Architecture Considerations for Advanced Earth Observation Application Systems

Hermann Ludwig Möller[1], Marcello Mariucci[2] and Bernhard Mitschang[3]

[1] European Space Agency, Directorate of Application Programmes, hmoeller@estec.esa.nl
Formerly ESA-ESRIN, Italy, since January 1999 affiliated to ESA-ESTEC, The Netherlands.
[2] Universität Stuttgart, Institut für Parallele und Verteilte Höchstleistungsrechner (IPVR), Marcello.Mariucci@informatik.uni-stuttgart.de
[3] Technische Universität München, Department of Computer Science, Bernhard.Mitschang@informatik.tu-muenchen.de

Abstract. Application systems in the earth observation area can be characterised as distributed, platform-inhomogeneous, complex, and cost intensive information systems. In order to manage the complexity and performance requirements set by these application scenarios a number of architectural considerations have to be applied. Among others the most important ones are modularization towards a component architecture and interoperation within this component model. As will be described in this paper, both are mandatory to achieving a high degree of reusability and extensibility at the component level as well as to support the necessary scalability properties. In our paper we refer to the state of the art in earth observation application systems as well as to a prototype system that reflects to a high degree the above mentioned system characteristics.

Key Words: Distributed Information Systems, Earth Observation Systems, Applications, Interoperability, Middleware, CORBA

1 Introduction

Since the early '70s an increasing number of satellites orbit our planet and make observation data related to sea, land, and atmosphere available globally. The data is used in support to a number of applications; the best known might be the daily weather forecast satellite maps and animations shown on the TV news programmes. Fleets of new satellites will produce about 1TByte of new data every day and soon the amount of data collected within a single year will equal the size of all acquired data of the last 25 years. With the increase in observation platforms also the number of applications is increasing. For example, since the early '90s, Europe has been exploiting its European Remote Sensing Satellite

[1] The contribution of the author is related in particular to his work during a one year research fellowship assignment at NASA/GSFC

ERS-1 and -2, e. g.for sea ice monitoring, oil-pollution monitoring or in support to disaster management. Earth observation (EO) satellites represent an investment of several hundred million ECU per space craft. In order to justify such investment, which still is largely public funded, new emphasis has been given in recent years to the development of ground segments with focus on exploiting the data streams received from the satellites for specified applications also beyond scientific use. The present paper focuses on architectural considerations of application-specific information systems for data exploitation.

2 Earth Observation Application Data and Information Management

EO systems distinguish between a space segment, comprising the space craft and associated command and control systems for flight operations, and a ground segment, comprising facilities for data acquisition, processing, archiving and distribution. In the following focus is given on the ground segment and in particular ground system elements that may help to relay observation information to its specific use in end-user applications. EO Ground Systems are defined by means of three levels:

- 'Data Level' (DL): large scale infrastructures primarily operated by space agencies and satellite operators as data providers (DP), handling the data acquisition from the EO satellite and the data handling for standard data processing and archiving;
- 'Information Level' (IL): infrastructures primarily operated by Value Adding Companies (VAC) and scientific institutions (SI) for creating higher level application specific information, e. g., through thematic processing, and used by Value Added Resellers (VAR) for distributiong such information;
- 'End-User Level' (UL): user access infrastructure, interface and local infrastructure serving scientists, governmental and commercial users, and the educational sector.

Traditionally, DL infrastructures interface directly with the user segment on UL, serving a multitude of user requests through a single system architecture. Search and order of standard data products are the main functions externally accessible in such systems. Figure **??** illustrates now the additional level, at present in prototyping stage, the IL. The IL constitutes an additional layer, interfacing upstream with the DL for the provision of standard EO products, and downstream interfacing with the UL for application-specific user access and distribution. The IL is not one single system but comprises a multitude of smaller systems each serving a different user domain and each interfacing the DL level separately. Some IL functions, in particular data ingest may be provided by space agencies and data providers in general. Other IL functions, e. g.thematic, application-specific processing, may migrate towards VAC or SI. Selected IL functions may be shared within a cluster of individual IL systems. IL functions such as archiving of thematic products may be covered by DP, thus reaching from DL into IL functions. End-users (EU), depending on the service support required may be associated either with an SI, a VAS or a VAR; even in some cases EU may be registered with more than one information service supplier.

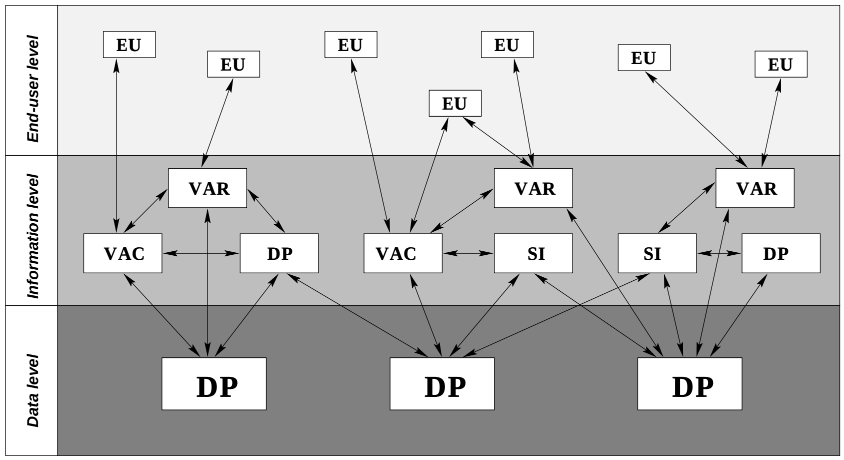

Fig. 1. Information Level

Prior to focusing on the IL, its functions, associated interoperation, technology issues, and federation concepts, an estimation of EO data volumes and access statistics shall help to understand the magnitude of the problem present day DL are facing. Estimates in the domain depend significantly on a number of assumptions, e. g., user behaviour and market evolution, and on the range of satellites considered, e. g., geostationary and low-orbiting satellites, meteorological and environmental satellites. However, the numbers provided will help in a first approximation to determine a number of requirements for the design of the IL infrastructure and IL system concepts.

2.1 Data Volumes

EO Ground Segments handle high volumes of data, both in terms of archived historical data and in terms of newly ingested data. Estimates of the total volumes considering all major LEO observation constellations world-wide point at more that 300 TByte of archived data in over 20 million inventory records, varying between a few kBytes and more than 2 GByte in size. Ingest of new EO data is estimated in the range of 1 TByte per day, with a corresponding user pull above 2 TByte in the next future. The actual user pull will depend on how successful the data exploitation will be performed. One important enabling element for such exploitation and global data usage are interoperable IL systems.

Existing Archives (1970's-present)			Number of Users	**Daily** Ingest from Satellites as from 1998/99	**Daily** User Pull as from 1998/99 (media&on-line)
Total Data Volume	Inventory Entries	Single Data Item Size			
> 300 TByte	> 20 Million	kByte – > 2 GByte	> 20000	500 GByte – > 1 TByte [a]	> 2 TByte

[a] Numbers derived from NASA, ESA and EUMETSAT estimates

Table 1. EO Information Systems - Volume Estimates - World-wide

2.2 Access Statistics

As an example, based on earlier NASA estimates for US based systems, Figure **??** provides estimates of category and number of users, and nature of electronic access to EO data and information, as expected for state-of-the-art DL systems becoming available in the near future.

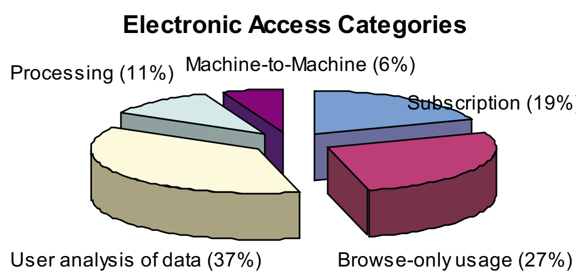

Fig. 2. EO Information Systems – Electronic Access Categories

According to the example, the scientific community with an estimated >12000 users globally represents the largest part of the user population. It is also within this community that 'user analysis of data' at UL is most prominent with 37%, as the scientific usage of data often requires in depth analysis of particular phenomena related to a data set. User requests leading to 'processing' represent another significant usage type and may include application specific processing of data, e. g., extraction of thermal fronts from sea surface temperature data along coast lines for detection of fishing grounds. Such thematic processing represents an estimated 11% of the accesses. It is also within this type of usage that user access may result in subsequent machine-to-machine interoperation hidden to the user, e. g.transfer of a particular data set from a remote archive to a thematic processing system.

A different utilisation is related to the users in the educational community, estimated at 70000 – 200000 users. Usage is focused on browse-only (27%), e. g., investigation of thumb-nail images without further data processing or analysis, and subscription services (19%), where users register for data provision related to specified locations, times, or events, e. g., the weekly image of their hometown.

3 Information Level Federation Concepts

3.1 Information Level

At present, DL systems need to handle very diverse usage profiles in a single system. However, government use and in particular the percentage of commercial usage are expected to grow significantly in the coming years as distinguishable application domains. VAC and SI emerge for serving these new user market segments. They will require automation. E. g., a high percentage of 'user analysis of data', today performed by experts at UL after consultation of DL, may need automation to serve these new user communities. To do this effectively, systems will need to be capable of handling application specific usage in a way today only available for expert users, for example a scientist with the knowledge where to find the data and which processing to apply in order to obtain desired information. Prototyping of such systems, constituting the forerunner of IL systems, is currently in progress in a number of selected application domains.

First results have shown that individual IL functions may be applicable to more than one application domain and system. E. g., advertisement of available services, which is today rather limited because services are already known to the scientists user group, becomes a challenge across IL systems in order to attract new governmental and commercial users to the information services. Interoperation of participating components for advertisement, storage, processing, workflow, and access is becoming essential. Proper partitioning of applications into a federation of individual IL systems is important in order to achieve the scalability and performance required by individual user groups. I.e., distribution and interoperation of data and functions and machine-to-machine interoperation is becoming an issue. Table **??** provides a comparison of complexity between individual data level systems vs. information level systems.

Individual IL systems distinguish themselves from large-scale DL systems in that they typically respond to the needs specific to a scientific discipline or particular application, e. g., flood monitoring or urban planning. Unlike satellite operator's large-scale data systems managing multiple TBytes, these systems are mostly concerned with a more limited amount of data, e. g., multiple GBytes, specific to a defined usage. The number of individual systems is expected to be one order of magnitude higher for IL than for DL. DL systems interoperate primarily internally, whereas IL systems may develop different degrees of interoperation in smaller clusters. User interfaces, rather uniform for DL may become application specific on IL, reflecting well-identified requirements of much smaller and defined user groups. This will lead to complementary interfaces to

	DATA LEVEL	INFORMATION LEVEL
Nature of System & Service	Universal ('query from hell')	Specific to application
Typical Data Volume per system/service	TBytes	GBytes
Number of individual systems/services globally	< 10	> 100 (estimate for 2005)
Level of interoperation	Primarily internal to individual system, Inventory external	Interoperation in a number of clusters (estimate > 20 globally) to varying degree depending on type of federation
User Type	Mostly scientific	Science, commercial, education/wider public
Number of users per system/service	> 10.000	± 10 government and commercial ± 50000 education/public
Processing requirements per system/service	Complex, across many scientific disciplines	Known with application, e. g.one algorithm per system
Access/Dissemination requirements	Universal	Tailored to user type

Table 2. Complexity: Data vs. Information Level

the general purpose data search and discovery interfaces dominant in today's Earth Observation User Information Systems on DL [MUIS].

The emergence of open, network-centric middleware, in particular the Object Management Group's Common Object Request Brokers Architecture (CORBA) [OMG98, Si96, OH97], and the wide availability of advanced communication networks, in particular the Internet, provide the essential elements for the required underlying infrastructure for distributed IL services, operating in a federation of individual systems.

3.2 Various Federation Scenarios

DL system developments inherently provide a high level of interoperation in a distributed environment as they are typically implemented as single large-scale projects, e. g., ESA's ENVISAT Payload Data Segment (PDS) [PDS], or NASA's Earth Observing System Data and Information System (EOSDIS) [EOSDIS] with its Data Model and Distributed Information Management (DIM). However, this interoperation is merely internal, i.e., interoperation options for external data providers and different DL systems either assume the adoption of the internal standards of a given DL system or external interoperation is limited or non-existent. The Committee on Earth Observation Satellites (CEOS) has made an attempt to achieve interoperation of DL systems for a limited functional scope, i.e., for queries of large scale inventories, and has specified the Catalogue Interoperability Protocol (CIP) [CIP97].

In contrast to the above, IL systems are expected to be developed in a multitude of smaller development projects. As individual projects may not have the resources to develop and serve all required service functions internally, a market opportunity for externally available, interface-compatible system functions and services beyond catalogue services exists. IL systems and services provided primarily by VAC and SI would orchestra in a federation characterised by

- The *level of distribution* of individual information service functions, like thematic image processing, and its allocation under the responsibility of different players in the service market, e. g.data providers, VAC, SI.
- The *level of interoperation and reuse* between different systems and distributed information service functions, like common advertisement services or sharing of a common data archive between IL systems under the responsibility of VAC or SI.

Four different degrees of federation are proposed in the following, ranging from a 'Non' Federation, with distribution and interoperation only within the DL, to a 'Full' Federation configuration where most functions are migrated to IL and distributed with a high degree of function interoperation.

'Non'-Federation The 'Non' federation (see Figure **??**) represents the *state-of-the-art* in EO information systems and in principle reflects today's DL-only configurations, i.e. without IL. All functions typically are provided through a single distributed DL system (per world-region) appearing to the user as a single system. DL sub-systems, e. g., for processing and archiving, interoperate according to a single DL-internal schema, called the common DL schema.

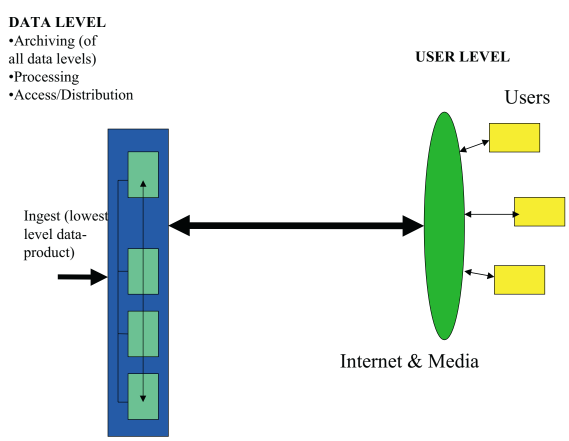

Fig. 3. 'Non'-Federation

This schema distinguishes a *user view* from a *conceptual data description* and a *physical data description*. The user view is a logical description that represents the end-user perception of the data. The conceptual data level describes the data assets, illustrating relationships among classes, e. g.defining the data input/output relations per sub-system, and specifies the attributes of the data. The physical data level refers to the platform dependent representation of the data as implemented using commercial database management systems. The common DL schema is essential for DL-internal interoperation. Different DL may offer catalogue interoperability services based on an agreed schema subset or a derived standard, limited to their inventory subsystems. Reuse of components within the DL is maximised. Partitioning of the DL into application-specific information system components is however strongly limited by the fact that individual components are typically designed to handle a broad range of different information. Therefore, DL systems as such are not easily adaptable or sizable to application-specific systems as defined in the IL.

'Processing' Federation The 'Processing' federation depicted in Figure **??** shows thematic processing as first functions migrated into a thereby created IL. Application specific information is typically generated under VAC or SI responsibility. Some VAC or SI may decide to share processing functions, i.e., an algorithm may be executable within a cluster of VAC IL systems. User access and distribution has not changed and is still achieved through the DL system. The IL typically builds on the DL schema. Although the 'Processing' federation improves the adaptability of the system to application-specific requirements and allows a first separation of service functions, its scalability and its reuse potential are still limited by the underlying DL.

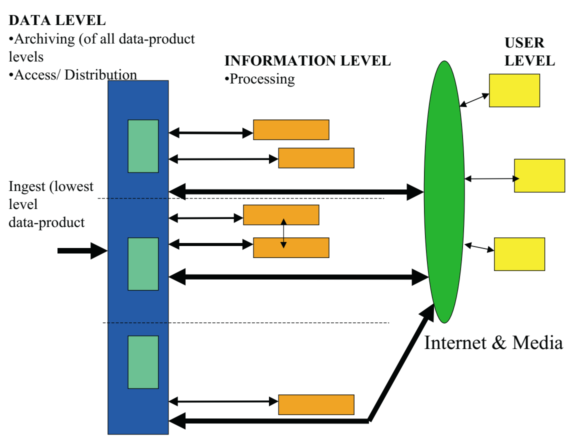

Fig. 4. 'Processing' Federation

'Processing and Access/Distribution' Federation The 'Processing and Access/Distribution' federation (in Figure **??**) is an essential step towards a federation allowing the VAC or SI to identify and distinguish itself visible to the user. The IL may be defined by different overlapping IL schemata, which for archival functions may still be those of the associated DL. A cluster of a few IL systems may decide to operate through a single user interface and provide a common user request management, e. g., forwarding of user requests and context information between their services. However, access functions are expected to serve as distinguishing feature for individual IL systems and will only show a limited level of interoperation. This interoperation may focus on directory and trading services for the advertisement of available IL services across a cluster. Interfaces to GIS for the provision of non-EO information as complementary data may be part of a VAC or SI offering. Standard data product archival is still performed at DL, management, storage and archival of thematically processed data may however be allocated to the IL. IL sub-system components, e. g.ingest module, may be reused from one IL system to another. With most functions migrated to the IL, the level of modularization of service functions, the system scalability, its reuse potential and adaptability are high. *Prototype systems* reflecting such a federation concept are currently being demonstrated (see chapter 4).

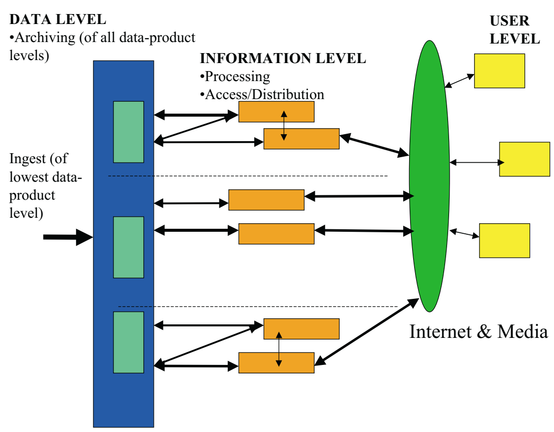

Fig. 5. 'Processing and Access/Distribution' Federation

'Full' Federation The 'Full' federation depicted in Figure **??** reduces the DL to the long-term archive for lowest level data products. Other archival functions have been moved into the IL segmented according to geographical regions and applications. Together with interoperable interfaces for archive services within a cluster, the number of interoperable, related processing services is expected to increase. User access/dissemination is expected to remain with a lower level of interoperation. Interoperation with traditional DL systems is no longer an issue for most VAC or IS as this is dealt with within the IL. The lack of any common IL schema maintained outside the IL provides a challenge to the interoperation within the IL. Leading federation members, e. g., those comprising the archival function, may provide a reference data model and serve as architectural model. Level of reuse, scalability, modularization, and adaptability is comparable to the 'Processing and Access/Distribution' federation.

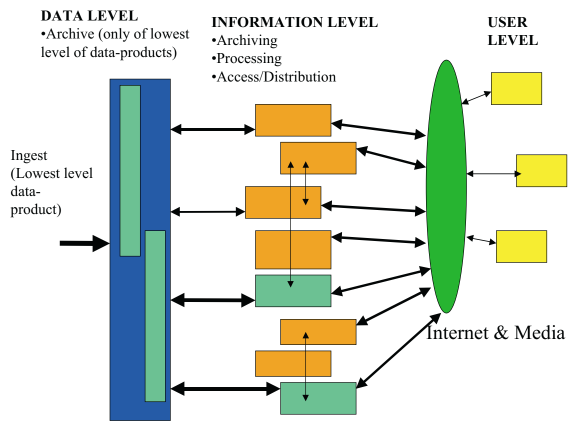

Fig. 6. 'Full' Federation

Outlook on Federations Depending on the success of DL systems in the different world-regions and the pressure from application markets, the degree of federation in the longer run may be different between regional markets. It may be established between the 'Processing and Access/Distribution' federation and the 'Full' federation, which most likely remains merely a long term perspective. Much will depend on the way in which new enabling technologies will be inserted into the development process. Leading IL service and system developments may set the pace for the federation, or a co-ordinating working group may provide recommendations for the IL.

3.3 Technology Considerations

During the last years system engineering and software design for federated information systems have undergone tremendous technological changes. Modern architectures are based upon components with well-defined interfaces and behaviour. This step enables reuse, extensibility, and scalability of systems or system components. Adoption of new technology or application requirements can be done by exchanging single components only. Furthermore, each module can be developed by separate companies with different expertise. In order to enable the usage of third-party software published without source code, several component models have been defined, e. g.Microsoft's DCOM or ActiveX [Se98], IBM's DSOM [La95], OMG's CORBA [OH97], etc.

Considering the great Internet wave starting in the early 90's, component-based design became even more important. Companies want to establish so-called virtual enterprises, accessing resources of their partners and offering e. g.online order facilities for their customers. Thus there was an increasing need for standardized components which can interoperate with other ones via Intra/Internet, no matter which programming language or operating system is used. The success story of Java and its well defined components started. Though Java offers several ways for component interoperabilty, it is still a particular programming language. In order to be extensible w.r.t. any kind of system aspects it is not appropriate to define interfaces of components in a specific and single programming language. For example, in independence w.r.t. interface definition can be easily achieved by means of CORBA's Interface Definition Language (IDL). It is independent of programming languages, but mappings exist or can be developed for any language as needed. In addition, arbitrary components can interact through ORBs and interfaces, and the behaviour of basic components are already standardised by the OMG.

Though we have just detected a suitable component model, federated systems generally raise another requirement: Data models of each participating partner have to be compatible. CORBA does not offer mechanisms to resolve this issue, but federated database technology (as e. g.provided by so-called database middleware [RH98] like IBM's Data Joiner [IBM97]) together with standardization endeavours (e. g.OGIS) may help. The different federation scenarios presented before differ in their degree of distribution of functions and degree of standardisation. From the database point of view, these differences translate to system transparencies for data storage and data access. 'One-stop-shopping' is the idea that a user can issue an information service request to a (logically) single system and is freed from possible query decomposition into multiple sub-queries to distributed data sources, issuing those queries, and perhaps finally integrating the results to be presented to the user. For higher level information it may be required to know how this information was derived. The ability of the system to support such questions is the issue of pedigree or data provenance. Clearly, less control over data and metadata model, and an increased degree of distribution make the necessary transparencies more difficult to be achieved.

4 A 'Processing And Access/Distribution' Federation Prototype

In order to prepare for the challenge described above, a number of prototype developments have been initiated, e. g.the Earth Science Information Prototypes (ESIPs) [ESIP] in the USA, or European projects funded by ESA and the European Commission. A European prototype, the Interactive Satellite Image Server (ISIS) project [COD, ISIS], and its successor project RAMSES, aim at defining and validating IL interoperation with emphasis on the distinct definition, distribution, and interoperation of functions such as data ingestion, processing, cataloguing, and user access as well as their implementation and validation on a common system backbone. More information on this prototype can be found in [RAMSES].

4.1 Distributed Functions

Figure **??** depicts the various IL functions identified and illustrates their interconnection through a common bus based on CORBA technology. The User Client interfaces refer to the UL, whereas the Data Ingest interfaces refer to the DL. This architecture has been applied by three IL system developments, each for a selected application: detection of fishing grounds, urban planning, and oil pollution monitoring. All components are defined through a well defined CORBA IDL interface. A high level of reuse of components for standard data ingest, image processing, and catalogue has been achieved between the systems. Although client and workflow functions are application specific, different modules on sub-component level are reusable also, e. g.image display and animation.

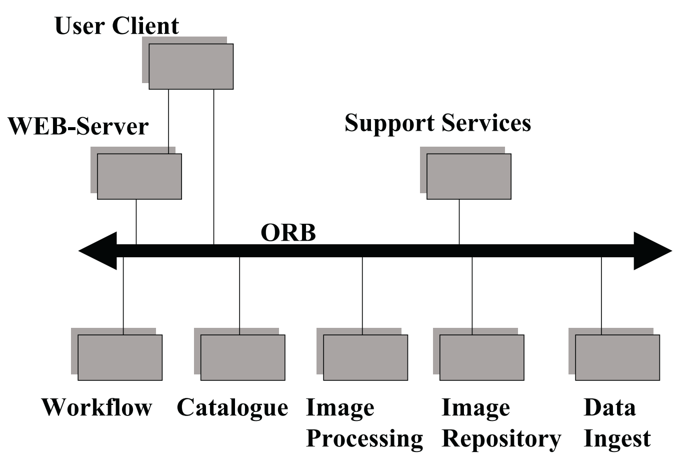

Fig. 7. Information Level System Components

Figure **??** presents the interactions between the different functions for a given application, here the monitoring of oil pollution through the analysis of radar imagery. The UL client's interface to IL functions is through a workflow module which has a priori knowledge of the user, its application, and the associated relevant catalogue entries and processing algorithms.

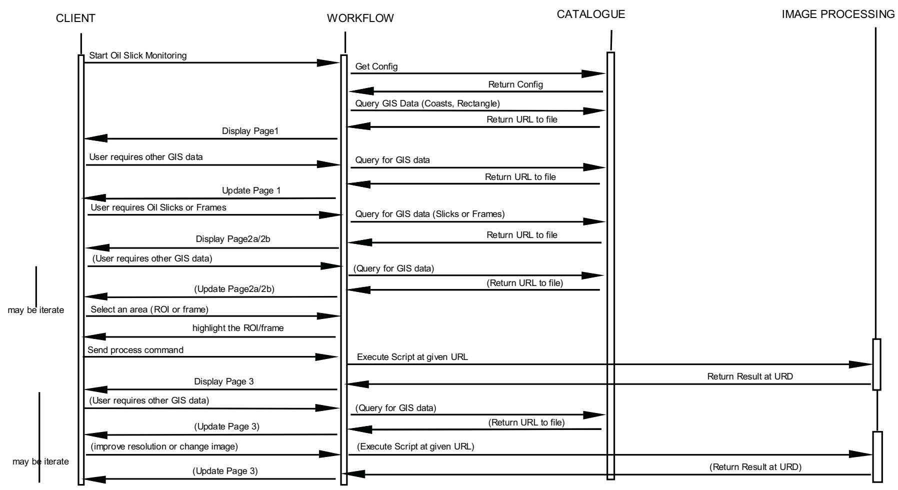

Fig. 8. Example: IL internal Interaction between IL components

The workflow manages the client requests and the catalogue queries, e. g., it automatically retrieves the user relevant catalogue information at the start of the user session. In the example it triggers the processing function once the user has confirmed a pre-selection of a suitable data set. Such data set is identified by the catalogue based on a number of parameters provided by the user. It is displayed as a vector or sub-sampled image on the user screen for selection.

The prototype makes use of the Internet InterORB protocol (IIOP) for client access and has been implemented based on Orbix products. The OpenGIS simple feature specifications are under consideration for interfacing external GIS for the provision of complementary, non-EO data. At present, this data is stored as vector data inside the catalogue component.

4.2 Application of CORBA Services

A number of commercially available CORBA service are suitable to be directly applied in support to earth observation data and information services. An example is the Trader service [Si96] which can be used as a directory function where different catalogues register with a description of the nature of available data sets, e. g., European radar imagery, ERS-2 available at the ESA-ESRIN catalogue. A catalogue query may thus identify and be routed to the adequate catalogue site without running a full query on all sites (see Figure **??**).

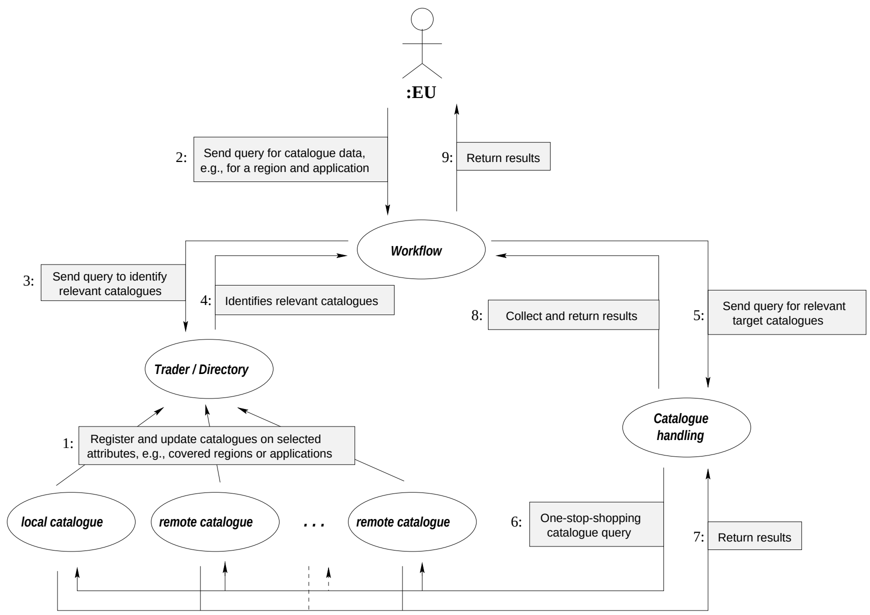

Fig. 9. Directory Service and CORBA Trader Service

Another example is the Event service [Si96] as illustrated in Figure **??**. Well specified events are triggered by the ingestion of new data sets into the system archive. Events specification may include the source of the data, information on its applicability for selected applications, or a first indication on geographical coverage. End users (EU) may register at the subscription service for a sub-set of such events, and in conjunction with a message box will be notified by the system in case such event occurs. The same events may be used for logging and accouting, or may act on workflow functions to automatically trigger processing functions on the newly ingested data.

5 Conclusion And Future Issues

The 'One-size-fits-all' approach of today's state-of-the-art EO information systems, i.e. large-scale DL systems, may be adequate for the handling of standard EO products. But it leads to overly complex solutions in view of the multitude of emerging, very different user communities and application domains. It also risks not to meet the adaptability requirements resulting from the future role of VAC, SI and VARs, which demands a higher degree of independence to distinguish their service offering. An additional system layer, the IL, may

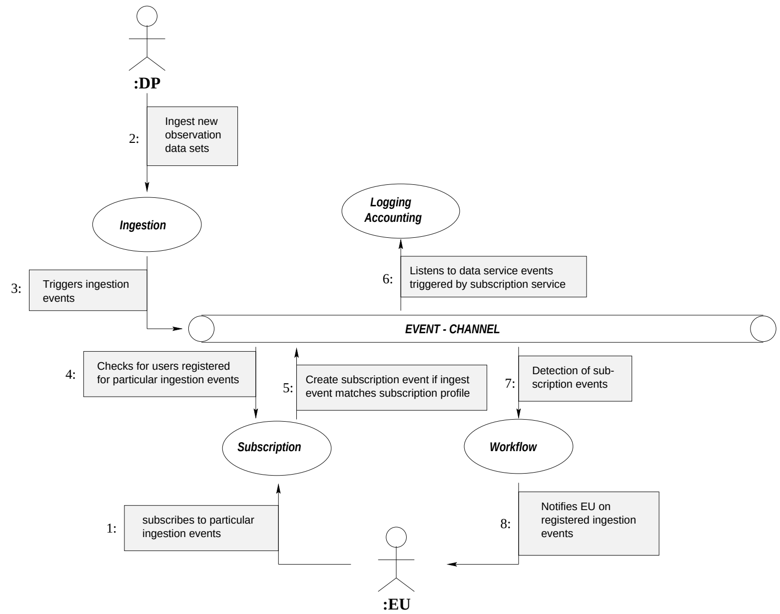

Fig. 10. Data Subscription Service and CORBA Event Service

provide the adaptability needed, balancing the complexity of individual, smaller systems against an adequate level of interoperation among such systems in a federation. Prototype demonstrations indicate that the evolution of the Internet, together with the already cheaply available processing and storage power, and the emergence of open middleware standards, provide the technological basis for federations of IL systems yet to develop.

A number of activities are underway worldwide with space agencies and related organisations to advance along this line. In particular, the authors are involved in the development of a 'Processing and Access/Distribution' prototype [RAMSES] and the ESA author has prepared a modelling study for better mapping CORBA and EO IL systems [MAAT] which initiated recently. This shall help to perform system verification in a real case of an application and user scenario (oil pollution monitoring) and in optimising the EO service component model to make best use of CORBA.

References

[CIP] Triebnig, G., et al.: Catalogue Interoperability Protocol Specification (CIP) - Release B, CEOS Working Group on Information Systems and Services, 1997.

[COD] Moeller, H.L.: European Earth Observation Application System Prototype, CODATA - Conference on Scientific and Technical Data Exchange and Integration, Bethesda, MD, USA, December 1997.

[EOSDIS] NASA Earth Observing Data and Information System EOSDIS, http://spsosun.gsfc.nasa.gov/New_EOSDIS.html.

[ESIP] Creative and Innovative Working Prototype Earth Science Information Partnerships in Support of Earth System Science, NASA CAN-97-MTPE-01, May 1997.

[IBM97] IBM Co.: Data Joiner: Administrator Guide and Application Programming, IBM Co., San Jose, 1997.

[ISIS] Moeller, H.L.: The Interactive Satellite Image Server Project, International Society for Photogrammetry and Remote Sensing - ISPRS Workshop 'From Producer to User", Boulder, Colorado, USA, October 1997.

[La95] Lau, Chr.: Object-Oriented Programming Using SOM and DSOM, Wiley & Sons, 1995.

[MAAT] Object Oriented Environment for Distributed Facilities, Earth Observation Applications Architecture, ESA Statement of Work, June '98.

[MUIS] Landgraf, G.: The ESA Multi-Mission User Information System (MUIS) and their evolution, Earth Observation & Geo-Spatial Web and Internet Workshop '98, Salzburg, http://earthnet.esrin.esa.it.

[OH97] Orfali, R., Harkey, D., Edwards, J.: Instant CORBA, Wiley Computer Publ., 1997.

[OMG98] Object Management Group: The Common Object Request Broker Architecture: Architecture and Specification, Revision 2.2, http://www.omg.org/corba/corbiiop.htm, Upd. February 1998.

[PATHEMA] Fischer, P., Gorman M.: Generation of Mutilayer Thematic Products for End Users - A WWW Interfacing System for Integration earthborne GIS and spaceborne EO Information, Earth Observation & Geo-Spatial Web and Internet Workshop '98, Salzburg, http://earthnet.esrin.esa.it.

[PDS] ESA: Envisat-1, Mission & System Summary, ESA Publication, March 98.

[RAMSES] Regional earth observation Application for Mediterranean Sea Emergency Surveillance, Project Description, July 1998, http://tempest.esrin.esa.it/ramses.

[Se98] Sessions, R.: COM and DCOM, Wiley Computer Publishing, 1998.

[Si96] Siegel, J.: CORBA: Fundamentals and Programming, Jon Wiley & Sons, 1996.

[SM98] Sellentin, J., Mitschang, B.: Data-Intensive Intra- & Internet Applications - Experiences Using Java and CORBA in the World Wide Web, in: Proceedings of the 14th IEEE International Conference on Data Engineering (ICDE), Orlando, Florida, February 1998.