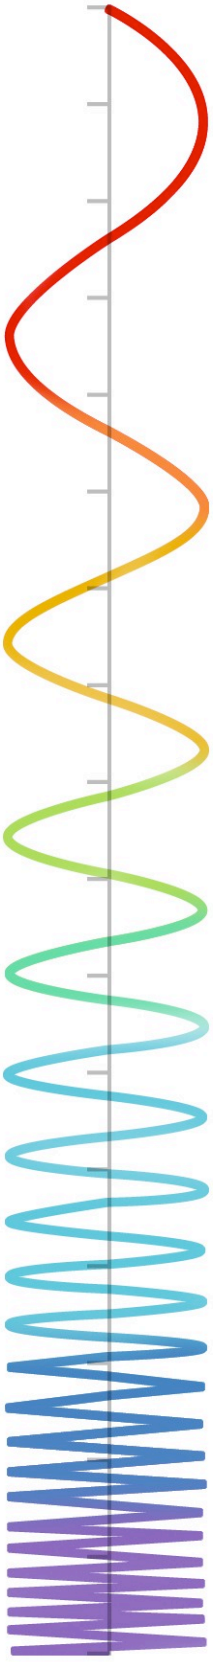

Quarterly Report
April – June 2008

Building the Framework for the National Virtual Observatory

NSF Cooperative Agreement
AST0122449

NATIONAL VIRTUAL OBSERVATORY

IVOA

INTERNATIONAL VIRTUAL OBSERVATORY ALLIANCE

Building the Framework for the National Virtual Observatory

NSF Cooperative Agreement AST0122449

Annual Report

Period covered by this report: 1 April – 30 June 2008

Submitted by: Dr. Robert Hanisch (STScI), Project Manager

Executive Summary

In coordination with our IVOA partners, the transition to V1.0 of the Registry Interfaces was implemented in April. The transition was carefully planned so that all groups deployed the new services in synchrony. Harvesting and replication are now working routinely, and the new registry is being used in support of the NVO data discovery portal.

All components of the data discovery portal have been tested separately and as working together, with results from one component being forwarded to others as appropriate. The new registry interfaces caused some of the components to break, owing to changes in the way VO services with multiple access methods are represented in the registry. These problems were largely resolved in this quarter, and a scientific testing team has been assembled.

Efforts within the team are increasing in the area of operations, with routine testing of registered VO services and feedback to service providers whose cone search or SIAP functions are failing or not fully compliant with VO standards. Service providers are typically very responsive in correcting problems.

A comprehensive review of all NVO software products was conducted this quarter and discussed at the NVO team meeting (June 30/July 1) in Estes Park, CO. All libraries, tools, and applications have been assessed for the quality of documentation, the extent of testing, and their utility for the VO in the long term.

As a result of extensive discussions at the IVOA Interoperability Meeting in Trieste, Italy in May, progress was made on the Table Access Protocol (TAP) and Astronomy Data Query Language (ADQL). TAP is a much more general table and catalog search interface than the simple cone search, and will provide much needed capabilities for querying and comparing complex databases. TAP will accept ADQL queries and will also have a parameter-based interface for simplicity of implementation and greater similarity to existing data access protocols. The ADQL specification underwent revision following the RFC period and is now pending IVOA Exec review. Two partial drafts of the TAP protocol will be merged, and the NVO team will be building several prototypes in order to validate the interface specification prior to the October IVOA Interop meeting.

Preparations for the September NVO Summer School, the fourth to be offered, are nearing completion. Applications were received from over 60 people and 50 were accepted

(the applicant pool was excellent this year). The overall program has been prepared and work was nearing completion on updating the software packages and tutorials.

A second Newsletter was distributed to the NVO community mailing list in June. Use of NVO tools in the astronomy community continues steady or increasing, with over 27 million hits on NVO websites and 168,000 user sessions thus far this year.

Activities by WBS

1 Management

1.1 General (planning, reporting, communications, team meetings, etc.)

Regular weekly telecons of the Technical Working Group (TWG) continue. Similarly, the Executive Committee also continues to meet weekly by telecon.

The summer NVO team meeting was held in Estes Park, Colorado on June 30/July 1. As the NVO project approaches its conclusion, we reviewed our team's software legacy and discussed how to best document and preserve all that has been developed.

1.2 Science

The NVO Executive Committee, and NVO Project Scientist D. De Young (NOAO), have been closely following the work on the data discovery portal (led by T. McGlynn, HEASARC). Progress on the portal was reviewed at the summer team meeting. R. Hanisch has assembled a team of scientist-testers to evaluate the portal once the development team has concluded its development and testing.

1.3 Technical (including standards, configuration management)

There was intense discussion of the Table Access Protocol (TAP) and the Astronomical Data Query Language (ADQL) standard at the VOA Interop workshop in Trieste in May. NVO management had several significant concerns about both protocols, e.g., that TAP did not include a parameter-based query mechanism and that ADQL region specifications were defined with no reference to the Space-Time Coordinates (STC) definitions that are already an IVOA standard. Our concerns were acknowledged and addressed.

1.4 Financial

Total project expenditures since inception now stand at $12,581,965 compared to a budget of $14,151,966, leaving a balance of $1,570,001. We have $318,221 in invoices received and pending payment, leaving a net of $1,251,780. Since January 2008 invoices from SAO were being sent to the wrong department at JHU; these have now been found and will further reduce the net once they have been reviewed and approved for payment. Average quarterly expenditures have been around $450k for the past year. Thus, we expect to end the fiscal year with a positive balance of as much as ~$750k. This includes the budget for the upcoming Summer School ($150k), however, leaving ~$600k. In July we submitted a no-cost extension request to NSF which, if approved, would allow the team to continue to function through at least the first quarter of FY 2009.

1.5 International coordination/collaboration

Senior members of the NVO project participated in IVOA Executive telecons and discussions. D. De Young (NOAO) continues as chair of the IVOA Executive Committee, and R. Williams (Caltech) chaired the IVOA Technical Coordination Group. At the May Interop meeting Williams concluded his term as TCG chair and was succeeded by C. Arviset (ESAC). T. McGlynn was appointed the new chair of the Applications Working Group. M. Graham and R. Plante continue as chairs of the Grid and Web Services Working Group and Registry Working Group, respectively. R. Hanisch continues as chair of the Data Curation and Preservation Interest Group, as a member of the Standing Commit-

tee on Standards and Process, and as chair of the IAU Commission 5 Working Group on Virtual Observatories, Data Centers, and Networks.

We learned with dismay of the decision in the UK to terminate the AstroGrid development efforts several months earlier than planned. This is despite generally positive statements about AstroGrid and a general recognition that UK astronomy has a clear need for VO infrastructure and support. Our AstroGrid collaborators assure us that they will continue their commitments to IVOA standards efforts as they regroup and pursue new sources of funding.

2 Science Requirements

2.1 Usage scenarios for all areas of astronomy research, including theoretical simulations

A use case based on comparing radio and x-ray properties of rich clusters of galaxies from the Abell catalog, which we have been using for some time to test the interoperability of components in the data discovery portal, is now being fully documented to serve as one of several "nature trails" demonstrating the use of the portal for scientists. The user discovers the location of the Abell catalog, selects a subset of the richest clusters, checks the inventory service to see which have radio and x-ray data, queries the relevant catalogs, and integrates the catalog information and imaging data together.

2.2 Requirements analysis

No activities to report this quarter.

2.3 Demonstration definition and review

Pre-release versions of components of the data discovery portal were demonstrated at the spring and summer team meetings and at the IVOA Interop meeting in Trieste in May.

3 Operations, System Integration, and Testing

3.1 Quality assurance and software engineering design

M. Nieto-Santisteban (JHU) led a major review of the status of NVO software systems. Over 60 software efforts were evaluated based upon information supplied by the responsible parties. A variety of software engineering and documentation practices were evaluated. Nieto-Santisteban reported at the summer team meeting and a number of recommendations for enhancing and standardizing practice within the NVO team were proposed.

The NVO Summer School used the new SVN repository to collect software to be used in the summer school. The organization used was similar to that in previous summer schools which used the CVS repository.

3.2 Facility operations

A final set of standard Web templates was developed and released by S. Emery Bunn (CACR) and were adopted by many of the NVO team sites including particularly the portal sites.

The NVO Twiki was moved to http://nvo-twiki.stsci.edu/twiki/bin/view/Main/WebHome. The Twiki was substantially upgraded using a new version of the underlying software and a new host machine. During a transition period both the new and old systems were maintained simultaneously.

All major NVO facilities continued standard operations during the quarter with no major interruptions of services were reported though a number of facilities experienced scheduled and unscheduled downtimes for periods of several hours.

New operational Web pages included a site released by M. Preciado (HEASARC) allowing NVO team members to note scheduled and unscheduled downtimes and including when the system was anticipated to recover.

S. Emery Bunn (CACR) installed VIM and the Portal Table Wizard into us-vo.org website. They had been hosted on site-specific machines earlier.

The SVN configuration management system remains fully operational at CalTech. This system is now being actively used by many development projects as their active repository and is central to the summer school software integraton.

The ticketing and tracking system (Trac) is now being extensively used in the Portal development and to a lesser extent in other NVO development projects. At the summer team meeting it was agreed to create Trac pages for those ongoing NVO projects that do not have them already in place. The Zoho management system developed by R. Hanisch (STScI) continues to be used to describe the progress of NVO tasks.

M. Preciado (HEASARC) has continued to operate and augment a suite of tests using the NAGIOS software system which checks every site providing VO services. Typically tests are run every hour. For sites which provide many services, only a representative subset of the services are checked to avoid undue use of the services. Currently about 100 services are tested including both NVO-hosted sites and sites maintained by other institutions around the world. The status of all tested sites is available at http://heasarc.gsfc.nasa.gov/cgi-bin/vo/monitor/monitor.pl.

Preciado reviews the status of all sites periodically (~daily) and sends notifications to sites where requests are failing. Critical sites are informed immediately while less critical and non-NVO sites are normally queried if they have been down for several hours. This continual monitoring of the health of the VO along with followup has kept the number of non-working VO sites relatively low.

A more comprehensive facility, which will test not merely whether facilities are up but also whether they pass validation, is under development and should be deployed before the end of the year. This involves the collaboration of NCSA, STScI, and HEASARC personnel.

Records of all correspondence sent to sites and any responses received are maintained and are available for review. The procedure for informing sites of problems is itself being reviewed and standardized. Approximately 2-3 issues are brought up each week and almost all are quickly resolved.

In this quarter, there has been significant progress in the VOEvent area, i.e., reporting of astronomical transients (A. Drake, M. Graham. R. Williams, CACR). New event streams have been added, including alerts from the MOA microlensing experiment (R. Bond, Massey University, New Zealand). The VOEvent activity has been feeding Google Sky with its up-to-the-minute event feeds, including gamma-ray bursts (Swift, Milagro, Integral), the two microlensing experiments (MOA and OGLE), the Catalina and Palomar-Quest optical synoptic surveys, and others. At the international standard level, the next

version (2.0) of the VOEvent specification is being developed, with support for Time Series, Orbital Elements, and Digital Signatures.

3.3 User support

All user queries were reviewed by S. Emery Bunn (CACR) and directed to appropriate NVO team members. Responses for all items were obtained.

M. Preciado (HEASARC) continues to develop Web pages designed to make it easier for end-user developers in the NVO. Many more examples of VOTable and FITS pages have been developed and an index page with clear links to the various test files is included. Links to external test files are also included. Facilities for easily downloading all or some subset of the test files have been provided.

S. Emery Bunn (Caltech) wrote and distributed the second NVO Newsletter subsequent to approval by the editorial board.

4 Registries

In this quarter, work on registries have been focused mainly in three areas: completing the transition to the latest publishing standards, advancing registry related standards through the IVOA process, and providing support to the project portal development.

4.1 Resource metadata

Two important standard development efforts have driven additional registry metadata modeling. The first is the Table Access Protocol (TAP, discussed in Section 6). Because we have in the IVOA several standards for accessing data and metadata in tabular form (including Simple Cone Search, OpenSkyNode, Simple Image Access, etc.), we had developed a common registry model for describing tables and realized this model as part of the VODataService XML schema. TAP represents a more general approach to accessing tables, and so the Data Access Layer (DAL) and the Registry working groups are now updating VODataService to accommodate the TAP data model. In the NVO, D. Tody (NRAO) and R. Plante (NCSA) have been the main contributors to this effort, and a proposed update was presented at the May IVOA Interoperability Meeting. Among the proposed additions is the ability to describe how a set of queryable tables relate to each other so that a user can readily understand how to form queries that join data from multiple tables.

Another important driver is the VO Standard Interfaces standard. One of the things this standard enables is the ability to pull metadata about service—in particular, information about its interface and behavior (referred to as its *capabilities*) and information about the tables the service interacts with—directly from the service itself. Not only will this be useful to service clients, it will also help make the process of registering services simpler.

There is a strong drive to finish these two specifications this summer; consequently, we have set our roadmap to complete this work by the end of July.

Plante continues to liaison with the separately funded Virtual Astronomy Multimedia Project (VAMP) that is developing a standard for tagging graphics images intended for outreach purposes with metadata, referred to as Astronomical Visualization Metadata (AVM). A presentation from the authors is planned for the IVOA Interoperability Meeting in the fall.

4.2 Resource metadata schema

Our current focus in this area is completing the standardization process for standard extensions to the VOResource XML schema, the core of the registry metadata. In addition to the work on standardizing the VODataService extension (described above), we are also working to complete the standards documents for the DAL service-specific extensions. Plante, collaborating with IVOA partners, is working on an omnibus document for the extensions that describe the currently support DAL services: namely Cone Search, Simple Image Access, and Simple Spectral Access. Plante and Tody are also assisting on the development of the TAP extension, and Graham has been working on an extension for VOEvent related services.

4.3 Publishing and harvesting protocols

The IVOA Registry Working Group set April 21 as the date for bringing on-line the majority of registries in the IVOA that support the latest specification for registry harvesting (Registry Interfaces v1.0). That date was chosen to allow several weeks of "shake-down" prior to the May IVOA Interoperability Meeting. Within a week of that date, all of the target registries were available (marked by their registration in the Registry of Registries, http://rofr.ivoa.net/), including the four publishing registries within the VO, and global exchange of resource records in the latest VOResource format commenced. A few interoperability issues arose in that shake-down period, but these were quickly addressed. As a result, the new NVO registry at STScI can now search nearly the entire global VO. The missing resources are those published by the ESAC registry in Spain, which has not yet completed its upgrade.

In this quarter, updating our registry publishing interface, the web service that data providers use to publish their resources into a registry, has been a major area of focus. Following the NVO registry team's design meeting last quarter, T. Dower (STScI) has been working on an improved version of the interface that supports the latest VOResource schemas, and she demonstrated the basic functionality at the summer NVO Team Meeting. This new design takes a wizard-like approach to registering a resource and takes advantage of the best features of all our previous implementations. One important design feature is that it can be configured and attached to any IVOA compliant registry. It also is designed to make it easy to support new kinds of resources in the future; this is enabled by the use of the XML XForms technology, pioneered by M. Graham via the Carnivore registry.

While the new publishing interface is in development our old publishing interface needed to be adapted to the new VOResource format. Plante completed a minimal upgrade to the current NCSA registry interface—just enough to support the existing user interface, and he deployed two versions: one for the NCSA registry and a second to support the STScI registry. Both will be replaced once the new publishing interface is complete.

4.4 Search protocols

With the upgraded harvesting processes in place, our focus has shifted to improving the search capabilities of the NVO registry at STScI, driven by the needs of the NVO portal. G. Greene (STScI) has been leading this effort with assistance from Dower and Plante. One innovation we incorporated into the registry last quarter was a custom interface that returns search results in VOTable format (in addition to the standard interface that returns VOResource). This was done to provide greater consistency with the other portal appli-

cations, most noticeable in our use of the NVO table browser that supports sophisticated sorting and filtering of tables all within a web browser. Also, VOTable is the format used to pass information between the different portal applications (which run on different sites). In May, Greene shared this work with our IVOA colleagues at the May Interoperability Meeting, and we found broad interest in standardizing the use of VOTable as a return format. Greene went on to propose a way to evolve the standard search interface in way that makes fuller use of the ADQL query language and is more consistent with the emerging TAP standard.

The spring upgrade of the IVOA harvesting network had two major impacts on our search interfaces, which required changes to the NVO portal. One impact relates to the ConeSearch services provided by CDS (France) that give access to thousands of catalogs available via their Vizier collection. While the resource records from the CDS Registry that describes these services are fully compliant, the underlying services they pointed to are not compliant with the Cone Search specification; in particular, the service (illegally) combines query results from multiple tables that make up a logical Vizier catalog into one query response. This causes certain clients to fail, including the search applications provided via the NVO portal. While we are working with the CDS to correct the problem on their end, Plante in the meantime provided a fix to the NVO registry at STScI to allow our portal to operate properly. The portal passes user registry queries into a custom search interface. As the results are returned to the user, the CDS records that describe the Cone Search services are transformed to split them into service interfaces, each returning a single table in compliance with the standard.

The second impact resulted in a change in the way resource metadata are now bundled in the registry, particularly as part of the results from a search. Prior to the upgrade, each "resource" listed in the search results tended to include only one service interface. The new VOResource schema required by the v1.0 Registry Interfaces standard allows a "resource" to list all the related service interfaces together. Thus, prior to the upgrade, resources provided an *interface-based* view of the metadata; afterward, they (intentionally) had more of a *data-collection-based* view. CDS Vizier took advantage of this revised view so as to collapse their approximately 20,000 records by a factor of about 3. While the data collection-based view is more useful to users who interactively search the registry directly, the rest of the portal was developed around the interface-based view. We were able to address this mismatch of views by inserting an XSL transformation of registry search results before they are transferred to other portal applications.

At the international level, more focus is being placed on the standard search mechanisms. In particular, Plante has been looking at registry compliance with the standard search specifications. He discovered that an error in the specification document that resulted in some non-interoperable keyword search implementations. He has since worked on a plan for migrating registries to a common, corrected standard in a manner that is backward-compatible to registry clients that already exist, and he updated the reference client library (ivoaregistry in the NVO repository) to verify this.

4.5 Replication, synchronization, maintenance, revision control, and curation

Plante continues to work with publishers to assist them with updating their records in the wake of the registry upgrade. Also this quarter, the registry team continues to migrate

more of its code into our shared code repository (http://trac.us-vo.org/nvo). This has been invaluable for our distributed team to share our work and coordinate efforts.

5 Data Models

5.1 High-level (image, spectrum, time series, event lists, visibilities, catalogs, simulations, data quality)

J. Cant (SAO) is updating the Java library that implements the Spectrum model. Changes to the external JAXB library that we use to map the schema to Java classes have required a significant rewrite of the library internals.

Progress in the area of mapping NED's SED data model into the VO SED data model, and providing access via a DAL, is waiting for revisions to the Spectrum data model (follow up to the meeting at IPAC in April 2008 which hammered out a preferred method to serialize "segments") and the updated Java package.

In the mean time, at the recent NVO Team Meeting there was an agreement that SAO/NED would work bilaterally to translate a sample SED from NED's current VO-Table format into a new XML structure that represents the SED essentially as an extension to the VO Spectrum data model.

J. McDowell is working on a translator from the current NED format.

5.2 Low-level (measurement, quantity, uncertainty, relationship)

No activity during this quarter.

5.3 Descriptors and ontologies (UCDs)

No activity during this quarter.

5.4 Space-Time and regions

A. Rots and McDowell discussed the STC/VOTABLE interaction with F. Ochsenbein and agreed on his method of representing STC information, but pointed out some potential problems.

Rots has communicated his strong reservations regarding the ADQL PR to the VOQL WG. This concerns in particular the handling of Regions.

Rots and D. Gunasegaran have begun work on a Java library to provide an interface to STC. We have been experimenting with JAXB, XML Beans and other methods, all of which have limitations in representing the full features of XML used by STC.

5.5 Standard schema

No activity during this quarter.

6 Data Access Layer

Work continues on the new Table Access Protocol (TAP), with draft specifications being completed during the quarter and discussed at the IVOA Interop workshop in Trieste in May. A design analysis for the second generation image access interface (SIAV2) was completed and discussed in Trieste, identifying the key areas where further work is needed to complete the initial specification. A first session was held on the planned footprint service to initiate work on this project within the IVOA. Work began on a conceptual design document for the application framework system interfaces. An effort has

begun to coordinate development of the application framework with related work within the AURA and AUI observatories as well as with OPTICON and Euro-VO in Europe.

6.1 Data access services (catalog, image, spectrum, time series, visibilities, ...)

Catalog access. The effort to define a new table access protocol (TAP) is being spearheaded through efforts in the US and in the UK and Europe. A preliminary draft specification for TAP edited by D. Tody, based mainly upon the tiger team meeting held at JHU in November 2007 was distributed in early May. A second specification edited by G. Rixon (AstroGrid) was also distributed. Both specifications were presented and discussed at the IVOA Interop in Trieste. It was agreed to work toward a single common TAP specification meeting the requirements of both groups, ideally with a first working draft available by the fall of 2008, although the details of how this effort will be coordinated at the IVOA level are not yet clear.

Despite the slow progress in achieving a joint IVOA TAP standard it is a high priority to have a functional TAP interface within NVO this year. The current plan is to proceed with prototyping of some initial TAP services within NVO while continuing to work with our IVOA partners towards a joint standard. NRAO, IPAC, STScI, and HEASARC plan to implement at least partial TAP prototypes this summer, with the goal of having something functioning by the fall.

The scope of these initial TAP prototypes is not yet certain but it is expected to include the following capabilities:

- Basic filter-type queries of a single table (most astronomical catalogs are single tables).
- Basic table metadata queries using the same query mechanism provided for data tables.
- Optimized, simple to use support for cone search and multi-position queries. Multi-position queries provide the first stage of a cross match, permit scaling up to large queries, and provide a more convenient interface for cases where a user wants to query against a list of objects.

Experimental support for ADQL-based queries, and integration of Grid capabilities for user table data management based upon VOSpace and asynchronous job execution based upon the UWS pattern, are also planned. These capabilities are needed to enable arbitrarily large queries.

Even a preliminary TAP interface will provide a uniform interface to table data from the major US data centers, which will allow us to begin implementing table-based client applications while the full IVOA TAP interface is being finalized.

Spatial Footprints. A session was held at the IVOA Interop in Trieste to kick off an effort to define an IVOA standard footprint service protocol. This will be used to access and represent the spatial footprint of complex datasets or entire data collections in various levels of detail, as well as provide a capability to perform region-based queries and arithmetic against spatial footprints. G. Greene, T. Budavari, and F. Bonnarel (CDS) presented materials in this session and led the discussion.

Spectral and time series data. Little was done on data access standards for spectral data in this quarter, although work continues to complete and verify the initial SSA services and integrate them with the new registry and client applications. J. McDowell is leading an effort define a data model for SED data. Once this gets to the point where we can interchange SED datasets a data access interface will be required. This should be similar to the current SSA interface.

Image data. Meetings to define the scope, concepts of operation, and interface for the second generation image access protocol (SIAV2), which includes support for spectral and time cubes, were held earlier this year and reported upon in earlier quarters. An IVOA Note written by D. Tody summarizing the conclusions of these meetings was written and distributed in early May. The results of this analysis were presented and discussed in a DAL session at the IVOA Interop in Trieste. F. Bonnarel presented a demo of an SIAV2 mock-up in the same meeting, using Aladin to demonstrate the query.

Much of the SIAV2 interface, being based upon the generic dataset model, is expected to be common with SSA. The main areas needing further work include the following:

- Extension of the FITS WCS and STC models to describe data with spatial, spectral, temporal, and polarization measurement axes. Currently FITS WCS provides limited support for time, and STC provides limited support for the spectral axis and polarization.
- How to specify the image world coordinate system in the query response, using the extended image cube WCS model described above.
- How to specify the parameters of the desired image world coordinate system in a SIA query.

The image WCS plays a central role in how we plan to access both 2-D and cube image data, as access is based mainly upon specifying the geometry and WCS of the output image to be generated by the service. The output image may be a sub-cube, a 2-D plane, a reprojection, or a 3-D slice, and may be dimensionally reduced along a particular axis, and may filter data along an axis.

SIAV2, along with TAP, will also be one of the first of the second generation DAL interfaces to provide Grid capabilities. These include use of VOSpace to store and transport data, support for asynchronous jobs to generate and stage data, and support for user authentication. As with TAP, support will be provided for multi-position queries to allow the service to scale up to potentially very large queries for large-scale analysis.

Complex data. The generic dataset query and generic dataset metadata are the basis for all the "typed" DAL queries (SIA, SSA, TAP, etc.), and define what is common to all these queries. A closely related issue is how to model and query complex data, which is a logical association of primary datasets used to model higher level data objects, such as a survey field consisting of several related primary data products all resulting from a single observation or processing run.

The issue of complex data and the generic dataset is being developed along with SIAV2, and was summarized in the IVOA Note also discussing SIAV2, prepared by D. Tody in early May. This was discussed in the IVOA Interop along with SIAV2. One preliminary conclusion from this analysis and discussion is that it may be possible to limit use of

ADQL within the DAL interfaces to the generic dataset query (except of course for TAP). The generic dataset query is closer to a conventional database query than the typed DAL interfaces as it does not have to deal with specification of virtual data.

6.2 Data representation (VOTable, etc.)

An issue we are currently looking at with SIAV2 is how to represent the image world coordinate system (WCS), both in the metadata returned by a service and as input parameters in the query. The greatest challenge here is actually the WCS data model, but a representation issue is how we deal with vectors since the WCS is multi-dimensional. When dealing with multi-dimensional data it is convenient to represent vectors as vectors so that the representation is the same regardless of the dimension of the data (WCS in this case).

Since the SIA query response is a table that can describe many objects, each with their own WCS, it is most convenient to map the fields of the WCS to columns of the table. This is done using UTYPE to map the table fields back to the WCS data model. However we would also like to have a clean mapping between the WCS defined by SIAV2, and the existing FITS WCS and STC data models. This is no problem for FITS WCS, but is an issue for STC, which still does not permit (formally at least) UTYPE references to the fields of the STC data model. What is needed here is such an "STC/U" mapping, similar to the STC/S and STC/X representations that are already defined. An objection is that a UTYPE-based mapping is not sufficiently general, but it is already the case that STC/S can only represent a simple subset of what STC defines.

6.3 Framework (mediators, components)

VOClient has been updated by M. Fitzpatrick to support the new registry. Other near term work planned for VOClient, and the DALClient Java library upon which it is based, is integration of support for SSA and the TAP prototype when it becomes available. R. Plante has been looking into VOClient integration into Python in connection with the NVO summer school.

M. Fitzpatrick has demonstrated an interesting approach to providing client access to VO data in a desktop environment. This uses the open source FUSE technology for user-space mounting of an external data source to a host operating system file system mount point. This can be used, for example, to make a VOTable or a VOSpace appear as a directory tree within the host operating system, allowing client software that does not know anything about these mechanisms to access data contained therein.

D. Tody has begun work on the DALServer service framework to add support for the simple cone search, TAP, and SIAV1 service protocols. The current version only supports SSA, for publishing spectra. These enhancements make it much easier for user sites to produce fully compliant VO protocols to publish their data.

D. Tody, M. Fitzpatrick, and B. Stobie attended the AURA science software workshop held in Hilo in June, representing NVO (as well as NOAO in the case of Fitzpatrick and Stobie). R. Hanisch and others also participated via videocon from STScI. While the primary goal of the AURA workshop was to determine what to do in the future to process science data from the AURA observatories, an important aspect of this project is coordination of AURA O/IR science software development with NVO and with related software development in other branches of astronomy such as radio. It makes sense to work

toward a broader integration of science software development within the US, both in a multi-wavelength sense, and in terms of integrating observatory data processing with VO.

As part of the ongoing collaboration between NVO and OPTICON in Europe, work has begun on a document presenting a conceptual design for the applications framework system interfaces (D. Tody and P. Grosbol will edit this document on behalf of the US and Europe, and D. Tody is the lead author). This document will go one step beyond the requirements and architecture documents already produced by this collaboration. The system interfaces document, and the first year of development planned for 2009, are well aligned in terms of both goals and schedule with science software development by both AURA and AUI.

6.4 Data provider/consumer implementations and end-to-end testing

Development of verification and test facilities for the operational NVO data services is well along at this point. The major groups involved in this are HEASARC (operational testing of the distributed NVO data system), NCSA (data service and registry verification), and NRAO (data service development and verification). T. McGlynn and M. Preciado are developing facilities for operational testing of data and other services. R. Plante supports service verification facilities for cone search and SIA services. D. Tody and J. Crossley are developing a load tester service that is expected to become available this summer.

Many groups currently have work ongoing to develop SSA services for accessing 1-D spectra. Completion of these services requires end-to-end testing involving client applications, the registry, and the service itself. This has been underway for some months but is ready for final end-to-end integration testing now that the new registry has come online.

The NED group at IPAC (J. Mazarella, R. Ebert, O. Pevunova, J. Jacobson) has begun work with the DALServer service framework (D. Tody) to develop SSA services for NED spectra. This will be extended later to support SEDs as well.

7 Query Language

7.1 Low-level: Astronomical Data Query Language

The Astronomical Data Query Language standard document became an IVOA Proposed Recommendation at the end of April 2008, and the RFC period for it ended five weeks later (in June). Issues were raised regarding STC metadata and function definitions during the RFC period, which included the IVOA Interop meeting in Trieste. The VO Query Language Working Group amended the document in response to comments received during the RFC period and at the Interop meeting.

7.2 Mid-level: VOQL and OpenSkyQuery/OpenSkyNode

JHU continues work on a cross-matching engine that will enable queries between very large datasets. The new hardware infrastructure for Open SkyQuery – Graywulf – has been installed and is being tested. Part of the GrayWulf cluster is being used as a sandbox for a partitioned Pan-STARRS database design. This uses an advanced Cluster Schema Manager (CSM) and a Microsoft scheduler and workflow system called Trident to implement a parallel data loader and fault-tolerant archive. Both of these technologies

will be incorporated into Open SkyQuery to enable it to handle arbitrary cross-matches of large datasets.

7.3 High-level: Complex queries
Nothing to report this quarter.

8 Web and Grid Services

8.1 Web Services (SOAP, WSDL, etc.)
VOSpace is the IVOA interface to distributed storage. Work this quarter has continued on the VOSpace 1.1 specification that extends the existing 1.0 recommendation to support containers, links between individual VOSpace instances, third party APIs and a "find" mechanism. Outstanding issues were resolved in discussion at the IVOA meeting in Trieste and are now propagating into reference implementations and the accompanying IVOA documentation–Working Draft, WSDL and schema. There has been some coordination of international effort in implementing VOSpace 1.1 based on the iRODS technology from SDSC, including a pure PHP version at UCSD. The JHU VOspace 1.1 implementation also continues to be developed.

There is interest from the LSST project in utilizing VOSpace.

8.2 Grid Services (OGSA)
No activities to report this quarter.

8.3 Computational resource management
An interface to manage access control has been developed and was presented at the IVOA meeting in Trieste. Although this arose out of work related to the Universal Publishing Interface in the NVO Registry WBS, it deals with an issue that affects several IVOA interest groups. An assessment is now required of the merits of this approach as opposed to using an infrastructure based around authorization assertion services.

8.4 Virtual data
No activities to report this quarter.

8.5 Application and service integration with Grid
NCSA has been working with portals that make use of NVO login services to upgrade them to use the production versions. The NOAO Archive portal and the Dark Energy Survey project portal have been converted so far.

9 Applications

9.1 Data location services
A demonstration of the NVO discovery portal was made at the IVOA meeting in Trieste in May 2008. All services are now available for use and operational.

A major issue in the portal development that has slowed the release of these services to the public is the upgrade of all portal elements to full support for the V1.0 registry interfaces. It had been anticipated that it would be possible for some services to use the older registry elements, but in practice major data providers, notably Vizier, have registered services in incompatible fashions in the new registry. This resulted in services that appeared to be available in some portal elements not working in others.

At STScI, T. Dowler and G. Greene have continued to upgrade the registry and responded to concerns of the other portal elements to provide simple VOTable format results from the registry. More interfaces to the new registry have been provided, and full support for keyword queries (as used in the portal) is available.

At IRSA, J. Good, has built two of the main components of the NVO Data Discovery Portal. To aid users in importing their own tables (usually of sources) for comparison to VO holdings, IRSA provides a Table Ingest service. This service automatically detects a wide variety of file types and formatting and translates the user data into standards such as VOTable (but also comma separated values, etc.). As needed, it also identifies the table columns best used for source location and adds RA, Dec columns (decimal degrees J2000) to be used by other VO services.

The VOInventory service is one of the cornerstones of the Portal effort and provides several important functions. Given a user source table (or region on the sky), the VOInventory will very quickly identify all registered datasets with overlapping data. This can be returned at various levels of detail (e.g. how many user sources are matched in each registered dataset; which sources in a specific dataset match which user sources; etc.) This can be used directly to obtain directly to return subsets of the real data or in conjunction with other Portal services (and VO-Registered Cone Search and SIA services) as part of the broader Portal capabilities.

At the HEASARC T. McGlynn and J. Wendt have fully implemented the table browser and also continue to upgrade the DataScope element. The table browser has been very useful in testing the availability of services.

Users of the Virtual Observatory have noted the difficulty of taking text-based tables for import into VO tools. The NVO Table Importer (CACR, J.McQuorquondale) is a pure JavaScript application that has been developed in the last few months. It takes a text representation of a table, asks straightforward questions of the user, and creates a VOTable, understanding sexagesimal coordinates, assigning datatypes to columns, allowing input of column names and UCDs. The Table Importer scales to very large tables, and allows direct copy/paste of source tables from PDF documents such as those from astro-ph or on-line journals.

M. Fitzpatrick (NOAO) has worked to provide VOClient support for the new V1.0 registries. An initial attempt to supported Vizier services but a broader implementation that supports all V1.0 registered services was to be released earlier in the next quarter (July 2008).

9.2 Cross-correlation services
No activities to report this quarter.

9.3 Image combination, registration
The Montage on-demand mosaic service has been in operation for the bulk of this fiscal year at http://hachi.ipac.caltech.edu:8080/montage/ and has proven highly successful. This service provides photometric, calibrated, science-quality mosaics of 2MASS, SDSS, and DSS data and is based on the Montage software suite (developed to support VO activities under a NASA grant) and the Request Object Management Environment (ROME) request management environment funded directly by NVO. To date, over 2500 mosaics

of various sizes have been constructed.

A new release of the *SkyView* stand-alone toolkit was made at the HEASARC. *SkyView* can now add and mosaic images weighting images either using constant exposures, or exposure files. *SkyView* can access image data from any source supporting the SIA protocol. This effort was not funded through the NVO, but involves members of the NVO team.

9.4 Visualization tools and services
R. Williams (CACR) continues to develop VIM both as a standalone tool and as part of the NVO portal. VIM is a tool for researching multiple sky positions simultaneously. Each source becomes a row in a table, with catalog, image cutouts, and spectral information. This information is drawn from all the published surveys of the astronomical literature.

VIM has advanced in this quarter from prototype to beta, with a release (us-vo.org/vim), substantial documentation, as well as many bug fixes for many of the backend services. VIM also has a new AJAX interface to show progress of longer-running data-assembly operations, as well as the new version of the Yahoo JavaScript kit, so now the tables have columns that can be moved, sorted, and resized.

9.5 Theory
R. Wagner (SDSC) attended the Trieste Interop meeting in May and was very involved in discussions of how to represent and publish theoretical simulations to the VO. The Simple Numerical Access Protocol is being divided into two pieces: Simulation Database (SimDB) and Simulation Data Access Protocol (SimDAP). SimDB is an online service offering query capabilities to a database containing metadata describing results of simulations and their post-processing as well as about the codes used in these algorithms. Currently the simulations are still supposed to be those that produce a representation of 3+1D space. SimDB has similar functionality to the VO Resource Registry, and further developments will be coordinated with the Registry WG.

9.6 Statistical analysis
No activities to report this quarter.

9.7 Data mining, outlier identification
No activities to report this quarter.

9.8 Interfaces to/from legacy software systems
The VOClient toolkit has been released and can be used from any scripting environment to efficient access VO data.

10 Community Engagement

10.1 Documentation
The second quarterly NVO Newsletter was distributed in June. Topics covered included "NVO Inside", the WorldWide Telescope, the Visual Integration and Mining (VIM) tool, and news about the Heterogeneous Telescope Network (HTN).

10.2 Web site
No activities to report this quarter.

10.3 Technical training initiatives

This quarter has seen much progress in the planning for the NVO Summer School 2008 in September. Advertising, e-mail, and posters were created and distributed, a web-based application form prepared, and 60 applications were received for 40 places at the School. The acceptance committee selected whom to accept, and those students have been informed. The quality of applications was very high, and given the lower costs of running the Summer School in Santa Fe, we decided to take on 50 participants.

The lecture schedule has been completed, with eleven faculty speaking over eight days on topics such as:

- Using the VO to Study Clusters of Galaxies
- The Power of VOTables
- Introduction to XML
- A Universal View of Spectra
- VOEvent: Rapid Reporting for Transients
- Using IDL & IRAF with VO
- Finding Services in the VO Registry
- Google Mashups and Worldwide Telescope
- Databases and SQL
- Astronomy with Crossmateched Catalogs
- VO for Mobile Devices
- Web Services with Apache and CherryPy
- Join the VO: Publishing and Serving Data

10.4 Advocacy

No activities to report this quarter.

11 Education and Public Outreach

11.1 Strategic partnerships

C. Christian, K. Borne, and J. Raddick (JHU) continued to present to EPO groups methods for enabling education with large data archives. They participated in the LSST-EPO working group meeting at the LSST team meeting. The purpose was to clarify what the data needs are for LSST EPO activities and how those cascade to requirements and specifications for the LSST Data Management Facility. The outcome is that the LSST Data Management team believes a lot of the value added materials (keywording, metatags, etc.) will come from the VAO or from the LSST-EPO group itself. The LSST-related EPO activity is being formulated for input to NSF within the proposal for the construction grant for the LSST.

11.2 Formal education

Three proposals were submitted to NSF or the Department of Education by external groups that contained a plan to use data from the VAO and/or contribute to the VAO. Unfortunately, none were funded, but reviewer comments for resubmissions will be evaluated.

11.3 Informal education

Nothing to report this quarter.

11.4 Outreach and press activities

Nothing to report this quarter.

11.5 Technical development

Collaboration on the AVM meta-tag project continue so that at least the EPO community implements these standards. C. Christian has been advocating strongly that science archives consider a mechanism for feeding metatags to services that search for science data.

Activities by Organization

Caltech–Astronomy Department and Center for Advanced Computational Research (CACR)
Activities at CACR this quarter included:

- Further development and release of the VIM application for the NVO Portal.
- Development and release of the Table Wizard for the NVO Portal.
- VOEvent has new surveys on line, a new interface to Google Sky, and a new VOEvent specification.
- Definition of VOEvent registry extensions (VOEventStream and VOEventServer).
- Progress on specification of VOSpace 1.1 and release of reference implementation.
- Experiments with alternate authentication methods, including openID and OAuth.
- New IVOA specification to manage access control.
- Worked with software inventory team on data input mechanism for NVO team, and evaluate results.
- Progress on NVO Summer School.
- NVO newsletter created and disseminated, IVOA newsletter under development.
- Carnivore registry passed IVOA compliance tests, and added to Registry of Registries.

Caltech–Infrared Processing and Analysis Center (IPAC)
R. Ebert attended the NVO Team Meeting held 30 June and 1 July 2008 in Colorado. He reported progress on the SSAP interface into the NED spectrum database using the DAL Server toolkit, and participated in the various discussions. O. Pevunova maintained the 19 NED services during recent changes to the NVO Registry infrastructure, and she maintained the automated generation of daily access logs. Pevunova updated the metadata in NED's image archive to make 29,000 new FITS images with world coordinate system information available through the SIAP interface. She also provided support for people utilizing the VOTable services. Pevunova participated in the weekly Technical Working Group telecons, and J. Mazzarella participated in the weekly Executive Committee telecons.

J. Good continues to chair the NVO Technical Working Group.

The Montage on-request image mosaic is available at: http://hachi.ipac.caltech.edu:8080/montage/.

High Energy Astrophysics Science Archive Research Center (HEASARC)
HEASARC efforts focused in the software development and operations areas. In software T. McGlynn coordinated the efforts of the portal team and he and J. Wendt continued to develop and augment the table viewer and DataScope tools. An initial release of the portal suite was made Trieste IVOA meeting May 2008. An interim Portal home page has been developed and is currently hosted at the HEASARC but should be moved to the NVO Web site.

M. Preciado has continued to augment software that monitors the health of all VO hosts. New sites have been added. He regularly reviews the status of services and reports to

users when services are unavailable. Approximately 2-3 contacts are made with external sites each week. Preciado has begun the development of a more complete validation service that will use validator software developed at NCSA to test all queryable VO services. This will be done in collaboration with other VO sites.

Preciado has released a service to allow users to notify the NVO of service issues.

T. McGlynn is a member of the NVO Executive committee. Wendt, Preciado and McGlynn both attended the regular NVO technical working group telecons and support the general activities of the NVO including the NVO team meetings. McGlynn attended the summer NVO Team meeting. T. McGlynn was selected as the head of the IVOA Applications Working group at the Trieste Meeting. His term will begin in July 2008.

J. Wendt has left the Goddard and a replacement is currently being sought.

Johns Hopkins University

M. Nieto-Santisteban continued to support and help users who reported problems with SkyNodes and OpenSkyQuery, continued her work on the management of very large astronomical databases and underlying infrastructure such as cluster management, and loading workflows. Nieto-Santisteban led the NVO Software Inventory Assessment effort. Nieto-Santisteban, in collaboration with S. Emery-Bunn (Caltech), M. Graham (Caltech), T. Donaldson (STScI), and A. Rots (SAO), collected information about NVO software and services. Applications were evaluated and a summary report was presented by Nieto-Santisteban at the NVO meeting held in Estes Park, CO on June 30-July 1. Nieto-Santisteban also attended the NASA AISRP workshop held in College Park, Maryland on 6 May 2008 and presented the talk "Large-Scale Cross-Matching with Open SkyQuery."

A. Thakar continues to maintain the NVO Weblogs and monitor the harvesting at JHU on a daily basis, and contributes log statistics and charts to quarterly reports.

Thakar is still interviewing candidates for a full-time NVO position at JHU to work on the OpenSkyQuery re-engineering. In the meantime, S. Carliles will take over the development and support until a new person joins the team. The GrayWulf hardware infrastructure, which will eventually host a redesigned cross-match facility, is in place and being tested. Part of the GrayWulf cluster will host a testbed for loading a simulated Pan-STARRS dataset.

Thakar continues to provide triage and support for technical problems with OpenSkyQuery and other VO services at JHU, delegating to developers of the services as necessary. He also attended the NVO Team Meeting (Jun 30-Jul 1) in Estes Park, CP.

A. Szalay, along with J. Vandenberg and A. Wonders (both JHU), are building the GrayWulf (in honor of Jim Gray) data cluster, which will host a Petabyte of astronomical data at JHU in SQL Server databases – one of the largest NVO holdings. Most of the cluster nodes have been configured, installed, and tested. The cluster will eventually host copies of the largest NVO datasets for cross-matching purposes.

J. Raddick continues to maintain the virtualobservatory.org public outreach website. He attended the summer AAS meeting in St. Louis, MO, where he gave a talk entitled "PocketVO: A Simple Tool for Viewing Images through the Virtual Observatory." The talk detailed the development work he did with R. Santos (INPE, Brazil) on PocketVO, a

tool to display side-by-side images drawn from multiple wavelengths through VO tools. Raddick also worked with D. Powell, a graduate student in science writing at JHU, to draw up preliminary plans for using PocketVO in education.

National Optical Astronomy Observatories (NOAO)
As part of the transition from development into an operational NVO facility, NOAO DPP Operations has been moving towards mirrored services to help better guarantee availability of NVO services, especially for critical services like the NVO Single-Signon (SSO) Service, which the NOAO VO Portal in particular uses to grant PIs access to their proprietary data. In June 2008 the NOAO NVO Portal in Tucson and La Serena was deployed into production using sso.us-vo.org as the sole interface for NVO authentication. To this end, I. Barg, working with C. Miller, N. Saavedra and M. Fleming (NOAO) and R. Plante and B. Baker (NCSA) have built VMware Linux servers to host an SSO mirror (nvoauth1.tuc.noao.edu). This is now a production-level service that is mirrored from NCSA and visible to the world at sso.us-vo.org.

Similar work has been done in collaboration with R. Williams (CACR), Warner, and R. Seaman (NOAO) for the VOEVent service at voevent.noao.edu, which is hosted on the same physical server as the SSO mirror using virtualization provided by VMWare. VMWare allows for the consolidation of several NOAO NVO services under one production-quality system, thus reducing server costs and maintenance overhead.

Warner has also continued working on the VOEvent services for the NASA-funded VO-GCN project. A release of the latest version has been delayed, but should take place within the next couple of months. Seaman has also served on the organizing committee for the HTN IV workshop to be held in Santa Barbara, CA July 30 through August 6, 2008. Seaman, Warner, and C. Smith will attend this workshop in the interests of integrating VOEvent standards into HTN operations. Seaman assumed the chairmanship of the IVOA VOEvent working group at the IVOA InterOp meeting in May. Top priority is finalizing v2.0 of the VOEvent specification.

R. Massad, Miller, A. Egana, and E. Fuentes have built a prototype testing framework which enables automated regression testing of web-based applications in the VO. This testing package has server-side and client-side components and allows for both “white box” (i.e., where the functional flow of the application is known and utilized) and “black box” testing (i.e., where no knowledge of the application's internal structure is assumed). On the server side for white box testing, the service and portal providers will create their own functional tests using a simple Domain Specific Language (DSL). These server-side functional tests do not use the web GUI, but instead make a series of direct calls to the portal/service via something like PHP “actions”, Ruby “controllers”, or simple HTTP posts. The service provider writes their functional regression tests in the DSL and runs the test framework that generates the reports. The client-side component integrates this standard framework with the “Selenium” (http://selenium.openqa.org/) test-tool for web browsers. This client-side integration allows for real-world actions like: “user clicks button”, “user uploads file”, etc. This client-side functionality will allow for easy user-based regression tests of VO web-based applications.

M. Fitzpatrick continued development of the VO-CLI command-line tools package, adding support for the new Registry interface as well as an option to call the IPAC Inventory

service. This work was done in support of the continued development of the NVO Portal project. The IRAF NVO package was updated to include these new features. Fitzpatrick began work to update the VOClient interface, the associated VO-CLI tasks, and the IRAF NVO package to support the new Registry interface.

Fitzpatrick also continued development of the new SAMP (Simple Applications Messaging Protocol) standard, authoring part of the document presented at the IVOA Interop meeting in Trieste and now released as a Working Draft. Several implementations of SAMP are now available and the standard is being refined to resolve outstanding issues with the hope of promoting this to a Proposed Recommendation at the Fall Interop meeting. M. Taylor (AstroGrid) and T. Bosch (CDS) are co-authors.

In other VO application work, Fitzpatrick updated the WCSFixer web application (http://iraf-nvo.noao.edu/wcsfixer) to provide better catalog support for infrared image data. He also implemented a proof-of-concept application called "VotFS" that uses the open-source FUSE (File System in User-space) kernel extension to make a VOTable appear to the operating system as a directory tree. Because the OS presents the data and metadata in the VOTable as a plain-text file, data can be accessed by any application without the need to explicitly parse the XML file or link in special-purpose libraries. An end-user application of this technology is under consideration.

B. Stobie joined the NOAO staff as Program Manager of the Data Products Program (DPP), and will take over day-to-day management of the NVO and other VO-related activities at NOAO. She has many years of experience in astronomical software development and management, and has worked with many of the NVO software development staff in the past.

The NOAO NVO team attended various VO and VO-related meetings. D. DeYoung, M. Fitzpatrick, and R. Seaman attended the IVOA Interop meeting in Trieste, Italy, May 19-23. M. Fitzpatrick, C. Smith, B. Stobie, and F. Valdes attended the AURA Scientific Software Workshop in Hilo, HI, June 16-18, while R. Seaman attended this meeting remotely via videocon. D. DeYoung, M. Fitzpatrick, and B. Stobie (with C. Smith via telecon) attended the NVO Team Meeting in Estes Park, CO, June 29-30.

National Radio Astronomy Observatory (NRAO)

NRAO participation in the NVO over the past quarter has emphasized ongoing IVOA standards development, particular for the new table access protocol and the second generation image access protocol, and organization of a collaboration to develop a new applications framework for processing and analysis of astronomical data, helping to integrate the next generation of observatory data processing with VO.

In the area of IVOA standards development, D. Tody led preparation of a draft specification for the table access protocol (TAP), and worked with IVOA partners to develop a major upgrade to the simple image access protocol (SIAV2), documented in an IVOA Note released in May. Both of these were discussed within the IVOA working groups at the IVOA Interop in Trieste in May.

About half of the NRAO effort this past quarter went to development of a conceptual design for the applications framework. D. Tody participated with other colleagues within NVO in the AURA science software workshop held in Hilo in May, representing the

NVO and Opticon efforts in this area, as well as longer term science software development within NRAO. W. Cotton and W. Young also contributed to this effort on the NRAO side, helping to ensure that the evolving applications framework design will be suitable for NRAO and ALMA data processing with CASA and AIPS.

J. Crossley is nearing completion of a first version of a loadtest Web application to be used to test NVO services, and also continues work with L. Sjouwerman on an effort to pipeline process legacy VLA observations to produce images for the VO.

Raytheon/ADC (University of Maryland and George Mason University)
George Mason University (GMU) staff K. Borne presented an invited talk at the Department of Energy Workshop on Mathematics for Analysis of Petascale Data in Rockville, Maryland in June 2008. The talk covered large astronomy surveys and the corresponding petascale data challenges. The talk title was "Data Science Challenges from Distributed Petascale Astronomical Sky Surveys." Borne submitted an extended abstract on the same topic for the IEEE workshop on Computing with Massive and Persistent Data, scheduled for September 2008 in Baltimore. His abstract was accepted for presentation as a talk at the workshop. Borne continued his collaboration with UMBC distributed data mining experts—their focus is on peer-to-peer applications of data mining on extragalactic sources within the combination of the SDSS and 2MASS catalogs. They are preparing two papers for submission to peer-reviewed journals, one in astronomy and the other in database science.

Smithsonian Astrophysical Observatory
G. Fabbiano, J. McDowell, and J. Cant attended the team meeting in Estes Park. Fabbiano has participated in NVO Executive telecons. She will give a talk on the VO at the meeting of the Brazilian Astronomical Society in September.

SAO has hired Dinesh Gunasegaran to join the VO programmer team.

J. Cant is implementing ConeSearch for some CfA data collections.

Space Telescope Science Institute
T. Dower wrapped up the framework for the registry publishing interface to be used at STScI, tested it as a live system with the help of R. Thompson, and used it to publish the first new record directly into the STScI registry. She gave a demonstration of the interface and received substantial useful feedback at the NVO team meeting. She also worked on basic registry maintenance and operations, particularly regarding interfaces for the data discovery portal project.

A suite of customized web services were developed and delivered to support the portal requirements to have unique service interface descriptions. The portal integration process produced a new awareness that will help with future development efforts in the understanding of Registry resources as composites of multiple service capabilities. The interfaces also were enhanced to support additional filtering on bandpass in addition to capability and SQL-like query inputs.

The registry source code maintenance in the NVO SVN repository was reviewed and had several issues that required maintenance and cleanup. As part of the ongoing effort to improve the operational code base for the VO applications, we will continue to improve

and monitor the repository usage of the registry to support more efficient release versions along with continued development.

K. Gillies continued development of an SIA service to provide footprint information to STScI's APT application in the format specified in the IVOA Footprint Overlay Specification Note used by Aladdin. Database issues have been resolved and he is working on finishing this. He completed a report on distributed version control systems, which was distributed to the NVO Technical Working Group.

R. Hanisch worked on preparations for the fall IVOA Interoperability workshop (26-31 October).

R. Hanisch, G. Greene, and T. Dower attended the summer NVO Team Meeting in Estes Park, CO (June 30/July 1). C. Christian joined by telecon.

R. Hanisch, G. Greene, R. Thompson, and A. Conti attended the IVOA Interop meeting in Trieste, Italy. Hanisch participated in the IVOA Executive meetings, the ADQL discussions (pushing for region specifications that are consistent with STC), and the TAP discussions (to make sure that parameter-based queries are acceptable as well as ADQL-type queries). Greene presented the NVO footprint standard service specification at a Data Access Layer working group session. The presentation included the STScI HLA footprint hierarchy as a reference implementation. Several offline discussions with CDS/Aladin developers and the Data Model working group resulted in a much clearer direction for moving forward in a unified approach to a standard footprint service development. The goal is for the development of a draft service proposal to present at the Fall IVOA Interop meeting. In the Registry working group sessions all projects reported positive results for the harvesting of resources between the standard publishing registries. The STScI NVO registry is the primary fully searchable registry using open web application technology in the IVOA. Feedback on the upgraded VO STScI registry performance from the external organizations was very positive. Hanisch also chaired a session of the Data Curation and Preservation Interest Group.

Greene is working with Microsoft developer J. Fay, to develop interfaces between the HLA footprint web services and the WorldWide Telescope. With the professional release of WWT Pro, Fay would like to provide high performance tiled rendering of HST footprints along with a tiled view of the All-Sky Guide Star Catalog 2. The overlays will utilize the HTM software for spatial indexing of the sky.

Hanisch gave presentations at a digital library conference at Purdue University, the astronomy department colloquium (about NVO) at Indiana University in Bloomington, and a talk about NVO at the NASA Applied Information Systems Research workshop at the University of Maryland.

C. Christian, K. Borne, and J. Raddick (JHU) continued to present to EPO groups methods for enabling education with large data archives. They participated in the LSST-EPO working group meeting at the LSST team meeting. The purpose was to clarify what the data needs are for LSST EPO activities and how those cascade to requirements and specifications for the LSST Data Management Facility. The outcome is that the LSST Data Management team believes a lot of the value added materials (keywording, metatags, etc.) will come from the VAO or from the LSST-EPO group itself. The LSST-related

EPO activity is being formulated for input to NSF within the proposal for the construction grant for the LSST.

Three proposals were submitted to NSF or the Department of Education by external groups that contained a plan to use data from the VAO and/or contribute to the VAO. Unfortunately, none were funded, but reviewer comments for resubmissions will be evaluated. Collaboration on the AVM meta-tag project continue so that at least the EPO community implements these standards. C. Christian has been advocating strongly that science archives consider a mechanism for feeding metatags to services that search for science data.

University of Illinois-Urbana/Champaign/National Center for Supercomputer Applications (UIUC/NCSA)

As a faculty member, R. Plante is participating in the planning for the 4th NVO Summer School. Much of the effort this quarter has been in updating summer school software, including the ivoaregistry library (see Section 4.4). He also added SSA support to the VOlib python library as well as made improvement to the VOClient python library.

R. Plante and B. Baker lead the development and support for secure VO logins (see section 8). Plante also leads the Registry WBS, working with team members at STScI (G. Greene and T. Dower) and Caltech (M. Graham). Plante continues to contribute to the development of our Year 7 software process, including assisting M. Nieto-Santisteban and S. Emory Bunn with the on-going software assessment effort.

Publications and Presentations

"Pocket VO: A Simple Tool for Viewing Data Through the Virtual Observatory," J. Raddick, R. Santos, & I. Fernandes, AAS Meeting 212, #73.03.

"NVO Newsletter No. 2", http://www.us-vo.org/newsletter/2/

"NVO Status Report," R. Hanisch, NASA AISRP Workshop, College Park, 6 May 2008.

"Large-Scale Cross-Matching with Open SkyQuery." M. Nieto-Santisteban, NASA AISRP Workshop, College Park, 6 May 2008.

"Digital Data in Astronomy," R. Hanisch, Purdue University, 12 May 2008.

"The National Virtual Observatory: Astronomical Research Enabled by Information Technology," R. Hanisch, Indiana University, 13 May 2008.

"Data Science Challenges from Distributed Petascale Astronomical Sky Surveys," K. Borne, Department of Energy Workshop on Mathematics for Analysis of Petascale Data, Rockville, Maryland, June 2008.

Virtual Observatory Articles in the Popular and Technical Press

"The Virtual Observatory Inside WorldWide Telescope," http://research.dynamic-patterns.com/ 2008/07/the_virtual_observatory_inside.html#000160

WorldWide Telescope, http://www.worldwidetelescope.org/experienceIt/ExperienceIt.aspx

Usage Logs

A number of NVO participating organizations have implemented standard interfaces to their web and service logs, and we collect these logs in order to track VO-enabled use of data and services. The main NVO web site continues to be used frequently, with an all-time peak in March 2008 of nearly 500,000 hits and over 470,000 hits in June 2008.

Monthly Usage of NVO Web Site (http://www.us-vo.org):

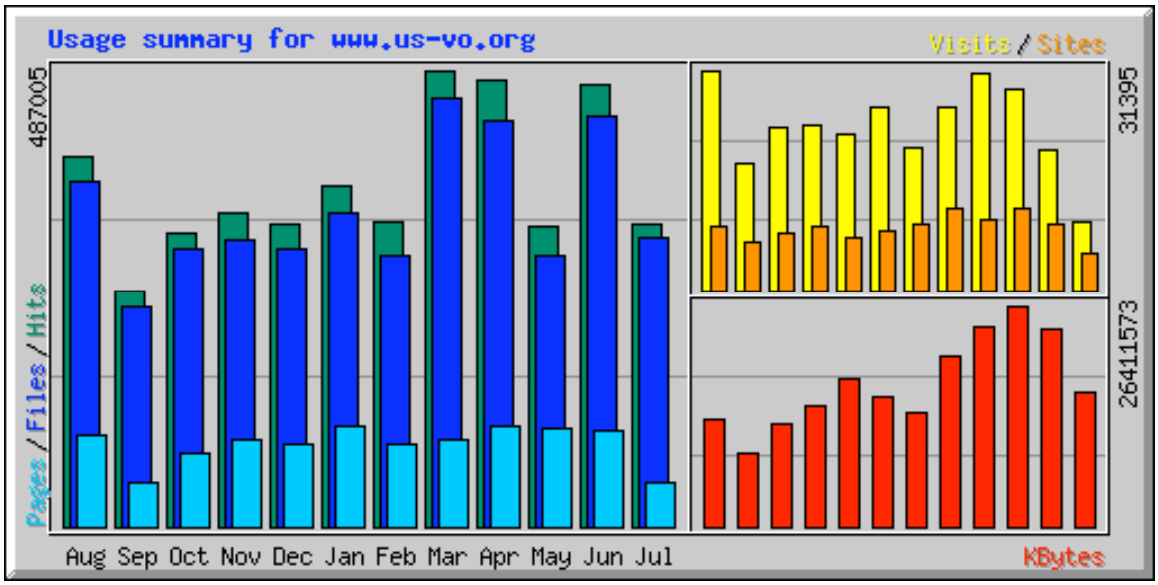

Monthly Hits at NVO Organizations Originating from NVO Applications:

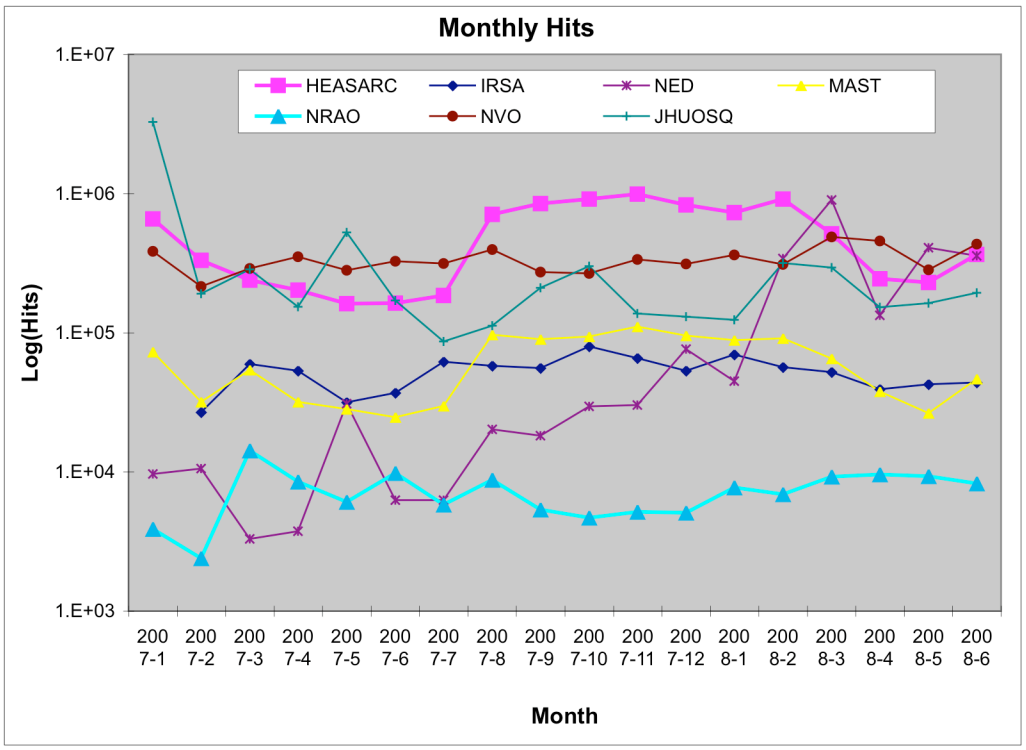

Monthly Users of Any NVO Application or Website:

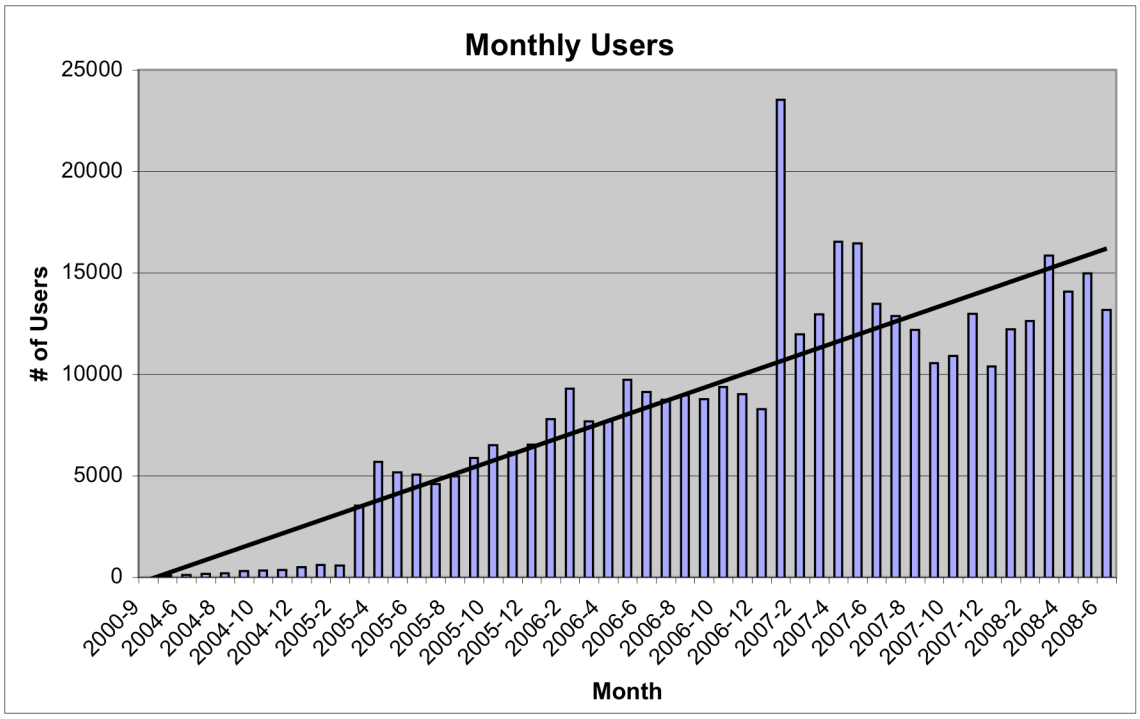

Cumulative Use of NVO Services:

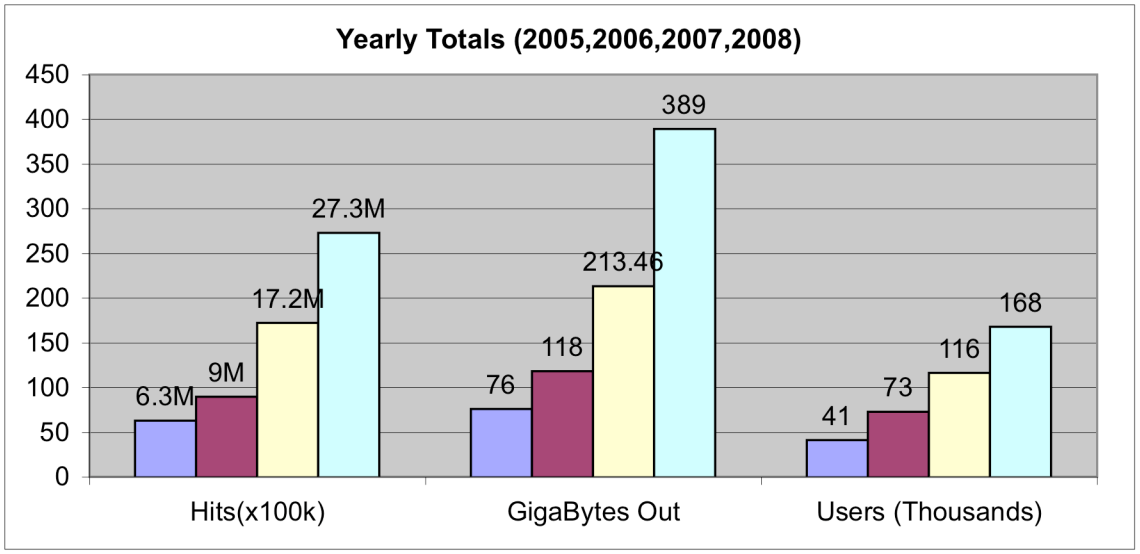

Acronyms

AAS	American Astronomical Society
ADC	Astronomical Data Center
ADEC	Astrophysics Data Centers Executive Committee (NASA)
ADQL	Astronomical Data Query Language
AIPS++	Astronomical Image Processing System++ (NRAO)
API	Applications Programming Interface
AVO	Astrophysical Virtual Observatory
CACR	Center for Advanced Computational Research (Caltech)
CADC	Canadian Astronomy Data Centre
CDS	Centre de Données astronomiques de Strasbourg
CMU	Carnegie Mellon University
CXC	Chandra X-Ray Center
CY	calendar year
DAG	Directed Acyclic Graph
DAGMan	Directed Acyclic Graph Manager (Condor)
DAML	DARPA Agent Markup Language
DARPA	Defense Advanced Research Projects Agency
DIS	Data Inventory Service
DM	Data Model
DOE	Department of Energy
DPOSS	Digitized Palomar Observatory Sky Survey
DTD	Document Type Description
EPO	Education and Public Outreach
ESTO	Earth Science Technology Office (NASA)
ESTO-CT	ESTO Computational Technologies (NASA)
FIRST	Faint Images of the Radio Sky at Twenty Centimeters
FITS	Flexible Image Transport System
FNAL	Fermi National Accelerator Laboratory
FTP	File Transport Protocol
FY	fiscal year
GB	gigabyte
GLU	Générateur de Liens Uniformes (uniform link generator)
GRB	Gamma Ray Burst
GriPhyN	Grid Physics Network
HEASARC	High Energy Astrophysics Science Archive Center
HTTP	HyperText Transport Protocol
IPAC	Infrared Processing and Analysis Center (Caltech)
IRAF	Image Reduction and Analysis Facility (NOAO)
IRSA	Infrared Science Archive (IPAC)
ISI	Information Sciences Institute (USC)
ITWG	Information Technology Working Group (NASA data centers)
iVDGL	International Virtual Data Grid Laboratory
IVOA	International Virtual Observatory Alliance
IVORN	International Virtual Observatory Resource Name

JDBC	Java Data Base Connectivity (Sun, Inc., trademark)
JHU	The Johns Hopkins University
MAST	Multimission Archive at Space Telescope (STScI)
MB	megabyte
MOU	Memorandum of Understanding
MWG	Metadata Working Group
NASA	National Aeronautics and Space Administration
NCSA	National Center for Supercomputer Applications
NED	NASA/IPAC Extragalactic Database
NESSSI	NVO Extensible Secure Scalable Service Infrastructure
NOAO	National Optical Astronomy Observatories
NPACI	National Partnership for Advanced Computational Infrastructure
NRAO	National Radio Astronomy Observatory
NSF	National Science Foundation
NVO	National Virtual Observatory
OAI	Open Archives Initiative
OASIS	On-line Archive Science Information Services (IRSA)
OGSA	Open Grid Services Architecture
OIL	Ontology Inference Layer
OWL	Web Ontology Language
PB	petabyte
PMH	Protocol for Metadata Harvesting (of OAI)
Q	quarter
QSO	Quasi-Stellar Object
RC	Replica Catalog
RDF	Resource Description Framework
REST	Representational State Transfer
RLS	Replica Location Service
ROME	Request Object Management Environment
SAO	Smithsonian Astrophysical Observatory
SAWG	Science Archives Working Group (NASA)
SAWG	System Architecture Working Group (this project)
SciDAC	Scientific Discovery through Advanced Computing (DOE)
SDSC	San Diego Supercomputer Center
SDSS	Sloan Digital Sky Survey
SDT	Science Definition Team
SIAP	Simple Image Access Protocol
SOAP	Simple Object Access Protocol
SRB	Storage Resource Broker
SSAP	Simple Spectral Access Protocol
STScI	Space Telescope Science Institute
SWG	Science Working Group
TAP	Table Access Protocol
TB	terabyte
UCD	Unified Content Descriptor
USC	University of Southern California

UDDI	Universal Description, Discovery, and Integration
UIUC	University of Illinois Champaign-Urbana
USNO	United States Naval Observatory
USRA	Universities Space Research Association
UWS	Universal Worker Service
VDL	Virtual Data Language
VDS	Virtual Data System
VO	Virtual Observatory
VO	Virtual Organization
VOQL	Virtual Observatory Query Language
WBS	Work Breakdown Structure
WebDAV	Web-based Distributed Authoring and Versioning
WSDL	Web Services Description Language
XML	Extensible Mark-up Language
2MASS	Two-Micron All Sky Survey