Infrastructure Design Study Report

(Released Jan 2007)

Table of Contents

Introduction

This is a preliminary report by the Design Study 3 (DS3) team of the VOTECH project. It covers the work undertaken by the DS3 team from 1 January 2005 to 31 December 2006 and outlines the plans for the remainder of the project. This is a *Preliminary Design Study Report*, covering the period to the end of 2006, which will be followed later by a final version of the *Study Report*. This has the advantage of apprising the community of the significant progress made by the DS3 team to date, and enabling their input into the detailed planning of the remainder of the DS3 workplan, while ensuring that the deliverables from DS3 record all its achievements, and not just those before the end of 2006.

Executive Summary

The aim of DS3 is to undertake preparatory work for the creation of a EuroVO infrastructure. Such an infrastructure would be a combination of a thin layer of new access protocols which can be easily implemented by existing data holders and new interoperability protocols to allow communication between astronomers' analysis and visualization tools. This preparatory work includes studies of existing software as well as the development of infrastructure components which will demonstrate the viability of the EuroVO mission. One of the most important aspects of DS3 is that work is informed by the development of standards in the IVOA: the EuroVO must be based upon IVOA standards to ensure that it can interoperate with the infrastructure implemented in other areas. That said, much of the work on EuroVO is very much in advance of that done by other VObs projects and so the DS3 team is very active in leading the development of new IVOA standards.

In the period to the end of 2006, the DS3 team has made significant progress on these goals:

- access to catalogues, images and spectra have been enabled with implementations of the IVOA standard access protocols to many European data holdings
- work on a security model for the VObs in the IVOA has been paralleled with the development of prototype code and the demonstration of secure data access between three locations (UK, ESO, US)
- prototypes have been built of gateway components to grid hardware and software
- the AstroGrid CEA mode of making applications VObs-aware has been adapted by the IVOA into the UWS standard now under development: it is hoped this will make it possible for server-side applications to be invoked from anywhere
- some effort has gone into developing workflow infrastructure but this relies on the completion of the UWS standard and has a lower priority
- a number of client-side tools have been developed which make it easy for astronomers to interact with VObs server-side components: this has been made simpler with the construction of a AstroRuntime code developed within DS3
- to enable these tools to communicate, a protocol called PLASTIC has been created within the DS6 team and prototypes developed in DS3/4/6 to demonstrate actions being passed from one tool to another: this too is now under discussion as an IVOA standard protocol
- another important component is the provision of virtual storage space (server-side storage available for intermediate results, data sharing etc.): a VOSpace protocol is under development in the IVOA based on work done in the DS3 team

For the future, the DS3 team will continue to explore how the EuroVO infrastructure can be realised. This will include completing the work on security, applications protocols and virtual storage alongside continuing development of the runtime tools and communications protocols. It still remains to be seen whether a standard for workflow construction is viable but investigations will continue in this area.

VOTech Project Background

Project Summary

A Design Study will be undertaken aimed at completing all technical preparatory work necessary for the construction of the European Virtual Observatory (EuroVO). The concept of the Virtual Observatory (VObs) is that all the world's data should feel like it sits on the astronomer's desk top, analysable with a user selected workbench of tools, made available through a standard interface. Internationally this is set to transform and re-structure the way astronomy is done. EuroVO is a specifically European implementation of this idea, and will produce a world leading infrastructure providing a unified virtual data resource and the ability to preform complex data discovery and manipulation tasks across the whole range of astronomy. Access to data and tools will be equally good across Europe, regardless of location. This will require establishing an alliance of data centres, and a VObs facility centre in support of the community, but crucially requires the construction of an infrastructural glue of software components, in the context of rapidly evolving background developments in IT and the grid.

The VO-TECH project aims specifically at feasibility studies and design work aimed at integrating such new technologies into the EuroVO. Key IT advances to build on are in intelligent resource discovery (ontology and the semantic web), data mining, and visualisation capabilities. These will be integrated via global astronomical interoperability standards coupled with the latest distributed grid computing services. Additionally this project covers design and preparatory work to ensure that data from the major European telescopes and facilities (as represented by the Opticon and RadioNet networks) is fully accessible through the EuroVO, and where required, is able to offload mass scale computational process onto the EGEE backbone.

In summary, VO-TECH will lay the technical foundations of the EuroVO, a European infrastructure to revolutionise the scientific process

Project Objectives

The top-level objective of the VO-TECH proposal is to complete all technical preparatory work necessary for the construction of the European Virtual Observatory.

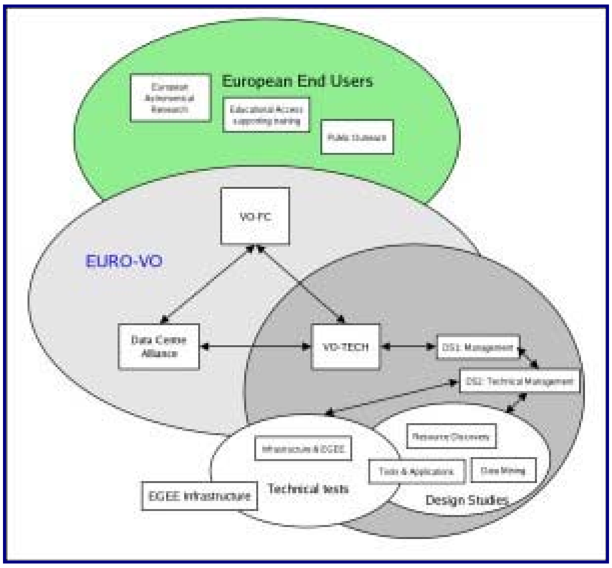

Figure 1: Click to see image full size

Context Figure 1 shows the conceptual structure of VO-TECH, and how it relates to EuroVO as a whole, the various classes of user, and the general astronomical infrastructure.

The VO-TECH preparatory work needs to link closely with the work of the Data Centre Alliance

(DCA) and VO Facility Centre (VOFC), with the final construction of EuroVO following the VO-TECH design study kept in mind. The work also takes place in the context of extensive developments - new algorithms, technologies, and protocols - in academic and commercial IT, and especially of course in generic grid middleware. This work will not be repeated. The job needed is to assess these developments and design astronomy-specific modules based on them. The working links of the project partners with this external world are excellent. Several of the VO-TECH partners have active working relationships with the academic and commercial IT communities, and the project will work with EGEE as an exemplar application area. It is also worth noting that Astronomy in general and the VObs projects in particular have attracted attention as leading edge but pragmatic exemplars of the new e-science approach - for example some of the VO-TECH co-Is have been invited to talk at Bio-Informatics meetings, as well as general Grid meetings.

Objectives of Project:

1. To assess new technologies and study the feasibility of their incorporation in EuroVO
2. To create designs of new infrastructure components based on those new technologies
3. To create designs of science user tools and datamining services
4. To develop trial versions of new infrastructure components, tools, and datamining services and to test them
5. To decide what new interoperability standards are required, and to define those standards with international partners
6. To liaise with the larger EuroVO structure, gaining refreshed versions of science functionality and architecture, and feeding back component test results, designs, and trial components for demonstration suites.
7. To liaise with computer science, IT industry, and related applications projects in order to mesh with larger standards and to save work wherever possible

Scope of DS3

DS3 covers the following areas:

- To design and deliver mature infrastructure components for EuroVO
- To ensure those components comprise a platform upon which the other Design Study can undertake their work
- To assess key new infrastructural technologies and incorporate them as appropriate
- To ensure interoperability and international integration

This Study aims at producing final designs of mature components, as well as assessments, designs and trials of new components that don't fit into the major categories of DS4-6. It is expected that the priority areas will be in:

- authentication/authorisation (and eventually accounting)
- an easy to use API for VO compliant services
- access to astronomical datasets
- work-flow
- distributed storage (VOSpace).

In addition, DS3 has a responsibility for considering interoperability, integration and testing within the context of the overall EuroVO architecture and hence liaising with the VOFC. This will also include full internationalisation of the EuroVO programme, designing customisation tools for deployment across Europe and mix-and-match integration with other projects. DS3 will also have the prime responsibility for liaising with the NA4 work-package of the FP6 programme Enabling Grids for E-Science in Europe (EGEE).

Web presence

VOTech uses its Internet presence for a number of activities:

- Provide a public introduction to the project and its progress (http://eurovotech.org/)
- Plan and monitor progress in the various stages and publish interim findings and reports (http://wiki.eurovotech.org/bin/view/VOTech/)
- Discussion of various issues and plans (http://forum.eurovotech.org/)

These services are hosted and maintained at the IoA Edinburgh.

Science Requirements of DS3

The science requirements of the infrastructure are exactly those of the VObs as a whole. Primarily, it is to make all the world's astronomical data accessible to the astronomer through a small number of simple protocols. Supplementary to this is the need for existing and new tools to be adapted to use these protocols. Secondary aims are that applications should be equally accessible and usable, that resources (data and tools, mainly) should be easy to locate and access, that an astronomer in a resource-poor location should still be able to access and use all these resources. A supplementary benefit to achieving these aims is that new and more powerful tools and techniques will emerge which will empower astronomers in their research.

Studies

Agent Architectures

The emerging VO architecture can leverage the peer-to-peer nature of agent based solutions to make the software smarter, more efficient and more responsive to the user. The purpose of an agent is to know what a user wants, even if they don't know themselves. Establishing architectures and building agent based solutions for the VO will flatten the network topology, enabling links between related datasets to form more quickly, and make the resulting network more robust.

Data Publishing: SkyNodes

The SkyNode has been developed in the frame of the NVO and is now implemented for a few catalogues. The aim of this work is to improve the SkyNode concept in the case of VizieR which is a very large collection of catalogues (>5000).

VizieR is a good test case for VOQL/ADQL, as well as for the scalability of the SkyNode protocol. We have tested the JVO and ESAC "SkyNode toolkit" and we have given SkyNode access to some VizieR catalogues using the ESAC toolkit. These first tests and evaluations will allow us to bring our feedback and suggest some improvements to the SkyNode protocol and to the implementation toolkits. A more exhaustive test (at least JVO, ESAC, NVO toolkits, performance comparison with CGI access (overhead, ...), etc.) is ongoing at CDS.

Infrastructure components

Registries

An IVOA-conformant Registry will allow an astronomer or service to be able to locate, get details of, and make use of, any resource located anywhere in the IVO space, ie in any Virtual Observatory. The IVOA defines the protocols and standards whereby different registry services are able to interoperate and thereby realise this goal. IVOA Registry WG has defined two types of standard:

- *metadata* - a set of XML schemas describing 'Resources', such as Services, Data Sets, Applications, and others.
- *interface specification* - which details how the registry can be searched and how registries must interoperate with other registries.

There are two types of Registries. A 'Publishing' registry will publish its data in accordance with the Interface specification, and a 'Full' Registry, which must be able to hold all 'Resource' records, must comply with the search interface and be able to harvest other registries.

Implementations

AstroGrid Registry is a software component which conforms fully with the IVOA specification and can be setup as either a 'Publishing' or a 'Full' registry. It contains an XML database backend for holding records and allows for extensions to the Resource Metadata. The Registry handles multiple versions of the 'Resource' metadata (0.10 and 1.0 versions).

CDS has set up a registry which publishes all VizieR tables. Currently, each VizieR table corresponds to a Resource (giving a total of over 13000 resources). This registry is compliant with the VOResource 0.10 schema. Each Resource contains the full description of the table, with details on the columns metadata (name, units, UCDs), and has type TabularSkyService. The CDS registry does not harvest from other registries, but can be harvested by other full VObs Registries through the OAI Protocol for Metadata Harvesting.

This registry uses the PERL library developed at Virginia Tech to build on the fly the resource descriptions from the VizieR metadata (extracted from both Sybase RDBMS and ascii ReadMe files for catalogues). The registry will evolve when VOResource 1.0 becomes an IVOA Recommendation, with im particular an improvement of the Resources description, and the use of UCD1+.

Provide access to SIAP, SSAP, ConeSearch, OpenSkyQuery etc.

A number of early studies were undertaken to evaluate how these protocols might best be supported from within the AstroGrid infrastructure. The DSA, CEA and JES components were all considered. This has resulted in recent work on the following:

- DSA: This component is now a family of components rather than a monolith. Thus we are working on DSA/Catalogue, DSA/SIAP, DSA/SSAP etc.
- The AR is now capable of access these and other IVOA compliant services. Existing AstroGrid components will be refactored over time to use the AR. Once this is complete, it will be possible to call all IVOA compliant services from within any AstroGrid component.

Astro Runtime

The Astro Runtime is a programming interface for any client code that wants to access Virtual Observatory services. It's facilities are exposed via a range of technologies - HTTP / XML-RPC / Java-RMI, making it easily accessible from almost all programming and scripting languages. It hides the complexities of the VO - security, configuration, service resolution - and lets developers get on with the interesting stuff.

The runtime simplifies access to

- IVOA standard services : Registries, SIAP, SSAP
- All AstroGrid services : MySpace, Workflow, CEA

- Other popular protocols: NVO Cone-search
- Useful one-off services - CDS Simbad, etc.

It also provides GUI components - simple dialogues that can be reused by client applications to perform common tasks (MySpace micro browser, registry browser, etc.). Other benefits of the runtime include single sign-on, single-configuration, and single cache service responses - making implementation of client-side applications simpler.

CDS has tested access from a client application (Aladin) to AstroGrid components via the ACR library. In that framework, following interactions have been developed:

- access to MySpace: allows one to download within Aladin data coming from its MySpace account
- access to AstroGrid registry : allows you to query the registry, for instance to look for SIAP (Simple Image Access Protocol) or SSAP (Simple Spectrum Access Protocol) services
- launch CEA applications : retrieves available CEA applications from the AstroGrid registry, and allows one to launch them with some user-defined parameters.

Those developments are available in the Aladin prototype version.

VOTechBroker

DS6 needs drove the development of interfaces from the AstroGrid infrastructure to Grid and HPC resources, in the shape of the the VOTechBroker (VOTB). VOTB builds on the GridSAM system developed by the London e-Science Centre, and offers users a way to submit jobs to a wide range of Grid resources using the same job description file, written in Job Submission Description Language (JSDL). The requirements for VOTB arose from a data mining application, which requires a large number (~10,000 or more) of small, parameter sweep jobs to be executed, but the system developed will be applicable to a wide range of situations, and it can provide a bridge between the Virtual Observatory and the Grid.

The VOTechBroker (VOTB) acts as a bridge for submitting parameter sweep computations from the Virtual Observatory to the Grid, and other distributed resources. It provides a number of features, including:

- Flexibility: Most existing projects provide interaction with a given Grid middleware only, for example Globus. In some cases interaction is tied to a particular version of Globus. Our approach is to utilise whatever middleware a site already has installed. For example if Globus or Condor are not available we can still execute the client's algorithm. We do not require system administrators to install additional software.
- Transparency: A standards based job description language abstracts clients from the syntax of underlying grid middleware. Conversion between this abstract language and middleware specific syntax is performed by drivers that plug in to the submission framework.
- Ease of use: Details of parameter sweep job submissions are hidden from the client. Issues such as location of resources, staging of application binaries and data, the monitoring (and restarting) of jobs, the staging of results are taken care of by the broker.
- Distribution: The broker does not assume a shared file system between resources. It is the responsibility of the information service to detect which endpoints have shared file systems and optimise file staging appropriately.
- Integration: The VOTechBroker provides a SOAP interface so that a diverse range of clients can submit jobs and monitor progress. Current clients include, CEA wrapped 'thin-clients' for integration with the AstroGrid Workflow, Java (Web start) applications and Web forms.

VO Data Storage

There is a strong desire within the Virtual Observatory to have a distributed storage mechanism, that will allow users to easily refer to a piece of data without necessarily needing to concern themselves with the physical location. The task of defining this system has been given to the IVOA Grid and Web Services WG This conceptual storage space is called VOSpace. It is envisaged that the VOSpace will manage references to the physical location of data. A primary design goal of VOSpace should be to allow easy integration of existing systems such as SRB or NGAS which have similar goals, but differing levels of abstraction and implementation

There are two primary use cases

1. A certain dataset is needed for analysis - for efficiency reasons it would be best if the data could reside "close" in network terms to the compute resource on which the analysis will be performed. In the ideal case this will involve the data being located on the analysis computer's hard disks. With potentially many analyses being performed on the same data at many locations, it is more efficient if the data are gradually replicated onto these locations
2. A user would like easily to publish and share his own data.

IVOA VOSpace standardization effort was split into two parts:

1. VOStore - this was intended to address the functionality of the "physical storage" domain of the logical architecture.
2. VOSpace - the metadata storage and user interface compontent. A VOSpace server was intended to control several VOStore servers.

It was decided that the definition of VOStore should be attempted first, leaving the more complex aspects of the VOSpace for later. The draft definition document is v0.18. The VOTech partners decided to attempt an interoperability demonstration at the 2005 ADASS conference by trying to implement this interface as a facade to their existing systems - in particular

- Astrogrid created a VOStore interface for MySpace
- ESO created a VOStore interface for NGAS

In addition Caltech (an IVOA partner) created a new implementation of the VOStore interface based on a WebDav server.

The demonstration was a success in that data could successfully be moved between the servers in a secure fashion, but it became clear that the separation of the system into the two VOStore and VOSpace levels did not map that well onto the existing implementations as they already contained elements of functionality that belonged to the VOSpace level, and this functionality was effectively being hidden by putting the services behind a VOStore server.

This has lead the VOTech members to propose a new architecture to the IVOA for VOSpace which essentially drops the VOStore layer and has each of the VOSpace servers acting as peers, with local data storage and all connected to a common metadata store. This architecture is now the basis of the latest IVOA standardization efforts.

Standards development and prototyping

Security

Components are being developed following the emerging IVOA architecture. These stand as reference implementations to validate the IVOA architecture and allow EuroVO to implement SSO in advance of IVOA’s final standard.

The AstroGrid components are as follows.

- Community services, including MyProxy, user management and the attribute service.
- Facilities in the Astro Client Runtime (ACR) to obtain and store credentials for use of a client application.
- Java classes to obtain credentials from a MyProxy service.
- “Security façade” of Java classes to:
    - sign outgoing SOAP requests;
- check signatures on incoming SOAP requests;
- handle authorization decisions based on attributes;
- participate in credential delegation

UWS

Application Support in the VO

Early efforts in the VO have been directed towards the definition of specifications for a "Data Grid", where the principle focus has been on how to provide a consistent description of data items and how to provide uniform interfaces to access that data. It is natural that effort should now move on to defining how to process that data using applications.

In the context of the VO an application is any process that can manipulate data in some fashion. This can be an extremely wide ranging definition that encompasses anything from calibration of raw data through to model fitting, image manipulation, calculation of observing parameters etc.

In the VO the "application" might run on the user's desktop or be a remote service and part of the purpose of the VO is to minimise the differences between these two modes of operation as far as the effort needed to invoke the application is concerned.

This work has the following general aims

- To create a uniform interface and model for an application and its parameters. This has twin benefits;
    1. It allows VO infrastructure writers a single model of an application that that have to code for.
- Application writers know what they have to implement to be compatible with a VO Infrastructure.
- To provide a higher level description than WSDL 1.1 can offer.
- Restrict the almost limitless possibilities allowed by WSDL into a manageable subset. There are too many ways that web service interfaces can be expressed in WSDL 1.1 - however, even when interoperability [WS-I] guidelines are followed, there are still many choices for expressing parameter values - we need a common way to do this in IVOA workflows.
- Provide specific semantics for some astronomical quantities.
- Provide extra information not allowed in WSDL - e.g. default values, descriptions for use in a GUI etc.
- Insulate upper software layers from the differences in WSDL styles, and possible grid implementations.
- To provide extensions with the VO Resource schema (See the IVOA WG) that can describe a general application
- To provide asynchronous operation of an application - This is essential as the call tree that invokes the application cannot be expected to be active for extremely long lasting operations - e.g. a user from a web browser invokes a data-mining operation that takes days
- Provide callback for notification of finishing.

- Provide polling mechanisms for status.
- Provide Job identification.
- To allow for the data flow to not necessarily have to follow the call tree. In a typical application execution the results are returned to the invoking process - In a VO scenario, it can be useful if the application can be instructed to pass the results on to a different location, especially as a staging point for the result of asynchronous calls. (cf the section on VOSpace)

The task can be split into two essential parts

1. The mechanics of calling the service asynchronously - which are covered by the SOAP WSDL at the basic level with WS-Resource and WS-Notification standards layered on top as usage patterns.
2. The description of the application or service. This is both a description of what a particular service does and how to call it. Ideally this description language can be standardized also and published in a registry to make it easier to connect the execution of different applications in a workflow and to make it easier to standardized GUIs.

AstroGrid has a prototype that covers both these aspects, though in a proprietary fashion, called the Common Execution Architecture (CEA). The intention is to evolve CEA application descriptions to become the IVOA Universal Worker Service (UWS) standard for application description.

IVOA developments

The Universal Worker Service (UWS) pattern is a work in progress by IVOA to standardize asynchronous executions. It is informed by AstroGrid CEA. Each detailed application of the UWS pattern defines a contract for an IVOA-standard service. UWS for Parameterized Applications (UWS-PA) is a proposed contact for a service that can replace AstroGrid's CEC. When UWS-PA is standardized by IVOA, AstroGrid expects to migrate its server components (the application servers and DSA) from CEC to UWS-PA. The migration will be gradual and, as always, the AR will hide the details from application code.

Although the UWS work in general is part of the IVOA Grid and Web Services WG, the VOTech members have been particularly active in this task. An XML schema for the description of an application and its parameters has been developed and should be included in the Version 1.0 release of the IVOA Registry schema.

PLASTIC

PLASTIC is a protocol for communication between client-side astronomy applications. It is very simple for application developers to adopt and is easily extended. Through PLASTIC applications can do tasks such as instruct each other to load VOTables, highlight a subset of rows or load an image of a particular area of sky. Although such operations are quite simple, they enable powerful collaborations between tools. The philosophy is that the astronomer should have a suite of interoperating tools at his disposal, each of which does one thing well and which can be composed according to need.

The DS6 effort provided the impetus for developing the PLASTIC interface when it was realised that client-side applications which were exploring and visualising data needed to be able to interact seamlessly. The specification of the PLASTIC interface was mainly developed as a DS6 task and its initial prototyping as a joint DS3/6 effort. While future development of the interface will be under DS3, it is fully expected that DS6 tool development will continue to drive the requirements for PLASTIC.

Other developments

Workflow

A workflow aims to accomplish a complex piece of work, for example an astronomical investigation. The workflow builder is designed to enable astronomers to design and develop these complex workflows in a simple and intuitive manner, whilst hiding much of the intricacies of the underlying XML document structure. The use of familiar drag and drop features, tooltips, examples and continuous error checking mean that a novice user can quickly produce simple workflows, and rapidly progress to ever more complex pieces of work.

At this stage there seems no comprehensive reason to pursue standards activities relating to workflow: it is ultimately reliant on the development of the UWS standard first and has, in the past, proven impossible to gain agreement across development communities.

Activities

The AstroGrid project has spent some time on developing workflow tools. Its Workbench product has an integrated workflow builder which incorporates complex logic and a scripting language along with the ability to add any data or application resource from the registry to a workflow. The workflow is enacted by a server-side Job Execution Service (JES) which keeps the user informed of progress via a VOLookout tool. The project is about to develop the Taverna tool as a replacement workflow builder and job execution service.

CDS has developed AÏDA (Astronomical Image processing Architecture) in the context of the MDA (Masses de Données en Astronomy) project (French ministry action). AÏDA can execute workflows and a client-side GUI workflow builder tool (JLOW, Java Libraries fOr Workflow) has been developed in the Cycle 2 of VOTech. In addition, a collaboration with ESSI to enable Grid access (EGEE, Grid5000) through a middleware component (MOTEUR) is ongoing at CDS.

Interim Conclusions

1. The EuroVO infrastructure **is** viable and scalable;
2. Early technologies can be deployed **now**;
3. More work is needed to meet the requirements of other environments;
4. More effort needed on: VObs-enabling tools and making data available;
5. More work is needed on linking to Grid and HPC resources.