

Federated data bases for the development of an operational monitoring and forecasting system of the ocean: the THREDDS Dataset Merger

S. Nativi¹, G. M. R. Manzella², F. Paolucci², P. Mazzetti¹, L. Pecci², L. Bigagli¹, and F. Reseghetti²

¹ Institute of Methodologies for Environmental Analysis, CNR-IMAA, Italy

² Italian National Agency for New Technologies, Energy and the Environment, ENEA, Italy

Received: 15 November 2005 – Revised: 26 January 2006 – Accepted: 5 March 2006 – Published: 6 June 2006

Abstract. During the last decade, operational monitoring and forecasting systems have been developed in all the European seas. The exchange of data and products and the development of services for a wide community of users pose some fundamental issues, whose solution has become a priority in integrated and GMES referring projects, such as the MERSEA European project. These projects aim to develop a European system for operational monitoring and forecasting on global and regional scales of ocean physics, bio-chemistry and ecosystems. GMES system and its operational projects need to federate resources and expertise coming from diverse organizations working on different Earth Sciences fields (e.g. satellite data processing, in situ observing systems, data management, ocean and ecosystem modeling, etc.). Therefore, it is required a Marine Information Management (MIM) system capable of facilitating the regular real-time exchange of high quality information, data and products. Moreover, MIM system must provide appropriate information for a wide range of external users both in real-time and delayed mode.

In this paper an architecture based on the OPeN-DAP/THREDDS technology is proposed as a solution for these operational systems. In this context, a catalog merging solution is introduced for the MIM system, which results in the design and development of the THREDDS Dataset Merger (TDM): a set of services meant to merge THREDDS Dataset Inventory Catalogs, so to achieve a unique catalog service for a whole database federation. TDM service merges distributed and autonomous THREDDS catalogs in order to work out a virtual merged catalog. The TDM service was extended in order to provide automatic catalogs synchronization. This service allows extending the pull-based TDM paradigm to support push-based applications. Some security issues are also considered.

Correspondence to: S. Nativi
(nativi@imaa.cnr.it)

1 Introduction

In the operational systems based on data and information coming from monitoring and forecasting of sea physical state, there exists the need to access datasets from many different data sources. Conversely to the meteorological weather prediction systems, where all data are conveyed to the forecasting centers through a unique standard service (the Global Telecommunication System – GTS), in the marine forecasting systems data are distributed in different data centers (i.e. satellites, in situ, meteorological forcing). Therefore, operational sea monitoring and forecasting systems are effectively modeled as Distributed Information Systems (DISs). Moreover, common requirements on Quality of Service (QoS) such as reliability, real-time support and high-volume data provisioning are difficult to achieve by means of centralized solutions (e.g. data warehouse), which have a single point-of-failure and pose scalability issues. On the other hand, the implementation of a DIS may support parallel data access and be more robust to malfunctioning, although at the cost of a more complex coordination among the subsystems (e.g., enforcing security policy and transactions).

DISs may be further classified into several categories, depending on the degree of heterogeneity and autonomy allowed, the overall provided functionalities, and the underlying technologies (Busse et al., 1999). For example, a cooperating collection of data base instances based on the same schema and interoperable technologies constitute a so-called *distributed database*.

In the above context, no assumption can be made on internal schemas and technologies of the individual subsystems, so a loosely-coupled approach is more appropriate, e.g. the Service-Oriented Architecture (SOA) (ISO/IEC 19119:2005, 2005) approach. Notably, Spatial Data Infrastructures (SDI) (GDSI Association, 2004), which may encompass sea monitoring and forecasting systems, are generally based on a SOA approach.

A desirable property of DIS is location transparency, that is the ability to interact with the DIS without knowing its physical distribution details (i.e. addresses). To implement location transparency, a DIS usually provides functionalities of data discovery, which support (automatic or user-driven) retrieval of the needed data access information.

In SOA terminology, we will refer to such data discovery functionality as to a *Catalog Service*. Catalog queries deal with information such as: human-readable description, temporal and spatial extent, access method, etc. In the context of SDI, examples of Catalog Services specifications include the OGC Catalog Service (Open Geospatial Consortium Inc., 2005), the FGDC Clearinghouse (FGDC), the Unidata Thematic Realtime Environmental Distributed Data Services (THREDDS) (Unidata THREDDS Group^a, Unidata THREDDS Group^b, Domenico et al., 2002).

The most important requirement for Catalog services is scalability: they usually feature hierarchical data models, with support to catalog nesting and, possibly, to distribute queries. In the last years many efforts have been dedicated to the development of complex cataloguing solutions for general federated information systems (SeaSearch.NET). Besides, when some assumptions can be made about the information system structure, tailored and more efficient solutions can be adopted. Indeed, when the hierarchical structure can be considered near-static and the needed scalability in terms of data sources is relatively small, an approach based on data aggregation can be suitable and profitable. In such a schema, upper level catalogs are built merging lower-level catalogs. This simple schema allows to reach good performances, avoiding to cascade queries through the catalogs tree. On the other hand, it results inappropriate when the system hierarchy is not stable or data sources increase, because a frequent aggregation becomes necessary.

1.1 The Marine Information Management System

Marine forecasting systems such as MERSEA (MERSEA Consortium), and their regional components have the strategic objective of providing an integrated service of global and regional ocean monitoring and forecasting to intermediate users and policy makers in support of safe and efficient offshore activities, environmental management, security and sustainable use of marine resources. One of the main goals for these operational systems is to facilitate the routine real-time exchange of high quality information, data and products within the project and the provision of appropriate information for a wide range of external users both in real-time and delayed mode. Since data and products will be generated in different institutions, efficient information and products provision from a single user portal requires a network of 'virtual data centers'. This objective is carried out conceiving and implementing a specific Marine Information Management (MIM) System.

In the framework of the European Global Monitoring for Environment and Security (GMES) programme, important operational system applications are:

- improvement of safety and efficiency of maritime transport and marine operations;
- support to safe and efficient offshore oil and gas industry activities;
- mitigation of effects of environmental hazards and pollution crisis;
- contribution to studies about ocean climate variability;
- advancements in marine research on global climate, the ocean and its ecosystems.

The development of such systems requires an unprecedented level of collaboration and coordination of multidisciplinary research and operational activities. For instance, the MERSEA project is federating the resources and expertise of diverse institutes, agencies and companies in the fields of satellite data processing, in situ observing systems, data management, ocean and ecosystem modeling, ocean, marine and weather forecasting and fisheries.

MIM is strategic for accessing data, model results and added value products residing in heterogeneous data bases located in different sites. Moreover, it is essential for providing users (both human and systems) with access to products and core services. A possible solution considered for implementing MIM systems could be based on OPeNDAP and THREDDS technologies (Unidata THREDDS Group, OPeNDAP Inc., Domenico et al., 2002).

A marine operational system deals with numerous activities, such as: data collection, scientific research, data assimilation in numerical models, information management, services and information provision for users and the public, respectively. These activities are carried out by providers such as: data centers, modelling centres, scientific teams, etc.

For MERSEA the two main categories of providers are observation data and models providers.

Observation data providers grant functionalities to access data obtained from satellites and in situ sensors, as well as from meteorological forcing fields. Models providers make assimilation model results available at global and regional scale. MIM systems are in charge of supplying functionalities to allow resource sharing and provide users (both human and systems) with data and services access.

Due to the amount of providers involved in operational projects such as MERSEA, a MIM system is necessary to support exchanges within the project and an effective relationship with users. As a matter of fact, the MERSEA MIM system federates the distributed global and regional observations and models providers, implementing a consistent access to their data and products.

1.2 MIM System Portals

In the MERSEA framework, a group of experts defined eight core themes. Table 1 reports selected themes: F1 to F3 deal with observation data (in-situ data, remote-sensed data and forcing field), while F3 to F8 concern model results for different ocean/marine areas. For each theme a set of core functionalities is required, including: data discovery, publishing and access. For distributed systems, data discovery and access services are based on an efficient catalog service.

A thematic portal (TEP) implements this set of functionalities. Therefore a TEP is responsible for providing services to users or other systems in order to discovery and access thematic data which are managed by distributed providers. A TEP implements a delegation mechanism for thematic communities. In this way, TEPs are in charge of:

- federating services checking their consistency and uniqueness;
- describing services and products;
- operating services in a homogeneous way;
- monitoring services and reporting (Key Performance Indicators).

Each TEP is designed as an access point to information about specific themes.

Actually, in service-oriented approaches TEPs services are built on the top of data and models providers services and therefore they can be considered as a higher level of services. In turn, the MIM system implements a general portal (i.e. the MIM Portal) which provides query and access services to all the different TEP's. Therefore it is a portal of portals implementing a unique access point to the MIM system for users (both human and other systems).

A portal stays at the top of this component hierarchy, giving access to the MIM System. This component is called MIM Portal and provides query services and links to all the different TEPs.

1.2.1 MIM catalogs hierarchy

In the previously described context, several catalogs have to be maintained at different hierarchical levels. Following a bottom-up approach, Regional Centers catalogs publish and give access to the collection of datasets managed by local data providers. Thematic catalogs provide access to datasets for thematic communities. Eventually the general catalog provide access to MIM system users. Figure 1 depicts the three-levels hierarchical structure of cataloguing. Since no more hierarchical levels are foreseen, this structure can be basically considered static. The number of thematic portals (and catalogs) is defined and there is no need for a substantial growth in the number of data providers. Anyway, for the sake of investments, a limited extensibility of the overall structure was taken into account.

Table 1. MERSEA selected teams.

Code	Short Description
F1	Global and regional ocean in-situ data
F2	Remote-sensed ocean products (sea level, sea surface temperature, sea spectral reflectance, sea ice)
F3	Numerical Weather Prediction (NWP) derived forcing fields
F4	Ocean Monitoring and Forecasting on Global Ocean
F5	Ocean Monitoring and Forecasting on Arctic Ocean
F6	Ocean Monitoring and Forecasting on Baltic Sea
F7	Ocean Monitoring and Forecasting on Atlantic North West Shelves
F8	Ocean Monitoring and Forecasting on the Mediterranean Sea

1.2.2 The main use-case

The requirements and technological choices are base on a typical use-case. In this use-case the user is a researcher who needs to query and browse observation data and models. He/she can access the system through either the main portal or the thematic portals. However, since the typical user is a member of one of the previously defined communities, the use of more than one portal in the same session is considered unusual. Thus, in the same session, the access to a single portal to perform many different queries is the typical user scenario.

1.3 THREDDs Dataset Merger (TDM)

Several data and models providers use Unidata's THREDDs and OPeNDAP technology. Thus, THREDDs can be considered a natural choice for MIM cataloguing service implementation. THREDDs Inventory Catalog data model supports distributed catalogs through the catalogRef construct, which directly supports browsing in a scalable fashion. However, a nesting approach based on remote references was judged non-optimal for the MIM use-case, characterized by searching, more than by browsing. In fact, this solution would imply frequent cascaded requests to the lower-level servers throughout the near-static service hierarchy as described in the use-case. Hence, a simpler catalog merging strategy was implemented: the lower-level server catalogs are replicated and cached on the upper-level aggregator servers. The system performance could be further improved using DBMS solutions for storing catalogs (namely, a Native XML Database would seem appropriate).

This solution, based on catalogs replication and aggregation, has the benefit of simplifying the search task avoiding query distribution. Anyway, it has important drawbacks mainly related to scalability issues. In particular, in a system where catalogs are frequently updated, or where hierarchy is complex or variable, the replication process can easily result too much demanding.



Fig. 1. Hierarchical structure of the MERSEA information system.

```
<?xml version="1.0"?>
<list-of-dic>
  <item>http://motherlode.ucar.edu:8080/thredds/catalog.1.0.xml</item>
  <item>file:///D:/xmlfiles/MerSea/catalog.xml</item>
</list-of-dic>
```

Fig. 2. Example of a list of Dataset Inventory Catalogs.

Although the solution is generally questionable, for our use-case, characterized by a simple and near-static hierarchy the trade-off was considered effective.

To implement this strategy, a solution based on the THREDDS catalog technology was designed and developed, namely: the THREDDS Dataset Merger (TDM) component and its related services.

The main TDM objectives are:

- To merge distributed and autonomous THREDDS Catalogs, at the dataset level, in order to work out a virtual merged catalog.
- To support an extensible hierarchical architecture of THREDDS Catalog servers.
- To aggregate and publish catalogs at different granularity levels, preserving catalog systems autonomy.
- To support both pull and push paradigms (that is the aggregation can be started by the Merger client itself, according to the so-called pull approach, or by the data provider, according to the so-called push approach).

This strategy is implemented by a couple of services:

1. The Merge service, which retrieves the specified catalogs and carries out their Datasets merging.

2. The Trigger service, which waits for specific event and activates the Merge service.

The Merge Service is the TDM core, while the Trigger Service allows to automate the merging task, enabling push-based interactions. Both services are implemented as servlets to work in a distributed environment based on Internet and Web technologies.

An experimental TDM Web service was developed and published in order to collect feedbacks from both THREDDS and GMES community (Nativi et al.).

1.3.1 MergeServlet

The MergeServlet merges local and remote THREDDS Dataset Inventory Catalogs (ICat) using a list of URIs encoded in an XML file, which is transferred as a parameter of a HTTP POST request. Both *Http* and *file* URI schemas are currently supported. For instance, *Error! Reference source not found. shows a list containing references to a remote and a local catalog.*

The merging process consists of the following steps:

1. create a blank new catalog;
2. retrieve the listed catalogs;

```

<?xml version="1.0" ?>
- <list-of-dc>
  <item>http://motherlode.ucar.edu:8080/thredds/casestudy/vgeeCatalog.1.0.xml</item>
  <item>http://motherlode.ucar.edu:8080/thredds/cataloggen/catalogs/uniModelsInvCat1.0en.xml</item>
</list-of-dc>

<?xml version="1.0" encoding="UTF-8" ?>
- <catalog xmlns="http://www.unidata.ucar.edu/namespaces/thredds/InvCatalog/v1.0" xmlns:xlink="http://www.w3.org/1999/xlink"
  name="Data for VGEЕ Curriculum">
  <service name="motherlode" serviceType="DODS" base="http://motherlode.ucar.edu/cgi-bin/dods/DODS-3.2.1/nph-dods/" />
+ <dataset name="ENSO Curriculum">
</catalog>

<?xml version="1.0" encoding="UTF-8" ?>
- <catalog xmlns="http://www.unidata.ucar.edu/namespaces/thredds/InvCatalog/v1.0" xmlns:xlink="http://www.w3.org/1999/xlink"
  name="Unidata IDD Model Data">
  <service name="latest" serviceType="Resolver" base="http://motherlode.ucar.edu:8080/dqcServlet/latestModel" />
  <service name="mlode" serviceType="DODS" base="http://motherlode.ucar.edu/cgi-bin/dods/DODS-3.2.1/nph-dods/" />
  <property name="CatalogGenConfigOrigURL" value="/thredds/cataloggen/config/catGenConfig.ModelsInvCat1.0en.xml" />
+ <dataset name="Unidata IDD Model Data" ID="ucar.unidata.idd.model">
</catalog>

<?xml version="1.0" encoding="UTF-8" ?>
- <catalog xmlns="http://www.unidata.ucar.edu/namespaces/thredds/InvCatalog/v1.0" xmlns:xlink="http://www.w3.org/1999/xlink"
  name="MerSea THREDDS-IDD Server">
  <service name="motherlode_1" base="http://motherlode.ucar.edu/cgi-bin/dods/DODS-3.2.1/nph-dods/" serviceType="DODS" />
  <service name="latest_2" base="http://motherlode.ucar.edu:8080/dqcServlet/latestModel" serviceType="Resolver" />
  <service name="mlode_2" base="http://motherlode.ucar.edu/cgi-bin/dods/DODS-3.2.1/nph-dods/" serviceType="DODS" />
  <property name="CatalogGenConfigOrigURL" value="/thredds/cataloggen/config/catGenConfig.ModelsInvCat1.0en.xml" />
+ <dataset name="ENSO Curriculum">
+ <dataset ID="ucar.unidata.idd.model" name="Unidata IDD Model Data">
</catalog>
  
```

Fig. 3. An example of catalog merging.

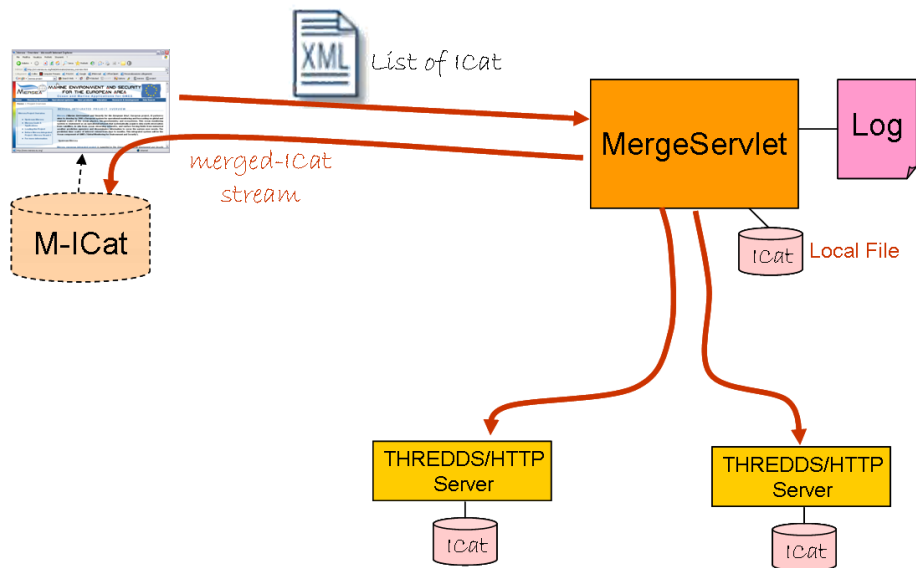


Fig. 4. TDM Collaboration/Deployment diagram.

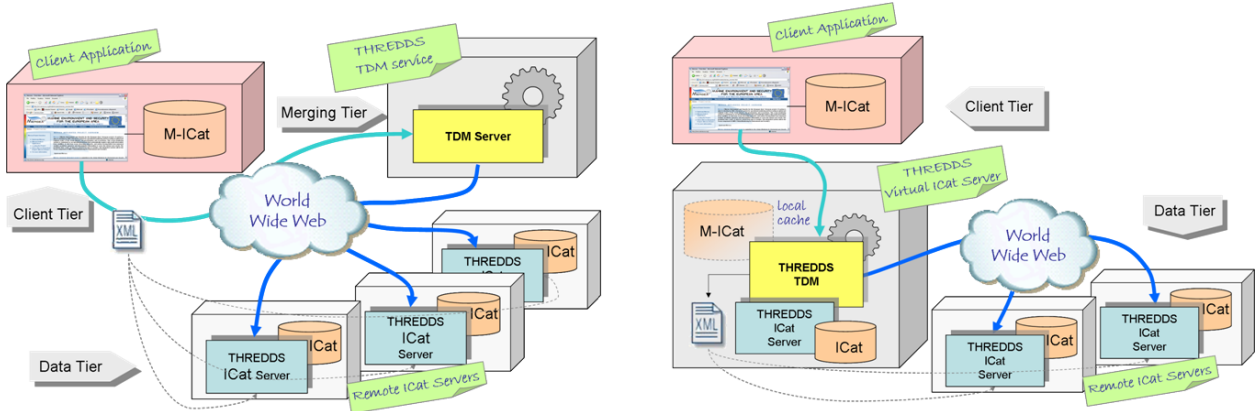


Fig. 5. The TDM as a catalog merging service and as an application, respectively.

3. examine the retrieved catalogs, for each Service tag found:
 - 3.1 create a new Service tag in the new catalog. In order to avoid name collisions, an identifier corresponding to the catalog retrieving order is added to the original name;
4. examine the retrieved catalogs; for each Dataset tag found:
 - 4.1 If the Dataset tag is not in the new catalog, create a new Dataset tag in the new catalog, changing the Service tag reference to the new identifier (the presence of the attribute ID enforces the uniqueness of Dataset tag reference);

In case of a URI pointing to a missing or unreachable resource, the event is logged, the missing catalog is skipped and the merging process continues.

Error! Reference source not found. shows an example of catalog merging. Two remote catalogs containing one dataset each are merged in a single catalog containing two datasets. The merging process is performed through an XSLT transformation.

Actually, the MergeServlet is sufficient to implement the pull strategy. In fact, *Error! Reference source not found.* shows the Collaboration diagram for the pull approach based on the MergeServlet component. An External Process calls the MergeServlet passing a list of catalogs to be aggregated. The MergeServlet retrieves the original catalogs accessing both THREDDSS server components and local filesystem. The merging result: the merged inventory catalog, called M-Icat is returned to the External Process.

This general Collaboration Diagram can be implemented according to a couple of deployment solutions: a catalog merging service or a client application component.

In the first deployment solution we can distinguish three different tiers. A Client Tier containing the Client Application which asks for and manages the merged catalog. A

Merging Tier which consists of a TDM Server that is an application server to run the MergeServlet. The last tier is the Data Tier and contains all the THREDDSS server which manage the original catalogs. The three tiers interoperate in a Web distributed environment. In particular, the Client Application passes the catalog list to the TDM Server as parameter of a HTTP POST request. The TDM Server retrieves the original catalogs using HTTP GET calls.

The second deployment solution implements a two-tier approach. This time the TDM Server is deployed on the Client application node; thus the Client Application and the TDM server which communicates by a local call (i.e. an internal HTTP call).

The two possible deployments are depicted in Fig. 5.

1.3.2 Trigger Servlet

The Trigger servlet (TrigServlet) triggers the catalog merging task allowing for the implementation of the push strategy. This is done by exposing a Web interface, implemented as a servlet, that external applications can call to initiate the aggregation task. The trigger reads a property file which stores the path of an XML document containing: the list of catalog URIs to be aggregated, the merged catalog path and the MergeServlet URI. The TrigServlet is in charge of launching the MergeServlet whenever a specified event occurs. *Error! Reference source not found.* shows a Collaboration & Deployment diagram of a synchronization and merging use-case. A remote catalog and its local copy are maintained synchronized by a specific external application (e.g. a SyncML-based application). Each time a change occurs in the remote catalog, the local copy is updated. To start the catalog merging task, a signal is sent to the Trigger Servlet as a HTTP GET call. The Trigger component reads the required information and calls the MergeServlet. This could be deployed either in the same computational node of the TrigServlet, or in a different one.

This approach avoids misalignment between the remote and local catalogs and allows running the merging process

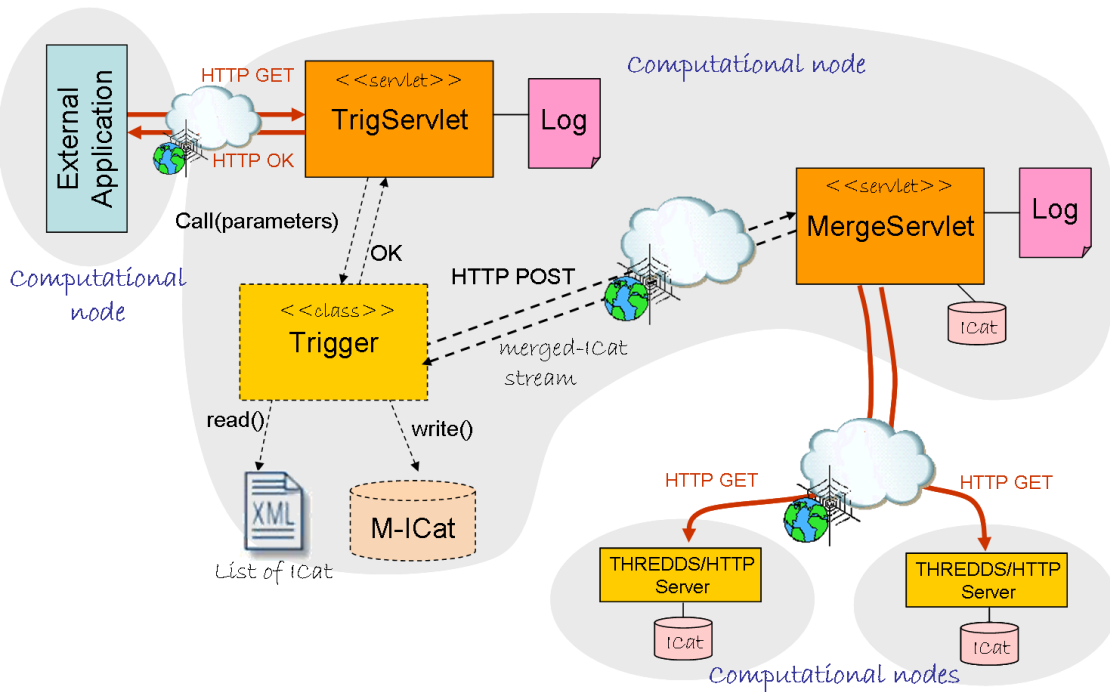


Fig. 6. Synchronization and merging use-case basing on the Trigger Service.

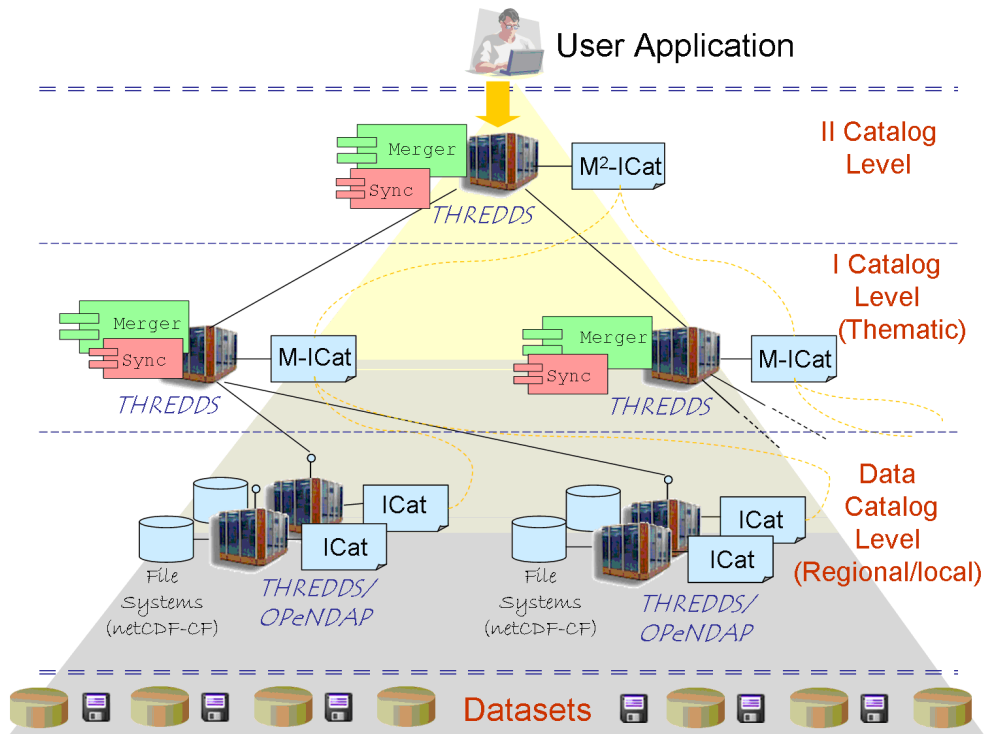


Fig. 7. TDM and synchronization: the MERSEA use-case.

only when really needed that is when a change in the remote catalogs occurs. Indeed this implements a push-based interaction paradigm.

1.4 The TDM application to the MIM System: a push-based implementation

In the MERSEA MIM framework, an experimental TDM service was published by IMAA-CNR and ENEA (Nativi et al., 2005). A three-levels merging chain was implemented, as shown in *Error! Reference source not found.*. At the bottom level there are the regional/local data catalogs; they are THREDDS ICats and describe MERSEA datasets which are netCDF-CF files (Unidata netCDF Group, Gregory, 2003) accessed through OPeNDAP servers. Since each provider should be free to make its own technological choices, no assumption was made on the adoption of particular OpenDAP servers' implementation. These ICat datasets are merged in a set of virtual ICats, called M-ICats. It is the first MERSEA aggregation level: the Thematic level. Thematic virtual catalogs are merged in turn at a second aggregation level: the general catalog level; this process works out the M²-ICat. User applications can access this top level virtual catalog as well as any thematic level catalog (i.e. M-Icat).

From the bottom, the first catalog type is the Data Catalog. At this level, each catalog describes datasets managed by a regional or local data/model provider. The second catalog type is the Thematic Catalog. Each of them aggregates information stored on the Data Catalogs, on a thematic base. The third catalog type is the Global Catalog (i.e. the MIM catalog). The global catalog aggregates thematic information supporting query and access services.

Each aggregating node is made of two components: the Merger and the Synchronizer communicating through the Trigger service. TDM provides the Merger and the Trigger components, while a specific Synchronizer module (Sync) was developed to support the push-based aggregation process. It maintains synchronized a local copy of each catalog to be aggregate. Thus, the aggregation can be performed using local copies instead of remote ones, avoiding connection related problems. The Sync module is implemented through a messaging queue which distributes the catalog update actions. On each remote node a specific application acts as a Producer sending a message whenever a catalog is updated. On the aggregating node, a specific application acts as a Consumer updating the local copies through the message content. The message exchange is based on JBOSS JMS Server which is in charge of handling queues. In the current implementation each time an update is performed, the Consumer retrieves the whole catalog. This approach is acceptable since the catalogs sizes are rather small (less than 1 MB) and their update rates are low (less than once a day). Future implementations will adopt incremental synchronization.

1.5 Security issues

In a context of operational services, different providers must share their products; whereby, data policy and security issues must be carefully considered. The required solution must provide security functionalities in terms of authentication, access control and logging. On the other hand, technological and organization constraints restrict the range of possible solutions; for example, THREDDS/OPeNDAP servers have very limited security capabilities. Hence, we designed an AAA (Authentication-Authorization-Accounting) architecture that is pluggable on the hierarchical structure of the previously described information system. It adopts a centralized-authentication and distributed-authorization schema. This solution implements a flexible approach that simplifies the authentication process. A unique authentication server implements all the required different schemas (e.g. username/password, digital certificate, etc.). However, each Authentication Manager allows providers to control the access to its own data, using the authenticated requestor identity. This enables the application of locally defined data policies.

As far as the MERSEA MIM framework is concerned, an Authenticator module must be deployed on the MIM portal server, while an Authorization module must be installed on each data server to control data access. Figure 8 sketches the described solution. For the sake of existing data servers autonomy, or for the extensibility limitations of used technologies a proxy-based approach was adopted. A transparent proxy equipped with the Authorization module is installed beside each data server. Therefore, the authorized requests are directed to the real data server, while unauthorized requests are blocked by the proxy.

A security system based on this architecture is under development. In this implementation the Authentication Manager and the Authorization Managers are realized as servlets running in a Tomcat application server.

2 Conclusions

In many Earth Sciences operational fields, it is necessary to implement performing distributed catalog services. A simple solution consists of merging catalogs. This is suitable for near-static hierarchy of catalog servers characterized by limited planned scalability. In the framework of a European GMES project (i.e. MERSEA), a solution based on THREDDS Dataset Inventory Catalog was designed and developed; it is called THREDDS Dataset Merger (TDM). TDM can be used to build hierarchical distributed systems of catalogs. TDM functionalities can be provided by either a locally deployed software component or a remote Web service. Both deployment approaches were presented. An experimental TDM Web service was developed and published

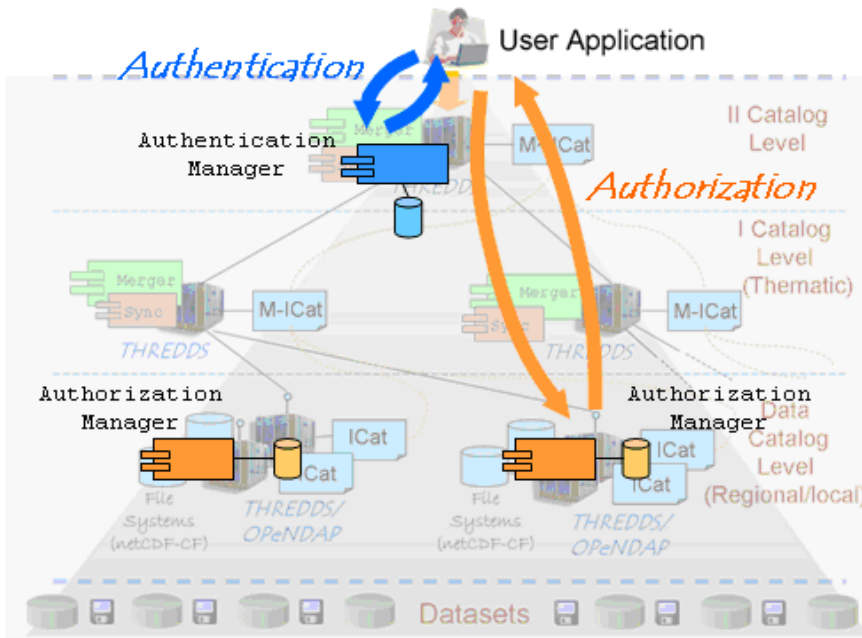


Fig. 8. Security architecture for centralized authentication and distributed authorization.

in order to collect feedbacks from both the THREDDS and GMES communities.

TDM has been experimented in order to implement the catalogue service of the MERSEA information system. Using a bottom-up description, the MERSEA information system is characterized by three catalog levels: a data, a thematic and a general catalog level. TDM provides all the components and services needed to implement all the three catalog levels.

Currently, an on-demand request-response message pattern and an automatic publish-subscribe pattern are both available. For the publish-subscribe pattern, a TDM synchronization service was worked out. Future development includes the deployment of a security system based on a centralized authentication schema along with a distributed authorization approach. These schema allows a full access control to cataloged data sources.

Acknowledgements. This work is a contribution to the MERSEA Project. Partial support of the European Commission under Contract SIP3-CT-2003-502885 is gratefully acknowledged.

Edited by: E. Cutrim, M. Ramamurthy, S. Nativi and L. Miller

Reviewed by: anonymous referees

References

- Busse, S., Kutsche, R.-D., Leser, U., and Weber, H.: Federated Information Systems: concepts, terminology and architectures, Technical Report Nr. 99-9, TU Berlin, 1999.
- Domenico, B., Caron, J., Davis, E., Kambic, R., and Nativi, S.: Thematic Real-time Environmental Distributed Data Ser-

vices (THREDDS): Incorporating Interactive Analysis Tools into NSDL, *J. Dig. Inf.*, 2, 4, Art. No. 114, 29 May 2002.

FGDC: Clearinghouse, [urlhttp://www.fgdc.gov/dataandservices/](http://www.fgdc.gov/dataandservices/) (last access: April 2006), March 2006.

Gregory, J.: The CF metadata standard, http://www.cgd.ucar.edu/cms/eaton/cf-metadata/clivar_article.pdf, 2003.

GSDI Association: Developing Spatial Data Infrastructures: The SDI Cookbook; version 2.0, GSDI publication, 2004.

ISO/IEC 19119:2005: Geographic information – Services, ISO International Standard, Geneva 2005.

MERSEA Consortium: MERSEA project web site, <http://www.mersea.eu.org/> (last access: April 2006), 2004.

Nativi, S. and Manzella, G.: THREDDS Dataset Merger: a push solution to aggregate and synchronize distributed Catalog datasets, Create your Merged Catalog, <http://athena.pin.unifi.it:8080/mersea/tdm> (last access: April 2006), 2005.

OPeNDAP Inc.: OPeNDAP web site, <http://www.opendap.org/> (last access: April 2006).

Open Geospatial Consortium Inc.: OGC™ Catalogue Services Specification, OGC™ project document: OGC 04-021r3, Ver. 2.0.0, 2005.

SeaSearch.Net: Sea Search web site, <http://www.sea-search.net/> (last access: April 2006).

Unidata THREDDS Group^a: netCDF Conventions, netCDF Users Guide, <http://www.unidata.ucar.edu/packages/netcdf/conventions.html> (last access: April 2006).

Unidata THREDDS Group^b: netCDF Usage, Unidata software:netCDF, <http://www.unidata.ucar.edu/content/software/netcdf/usage.html> (last access: April 2006).

Unidata THREDDS Group: THREDDS project web site, <http://www.unidata.ucar.edu/projects/THREDDS/> (last access: April 2006).