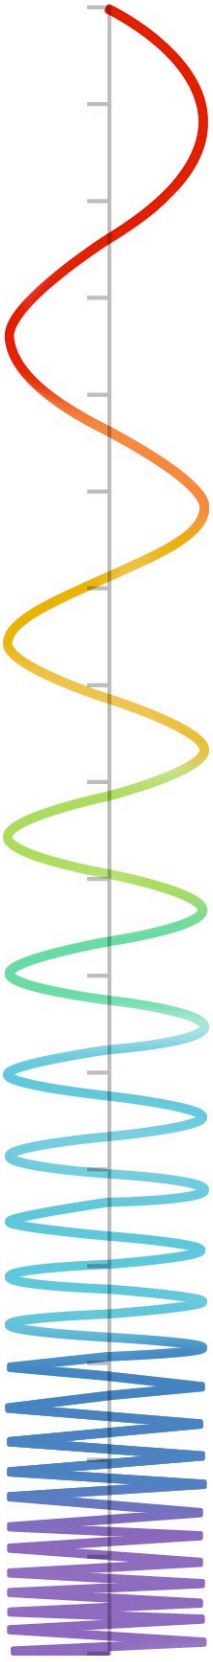

Quarterly Report
October – December 2008

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 October – 31 December 2008

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

Executive Summary

In October 2008 the NVO Project began operating solely on carry-forward funding from the previous fiscal year. Our budget estimates indicate that we will be able to continue work at a reduced level for approximately six months. We have reprioritized work and attempted to match development goals to the remaining resources. Our top priorities are to complete work on the NVO data discovery portal (essentially done, with the tools and interfaces undergoing testing as the quarter came to a close), follow through on commitments to the international VO standards process, and support and maintain the software and tools in use. To be frank, however, it is a very difficult period for the project team. The uncertainty surrounding future support is creating morale problems, and management is concerned that experienced people will be lost from the project team. Several organizations have prepared, or already had to carry out, staff reductions owing to the severely constrained budget. Efficiency and productivity have decreased as team members have had to spend more of their time on other projects.

The Astronomical Data Query Language (ADQL) specification was finalized at the Baltimore Interop meeting and endorsed as a full Recommendation by the IVOA Executive.

Substantial technical progress was made on the Table Access Protocol (TAP) specification, which was discussed in detail at the fall IVOA Interoperability meeting in Baltimore. Unfortunately, little progress has been made since early December. Senior members of the NVO team are working with our IVOA partners to help get the process going again.

The IVOA Grid and Web Services Working Group completed RFC periods for the VOSpace V1.1 specification and the Credential Delegation Protocol. VOSpace V1.1 supports transparent access to distributed storage systems, and the Credential Delegation Protocol supports hand-off of user authentication from one VO service to another.

The major development effort in the past quarter focused on completing and testing the data discovery portal components, their integration, and their presentation via a new home page for NVO. Astronomers outside the core NVO team were asked to evaluate the tools and their organization, and their feedback was immensely helpful. The new tools and website will be made available to the astronomy community in the next quarter.

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. Both groups decided to continue these regular telecons in order to maintain as much continuity and team unity as possible despite operating solely on carry-forward funding. No fall team meeting was held, however, in order to conserve funds. Several internet-based virtual meeting tools were tested, particularly, the EVO system at Caltech (NSF sponsored, used widely in high energy physics). No facility has been found, however, that deals well with more than a handful of video connections.

The Annual Report for FY2008 was completed and submitted to NSF on 31 October.

An updated project plan was prepared in October reflecting top priorities and available resources. Our top priorities are to complete work on the NVO data discovery portal (essentially done, with the tools and interfaces undergoing testing as the quarter came to a close), follow through on commitments to the international VO standards process, and support and maintain the software and tools in use.

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). R. Hanisch and P. Fabbiano recruited a team of scientist-testers to evaluate the portal, and their feedback helped us to identify a number of improvements.

1.3 Technical (including standards, configuration management)

Substantial progress was made on the Table Access Protocol (TAP) during the quarter, thanks to a core technical team composed of G. Rixon (AstroGrid), P. Dowler (Canadian VO), and D. Tody (NVO/NRAO). Disagreements on process, however, led Rixon to remove himself from the work, and K. Noddle (AstroGrid), the chair of the IVOA Data Access Layer Working Group, assumed editorial responsibility for the standard. This has failed to produce any further progress, however, and several technical issues remain to be resolved and prototyping activities are thus stalled. The NVO management continues to seek solutions with our IVOA partners.

During the October IVOA Interoperability meeting, held at STScI/JHU, agreement was reached on the Astronomical Data Query Language (ADQL) standard. This was quickly approved as a Recommendation by the IVOA Executive.

1.4 Financial

The NVO Executive Committee agreed with a plan proposed by the Program Manager to maximize participation of all key team members by reapportioning the remaining project funds based on the steady-state levels of effort that prevailed at the beginning of FY2008. JHU issued subaward modifications to all organizations, some of which decreased spending authority and others of which increased it. This was particularly painful for those or-

ganizations that had carefully conserved funds, but ultimately all groups and managers agreed that this approach was the best overall for the project team.

Total project expenditures since inception now stand at $13,486,399 compared to a budget of $14,151,966, leaving a balance of $683,567. We have $222,405 in invoices received and pending payment, leaving a net of $461,162. Total spending this quarter, including pending invoices, is $486,862. We expect our remaining funds to be essentially exhausted in the second quarter.

1.5 International coordination/collaboration

Senior members of the NVO project participated in IVOA Executive telecons and discussions. D. De Young (NOAO) completed his term as chair of the IVOA Executive Committee, and was replaced by F. Pasian (VO-Italy).

STScI and JHU co-hosted the fall IVOA Interoperability meeting. AURA, AUI, and Computer Sciences Corporation provided financial support. About 100 people attended and the usual Working Group and Interest Group sessions were held. The plenary meetings were webcast and archived; those interested may view them at

http://www.stsci.edu/institute/itsd/information/streaming/archive/IVOAInteroperabilityMeeting

VO-related presentations and discussions continued the following week at the ADASS Conference in Quebec, Canada, and a number of NVO personnel attended and gave talks or posters on VO projects.

De Young attended the EuroVO Science Advisory Committee meeting in December in Strasbourg.

S. Emery-Bunn (Caltech) prepared the first issue of the IVOA Newsletter for publication on the web and notified the international VO user community of its availability through IVOA and NVO mailing lists. The Newsletter has been accessed over 1000 times. See http://www.ivoa.net/newsletter/001/ .

We anticipate convening the IAU Virtual Observatories, Data Centers, and Networks Working Group during the IAU General Assembly in Buenos Aires, August 2009.

2 Science Requirements

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

No activities to report this quarter.

2.2 Requirements analysis

No activities to report this quarter.

2.3 Demonstration definition and review

No activities to report this quarter.

3 Operations, System Integration, and Testing

3.1 Quality assurance and software engineering design

M. Preciado (HEASARC) prototyped a comprehensive validation framework for all VO services as a collaboration between the HEASARC, NCSA and STScI. Using the valida-

tion tools developed by R. Plante (NCSA) this validation framework will check all registered VO verifiable services (registries, cone search, SIA, and SSA services). These services will be validated on an anticipated one-month cadence and the results of all validations will be maintained in a user queryable database. Registries may then use this database as one basis for settings their validation status flags for VO services.

3.2 Facility operations

M. Preciado (HEASARC) continues to operate a suite of tests using the NAGIOS software system which checks every site providing VO services. Typically tests are run every hour. For sites that provide many services, only a representative subset of the services are checked to avoid undue use of the services. Currently about 100 tests are made on servers that host about 7,000 distinct queryable services. All portal elements are tested.

Preciado reviews the status of all sites on a 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 follow-up has helped to significantly decrease the number of sites that failed testing over the course of the year.

A record of all issues that come up and their resolution is maintained and available for review. The procedure for informing sites of problems is itself being reviewed and standardized. On average 2-5 sites are notified of problems each. In most cases the problem was resolved quickly but in some cases the resolution requires multiple iterations.

Aside from occasional service downtimes there were no major interruptions to services at the NVO's major web sites at CACR, NOAO, CfA, HEASARC, IRSA, NCSA, STScI, or JHU.

VOEvent services continued nominal operations at the CACR (A. Drake, M. Graham. R. Williams).

Montage services continued nominal operations at IRSA.

The TRAC services continued to be used in the Portal development. Initial discussions have begun on the distributed configuration management systems that might supersede the existing SVN tools.

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.

S. Emery-Bunn (CACR) made preparations for the next newsletter (anticipated 1Q 2009).

4 Registries

In this quarter, the bulk of our effort has gone into testing the new registry as part of the new NVO portal and addressing issues arising as a result. We also deployed a version of the new registration interface though development continues.

4.1 Resource metadata

Nothing to report this quarter.

4.2 Resource metadata schema
We continue to refine the metadata model for describing tables. In particular, a consensus has emerged on how to describe foreign keys as an aid for creating joins within a complex set of catalog tables. This model is being incorporated into what will be the VODataService standard schema.

4.3 Publishing and harvesting protocols
As part of our effort to ready the new registry interface for the release of the NVO portal, T. Dower (STScI) has deployed our latest version of the Universal Publishing Interface at the project registry at STScI. This has allowed beta-testing of the interface with real users—in particular, users that have registered with the STScI registry and wish to update their records. Feedback has been good, so we are including it as working service that will be part of the portal release. Work continues on the interface, particularly in better management of access rights.

R. Plante (NCSA) continues to edit the Registry Interfaces document as part of the IVOA standardization process. .

4.4 Search protocols
Extensive testing of the new registry web interface to searching and its role in the NVO portal has been a major part of our effort this quarter. G. Greene and T. Dower (STScI) have been refining the presentation based on feedback from users.

Keeping VO jargon from confusing users continues to be a challenge. In particular, we have found the term "registry" to be fairly mysterious to new users, which can lead to incorrect expectations of what the registry can do. In response, the portal team recommended that within the NVO portal, we refer to the web interface to the registry as the "directory," a term that we hope will resonate better with astronomers.

One common misconception about the registry interface is that it will resolve and respond effectively to things like object names or return such fine-grained products such as individual images. We have thus improved the labeling of the interface to redirect users who want to do such searches to the more appropriate tools.

4.5 Replication, synchronization, maintenance, revision control, and curation
Our testing continues to uncover various metadata record curation issues. We are particularly concerned about Vizier tables at CDS as they represent the vast majority of available catalogs. These registry records are maintained at CDS. R. Hanisch (STScI) and R. Plante (NCSA) have continue to communicate with CDS staff to address issues of compatibility and support for the relevant standards.

5 Data Models

5.1 High-level (image, spectrum, time series, event lists, visibilities, catalogs, simulations, data quality)
J. Cant continued work on updating the Java library that implements the Spectrum model. Cant and J. McDowell are prototyping the NEDCOPY application which uses the library to convert NED's SED data model into the VO SED data model. This has resulted in improvements to the design of the library, which is being prepared for re-release.

5.2 Low-level (measurement, quantity, uncertainty, relationship)
No activity during this quarter.

5.3 Descriptors and ontologies (UCDs)
NVO team members M. Graham and K. Borne are on the Scientific Organizing Committee of the conference "Practical Semantic Astronomy," which will be held in Glasgow, Scotland, 2-5 March 2009.

5.4 Space-Time and regions
D. Gunasegaran finished a prototype STC-X to STC-S converter. This is the start of an STC Java library.

A. Rots is looking into adapting the HLA footprint service developed by G. Greene (and based on the libraries developed by T. Budavari) to the Chandra Data Archive.

A. Rots continued discussions in the VO on STC issues. As to the latter, those were specifically with D. Berry (UK), who is working on updating his STC-compliant coordinate transformation application that is based on the AST library; and with F. Ochsenbein (CDS) concerning the specification of Julian and Besselian epochs in STC.

5.5 Standard schema
No activity during this quarter.

6 Data Access Layer

A first draft of an IVOA specification for the Table Access Protocol (TAP) was presented and discussed at the IVOA interoperability workshop held in Baltimore in late October. Key issues relating to the second generation IVOA image access protocol, which includes support for remote access to large spectral data cubes, were also discussed. A revision of the TAP specification was produced following the Interop, and a round of prototyping is planned for both data access protocols in the months leading up to the next IVOA Interop to be held in Strasbourg in May 2009. Plans for a multi-observatory applications framework integrating observatory computing with VO were discussed in a BOF session at the ADASS, followed by prototyping and completion of a whitepaper and a conceptual design document in early January.

6.1 Data access services (catalog, image, spectrum, time series, visibilities, ...)
Catalog access. An international team (P. Dowler, G. Rixon, D. Tody) produced a first draft of an IVOA specification for the Table Access Protocol in late September. Related prototypes were completed within the US (IRSA, STScI, HEASARC, NRAO) and at the CADC (P. Dowler). The draft TAP specification and prototypes were discussed within a session of the DAL working group at the IVOA Interop meeting in Baltimore in late September.

A second revision of the TAP protocol was produced following the Baltimore Interop, however the core TAP design team was dissolved late in the year due to a lack of consensus within the IVOA community on how to proceed with TAP. NVO is continuing its efforts to move the TAP documents forward within the IVOA context and to resolve any conflicts in that environment. If this is not successful, NVO may be forced to separate its efforts on TAP from the overall IVOA context in order to preserve the momentum and progress on this crucial capability. A round of prototyping is planned for both the ADQL and paramater-based access protocols in the months leading up to the next IVOA Interop to be held in Strasbourg in May 2009.

Spatial Footprints. G. Greene (STScI) worked in coordination with the IVOA Data Model group to generate the footprint service use case for Observational data. This work was presented as part of the DAL working group footprint service session at the Baltimore Interop meeting along with the footprint standard specification effort..

Spectral and time series data. No activity this quarter.

Image data. An updated version of the DAL service architecture and standard profile (DAL2) document was prepared for the IVOA Interop meeting in Baltimore, and discussed in the meeting. This defines the linkage between the SIA V2 interface, SSA, and other second generation DAL interfaces such as TAP, and is critical to integrating all these second generation, production IVOA DAL interfaces.

An updated version of the SIAV2 conceptual design document was prepared and distributed, incorporating comments from F. Bonnarel since this document was first distributed and discussed in connection with the May 2008 IVOA Interop meeting in Trieste. Key issues relating to the SIAV2 image access protocol, which includes support for remote access to large spectral data cubes, were discussed in the fall Interop in Baltimore. Good progress has been made on how to represent and deal with polarization, and integration of multidimensional WCS and STC. The major concern expressed in the Baltimore meeting was the complexity and scope of SIAV2. Some wish that it could be simplified, but the advanced applications required for the facility phase of VO plus support for advanced data types such as the spectral data cubes produced by modern radio instruments and IFUs, require functionality beyond what is currently available.

SIAV2 remains a high priority for development over the next two years, second only to TAP among the DAL interfaces. The next steps will include specification of a preliminary interface, and prototyping of SIAV2 including the updated query interface, integration of grid functionality, and support for expanded image data access such as spectral and time cubes.

Complex data. Work on modeling complex data associations continues mainly in connection with the development of SIAV2. A discussion of complex data was included in the DAL2 architecture document released in advance of the fall 2008 IVOA Interop workshop in Baltimore. The related *generic dataset* concept is being investigated as the basis for modeling both complex data as well as to provide the basis for global data discovery in the VO, a high priority capability planned for the facility phase of VO.

6.2 Data representation (VOTable, etc.)

The UTYPE mechanism, used to tag data model fields to enable reconstruction of the logical content of a data model independent of how a data model is physically represented, was the topic of spirited discussion at the Baltimore Interop. While UTYPE is heavily used within the second generation DAL interfaces, some details such as how namespaces are identified and interact require further definition. The basic concept and use of UTYPE is still not fully accepted by some, with some preferring fixed representations such as structured XML. The problem with limiting data models to a single representation such as XML is that this would prohibit the use of many languages and storage formats in common use, including such VO standards as VOTable and FITS.

While VOTable is one of the most notable successes of VO it remains controversial as well. Some would prefer to do everything with XML. This approach would have some advantages as a broader non-astronomy standard, but XML requires use of a custom hierarchical data structure for every individual data model, whereas VOTable provides a standard table-oriented container for storing data, allowing data models which are simple enough to be represented as a set of keyword value pairs (UTYPEs with associated values) to be stored and manipulated as tables. VOTable as well as XML are also relatively inefficient, requiring a large amount of text to represent a given amount of information, compared to alternative, more compact text-based serializations such as JSON. While not a problem for most scientific analysis this can be an issue for interactive applications that transfer information over a relatively slow Internet connection.

6.3 Framework (mediators, components)

A special session on the applications framework project (an effort to develop a common framework for processing of data from the major observatories while integrating observatory data processing with VO) was held at the ADASS conference held in Quebec in early November. Representatives from AUI, AURA, NVO/VAO, and OPTICON (R. Hanisch, J. Kantor, P. Grosbol) chaired the session, which had over one hundred attendees, producing a lively and supportive discussion. A white paper summarizing the goals and architecture of the planned system has since been produced by a joint effort by representatives of NVO and OPTICON. The AURA observatories are producing a similar whitepaper to guide future software development at AURA. A conceptual design for the applications framework, first drafted in mid-2008, is in the process of being updated, in preparation for an initial round of prototyping planned to take place over the next year.

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

The DALServer service framework, used to implement data services, is now stable at version 0.3. The current version includes support for spectral data (SSA), simple cone search, and simple image access (SIA V1). Prototype support for TAP and SIAV2 is planned. NRAO (D. Tody), the NED group at IPAC (J. Mazarella et. al.), and ESO (A. Micol) are currently collaborating on DALServer development as well as using DALServer for local VO data services.

A validator for SSA services is under development at NCSA (R. Plante). This will augment our current suite of DAL validators, along with the cone search and SIA validators already available.

7 Query Language

7.1 Low-level: Astronomical Data Query Language

The Astronomical Data Query Language standard document became an IVOA Recommendation at the end of October 2008, after a few remaining issues about STC metadata and function definitions were resolved during the IVOA Interoperability meeting in Baltimore.

7.2 Mid-level: VOQL and OpenSkyQuery/OpenSkyNode

Nothing to report this quarter.

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. VOSpace 1.1 has finished its RFC period and awaits a final version and TCG approval. There was discussion about VOSpace 2.0 in Baltimore. There was also discussion about RESTful approaches to web services and whether the IVOA should be advocating particular toolkits or providing representations of IVOA data standards that are known to serialize with the most common web service frameworks.

8.2 Grid Services (OGSA)

No activities to report this quarter.

8.3 Computational resource management

There was discussion about the pros and cons of centralized authorization services at the Interop meeting in Baltimore.

8.4 Virtual data

No activities to report this quarter.

8.5 Application and service integration with Grid

The Credential Delegation Protocol specification has finished its RFC period and a final edit is underway before being forwarded to the IVOA Technical Coordination Group for review.

9 Applications

9.1 Data location services

Significant development in the data discovery portal services continued during this quarter. R. Hanisch (STScI) initiated a comprehensive review of portal services by members of the NVO team outside the portal and by interested outsiders. Results of this review will be collated and acted upon with the intent to release the portal on 1 Feb 2009.

Changes in the last quarter focused primarily on ensuring that the tools worked together. A wizard describing how to apply the portal to various science questions that had been developed by T. McGlynn (HEASARC) was adapted by S. Emery-Bunn (CACR).

Emery-Bunn and Hanisch developed a revamped NVO home page that highlights portal services.

M. Fitzpatrick (NOAO) began working to incorporate SSAP services within the VOClient framework. He has began incorporating the Simple Applications Messaging Protocol into this software.

The Portal Working Group met regularly to discuss issues with regard to the development of the portal.

9.2 Cross-correlation services

The data discovery portal working group decided to include Open SkyQuery as a core component. JHU was asked to update the web appearance of OSQ to follow the standard NVO application template.

9.3 Image combination, registration
Nothing to report this quarter.

9.4 Visualization tools and services
Nothing to report this quarter.

9.5 Theory
Nothing to report this quarter.

9.6 Statistical analysis
Nothing to report this quarter.

9.7 Data mining, outlier identification
Nothing to report this quarter.

9.8 Interfaces to/from legacy software systems
Nothing to report this quarter.

10 Community Engagement

10.1 Documentation
A draft of a third NVO Newsletter was prepared, but distribution is being postponed until the final release of the new data discovery portal.

10.2 Web site
In support of the data discovery portal deployment, the NVO web site has been completely redesigned and reorganized. The web site is now totally focused on the astronomer-user and emphasizes immediate access to and use of NVO tools. Background information and technical details have been moved to other pages, "behind the scenes". This new design is being evaluated as part of the data discovery portal testing. The current web page is shown here, followed by the new version.

The new version emphasizes the tools of immediate interest to astronomers. Text is minimized, and icons with a common graphical design are used to represent the tools on the home page, and then throughout the other pages that astronomers visit. User reaction to the new design has been extremely positive. S. Emery Bunn (Caltech) led the web site redesign effort in consultation with R. Hanisch (STScI).

10.3 Technical training initiatives

No activities to report this quarter.

10.4 Advocacy

No activities to report this quarter.

11 Education and Public Outreach

EPO activities are dormant in the no-cost extension phase of work.

Activities by Organization

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

- Working to bring the NVO Data Discovery portal to the alpha-testing milestone. Running tests and use cases on Vim, running interoperability tests.
- Wrap-up for summer school.
- Working on NVO portal.
- Redesign and rewrite of the NVO website.
- Evaluating EVO.

Caltech–Infrared Processing and Analysis Center (IPAC)
J. Good continued to lead the Technical Working Group.

IPAC upgraded the inventory and table converter portal services in response to defect reports and user evaluation.

O. Pevunova participated in the Technical Working Group telecons and J. Mazzarella participated in the Executive Committee telecons. Work on the SSAP DAL interface to the NED spectral database continued at the level of revisions and updates to the metadata. Deployment of the prototype SSAP service, originally planned for December 2008, has been delayed to Spring 2009 due to limited availability of our part-time Java developer. O. Pevunova and other IPAC staff spent time testing the NVO portal and submitted problem reports and suggestions.

High Energy Astrophysics Science Archive Research Center (HEASARC)
D. Hinshaw joined the HEASARC in September 2008 and is providing programming support for NVO development efforts.

T. McGlynn led the portal development team, which includes members at CACR, STScI, NOAO, and IRSA. During the past quarter all of the portal software has been released to public sites, but not explicitly linked to the NVO home pages. McGlynn continued to enhance and refine the Simple Query and Table Viewer tools and build a wizard to help users find appropriate Portal services. Work resumed on integrating DataScope into the new registry V1.0 environment.

M. Preciado periodically reviews the status the health of VO services each day. When a failing service is discovered a dialog with the responsible parties is initiated until the issue is resolved.

M. Preciado continued to augment software that monitors the health of all VO hosts. A simple report indicating the health of the VO is available through the NVO internal pages. A major focus of the effort in the past quarter has been the development of a framework that will validate all registered VO services using tools developed at NCSA. Since validation requires substantial resources on both the resource and validation site, only a few percent of all services will be validated on a given day. All VO services will be validated over a regular (~monthly) cycle.

McGlynn is the Operations WBS lead for the NVO and a member of the NVO Executive committee. He also chairs (since July 2008) the IVOA Applications Working Group. He leads the ad hoc portal working team. An RFC for the Simple Applications Messaging Protocol was initiated by McGlynn in this quarter and finished early in the following quarter. While some useful comments were made, no major issues have arisen during the RFC.

In addition to these major focuses the HEASARC has participated in many other elements of the NVO development. McGlynn developed a prototype of the TAP interface as a layer over the existing native HEASARC interfaces and is involved in discussion of TAP, VOQL, and other data access protocols.

HEASARC staff also participated in the NVO team meetings, IVOA meetings, the weekly technical telecons, and IVOA group discussions as appropriate.

Johns Hopkins University
M. Nieto-Santisteban attended the IVOA meeting in Baltimore in October, and ADASS in Quebec in November.

A. Thakar continues to maintain the NVO web logs and monitor the harvesting at JHU on a daily basis. He contributes quarterly and annual usage stats and charts for the NVO quarterly and annual reports. Thakar also presented the status of NVO logging at the IVOA Interop in Baltimore, and urged the VO community to develop a standard for service logging. A prerequisite for this is finalizing a format for the RunID that will be used to track service requests through various stages of execution.

Thakar continued to provide triage and support for technical problems with Open SkyQuery and other VO services at JHU, delegating to developers of the services as necessary. Thakar deployed the DR7 VO services for SDSS when DR7 became publicly available (10/31/08).

A. Szalay, along with J. vandenBerg and A. Wonders, continued to build the GrayWulf (in honor of Jim Gray) data cluster. The GrayWulf system won the Storage Challenge at the Supercomputing 2008 conference (http://www.jhu.edu/news_info/news/home08/dec08/storage.html).

J. Raddick continued to maintain the virtualobservatory.org public outreach website.

S. Carliles has developed a prototype TAP interface to OpenSkyQuery, which is currently available for testing.

National Optical Astronomy Observatories (NOAO)
M. Fitzpatrick continued participation in the Portal Working Group, making small modifications to the VOClient tools in response to user testing. Fitzpatrick continued work on the Simple Applications Messaging Protocol (SAMP) specification, which is now under consideration for promotion to Recommendation status. Work was begun to implement interfaces to Simple Spectral Access and the new SAMP protocol within VOClient, which is expected to be released in the near future. The WCSFixer tool was used over 2,000 times during the quarter, solving approximately 70% of submitted images. This is a four-fold increase in usage with an overall improved success rate from past usage. M. Fitzpatrick, R. Seaman, D. DeYoung, C. Miller, and A. Egana attended the IVOA Interoperability meeting held in Baltimore 26-31 October 2008.

NOAO's Data Products Program was reorganized and rescoped this quarter. It was given a new name, Science Data Management, to reflect the short-term mandate (next 18-24 months) from NOAO management to meet the immediate data management needs of NOAO and its community. The group is now slightly more than half its original size.

National Radio Astronomy Observatory (NRAO)
The primary VO activity at NRAO currently continues to be in the area of VO standards and infrastructure development. D. Tody led development of the TAP and SIAV2 standards within NVO, and participated in the workgroup discussions at the IVOA interoperability workshop in Baltimore.

Phase I of the NVO/OPTICON applications framework project was completed in early January. This project has been expanded and now includes AURA and AUI within the US, in addition to the work already being done by NVO in the US and OPTICON in Europe. An architecture whitepaper (D. Tody, P. Grosbol, et. al.) was completed in early January as part of this effort. The applications framework concept document drafted in mid-2008 is in the process of being updated. This will serve as the basis for applications framework prototyping as this goes forward in 2009.

A VO application, KMLNow!, begun at the NVO summer school in the fall of 2008 is now functional (J. Crossley, R. DuPlain, N. Radziwill). J. Crossley attended the January AAS to present a poster paper on this project. An SPIE paper on the project is in preparation.

A new project has begun at NRAO to produce a catalog of all VLA observations and publish this to the VO, intially via a simple cone search interface. This will be done with the DALServer framework, allowing addition of a prototype TAP interface later once this becomes available. The observation catalog will also be integrated with the associated catalog of images produced from selected VLA observations.

Raytheon/ADC (University of Maryland and George Mason University)
George Mason University (GMU) staff K. Borne presented an invited talk in November 2008 at the University of Notre Dame on NVO-enabling data-intensive astronomy. Borne presented a contributed talk at the AGU meeting in December 2008 on early science results from P2P (peer-to-peer) data mining, which addressed the use of the different VO Grid and Web Services (GWS) protocols required for the implementation. Borne presented essentially the same talk at the GWS working group breakout at the Baltimore IVOA Interop in October 2008. Borne is serving on the Scientific Organizing Committee for the March 2009 international conference “Practical Semantic Astronomy” in Glasgow Scotland; most of the astronomy talks at the conference are drawn from the IVOA semantics working group or from IVOA metadata development activities. Borne is co-author on a paper that was accepted for the peer-reviewed conference proceedings of the 2009 SIAM International Data Mining Conference; the paper used NVO-accessed distributed data on elliptical galaxies to study the fundamental plane of galaxy parameters. The paper is part of a broader distributed data mining research program using distributed scientific databases, including VO-accessible data, registries, and resources. The paper is titled “Scalable Distributed Change Detection from Astronomy Data Streams using Local, Asynchronous Eigen-Monitoring Algorithms.” The peer review was intense in this

case—only about 20 out of 150 papers were accepted for publication in the peer-reviewed proceedings.

Smithsonian Astrophysical Observatory
G. Fabbiano continued to attend telecons, planned budget to stretch SAO effort to the spring, continued serving in the Interagency Working Group for Digital Data, and the Smithsonian Institution Digitization Strategic Planning Committee. In both groups the VO is recognized as an important groundbreaking effort.

Space Telescope Science Institute
In November and December, T. Dower's work on the publishing interface shifted from basic development to support as several users began publishing to the new registry and required not only editing capabilities for a wide range of resources from the old registry, but support for Open Sky Nodes, a previously unsupported resource type. Basic user management functionality was also added. Support for the registry itself was improved with ongoing development in an advanced search page and with the introduction of weighted search results and a caching mechanism resulting in better results with a 50% improvement in search time.

Several STScI local registry resources were updated by G. Greene to correct accessURL locations. This activity included beta testing the new STScI VO Registry publishing Interface. There were several software issues resolved in the process. The publishing Web site will be opened to the public in January for updating locally managed VO registry resources as well as publication of new resources. This site is still in beta testing and users will be notified as such.

K. Gillies added a few new features to the SIA service that provides footprint information to STScI's APT application and Aladin. The tables needed for this application have now been migrated to the 2.5 HLA production system. The last step is to deploy the service, and that is taking place now. Development of this service will continue as it begins to provide features needed by the HLA.

R. Hanisch made major revisions to the IVOA Document Standards process document in response to RFC comments. The changes were sufficiently major that the document will go through the review and promotion process again. The revisions include an update to the IVOA document numbering system.

Hanisch reviewed the Table Access Protocol V0.30 and V0.31 documents and compiled extensive comments on the drafts. The chair of the IVOA Data Access Layer Working Group decided to assume editorial responsibilities for the TAP document himself, a move that appears to have brought progress to a halt. Hanisch continues to work with senior IVOA colleagues to try to get work on TAP back on track. Hanisch also reviewed the Simple Application Message Protocol (SAMP) specification RFC and sent extensive comments to the editor.

Hanisch worked with S. Emery Bunn (Caltech) to finalize the data discovery portal web site, which will eventually replace the current NVO home page. We have adopted a much simplified, iconic view of the portal components and redesigned the site to support the needs of the research community using the VO.

Hanisch organized a science testing team for the NVO data discovery portal, and invited them to evaluate the new services. Testing began over the holiday period and continued into January.

Hanisch compiled a history of IVOA events and key appointments and posted this on the IVOA Executive Committee's web site.

Greene, Dower, and Hanisch attended the ADASS Conference in Quebec. Dower presented a poster on the general workings of the new registry and the publishing interface, and Greene gave a poster on the footprint services for the Hubble Legacy Archive. Hanisch also participated in an AURA software planning discussion.

STScI and JHU co-hosted the IVOA Interoperability meeting, 27-31 October. About 100 people participated, with representatives from almost all of the IVOA partner organizations. In preparation for the Interoperability meeting, Greene worked in coordination with the IVOA Data Model group to generate the footprint service use case for Observational data. This work was presented as part of the DAL working group Footprint service session at the IVOA meeting along with the footprint standard specification effort.

Greene and Dower created a prototype TAP (Table Access Protocol) interface for the STScI registry. Greene presented a proposal for the Registry Interface standard specification using TAP as a method for supporting VOTable and simple query parameter sets. This would enable the IVOA standards to move forward in the direction of supporting programmatic interfaces that are more aligned with the current NVO Discovery Portal. The Registry interface design for the Discovery portal was also presented at the IVOA as a "lessons learned" with registry standards for application client development.

Hanisch participated in discussions of the ADQL (Astronomical Data Query Language) specification, which was being brought forward for final review and approval by the IVOA Executive in conjunction with the Interop meeting. Hanisch, A. Rots (SAO), and A. Szalay (JHU) had identified some serious concerns. Fortunately we were able to resolve these issues with some face-to-face discussions during the Interop meeting, and later in the week the IVOA Executive approved promotion of the document to Recommendation.

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

R. Plante continues to chair the IVOA Registry Working Group. The main efforts in this area are the Registry Interfaces document and the VODataService metadata schema extension. Plante also has contributed to the discussions of the Table Access Protocol standard.

B. Baker continues to support single sign-on services for the NVO. Our most active supporting portals today are the NOAO NVO Science Portal (nvo.noao.edu) and the Dark Energy Survey Project Team Portal (des.ncsa.uiuc.edu). Baker has been prototyping an OpenID-compatible version of our login services.

Publications and Presentations

"A Cost-Benefit Study of Doing Astrophysics on the Cloud: Production of Image Mosaics." B. Berriman, E. Deelman, G. Singh, M. Livny, & J. Good. ADASS XVIII Quebec, November 2008.

"An Astronomical Data Resource Registry for the Virtual Observatory." T. Dower, G. Greene, R. Plante, & M. Graham. ADASS XVIII Quebec, November 2008.

"Cross-Matching Astronomical Image Footprints." A. Egana & C. Miller. ADASS XVIII Quebec, November 2008.

"Astronomy 2020—A Pragmatic Approach." M. Graham (invited talk). ADASS XVIII Quebec, November 2008.

"Footprints for the Hubble Legacy Archive." G. Greene, S. Lubow, A. Szalay, & T. Budavari. ADASS XVIII Quebec, November 2008.

"Software Design, Process, and Testing for NOAO VO Portal." C. Miller, A. Egana, R. Massad, & E. Fuentes. ADASS XVIII Quebec, November 2008.

"Mixing the Public and the Private: Security in the Virtual Observatory." R. Plante (invited talk). ADASS XVIII Quebec, November 2008.

"The NOAO Data Products Program: High Quality Science Data and Technologies to Enable the VO and the US OIR System." E. Stobie & R. Seaman. ADASS XVIII Quebec, November 2008.

"An Applications Framework for Astronomical Data Processing and Analysis." D. Tody, P. Grobol, W. Cotton, M. Currie, M. Fitzpatrick, L. Paloro, & C. Surace. ADASS XVIII Quebec, November 2008.

"VOEventNet: A Portal to Astronomical Transients." R. Williams, A. Drake, M. Graham, R. Seaman. ADASS XVIII Quebec, November 2008.

Birds-of-a-Feather Session, "Future Astronomical Data Reduction and Analysis Software." Moderated by R. Hanisch. ADASS XVIII Quebec, November 2008.

The following poster papers were presented at the January 2009 AAS meeting in Long Beach, CA by the NVO Summer School science and technology prize winners:

"Using the NVO to Measure the Distance to Planetary Nebulae from Interstellar Reddening." Kristen A. Larson & S. G. Navarro. AAS 213, #487.07.

"Is the Na D Absorption Line Useful For Integrated Light Stellar Population Studies In Galaxies?" Marcel Bergmann & B. Milvang-Jensen. AAS 213, #444.11.

"Using Virtual Astronomical Observatory Tools for Astronomy 101." Kenneth J. Mighell, K. Garmany, K. Larson, & K. D. Eastwood. AAS 213, #462.03.

"Automating Image Import for Google Sky Using Virtual Observatory Tools." Jared H. Crossley, R. DuPlain, & N. M. Radziwill. AAS 213, #473.04.

"Probing the Quasar Distribution within the Virtual Observatory." Raffaele D'Abrusco, G. Barentsen, O. Laurino, P. Nayak, K. Borne, & G. Longo. AAS 213, # 332.04.

Virtual Observatory Articles in the Popular and Technical Press

Nothing this quarter.

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 a new all-time peak in October 2008 of nearly 662,000 hits and over 548,000 hits in November 2008. These peaks are probably a result of use and take-up of VO tools during and after the NVO Summer School. Web site usage declined in December and January to an unusually low level.

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

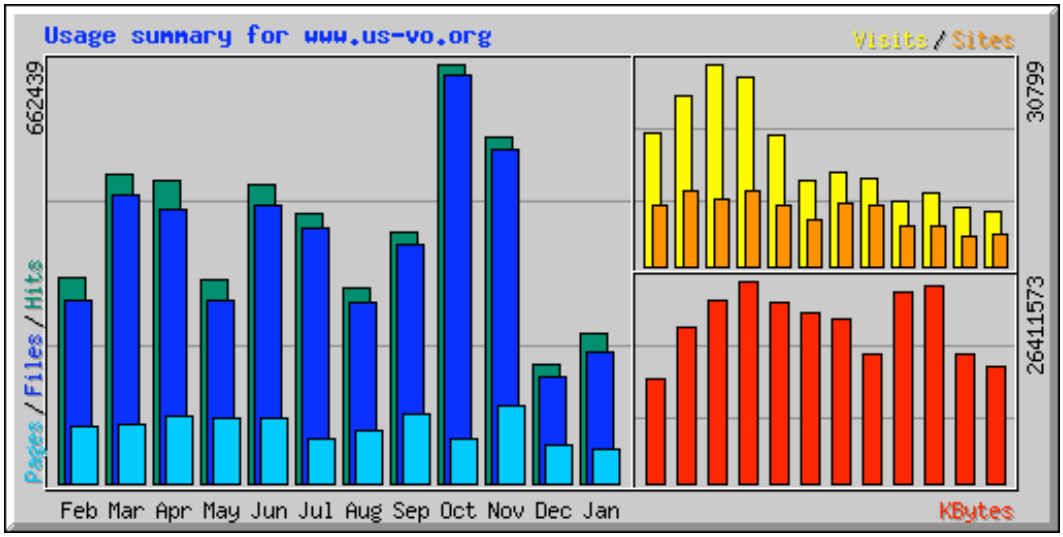

Monthly Hits at NVO Organizations Originating from NVO Applications:

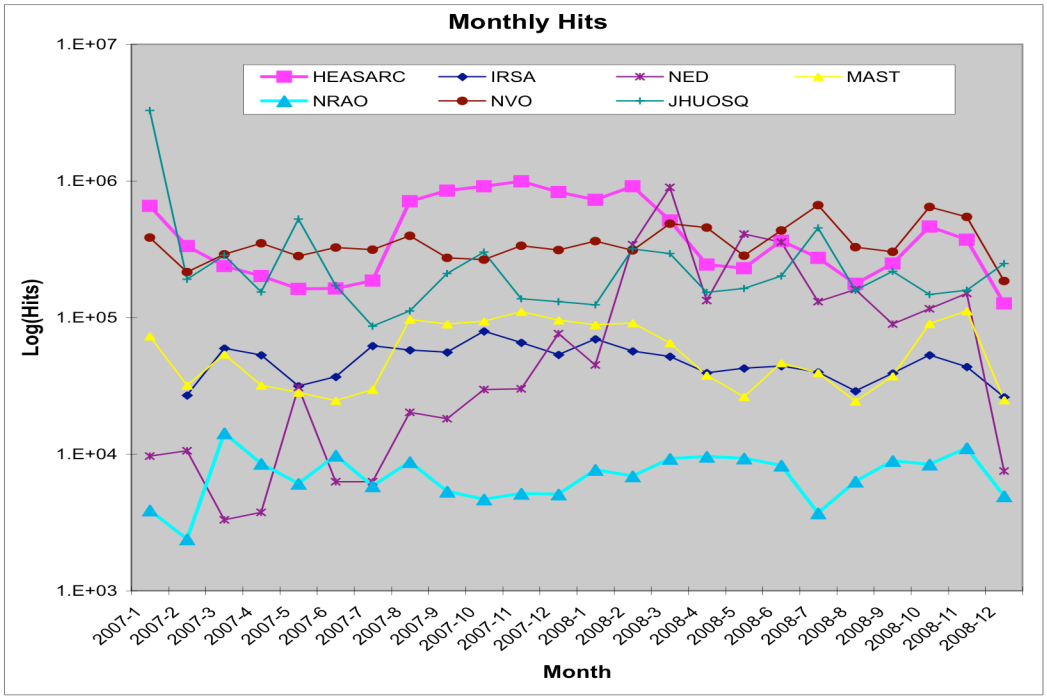

Monthly Users of Any NVO Application or Website:

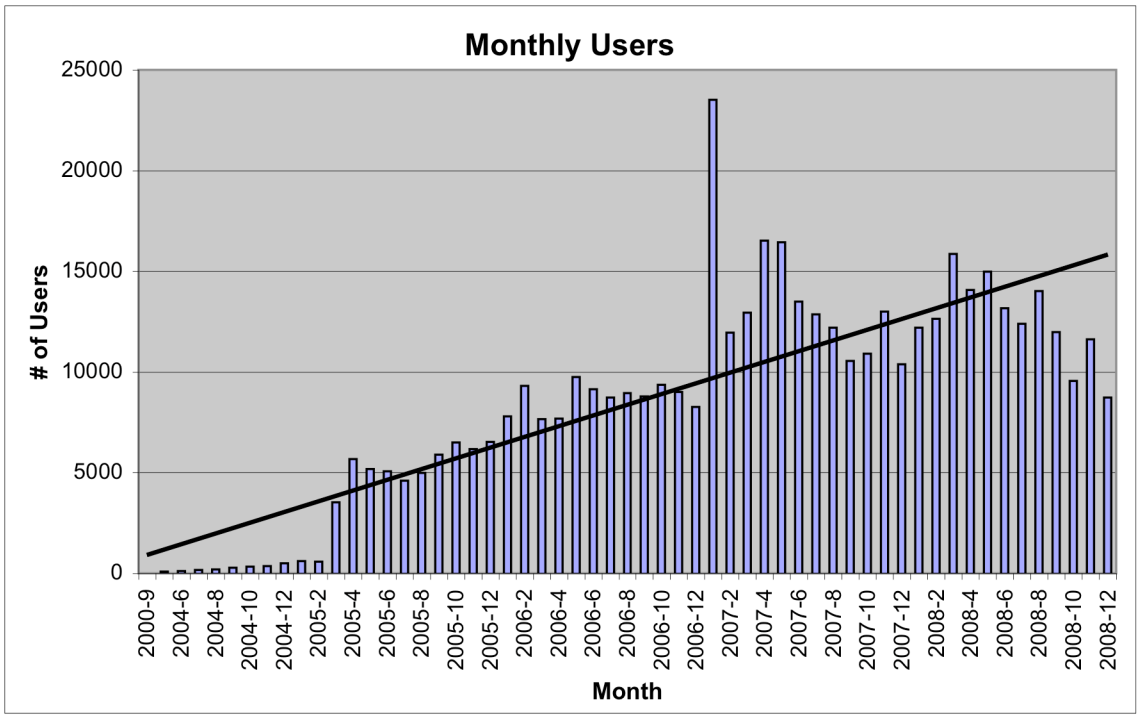

Cumulative Use of NVO Services:

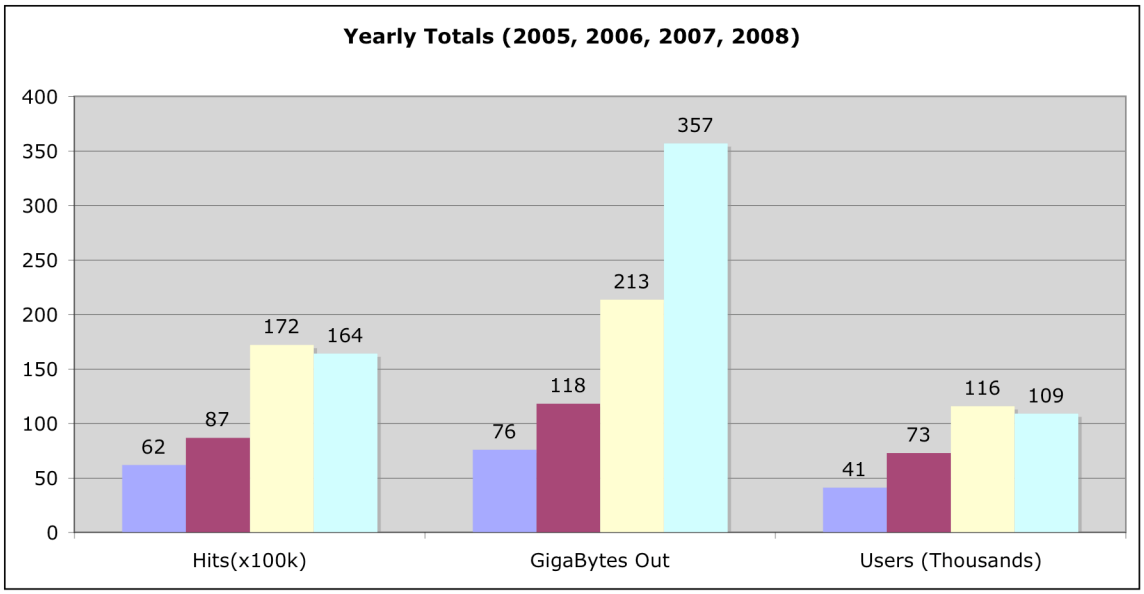

The number of people using NVO tools and services seems to have leveled off, but those users are making substantially more VO-enabled data requests than in previous years (an 80% increase from 2006 to 2007, and a further 67% increase from 2007 to 2008).

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