Review of the Heterogeneous Mission Accessibility Project

Bryan Lawrence, Matt Pritchard, Andrew Woolf
March 13, 2007
Document version 1.0
Distribution: Public.

This document has been commissioned by the British National Space Centre with the twin aims of providing a technical review of the HMA project for the benefit of the HMA team, and providing a high-level overview of key technologies to help engender future British involvement. The analysis was based on material from the Acceptance Review in late 2006, and attendance at the Final Presentation workshop in Frascati, 2007-2008.

Abstract Summary

From all perspectives, the technical opportunities for involvement in HMA in the future are good: the underlying technology is being developed in a public manner, a testbed and service validation system is planned, there is considerable scope for expansion (both by adding services and data products to the DAIL layer), and an investment in HMA technology is likely to have payoffs in the wider deployment of geospatial services (including commercial deployment).

Technical Summary

1. The HMA project is being developed with methodologies based on the ISO and OGC specifications. To understand the HMA, data providers and data consumers will need to be familiar with those specifications.
2. Not only are the baseline specifications in the public domain, but many of the HMA architectural specifications are in the public domain in form of OGC documents – so the only barrier to uptake on HMA technology is appropriate funding.
3. The project is on target to deliver new functionality based on the existing SSE toolkit and a Data Access and Interoperability Layer (DAIL).
4. The functionality that will be delivered by the DAIL will be limited by design decisions that have been made for pragmatic reasons in a changing landscape of what should be reliable inter operable web-service technologies.
5. As the underlying technologies change, and as the requirements of the HMA are driven by the wider GMES project, it is inevitable that changes in the DAIL (and associated toolkits) will be required. This is recognized by the establishment of a HMA project Architectural Board (HAB).
6. The membership of HAB may need to be reviewed to ensure it is forward looking and not limited to just the existing HMA partners (it may not be enough to have mechanisms for adding new members as new missions are added). HAB deliberations should be public (although obviously individual mission implementation timescales should remain confidential if desired).
7. There is considerable prospect for expansion of the HMA into other ESA activities.
8. There are some issues associated with identity management technologies which may slow progress with moving from prototypes to implementation. These are compounded by a potential lack of trust by data providers in (1) the ability of the DAIL to protect information about data/service use by individual users, and (2) the protections that their IPR has within the SOA. (The latter being unfounded in our opinion).
9. While the project has made good use of OGC specs for metadata management, service description and control, there has not been any significant data modelling, and that will limit

the use that can be made of OGC web services for data consumption, either within DAIL services, or by DAIL consumers.

10. The current development is based around layers; instruments that provide atmospheric profiles will not be well supported in the initial phases. (This is a consequence of the lack of data modelling and consequential lack of feature-type definition beyond the implicit assumption that the data consists of layers).
11. The permanent testbed to be created as part of the HMA-T project should ease development of HMA compatible services (both those which consume services via the DAIL and those which expose services via the DAIL).
12. The proposed OGC pilot project should expose the HMA technologies for wider constructive critique, and this will be of significant benefit both to GMES and the wider community.

1. Introduction

The Heterogeneous Mission Accessibility project is an ambitious attempt to deploy a service orientated architecture (SOA) to provide a portal entry-point to multiple missions in the context of the European Union's Global Monitoring for Environment and Security (GMES). As much as possible the work is following the methodologies for handling spatial data which are outlined in the series of standards issued by the TC211 committee of the International Standards Organisation (ISO), and using specifications for services defined by the Open Geospatial Consortium (OGC). The project is relatively young, aiming to complete deployment of a prototype in mid-2007, with full implementation at least eighteen months later (following tender specification and completion of contracts to deliver the final product, both of which may introduce further delays).

Currently the project partners are essentially the contractors building the HMA technology which are EADS Astrium, Spacebel, Siemens Austria, Scysis, Datamat and Spot Image and the missions involved: ASI and Alcatel Alenia (performing the technical work on behalf of ASI), CNES, CNES/Spot Image, DLR, EUMETSAT , the Canadian Space Agency and MDA (performing the technical work on behalf of the Canadian Space Agency) , ESA.

While the current scope of the project is to support GMES only, the project is under the scope of the Ground Segment Coordination Body, and should it be successful it would be hard to believe that ESA would not wish the architecture to be of wider use. In particular there are two immediate areas for further exploitation:

- Individual ground segments could be decomposed into services, allowing mission ground segments to incorporate common components. This would potentially allow more cost-effective and robust solutions for ground segments. In particular, this would allow ground segments to exploit more recent (and therefore efficient) technologies since components could be integrated far later in the implementation process than is possible with current design-implementation lifecycles. However, the current funding paradigms (characterised in the main by ESA returning funding to national missions which then look to support their local infrastructure) would work against this.
- Non-GMES missions could be supported, as the architecture itself is agnostic as to what the purpose of the data usage, and could support charging if necessary.

Even within the scope of GMES there is the expectation that future (“third-party”) missions will be added to the backend of the DAIL.

Because there are obvious extensions required to the HMA (to support new missions), and the possibility of further extension (outside GMES and within mission ground segments), the ability to evolve HMA is important. ESA are clearly thinking this way too, as there is an upcoming OGC pilot project based around the HMA intended to broaden participation through the OGC activity which would be subject to a parallel ESA and OGC ITT. ESA is also establishing an HMA architecture board which will be explicitly manage HMA evolution.

The current state of the project is summarized in five key documents and a skeleton software structure which supports the prototypes. The five key documents should be available from the portal at http://hma.eoportal.org:

1. The HMA Architecture Technical Note (currently at rev 1.4)
2. The GML Application Schema for Earth Observation Products (currently at rev 0.1.4)
3. The OGC Catalogue Specification 2.0 extension package for ebrim (currently at rev 0.1.0)
4. The Ordering Services for Earth Observation Products Specification (currently at 1.2.0)
5. The OpenGIS Sensor Planning Service Application Profile for EO Sensors (currently at 0.9.2)

The remainder of this document consists primarily of backup material for some, but not all, of the points made in the summary. It should be understood that because the authors are not members of the project, there is every possibility that some of our interpretations of the (voluminous) project material are incorrect.

2. Basic HMA Architecture

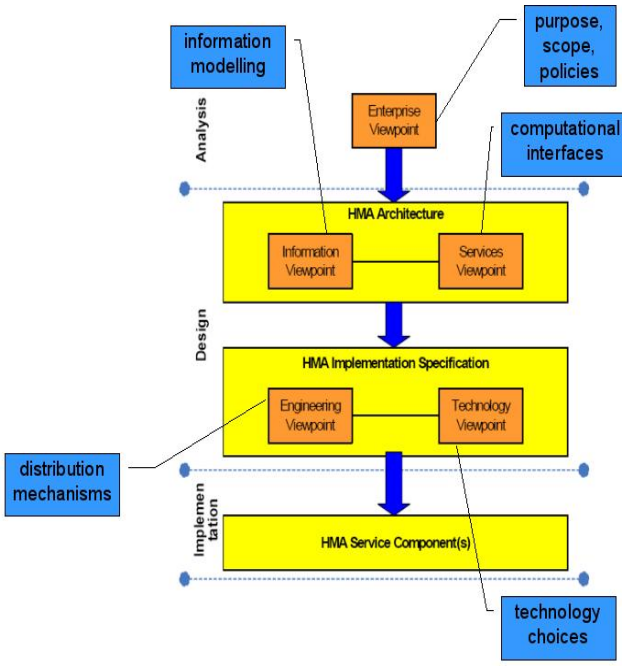

The basic HMA architecture is discussed in the HMA architecture design technical note (which is, or will be, publicly available). The architectural design is based on the Reference Model for Open Distributed Processing, and outlines a Service Orientated Architecture based on SOAP web services.

Information Viewpoint: Standards-base model driven approach (UML to XML).

- Service Metadata
  - low-level (for invocation/orchestration): UDDI model
  - high-level (for human-readable discovery): minimal ISO 19119
- Collection Metadata
  - ISO 19115
- Product Ordering
  - e.g. scene selection, delivery options, etc.
- Programming
  - e.g defining potential satellite tasking (actual tasking will be under individual ground segment control); similar to ordering information, but the information doesn't exist until the data is acquired.
- User Management

- user id, contact & organisation details, allowed services, default delivery and billing information

A key component of the development is the "Application schema for Earth Observation Products", which is published via OGC (OGC document 06-080), to define the product metadata. This exploits GML 3.1.1, and is primarily aimed at supporting cataloging satellite products for the catalog service. Although a WFS can be used to interrogate the metadata, a lack of data modelling (discussed below) limits the functionality of the schema and the use systems such as WFS for data delivery per se.

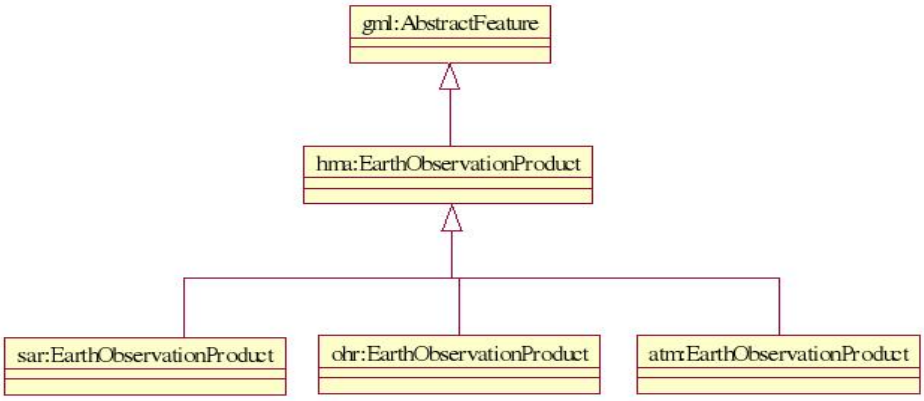

Services under discussion are classified according to the ISO19119 taxonomy.

- Architecture services include: Service Registration Service, Orchestration Service, Monitoring and Control Service, Service Configuration Discovery Service, User Management Service
- Application services include: Collection Discovery Service, Service Discovery Service, Catalogue Service, Programming Service, Order Service, Mission Planning Service, Data Access Request Service, Processing Service, Help & Documentation Desk Service
- Based on OGC Catalogue Service Spec: Collection Discovery Service (for product collections in HMA), Service Discovery Service (for 'human readable' services), Catalogue Service (metadata and browse images – EO App. Profile)
- Order Service: for placing orders for catalogued EO products, compliant with OGC OWS
- Programming Service: for placing requests for future EO products, based on SPS
- Service Registration Service: for deploying new services
- Orchestration Service: for designing and executing (BPEL), workflows
- Monitoring and Control Service: for logging and monitoring transaction traffic across HMA

The engineering viewpoint describes services allocated between Ground Segments and the Data Access and Integration Layer (DAIL). The latter includes: Service Container (runtime), Database, Workflow Engine, User Profile Repository (inc. authentication info), Collection/Service Discovery Servers, Service Description Registry, Catalogue Service, Order Service & Programming Service.

For user identity management, the DAIL handles user identity, retains federated GS policy stores; mission Ground Segments handle policy enforcement. DAIL ↔ GS interaction uses WS-Security.

Key technologies are

- SOA: XML, SOAP, WSDL, BPEL, UDDI, WS-Addressing
- Security & Identity Management: SAML, LDAP, WS-Security
- Services: OGC (GML, W*S, CSW, SWE: SPS), ebRIM

3. Toolkits

Much of the functionality that currently works is built around modifications and extensions to the ESA Service Support Environment. Service providers and consumers familiar with the SSE should find it relatively straightforward to work with the new toolkits.
The SSE toolbox is

- a JSP-based web application run by a service provider
- tool box supports development by
  - enabling code-free development of services compliant with required Interface Control Documents (ICDs)
  - providing configuration tools
  - providing test features
- it supports monitoring and logging of messages.

However, any web service engine capable of exposing services compliant to the ICDs should be usable in the HMA context.

There is a dedicated discovery service client:

The prototype also includes a Business Process Execution Language (BPEL) console which provides a workflow engine for orchestration and monitoring of services, as well as a test client.

Currently the HMA software stack has a heavy dependence on Oracle products (ranging from the Oracle BPEL process manager through the Oracle Application Service web service stack, Oracle LDAP, etc), and this has to be a risk for the HMA project: not all current and future HMA partners will be happy with a one-supplier solution. Other key technologies used in prototyping include various Apache components as well as the Oracle stack.

The existing prototype does show new functionality that will be of significant assistance in meeting GMES aims.

4. The HMA Project Architecture Board (HAB)

Role: The main roles of the HAB will be to provide

- long range planning and coordination between different areas of interoperability within the

GMES Space Component, and
- provide to GSCB oversight of the overall HMA architecture and of the protocols and standards used by the HMA Projects, and of the process leading to the protocol modification.

Specifically, the HAB is chartered as a monitoring, coordination and advisory body of the ESA HMA
project, and responsibilities include:

- HMA architectural oversight,
- Standards selection and modification oversight,
- Implementation and Coordination
- Verification and Harmonisation
- Standardisation

Membership : The Heterogeneous Missions Accessibility Architecture Board (HAB) shall consist of at least 4 (four) full members, composed of the Project Manager of the Heterogeneous Missions Accessibility Project (ESA) and of the HMA Project Managers of the GMES Participating Missions: Cosmo-Skymed (ASI); ENVISAT (ESA); Pleiades (CNES); Radarsat-2 (CSA – represented by MDA)'; Spot (CNES/SPOT IMAGE); Tandem and Terrasar-X (DLR); and Meteorological Missions (EUMETSAT).

The Secretary of the Ground Segment Coordination Body is member ex-officio of the HAB. Ex-officio and liaison members of the HAB may also attend HAB meetings but shall not participate in determination of decisions affecting the implementations either at ESA or on the Ground Segments of the GMES Participating Missions.

New missions contributing to GMES SC Phase-1 are identified by the ESA EOP-G (Ground Segment Department). Each new mission will have the possibility to nominate an ex-officio member to the HAB. Ex-officio members became full members when the relevant agency or organization signs an agreement or a contract with ESA for an HMA ground segment interface implementation.

Ex-officio and liaison members of the HAB have no standing to participate in HAB decisions but are expected to participate in HAB discussions as appropriate to their roles. However, an ex-officio position may be held by a full member, who does not thereby lose his or her standing to participate in HAB decisions.

Comment: While we understand the motivation that the HMA architecture board should provide inertia for change which affects existing ground segments, we are concerned that

1. All HAB deliberations should be public, as decisions made may affect future missions, so it would be desirable for candidates for future missions to at least have input (with or without liaison status, and
2. That further permanent (liaison) members of the HAB be assigned to represent communities that are very likely to interact significantly with the HMA in the future.

5. Software Engineering Constraints

The project is being built to deploy a service orientated architecture (SOA) based on SOAP systems and the WS-* hierarchy of standards. It is not obvious that all partners are fully aware of the immaturity of the relevant standards, and there is in our opinion an unrealistic expectation that commerical off the shelf software (COTS) will be available to implement the software stack (and an even more unrealistic expectation that if such software can be found, that it will be interoperable with all or part of other vendors WS-* SOAP stacks).

5.1 Identity Management and information Security

This part of the project is relatively immature, although it will be crucial to the final acceptance by the data and service providers as well as the funders.

The project has carried out a comprehensive analysis of the information security requirements, and is developing policies based on a series of standards: ISO/IEC17799 which is a code of practice for information security management); ISO27001 (formerly BS7799) which provides a list of criteria which can be used to evaluate information security; and ISO 15408 (also known as the "Common Criteria") which provides a framework for specifying and evaluating the software security requirements, The resulting policies are encoded in the "ISO27001 tailoring", although resulting technical software requirements do not yet appear to be fully and transparently documented.

The project carried out a review of possible identity management strategies and codified these in a small selection of abstract scenarios, for which various pros and cons were discussed. We were not convinced that the full range of access control paradigms were adequately explored (in particular there was no discussion of using proxy certificates to meet the goals of non-repudiation in a browser and service based environment), nor that all the statements[1] about the various options and their relative complexity and reliability in an SOA case would stand up to further independent scrutiny (Indeed we thought that combinations of the abstract models might have brought clarity to the situation, and shown some clearer paths through the options).

Nonetheless, given that the current version of the HMA SOA consists essentially of a DAIL portal orchestrating a closed family of services, the final recommendation to use their "Distributor" model seems appropriate, particularly given the time constraints. However, it will limit the options for building a more complex SOA based on multiple portals and more loosely coupled services (which would seem to be a obvious growth strategy).

The distributor model is depicted in the following figure:

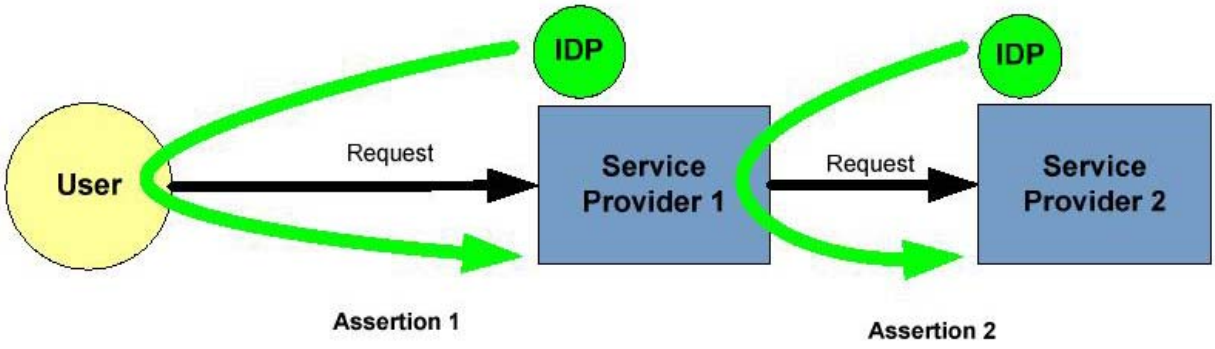

The key concept is that the user authenticates with SP1, and then SP1 invokes services in their own right with SP2 (they do not impersonate the user, but they can provide details in their own request as to who the original user was). This model relies on both the user and SP2 trusting SP1.

The HMA intends that all SOAP transactions will be secured using WS-Security, however it would appear that there is little practical experience amongst the HMA team of using WS-Security, and there is little resource available for experimentation. There is a strong risk that the reliance on one vendor to provide a WS-Security implementation will lead to non-interoperable solution without considerable further work codifying a profile of how WS-Security is to be used. (The main issues here are what parts of a message need to be signed, how and by whom, and what parts need to be encrypted: different vendors do not necessarily make the same decisions about defaults and mutually compatible configuration can be difficult if not impossible to establish. Further, we would assert that the expectation that WS-Policy or WS-PolicyAttachment are mature enough for use in

1 Assertions made in the HMA_DD_SIE_UM_001_1.0c document.

the HMA is not yet tenable). Accordingly, movement from a prototpye to implementation may be difficult in practice.

5.2 Trust

There are two issues associated with trust in a distributed enterprise: do the partners understand each others trust paradigms and can they agree on some level of unilateral or bilateral trust; and can the software architecture implement the agreed trust paradigms? While the former is not a software issue, we discuss it briefly here as it is very relevant to the success of the HMA.

The existing HMA ground segments are steeped in a commercial environment, with existing customers, and a significant investment in serving the types of requests they create. There is some unease (amongst more than one of missions) that it will be possible for commercial competitors in another mission to obtain via the DAIL either information about which customers use their services and/or (to some extent worse) what those customers are doing with their data/services. From a consumer point of view this is not a problem: it should allow competitive evolution of service quality; however, from a mission perspective this is a risk that may preclude full involvement. It will therefore be important that the information security and identity management policies and implementation specifications address these issues directly and then become more widely understood within the HMA team, so that reservations can be addressed as soon as possible.

There also appears to be some confusion amongst participants between the roles that User License Agreements and Copyright hold in protecting their IPR in the context of HMA where they appear to see some additional risks associated with data being used in external services before being provided to end users. It is our opinion that these additional risks are more likely to be managed better in the context of the information management within the DAIL, than they are by customers outside the DAIL scope, and so the DAIL does not provide any additional risk in this area. However, there needs to be clarity in how IPR propagates within the DAIL, and what licenses are provided to DAIL consumers.

5.3 Data Modelling

Thus far the project has been concentrating on the use of metadata for describing and orchestrating services and what we would categorize as “browse” information about data (that is data metadata which doesn't explicitly describe the semantics and syntax of the data itself).

However, it is clear that there is a desire to move onto building services which consume data products, and these will need to be aware of both the underlying syntax of the data (formats, layouts etc) as well as the semantics (phenomenon definitions, grid descriptions, sampling paradigms etc). Thus, the services will need explicit models of the data itself. In principle, such models can be constructed as application schema of GML, and exposed using the WFS[2]. When this is done, clients can interrogate a WFS for semantic information about the payload presented by a WCS, and make meaningful use of binary data thus exposed.

(We say “in principle” because we believe that it is likely that for a full range of appropriate EO data, some extensions to GML would be required, as was found in the development of the Climate Science Modelling Language - CSML[3] - itself an application schema with similar scope.)

Even without formal data modelling it is clear that the underling assumptions of the DAIL thus far is that the data objects are instances of images (or layers). It is not clear what the consequences of considering atmospheric profile instruments would be on the architecture.

---

2 It should be stressed that such application schema would be very different from the application schema currently being proposed for OGC standardization by the the HMA.

3 http://ndg.nerc.ac.uk/csml

5.4 Discovery and Search

The basic concept of the HMA catalogue interface is that both datasets and services are catalogued and discoverable via the same interface. The project initially started work based on the OCG CSW specification, and both built and implementation based on that, and developed a formal profile which was provided to the Open Geospatial Consortium. However, they have recently moved to catalogues based on the ebRIM formalism, and are about to issue a contract to establish a formal extension to the OGC ebRIM CSW specification for EO data.

As currently built, and perhaps based on their CSW heritage there appears to be an expectation that the formal association between datasets and services described in the following UML, is encapsulated in the metadata records themselves (in this case the service metadata). (At least this is what a number of speakers explicitly discussed in the meeting).

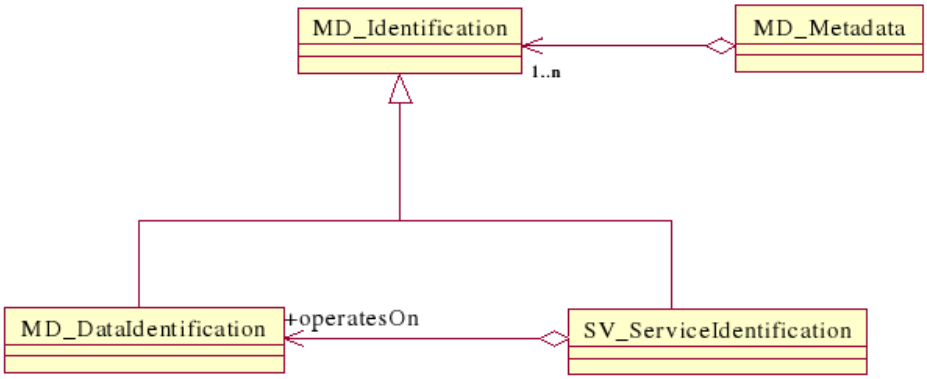

Such an approach precludes late-binding, and implies that the author of service metadata has full knowledge of all possible datasets. While this is probably pragmatic for services which directly manipulate mission data, this will be a fundamental limitation on the development of a complex service infrastructure (and economy) built on the HMA DAIL layer. Ideally the ebRIM formalism should be used to expose the operatesOn association, and the clients should use this rather than the contents of the metadata records themselves to present service capabilities etc. This will of course rely on the data models themselves as well as object instances and services being discoverable and identifiable.

While it is clear that some project participants understand this issue, the project documentation that we reviewed did not make this clear. It will be important that this issue is resolved as early as possible, and it will have implications for how client software interacts with the registry and the metadata contents.

6. HMA Testbed and External Pilot Projects

One of the main areas of concern amongst organizations not part of the existing HMA team is how they will avoid being “second-class” citizens because they cannot bootstrap the necessary expertise to develop services and/or will not be able to validate prospective services for inclusion in or interoperability with, the DAIL. To that end there are two significant initiatives which ESA have engendered: the HMA-T testbed project and an OGC pilot experiment.

6.1 HMA-T

A new project, HMA-T kicked off in early 2007, aimed at providing a persistent testbed so that new service and data providers, along with new service consumers, would be able to test their software against the DAIL. Limited details are available at this time, but the project will run for at least 30 months, and will consist of three phases.

- The first phase is aimed at improving aspects of the DAIL technologies, in particular the ebrim support. This will include both improving underlying specifications to explicitly support ISO metadata for dataset collections (19115) and service descriptions (19119) as well as the EO products; and in developing implementations, ideally COTS and open-source (probably in the geonetwork-opensource tool). Additionally, a conformance test service will be established, probably at ESA, and probably based on the the OGC compliance and interoperabilty test evaluation (CITE) environment (which itself is highly likely to migrate to the opensource Teamengine product). Such an environment would be capable of exercising clients and servers against both OGC specs and HMA specs (although the latter are expected to be incorporated in the OGC stack anyway).
- Further phases will explicitly support new service suppliers, and ESA will invite tenders to an ITT to be announced.

While the HMA-testbed project is currently a thirty month project, we strongly recommend that the testbed and evaluation environment become a persistent feature of the HMA environment: this should assist further uptake by both new producers for, and new consumers of, the DAIL.

6.2 An OGC Pilot Project

ESA and OGC are currently in the planning phase of an OGC Pilot Project, tentatively titled FedEO, which will provide an opportunity for even wider participation and testing of the underlying technologies and specifications. In the context of the HMA-T project ESA will issue two ITTs (one in 2007 and one in 2008), of which the one to be issued by mid 2007 will match the OGC one scheduled before summer 2007. If successful - the outcome of the FedEO Pilot will be presented at the GEO Plenary in November 2007.