reflect the CMCC organization in terms of divisions (*/cmcc/sco, /cmcc/ans, /cmcc/cip*, etc.*)*

Additional security services/components that are part of the CMCC Infrastrcture are listed below:

– *Certification authority:* the CMCC Certification Authority (CMCC CA) manages user, host and service certificates. It issues X509v3 digital certificates for all of the involved actors. Using grid certificates and the de-facto standard GSI (Tuecke 2001), we provide a secure access to the entire CMCC data/metadata management infrastructure.

– *Virtual Organization Membership Service*: VOMS (Alfieri et al. 2003) is a system for managing authorization data within multi-institutional collaborations. It manages user roles and capabilities providing specific authorization-based grid proxy extensions. In the CMCC context, it allows distributed sites to centrally manage user roles and capabilities. To address fault tolerance, a replica of the main VOMS service installed in Lecce (Tier0) will be managed in Capua (Tier1). At the initial step the consistency management between the two replicas will be based on Mysql support for one-way, asynchronous replication.

After having introduced the overall system in the large, let us now delve into details of the following three key aspects: *metadata management*, *data management*, *user support*.

Metadata management

Metadata management plays a critical role in such a distributed environment, since it (i) enables search and discovery activities, (ii) allows describing and cataloguing datasets, it (iii) makes data effectively accessible and shareable by the scientific community.

Several aspects have been considered here:

– *CMCC Metadata Agreement* and related schema implementation;
– CMCC metadata handling architecture and infrastructure;
– Metadata database;
– Distributed metadata search;
– Security-metadata grid services.

In the following subsections we discuss in detail each of these strategic issues.

Metadata agreement and schema implementation

The CMCC Metadata Agreement is a schema collecting and describing all of the metadata needed by the target scientific community. The aim of the schema is to classify the CMCC data production (input, intermediate data and output of the experiments), models, services, etc. Since the beginning of the project it was established an internal and interdisciplinary working group to properly address this issue. Basic requirements of this activity was mainly three: *interoperability*, *rich schema*, *light metadata publishing process*. In order to address the proposed requirements we started our research activity from existing and well-known standards (proprietary approaches were not considered). In particular the consolidated *ISO 19115* (Geographic Information Metadata) and *ISO 19139* (Geographic MetaData XML encoding, an XML Schema implementation derived from ISO19115) standards were taken into account as the basis for our schema.

The first outcome of the working group has been the CMCC Metadata Agreement v1.0. It is a subset of ISO19115 and it represents the best tradeoff (for CMCC purposes) between the need to fully describe climate datasets, scientific experiments/projects, etc. and to have a light metadata publishing process for data providers. In the future, additional scientists' requirements about data description will lead to new, refined and extended versions of our schema.

Most important classes of information we considered into our design provide appropriate information to fully

characterize and locate geographic data; simplify the organization/management of metadata; ease data discovery, retrieval, purchase, etc.

After designing the CMCC Metadata Agreement, we moved to the related schema implementation. ISO defines its XML implementation of the ISO 19115 standard (ISO 19139). Starting from our subset of the ISO 19115 we inferred the related XML schema implementation from ISO 19139.

The adoption of such a schema provides a high level of interoperability among the metadata describing datasets produced at CMCC and the ones available at other international centres. Moreover, ISO 19115 is even more becoming a widely adopted standard for metadata management in several climate research centres.

CMCC grid metadata handling system

The CMCC Grid Metadata Handling System is the infrastructural part of the CMCC Data Grid aiming at supporting the distributed metadata management.

Each CMCC site (both Tier0 and Tier1) has both *data* (storage, replica services, DBMSs, etc.) and *computational components* (parallel vector and parallel scalar machines to run models, post-processing algorithms, etc.). Metadata are locally managed at each site using both relational and XML back-end systems. The provided grid metadata architecture must be able to perform both access to and integration of metadata stored in different and geographically spread CMCC metadata sources by addressing *scalability, efficiency* and *transparency* (regarding data access) and taking into account *local autonomy* of each site (in terms of site policies, locally adopted metadata back-end system, etc.) as well as *fault tolerance*.

The adopted metadata solution exploits the GRelC DAIS, a grid metadata service providing access, management and integration functionalities concerning relational and non-relational (i.e. XML databases) data sources. It provides two levels of virtualization since it offers both a unique front-end to manage and access data sources in grid and a decentralized and scalable architecture to integrate geographically spread data sources. The CMCC grid metadata handlyng system is a P2P network (connected graph) of GRelC DAISs exploiting a super-peer model.

It is important to remark that there is no single point of failure and no centralized management for the adopted solution. Moreover, *local autonomy* (in terms of metadata management support in case of network interruption) is also preserved. When a site is in downtime for some reason, the remaining GRelC DAISs continue to work as well as when a site is disjoint from the GRelC DAIS backbone it can continue to locally provide metadata services/functionalities to site users. Two different schemas (see Fig. 4) have been designed: the first one (Fig. 4a) at the initial stage (pre-
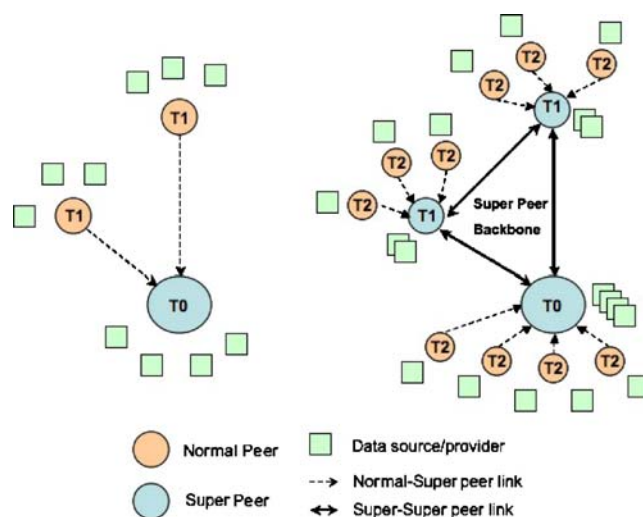

Fig. 4 Schema 4.a Initial stage – 4.b Production stage

production, 2008/2009) involving just Tier-0 and Tiers-1 (a single super-peer node is managed). The second one (Fig. 4b) in a mature production stage (from 2009/2010) involving Tiers-2 as well and the P2P backbone among Tier-0 and Tiers-1 super-peers.

In the second schema, the super-peer backbone among Tier0 and Tier1 is fully connected.

We chose the GRelC DAIS for several reasons. Among the others it provides:

- metadata integration and query forwarding capabilities for distributed searches;
- uniform access to several back-end systems (RDBMS such as Oracle, MySQL, Postgresql, etc. and XML-DB engines such as Xindice, eXist);
- different data access capabilities considering heterogeneous data models (relational, hierarchical);
- both user and VO centric authorization (even combined) via ACL and VOMS roles management;
- several kinds of queries and data delivery mechanisms with high levels of performance (in terms of efficiency) as described in previous works (Fiore et al. 2008a, b).
- full support for GSI (even including user credentials delegation);
- compatibility with the most important grid middleware (Globus, gLite, Unicore) (Fiore et al. 2007a, b)
- standard WS-I compliant interfaces already in production as well as an OGF compliant WS-DAIR and WS-DAIX interface in pre-production;
- a non invasive approach with regard to the managed data sources, the database query languages (including DBMS-specific extensions), the back-end system, etc.

Other works in the same area such as OGSA-DAI (Antonioletti et al. 2005), AMGA (Santos and Koblitz

2005), G-DSE provide similar capabilities even if for several reasons, they do not fit exactly the metadata solution designed at CMCC. For instance AMGA does not provide native XML support which is fundamental in our context, as well as full SQL support (native SQL99, DBMS specific SQL extensions, i.e. PostGIS), etc. Moreover the GRelC solution (with regard to OGSA-DAI and G-DSE) provides a combined authorization mode (GRelC DAIS and VOMS based) and a strong client side support (Command Line Interface, GRelC Portal, iGRelC Dashboard that can be easily customized and integrated in specific application-domain contexts such as in the case of CMCC). Finally, GRelC DAIS offers a P2P/grid support for data integration which is not currently provided by any other project in the same area.

Metadata database

The metadata database must reflect all of the information included in the CMCC Metadata Agreement. It is automatically ported in grid by the local (to the site) GRelC DAIS through a database registration process.

Each metadata database is basically made up of two parts: one *relational* and another one *XML*. It is woth noting here that the two parts play a different role in our metadata database, since the former acts as an index table for the local XML DB and the latter physically contains the entire set of scientific metadata. The relational part (also called *index database*) contains just a small set of information about the considered datasets. For instance the most important ones are: abstract, author, temporal and geographical extent, keywords, link to the XML document describing the entire dataset, etc. Moreover the links among projects, experiments and datasets are also modeled in the conceptual schema (Entity/Relationship model). The XML part is a collection containing a set of XML files describing all of the datasets managed at the related site. Each XML file is based on the ISO 19139 standard (a subschema) as described in a previous section. The XML part fully describes a whatever dataset and can be accessed, downloaded, completed, updated, etc. All of these access activities can be carried out through the available GRelC DAIS interfaces.

Distributed search

In the proposed data grid architecture, the metadata search process plays a significant role since it allows discovering datasets stored at the Centre.

First of all, the user has to create a valid proxy, starting from her user X509v3 digital certificates using the *grid-proxy-init* or the *voms-proxy-init* command.

After that, she has to submit a metadata query selecting an agent node (a super-peer) and specifying user search constraints. This operation can be carried out through the available Command Line Interface or the CMCC Data Distribution Centre (see *Client Support* Section). Each query is then forwarded to the peers (other super-peer GRelC DAISs) directly connected to the agent node (namely neighbours). This step is then recursively repeated (exploiting delegation of users' credentials) in order to securely spread the query across the network. Each peer is able to locally submit the query on the relational part of the metadata database, join partials results coming from the other peers, discard duplicated queries, avoid cycles, manage hops and time to live to limit the query with regard to space and time user constraints.

The user can check the status of the query (running, aborted, deleted, done, etc.), perform delete and abort actions, retrieve the entire output result when available (status equal to done), change/set a lifetime for the query result on the GRelC DAIS side. The search process is asynchronous and allows the user to obtain a list of climate datasets satisfying scientific requirements; in the provided list the user can read a small set of information (taken from the index database), even including the link to the XML dataset descriptions.

Security - metadata grid services

Concerning the metadata grid infrastructure, we decided to adopt the widely accepted and de-facto standard Globus GSI. Authorization leverages on a combined mode exploiting both a local authorization step on the GRelC DAIS side and a global authorization one on the CMCC VOMS server side. This mechanism allows great scalability, flexibility at a large scale and preserves local autonomy (in terms of policies/role management) on each site.

The GRelC DAIS is able to manage a wide set of data access policies (that is, specific user privileges) to update, read, delete, etc. metadata information. Groups of privileges are then mapped on specific roles (scientists, administrators, guests, etc.) on the VOMS server side. Main roles will concern with administration (/CMCC/Role=Administrator), CMCC researchers and scientists (/CMCC/Role=CMCCUser), metadata publishers (/CMCC/Role=MetaPublisher), data providers (/CMCC/Role=DataProvider), etc.

Each actor (user, service, host) is identified by an X509v3 digital certificates issued by the CMCC Certification Authority and each couple of actors interacting in the grid metadata handling system is mutually authenticated.

Data management

Data management is the second pillar we have considered in the CMCC data grid environment.

It enables users to (i) efficiently transfer and replicate data, (ii) perform data aggregation and sub-setting activities

on distributed data servers, (iii) makes data *efficiently* accessible by the scientific community.

Several aspects have been considered here:

– Data transfer/access/replication/management;
– Grid File Catalog;
– Security–data grid services.

In the following subsections we discuss in detail each of these strategic issues. While for the metadata management part several technical choices have been already done, for the data management one some services are still under evaluation.

Data transfer/access/replication/management

In order to allow an efficient data transfer, the three main CMCC sites will be able to efficiently move available datasets among them, by exploiting existing file transfer protocols and services. In particular, the proposed data grid environment will adopt GridFTP as standard data transport protocol. GridFTP fully supports GSI and provides mutual authentication based on X509v3 certificates. Moreover, it has several performance enhancing functionalities over FTP (i.e. parallel streams, retries and restarts supported, third party transfer option, etc.).

GridFTP will be used as the transfer protocol for data accessed through the OPeNDAP-g service. OPeNDAP provides, among the others, subsetting support for netCDF files and several filters for various kinds of data formats (netCDF, MATLAB, HDF, etc.) as well as format translation. Anyway, OPeNDAP relies on HTTP, does not support GSI and it is not suitable for large movement of data. For all of these reasons, in the CMCC context it is planned to adopt the related grid-enabled version (OPeNDAP-g) which has been widely and successfully used in the Earth System Grid (Middleton et al. 2006, Bernholdt et al. 2007) project, and supports both GSI and GridFTP.

As stated before in the overall architecture description, to manage the available storage resources, SRM services will be evaluated and integrated into the data grid infrastructure.

Moreover, to increase data availability (Madduri et al. 2002) a reliable grid replication service will be deployed between the Tier-0 and Tiers-1. At the moment two services, related to the reliable file transfer, have been evaluated: (i) Globus RFT which supports only GridFTP, is web service based and stores information in MySQL or Postgresql relational back-ends; (ii) gLite FTS which is a gLite/EGEE software, web service based and supporting SRM for data transfer. Right now, gLite FTS represents our first choice, since the SRM support (which is needed in our infrastructure) and the monitoring one represent interesting features that we plan to widely exploit.

Grid File Catalog

The CMCC Data grid environment must be able to provide the user with a global view of the entire grid as a single logical storage device. To achieve this goal, the LCG File Catalog service will be adopted. It is a gLite/EGEE service recording several information for each file (even including the locations of its replicas). Using the lfc-* commands we can query and update (i.e. the owner, permissions, group of a LFC file/directory) the LFC, whereas the lcg-utils move data into and out of storage elements maintaining consistency with the LFC. By exploiting this service, the file names appearing in the */cmcc* structure are called LFNs (Logical File Names), but a SURL (Storage URL) is needed to access data. Yet, the LFC provides a look-up to convert between LFN and SURL. It is worth noting here that, this is not necessarily a 1:1 mapping. If data is replicated in several grid storage resources, a single LFN can map to multiple SURLs.

Such a component is strongly adopted in the EGEE grid (expecially for LHC experiments) and it can be easily and straightforwardly integrated in the CMCC data grid infrastructure for grid file system/catalog purposes. The LFC service will provide the right level of data virtualization (naming transparency) concerning the CMCC data distribution.

Security - data grid services

Concerning the data grid infrastructure, all of the aforementioned services provide GSI support and so they adopt X509v3 digital certificates. As needed by design, the security framework is uniform for all of the data/metadata management activities in the CMCC data grid. Moreover, this element is coherent as well with the computational part (which is out of the scope of this work) at the CMCC.

Client support

In the CMCC context it is very important to have a user-friendly and complete client suite dealing with (i) metadata search, discovery, browsing, annotation, etc. (ii) file transfer, replica management, aggregation/sub-setting, etc. (iii) administration tasks, (iv) data visualization, etc.

In the following, we analyze the relevant client side topics discussing:

– CMCC Data Distribution Centre;
– Command Line Interface;
– CMCC Dashboard.

In the next subsections we cope with each of them trying to give to the reader a complete understanding of the proposed/adopted client-software solutions.

CMCC data distribution centre

The CMCC Data Distribution Centre (see Fig. 5), is the primary entry point (web gateway) to the CMCC. It is a Data Grid Portal providing a ubiquitous and pervasive way to ease data publishing, climate metadata search, datasets discovery, metadata annotation, data access, data aggregation, subsetting, etc. It does not centralize any functionality concerning authorization or authentication and for fault tolerance reasons it will be mirrored at several sites (Tier-1) providing additional web entry points to the CMCC system. The grid portal security model includes the use of HTTPS protocol for secure communication with the client (based on X509v3 certificates that must be loaded into the browser) and secure cookies to establish and maintain user sessions.

The CMCC DDC is now in a pre-production phase and it is currently used only by internal users (CMCC researchers and climate scientists). Right now it offers just some of the aforementioned functionalities (others will be available in a few months). The most important component already available in the CMCC DDC is the Search Engine which allows users to perform, through web interfaces, distributed search and discovery activities by introducing one or more of the following search criteria: horizontal extent (which can be specified by interacting with a geographic map), vertical extent, temporal extent, keywords, topics, creation date, etc.

By means of this page the user submits the first step of the query process on the distributed CMCC metadata DB (*relational part*, see Fig. 6). Then, she can choose one or more datasets retrieving and displaying the complete XML metadata description (from the browser). This way, the second step of the query process is carried out by accessing to a specific XML document of the CMCC metadata DB (*XML part,* see Fig. 7). Finally, through the web interface, the user can access to and download (partially or totally) the data stored on the storage device accessing to OPeN-DAP servers and to other available grid storage interfaces.

Requests concerning datasets stored in deep storage will be served asynchronously.

Command line interface

The data and metadata grid services can be managed through the classical command line interface. The CLI is already part of the adopted services and represents the basic way (less intuitive than a web-based solution) to manage services, transfer files, create a replica, store/access datasets, etc. for administrators, end-users, etc. Each data/metadata grid service adopted at the CMCC has its own CLI.

CMCC dashboard

The CMCC provides a proprietary tool (namely CMCC Dashboard) to monitor all of the data grid services that have been deployed at the Centre. This integrated approach (called Dashboard) is able to retrieve, process and display information coming from different data

Fig. 5 CMCC Data Distribution Centre

Fig. 6 CMCC DDC - Search Page

sources (both relational and non-relational) related to the aforementioned data grid services. The proposed dashboard is a *customizable* framework in which different users can choose a different set of panels and charts about the monitored resources, services, databases, machines, etc.

LFC, GRelC DAIS, SRM, FTS, etc. which manage their own metadata information into local catalogues can be monitored via this component. In all of these cases data is already available (including historical data) in relational or hierarchical databases since they are continuously produced by the related data grid services. For instance, concerning

Fig. 7 CMCC DDC - Metadata Browsing

LFC service, we set up both global and user-specific views about files distribution (in terms of number of files and in terms of total size) per storage element; number of created and updated files per storage element; replica distribution per file and storage element; etc.

Finally, it is important to remark that the CMCC Dashboard has been included into the CMCC environment since it provides the user with: (i) a global view of the CMCC data grid environment usage and status, including past and present, (ii) a valid solution to the monitoring of several grid activities inside the Centre as well as (iii) detection and diagnosis tools, (iv) advanced monitoring charts and (v) user/VO-centric reports about the CMCC activities. Additional info about the Dashboard approach/architecture that has been adopted at the CMCC, as well as several related works can be found in (Fiore et al. 2008a, b).

Related work

In the last years, other projects addressed similar issues at an international level (Earth System Grid, C3-Grid, Nerc DataGrid), with important differences (with respect to the proposed CMCC initiative) from the *middleware, metadata schema* and *metadata handling system* points of view.

The Earth System Grid (ESG) (Bernholdt et al. 2007) integrates supercomputers with large-scale data and analysis servers located at numerous national labs and research centers to provide a seamless and powerful environment that enables the next generation of climate research. The ESG project concentrates a lot of emphasis has been concentrated on the infrastructural part; it exploits *Globus middleware* and concerning the data grid part it leverages *GridFTP services*, *Storage Resource Manager*, *Replica Location Service* and *Opendap-g* servers. Concerning metadata management, ESG adopts a *centralized relational database* deployed at NCAR (directly queried by the portal) for descriptive or logical metadata which accurately describes a climate model experiment by means of a *Climate Model Metadata* (CMM). Concerning location or physical metadata (for replica management), ESG adopts a hierarchical and distributed framework based on Replica Location Services.

The C3Grid (Schindler et al. 2007; Kindermann 2006) project has been set up to enable an easier and more efficient resource management for the climate community, in order to improve the efficiency of scientific work both in terms of data storage and of computing. C3Grid strongly addresses data processing and data reuse through (i) portal integration of data processing workflows, (ii) grid workspace with data/job co-scheduling and (iii) metadata generation part of workflows. It offers an interoperable

framework able to deal both with *gLite* and *Globus* based environments. Moreover, C3Grid provides an uniform discovery of German climate data related providers (DKRZ, WDC Climate, IFM-Geomar, PIK, GKSS) through: (i) *ISO 19115/19139 metadata* based profile, (ii) OAI-PMH harvesting of metadata and GridSphere based Portal. Finally a *central metadata index* is used for metadata search from the C3Grid portal.

The NERC DataGrid (NDG) is a UK e-Science project that provides discovery of, and virtualised access to, a wide variety of climate and earth-system science data. Climate Science Modelling Language information Model (CSML) has been developed by the NDG project as a standards-based data model and XML markup for describing and constructing climate science datasets. It uses conceptual models from emerging standards in GIS to define a number of feature types, and adopts schemas of the Geography Markup Language (GML) where possible for encoding. In the NDG project a lot of emphasis has been devoted to metadata model (O'Neill et al. 2003), approach to discovery and use of data (O'Neill et al. 2004), data interoperability in the climate sciences (Woolf et al. 2004), NDG security (Lawrence et al. 2007), rather than the data grid infrastructural part (from a grid middleware point of view), compared with the other projects we mentioned before. Initial delivery services did not conform to any standard, de facto or otherwise. Concerning distributed climate metadata search, NDG Discovery is now based on Open Archives Initiative Protocol for Metadata Harvesting.

Conclusions and future work

This work presented a complete overview of the CMCC Data Grid environment, a distributed system aiming at managing tens/hundreds of petabytes of climate data for the scientific community. We highlighted several issues, concerning both data and metadata management as well as user support.

Concerning metadata management we discussed about the CMCC metadata agreement and related schema implementation, the CMCC metadata handling architecture and infrastructure, the metadata database and the distributed metadata search. On the other side for the data management part, we dealt with data transfer, access, replication and management services/issues and the Grid File Catalog. Security issues were also discussed in detail.

We also presented the available user support talking about the CMCC DDC, the Command Line Interface and the CMCC dashboard.

Future work will be related to (i) the enhancement of the infrastructural part (in particular the P2P metadata handling

system) and to the extension of the CMCC DDC. The CMCC dashboard will be further optimized and extended. The CMCC metadata agreement will be extended and completed taking into account (when needed) new scientific requirements. Data publishing will be strongly supported through portal web pages. Automatic ingestion of metadata (when possibile) will reduce time wasting for this step. For instance automatic registration of new data sources OPeNDAP-based (that means importing of metadata information about available NetCDF-CF datasets, involved variables, etc.) will ease indexing/registering new datasets into the metadata system.

Moreover in a few months the CMCC will start the pre-production phase. Complete suite of tests will allow to check and improve the global environment, the service integration and cooperation. Preliminary results and additional details about the involved services will be published in future works, when the entire CMCC infrastructure will be ready for the production phase.

In the next future, new tools will be considered to model scientific data-oriented workflow processes, (ii) provide easy to access visualization and post-processing tools, (iii) manage replicated datasets, (iv) increase the sharing of scientific results and new knowledge, (v) increase the scientific cooperation and collaboration, (vi) improve the resulting CMCC collaborative environment.

Finally the proposed CMCC technologies (GRelC DAIS metadata handling system and CMCC DDC) are currently tested as well, into the Climate-G testbed, an international research effort involving CMCC (Italy), IPSL (France), Fraunhofer-SCAI (Germany), NCAR (USA), University of Reading (UK), which provides a proof of concept about a large scale data environment for climate change. This international research effort (which is out of the scope of this work), will allow climate researchers and scientists to carry out geographical and cross-institutional data discovery, access, visualization and sharing of climate data.

References

Allcock W, Bester J, Bresnahan J, Chervenak A, Liming L, Tuecke S (2001) GridFTP: Protocol Extensions to FTP for the Grid. http://www-fp.mcs.anl.gov/dsl/GridFTP-Protocol-RFC-Draft.pdf Accessed 10 Sept 2008

Alfieri R, Cecchini R, Ciaschini V et al (2003) VOMS, an authorization system for virtual organizations. European Across Grids Conference 2003:33–40

Aloisio G, Cafaro M, Fiore S et al (2007) GRelC data gather service: a step towards p2p production grids. ACM SAC 2007:561–565

Aloisio G, Cafaro M, Fiore S, Mirto M (2005) The Grid Relational Catalog Project. In: Advances in Parallel Computing, Grid Computing: The New Frontiers of High Performance Computing. Elsevier, pp 129–155

Antonioletti M, Atkinson MP, Baxter R, Borley A, Chue Hong NP, Collins B, Hardman N, Hume A, Knox A, Jackson M, Krause A, Laws S, Magowan J, Paton NW, Pearson D, Sugden T, Watson P, Westhead M (2005) The design and implementation of grid database services in OGSA-DAI. Concurrency and Computation: Practice and Experience 17(2–4):357–376

Berman F, Fox G, Hey T (2003) The grid: past, present, future. In: Berman F, Fox GC, Hey AJG (eds) Grid computing. Wiley, New York, pp 51–63

Bernholdt D, Bharathi S, Brown D et al (2007) The Earth System Grid: Supporting the Next Generation of Climate Modeling Research. Available via arXiv.org. http://arxiv.org/pdf/0712.2262v1 Accessed 12 Sep 2008

Fiore S, Negro A, Vadacca S et al (2007a) Advanced grid database management with the GRelC data access service. International Symposium on Parallel and Distributed Processing and Applications LNCS 4742:683–694

Fiore S, Cafaro M, Negro A et al (2007b) GRelC DAS: a grid-DB access service for glite based production grids. International Workshop on Emerging Technologies for Next-generation GRID 2007:261–266

Fiore S, Negro A, Vadacca S et al (2008a) Advances in the GRelC Data Access Service, IEEE, to appear in Proceedings of ISPA2008

Fiore S, Negro A, Vadacca S et al (2008b) iGRelC: a Dashboard Implementation for Grid Environments, IEEE, to appear in Proceedings of PDCAT2008

Foster I (2005) Service-Oriented Science. Science Mag . doi:10.1126/science.1110411

Foster I, Kesselman C, Tuecke S, (2001) The Anatomy of the Grid: Enabling Scalable Virtual Organizations. International J. Supercomputer Applications, 15(3)

Gallagher J, Potter N, West P, Garcia J, Fox P (2006) OPeNDAP's Server4: Building a High Performance Data Server for the DAP Using Existing Software. AGU Meeting in San Francisco

Kindermann S (2006) Climate Data Analysis and Grid Infrastructures: Experiences and Perspectives–GELA Workshop. IEEE, Paris, France

Lawrence BN, Kershaw P, Blower J (2007) Practical access control with NDG-security. AHM 2007

Madduri R, Allcock W, Hood C (2002) Reliable File Transfers in Grid Environments Proceedings of the the 27th IEEE Conference on Local Computer Networks. p737–738

Middleton DE, Bernholdt D, Brown D et al (2006) Enabling worldwide access to climate simulation data: the earth system grid (ESG). Scientific Discovery Through Advanced Computing. J Phys 46(2006):510–514

O'Neill K, Cramer R, Gutierrez M, Kleese van Dam K, Kondapalli S, Latham S, Lawrence B N, Lowry R, Woolf A (2003) The Metadata Model of the NERC DataGrid, Proceedings of the U.K. e-science All Hands Meeting. Ed. S.J.Cox ISBN 1-904425-11-9

O'Neill K, Cramer R, Gutierrez M, Kleese van Dam K, Kondapalli S, Latham S, Lawrence B N, Lowry R, Woolf A (2004). A specialised metadata approach to discovery and use of data in the NERC DataGrid, Proceedings of the U.K. e-science All Hands Meeting. Ed. S.J.Cox, ISBN 1-904425-21-6

Santos N, Koblitz B (2005) Metadata services on the grid, in Proceedings of Advanced Computing and Analysis Techniques (ACAT'05). Zeuthen, Berlin

Schindler U, Bräuer B, Diepenbroek M (2007): Data Information Service based on Open Archives Initiative Protocols and Apache Lucene, German e-Science Conference 2007, Baden

Shoshani A, Sim A, Gu J (2004) Storage Resource Managers: Essential Components for the Grid. In: Grid Resource Management: State of the Art and Future Trends. Kluwer Academic Publishers, Norwell, pp 321-340

Stewart GA, McCance G (2006) Grid Data Management: Reliable File Transfer Services' Performance. CHEP06, Mumbai, India

Tuecke S (2001) Grid Security Infrastructure (GSI) Roadmap. Available via ggf1 website. www.gridforum.org/security/ggf1_-200103/drafts/draft-ggf -gsi-roadmap-02.pdf Accessed 12 Sep 2008

Woolf A, Cramer R, Gutierrez M, Kleese van Dam K, Kondapalli S, Latham S, Lawrence B N, Lowry R and O'Neill K (2004), Geophysical Research Abstracts, EGU 1st General Assembly-Standards-based data interoperability in the climate sciencese-science All Hands Meeting. Ed. S.J.Cox, ISBN 1-904425-21-6