Description of document:	National Science Foundation (NSF) Operational Review of the Virtual Astronautical Observatory, 2014
Requested date:	01-August-2014
Release date:	26-August-2025
Posted date:	15-September-2025
Source of document:	FOIA Request U.S. National Science Foundation Attn: FOIA Officer 2415 Eisenhower Avenue Alexandria, Virginia 22314 Fax: (703) 292-9041 Email: **foia@nsf.gov** FOIA.gov

NATIONAL SCIENCE FOUNDATION
2415 Eisenhower Avenue
Alexandria, Virginia 22314

OFFICE OF THE GENERAL COUNSEL

August 26, 2025

Via email

Re: NSF FOIA 2014-257

This is in response to your August 1, 2014, Freedom of Information Act (FOIA) request for a copy of the most recent operational review of the Virtual Astronautical Observatory (VAO).

Enclosed are the records responsive to your request. We have conducted a foreseeable harm analysis pursuant to the FOIA Improvement Act of 2016, and we reasonably foresee that disclosure would harm an interest protected by one or more of the nine exemptions to the FOIA general rule of disclosure or disclosure would be prohibited by law. Proprietary information (confidential commercial information) has been withheld under the provisions of exemption (b)(4) of the FOIA. Personal information such as biographical data has been withheld wherever it appears under the privacy protection of exemption (b)(6) (privacy interest) of the FOIA.

Your right of administrative appeal is set forth in Section 612.9 of the NSF FOIA regulation, 45 CFR 612.9, which can be found online at https://www.ecfr.gov/current/title-45/section-612.9 . Your appeal must be postmarked or electronically transmitted within 90 days of the date of this response to your request. You may appeal the decision to deny any portion of your request in writing to: U.S. National Science Foundation, Attn.: General Counsel, 2415 Eisenhower Avenue, Alexandria, Virginia 22314, or by email to foia@nsf.gov.

If you need any further assistance or would like to discuss any aspect of your request, please do not hesitate to contact our FOIA Public Liaison, Ms. Angela Nelson, in any of the following ways:

U.S. National Science Foundation
Attn.: FOIA Public Liaison
2415 Eisenhower Avenue
Alexandria, VA 22314
E-mail: foia@nsf.gov
Telephone: 703-292-8060

Additionally, you may contact the Office of Government Information Services at the National Archives and Records Administration to inquire about the FOIA Mediation Services they offer. The contact information is as follows: Office of Government Information Services, National Archives and Records Administration, 8601 Adelphi Road-OGIS, College Park, Maryland 20740-6001, email address: ogis@nara.gov; telephone: (202) 741-5770; toll free number: 1-877-684-6448; fax: (202) 741-5769.

There is no fee for FOIA services in this instance, in accordance with 5 U.S.C. 552(a)(4)(A)(i) et seq.

Sincerely,

Spencer D. Christian
FOIA/Privacy Act Officer

Enclosures:

145-pages

Project Report Printer Friendly Version

Click here to print

Preview of Award 0834235 - Annual Project Report

Cover

Federal Agency and Organization Element to Which Report is Submitted:	**4900**
Federal Award or Other Identifying Number Assigned by Agency:	**0834235**
Project Title:	**Management and Operation of the Virtual Astronomical Observatory**
PD/PI Name:	**Ethan Schreier, Principal Investigator** **David De Young, Co-Principal Investigator** **Giuseppina Fabbiano, Co-Principal Investigator** **Robert Hanisch, Co-Principal Investigator** **Alexander S Szalay, Co-Principal Investigator**
Recipient Organization:	**Virtual Astronomical Observatory, LLC**
Project/Grant Period:	**05/15/2010 - 04/30/2015**
Reporting Period:	**05/01/2013 - 04/30/2014**
Submitting Official (if other than PD\PI):	**Robert Hanisch** **Co-Principal Investigator**
Submission Date:	**05/09/2014**
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	**Robert Hanisch**

Accomplishments

*** What are the major goals of the project?**

Develop, operate, and disseminate information about the Virtual Astronomical Observatory

See attached report for full details.

*** What was accomplished under these goals (you must provide information for at least one of the 4 categories below)?**

Major Activities:	See attached report.
Specific Objectives:	See attached report.
Significant Results:	See attached report.
Key outcomes or Other achievements:	See attached report.

*** What opportunities for training and professional development has the project provided?**

Nothing to report.

*** How have the results been disseminated to communities of interest? If so, please provide details.**

Website, conference presentations, youtube video tutorials, twitter and facebook. Details in attached report.

*** What do you plan to do during the next reporting period to accomplish the goals?**

Conduct final close-out review, identify any outstanding issues/deliverables, and wrap up technical work by 9/30/2014.

Supporting Files

Filename	Description	Uploaded By	Uploaded On
VAO-AR-April2014.pdf	Annual Report for the period May 2013-April 2014.	Robert Hanisch	05/02/2014

Products

Books

Book Chapters

Inventions

Journals or Juried Conference Papers

Licenses

Other Conference Presentations / Papers

Other Products

Other Publications

Patent Applications

Technologies or Techniques

Thesis/Dissertations

Websites or Other Internet Sites

Virtual Astronomical Observatory

Website for all software products and services provided by the VAO.

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
	PD/PI	(b)(4); (b)(6)
	Co PD/PI	
	Co PD/PI	

Name	Most Senior Project Role	Nearest Person Month Worked
	Co PD/PI	(b)(4): (b)(6)
	Co PD/PI	

Full details of individuals who have worked on the project:

Ethan Schreier
Email: ejs@aui.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: (b)(4): (b)(6)
Contribution to the Project: member of VAO Board of Directors
Funding Support: none drawn from this project
International Collaboration: No
International Travel: No

David De Young
Email: deyoung@noao.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4): (b)(6)
Contribution to the Project: none; Dr. De Young passed away two years ago.
Funding Support: none
International Collaboration: No
International Travel: No

Giuseppina Fabbiano
Email: gfabbiano@cfa.harvard.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4): (b)(6)
Contribution to the Project: chair of Science Council
Funding Support: none from this award
International Collaboration: No
International Travel: No

Robert Hanisch
Email: hanisch@stsci.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4): (b)(6)
Contribution to the Project: Director
Funding Support: this award
International Collaboration: No
International Travel: Yes, Germany - 0 years, 0 months, 5 days; Ireland - 0 years, 0 months, 3 days

Alexander S Szalay
Email: szalay@pha.jhu.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4): (b)(6)
Contribution to the Project: technical advisor
Funding Support: this award
International Collaboration: No
International Travel: No

What other organizations have been involved as partners?

Nothing to report.

What other collaborators or contacts have been involved?

NO

Impacts

What is the impact on the development of the principal discipline(s) of the project?

See attached report.

What is the impact on other disciplines?

See attached report.

What is the impact on the development of human resources?

Nothing to report.

What is the impact on physical resources that form infrastructure?

See attached report.

What is the impact on institutional resources that form infrastructure?

See attached report.

What is the impact on information resources that form infrastructure?

See attached report.

What is the impact on technology transfer?

Nothing to report.

What is the impact on society beyond science and technology?

Nothing to report.

Changes/Problems

Changes in approach and reason for change

Nothing to report.

Actual or Anticipated problems or delays and actions or plans to resolve them

Nothing to report.

Changes that have a significant impact on expenditures

Nothing to report.

Significant changes in use or care of human subjects

Nothing to report.

Significant changes in use or care of vertebrate animals

Nothing to report.

Significant changes in use or care of biohazards

Nothing to report.

Special Requirements

Responses to any special reporting requirements specified in the award terms and conditions, as well as any award specific reporting requirements.

Nothing to report.

Project Report Printer Friendly Version

Click here to print

Preview of Award 0834235 - Final Annual Project Report

Cover

Federal Agency and Organization Element to Which Report is Submitted:	**4900**
Federal Award or Other Identifying Number Assigned by Agency:	**0834235**
Project Title:	**Management and Operation of the Virtual Astronomical Observatory**
PD/PI Name:	**Ethan Schreier, Principal Investigator Alexander S Szalay, Co-Principal Investigator Robert Hanisch, Co-Principal Investigator David De Young, Co-Principal Investigator Giuseppina Fabbiano, Co-Principal Investigator**
Recipient Organization:	**Virtual Astronomical Observatory, LLC**
Project/Grant Period:	**05/15/2010 - 04/30/2015**
Reporting Period:	**05/01/2014 - 04/30/2015**
Submitting Official (if other than PD\PI):	**Ethan Schreier Principal Investigator**
Submission Date:	**07/30/2015**
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	**Ethan Schreier**

Accomplishments

*** What are the major goals of the project?**

Develop, operate, and disseminate information about the Virtual Astronomical Observatory

See attached report for full details.

*** What was accomplished under these goals (you must provide information for at least one of the 4 categories below)?**

Major Activities:	See attached report.
Specific Objectives:	See attached report.
Significant Results:	See attached report.
Key outcomes or Other achievements:	See attached report.

*** What opportunities for training and professional development has the project provided?**

See attached report.

*** How have the results been disseminated to communities of interest? If so, please provide details.**

See attached report.

Supporting Files

Filename	Description	Uploaded By	Uploaded On
VAO-Final Report 07-30-2015.pdf	VAO Final Report	Ethan Schreier	07/30/2015

Products

Books

Nothing to report.

Book Chapters

Nothing to report.

Inventions

Nothing to report.

Journals or Juried Conference Papers

Nothing to report.

Licenses

Nothing to report.

Other Conference Presentations / Papers

Nothing to report.

Other Products

Other Publications

Nothing to report.

Patent Applications

Nothing to report.

Technologies or Techniques

Nothing to report.

Thesis/Dissertations

Nothing to report.

Websites or Other Internet Sites

Nothing to report.

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
	PD/PI	(b)(4) (b)(6)
	Co PD/PI	
	Co PD/PI	
	Co PD/PI	
	Co PD/PI	

Full details of individuals who have worked on the project:

Ethan Schreier
Email: ejs@aui.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: (b)(4) (b)(6)
Contribution to the Project: See attached report.
Funding Support: See attached report.
International Collaboration: No
International Travel: No

David De Young
Email: deyoung@noao.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4) (b)(6)
Contribution to the Project: See attached report.
Funding Support: See attached report.
International Collaboration: No
International Travel: No

Alexander S Szalay
Email: szalay@pha.jhu.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4) (b)(6)
Contribution to the Project: See attached report.
Funding Support: See attached report.
International Collaboration: No
International Travel: No

Robert Hanisch
Email: hanisch@stsci.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4) (b)(6)
Contribution to the Project: See Attached Report
Funding Support: See Attached Report
International Collaboration: No
International Travel: No

Giuseppina Fabbiano
Email: gfabbiano@cfa.harvard.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: (b)(4) (b)(6)
Contribution to the Project: See Attached Report
Funding Support: See Attached Report
International Collaboration: No
International Travel: No

What other organizations have been involved as partners?

Nothing to report.

What other collaborators or contacts have been involved?

See attached report.

Impacts

What is the impact on the development of the principal discipline(s) of the project?

See attached report.

What is the impact on other disciplines?

See attached report.

What is the impact on the development of human resources?

See attached report.

What is the impact on physical resources that form infrastructure?

See attached report.

What is the impact on institutional resources that form infrastructure?

See attached report.

What is the impact on information resources that form infrastructure?

See attached report.

What is the impact on technology transfer?

See attached report.

What is the impact on society beyond science and technology?

See attached report.

Changes/Problems

Changes in approach and reason for change

See attached report.

Actual or Anticipated problems or delays and actions or plans to resolve them

See attached report.

Changes that have a significant impact on expenditures

See attached report.

Significant changes in use or care of human subjects

Nothing to report.

Significant changes in use or care of vertebrate animals

Nothing to report.

Significant changes in use or care of biohazards

Nothing to report.

Special Requirements

Responses to any special reporting requirements specified in the award terms and conditions, as well as any award specific reporting requirements.

Nothing to report.

VAO

For the Period:

1 May 2013
through
30 April 2014

Associated Universities, Inc

The Virtual Astronomical Observatory (VAO) is managed by the VAO, LLC, a non-profit company chartered in the District of Columbia and established as a partnership of the Associated Universities, Inc. and the Association of Universities for Research in Astronomy, Inc. The VAO is sponsored by the National Science Foundation and the National Aeronautics and Space Administration.

Submitted to the National Science Foundation
Pursuant to Cooperative Agreement
AST-0834235

Contents

Executive Summary

The VAO has focused its activities on executing the VAO Project Execution Plan (PEP) for the period May 1, 2013 and through September 30, 2014, approved on June 8, 2013. Briefly, the plan calls for completion of all deliverables already underway, continued monitoring of services, preparation to deliver VAO assets to the community and delivery of a protocol and data discovery and data access service for multi-dimensional data sets.

The VAO hosted the semi-annual interoperability workshop in Waikoloa, HI, consecutively with the annual ADASS conference, and has developed an agenda for the VAO close-out review, scheduled for July 10 and 11, 2014 in Pasadena, CA, and stakeholders in the project have been invited to attend. VAO management consulted with NASA archive staff to support their proposal to sustain the VAO infrastructure.

VAO staff has been active in the community: they have made presentations at conferences, and have participated on panel discussions and workshops. VAO staff holds leadership roles in the IVOA, and have contributed to the development of five IVOA standards. VAO personnel have also been active in the Research Data Alliance (RDA), an international initiative to foster data access and exchange throughout the research community.

The final versions of the Data Discovery tool, Iris (the Spectral Energy Distribution service) and the Cross-comparison service have been delivered. These releases offer new functionality for users and correct several defects that have been reported. The new releases include updated user documentation and video tutorials.

Repositories have been set up to host the VAO code base and documentation with Google's free cloud-based services. The code repository has been populated with essentially all versions of software components delivered to date, and new components are transferred there each month as they are delivered. The documentation repository has been populated with user guides and tutorials, project reports, professional outreach materials, and a list of technical publications and presentations by team members.

The VAO has begun the transfer of IVOA communication assets (web site, wiki and mailing lists) to VOIndia and VOItaly. The VAO Program Manager is coordinating these efforts on behalf of the IVOA Executive Council.

VO science services were available 98% of the time, even with the shutdown of the federal government. The VAO validation services detected a dramatic decrease in validation statistics, which was attributed to changes in Vizier data services. The VAO informed Vizier staff and the problem was corrected. VO infrastructure requests remained steady during 2013 and early 2014, except for a surge in Registry queries in late 2103. Between March 30, 2013 and April 11, 2014, 327 tickets were filed with the JIRA ticket system and 245 were closed.

A major effort for the VAO is support for discovery of and access to multidimensional data sets ("data cubes'), identified as a key science initiative by the IVOA. The VAO has been providing leadership in this endeavor. A white paper delivered in May 2013 described the mechanisms and science drivers for serving data cubes, and led to the development of a prototype data access that served data from a range of projects. Concurrently, the VAO proposed modifications of existing standards and data models that were mandated by this service, and a successful demonstration at the Winter IVOA meeting led to agreement on the necessary requirements to support data cube. VAO staff is working with international partners to present mature versions of the Image Data Model, Simple Image Access Protocol v2, and DataLink at the Spring 2014 meeting.

The VAO established a partnership with the NRAO to support scalable access to the large datasets from ALMA and JVLA, and to support seamless access to their visualization toolkit, CASA. To date, the VAO has delivered an enhanced DALServer to provide access to data cubes, and NARO has integrated the VOClient tools into the CASA software.

The SciDrive data sharing and publication service has been made accessible to users for evaluation. It provides 100-TB storage with triple redundancy, and includes the Single Sign On service. Upgrades to the registry and data publishing tool are being readied for release.

A production release of the VOClient, which offers command line access to the VO, was deployed in January 2014, and two beta releases of the native Python package, PyVO, were released through the Astropy community-supported platform.

1 Management

1.1 Reporting and Communications

Regular weekly team telecons (Fridays) and weekly management telecons (Mondays) continued throughout the year, as well as the weekly telecons for each of the major work areas (Operations, Documentation Services, Science Applications, and Standards and Infrastructure).

R. Hanisch and B. Berriman represent the VAO at monthly telecons of the Astronomy Data Centers Executive Committee (ADEC), which coordinates activities among the NASA-funded data centers.

VAO Team Members participate in a biweekly telecon with members of the Atacama Large Millimeter/submillimeter Array (ALMA) software development and archive development team at the National Radio Astronomy Observatory (NRAO) to review progress in the development of an operational service that supports discovery of, and access to, multidimensional data sets from the Jansky Very Large Array (JVLA) and ALMA within their Common Astronomy Software Applications (CASA) software.

Requirements and design reviews for project deliverables are generally carried out by WebEx. These reviews are led by the development leads and are attended by representatives of management, User Support and Operations, as well as by the development team itself.

The Configuration Change Board (CCB) meets biweekly by telecom or WebEx as needed to review requests for changes to released software and documentation. The CCB is chaired by the Program Manager, and the board consists of the task leads of the work areas. All members of the VAO team are welcome to attend the CCB meetings.

1.2 Meetings

The project held two team meetings. The first meeting was held on June 10 and 11, 2013 in Annapolis MD. The second meeting was hosted by NOAO on January 30 and 31, 2014, in Tucson, AZ.

In their capacity as members of the Advisory Board of the Astronomy Source Code Library, R, Hanisch and B. Berriman were co-organizers of, and presenters at, a special session at the 223rd AAS meeting (January 2014) meeting, "Astrophysics Code Sharing II: The Sequel." Berriman, G. Djorgovski, C. Donalek, and M. Graham (Caltech), and A. Szalay and T. Budavari (JHU) gave presentations at the AstroInformatics conference in Sydney, Australia, December 9- 12.

Members of the VAO team presented four posters at the 223rd AAS meeting: "Virtual Astronomy: The Legacy of the Virtual Astronomical Observatory," "Accessing Multi-Dimensional Images and Data Cubes in the Virtual Observatory," "Data publication and Sharing Using the SciDrive Service," and "The Virtual Observatory for the Python Programmer."

Alex Szalay gave an invited presentation entitled "From AISR to the Virtual Observatory" at the session "Building the Astronomical Information Sciences: From NASA's AISR Program to the New AAS Working Group on Astroinformatics and Astrostatistics."

Hanisch and Berriman attended the workshop on "Metadata Semantics and Data Formats," organized by the AAS and held concurrently with the AAS meeting. They emphasized how VO services can support publishing of data through peer-reviewed journals.

In their capacity as members of the Advisory Board of the Astronomy Source Code Library, R, Hanisch and B. Berriman were co-organizers of a Birds Of A Feather session at ADASS XXIII entitled "Ideas for Advancing Code Sharing." Hanisch attended the Second Plenary Meeting of the Research Data Alliance, September 16-18, 2013, at the National Academy of Sciences in Washington, DC, and the Third Plenary Meeting of the RDA, March 26-28, 2014, in Dublin, Ireland. Hanisch now co-chairs the RDA Interest Group on Domain Repositories and thus has become a member of the RDA-US Advisory Committee.

R. Hanisch and J. Lazio attended *Radio Astronomy in the Era of LSST*, a meeting held at NRAO-Charlottesville, May 6-8, 2013, where Hanisch gave a presentation "Virtual Astronomy in the Era of LSST."

Hanisch and A. Szalay were invited to a workshop, *Sustaining Domain Repositories*, sponsored by the Sloan Foundation and held at the University of Michigan, Ann Arbor, June 24-25, 2013.

G. Djorgovski and A. Mahabal gave presentations at the conference *Interface 2013* held at Chapman University, Orange, CA, April 4-6, 2013.

R. Plante gave a presentation at the SciPy 2013 conference in Austin, TX, June 26, 2013.

Hanisch and Plante gave presentations to the NanoGrav consortium on November 6, 2013 describing the VO and the Python interfaces.

Hanisch described the VO infrastructure in a presentation to a small group of theoretical physicists and neuroscientists organized by the Kavli Foundation (Santa Monica, CA, February 19, 2014).

1.3 International Collaboration

The VAO hosted and participated in the IVOA Interoperability Workshop, held in Waikoloa, Hawaii, just prior to the annual ADASS Conference. We also supported publication of the IVOA Newsletter, distributed on September 18, 2013. A major element of that meeting was presentation of the VAO data cubes data access service, and a review of the updates to the Simple Image Access Protocol v2 (SIAPv2) protocol and the Image Data Model.

Members of the VAO collaboration attended the International Virtual Observatory Alliance (IVOA) Interoperability Meetings ("Interop") in Heidelberg, Germany, held May 12-17, 2013, and hosted by the German Astrophysical Virtual Observatory (GAVO). The Interoperability meetings are the focal point for discussion of the international VO standards and increasingly a venue for collaboration and consultation on VO-enabled research tools. Several VAO team members have leadership roles on IVOA Working Groups and Interest Groups. Hanisch, Lazio, and Fabbiano participate in the meetings of the IVOA Executive.

Work has begun on transferring the IVOA web site, document repository, Wiki, and e-mail distribution lists to other IVOA partners. VO-India has agreed to take on primary responsibilities for the document repository and web site, and VO-Italy will maintain the wiki, e-mail lists, and the IVOA domain name. SAO will be responsible for assembly and editing of the IVOA Newsletter. B. Berriman is coordinating this work on behalf of the IVOA Executive Council. He has developed a schedule that delivers the assets to the partners in time for the VAO close-out meeting, and he delivers bi-weekly reports to the Chair of the IVOA Executive Council.

1.4 Project Schedule and Milestones

The table below summarizes the milestones planned and accomplished in the reporting period..

Milestone	Delivered	Scheduled
Management		
1st Quarterly Report	July 2013	July 2013
2nd Quarterly Report	October 2013	October 2013
3rd Quarterly Report	January 2014	January 2014
Annual Report (this document)	April 2014	April 2014
Close-out Review Agenda	April 2014	April 2014
Standards and Infrastructure		
Spring IVOA Meeting	May 2013	May 2013
Winter IVOA Meeting	October 2013	October 2013
Python API	February 2014	December 2013
Upgrade Registry	October 2013	October 2013
Sci Drive released to astronomers at JHU	October 2013	July 2013
Command Line VO Services	December 2013	October 2013

Milestone	Delivered	Scheduled
PyVO Release	January 2014	October 2013
VO Client Implementations in C Test Release	January 2014	October 2013
NRAO Review of Client API for Access To Data Cubes	December 2013	December 2013
White Paper: Mapping Use Cases On To Proposed Standards	May 2013	May 2013
Demonstration data access and discovery service for data cube	September 2013	September 2013
VAO Single Sign On Services	May 2013	April 2013
Deliver Registry Update to Test	May 2013	April 2013
Deliver Publishing Interface To Test	May 2013	April 2013
VO Client C Delivered to test	July 2013	May 2013
Science Applications		
Archive software in repository	February 2014	December 2013
Data Discovery Tool v 1.5	June 2013	May 2013
SED Builder (Iris) v.1.2	July 2013	May 2013
Operations		
Monthly Operations Report	Monthly	Monthly
Deploy Close-out Code Repository	June 2013	May 2013
Archive SVN Repository	July 2013	June 2013
Documentation Support		

Milestone	Delivered	Scheduled
Updated videos for the Science Applications	May 2014	December 2013
Identify IVOA Partner to Curate Assets	February 2014	December 2013
Include initial set of documents in repository	July 2013	July 2013
Transfer Material from Presentations, Community Days, etc. to Repository	November 2013	November 2013
Deploy Close-out Documentation Repository	June 2013	May 2013

1.5 Close-Out Activities

Formal approval of the VAO Project Management Plan (PEP) for the period May 1, 2013 through September 30, 2014, was received on June 8, 2013.

NASA directed the archives at STScI (MAST), GSFC (HEASARC) and IPAC (IRSA and NED) to develop a proposal to support the essential components of the VAO-developed infra- structure as part of their core ("in-guide") activities, beginning in October 2014. This proposal was submitted to NASA on November 8, 2013, and a technical and managerial response to the review of the proposal was submitted on April 4, 2014. The partners in the proposal consulted with Hanisch and Berriman on the essential VAO components that can be supported within their core budgets, and have developed a transition plan to transfer the necessary VAO assets to them. The transition plan is straightforward: the VAO web page will be transferred from CACR to IPAC, whereas the Registry and the Operational Services will be sustained *in situ* at STScI and HEASARC by staff who currently support the VAO. JHU has agreed to support SciDrive, and the VAO is negotiating with JHU to deploy a mirror of the VAO registry.

A formal close-out review and final team meeting have been scheduled for July 10-11, 2014 at IPAC/Caltech. NSF and NASA representatives and members of the VAO Board have been invited to participate.

1.6 Personnel

There have been no major changes in VAO key personnel during the past year. The key personnel are:

Director	Robert Hanisch (STScI)
Program Manager	Bruce Berriman (IPAC)
Project Scientist	Joseph Lazio (JPL)
Technology Advisor	Alexander Szalay (JHU)
Business Manager	Maricel Claro (AUI)
Operations Lead	Thomas McGlynn (HEASARC)
Documentation Services	Sarah Emery Bunn (Caltech)
Science Applications Lead	Janet Evans (SAO)
Standards and Infrastructure Lead	Ray Plante (NSCA)

R. Hanisch will be leaving the project at the end of July 2014 to assume the position of Director of the Office of Data and Informatics, Materials Measurement Laboratory, National Institute of Standards and Technology in Gaithersburg, Maryland. He has arranged with NIST and cleared with federal ethics lawyers the ability to carry out fiducial responsibilities of the VAO, LLC as needed. E. Schreier (AUI) will be the formal representative of the VAO project to NSF effective July 27, 2014.

Over the past year approximately 12 FTE was spent on VAO work (NSF and NASA combined); we estimate another 1-2 FTE in in-kind effort. 70% of the effort was from men and 30% of the effort was from women.

1.7 Financial Status

For the period May 2013—March 2014 the VAO had expenses of $1.82M (NSF funding), with average monthly spending of $165K. As of March 2014, VAO has a balance of $1.720k (NSF funding); this includes the release of the final project funding of $1M from NSF on January 30th, 2014 and a budget carry-forward of $538K from prior years. Funding will be sufficient to complete all planned technical activities and to carry out all administrative matters associated with the close-out of the Cooperative Agreement with NSF.

The VAO, LLC completed its annual external audit of financial statements and A133 compliance for FY2013, which was conducted by BDO USA, LLP and presented to the VAO Board Finance and Audit Subcommittee on January 28, 2014, and approved by the full Board on April 8, 2014. BDO issued an unqualified opinion with no findings or questioned costs.

2 Operations

2.1 Highlights

VAO operations turned towards the establishment and population of facilities that will enable the legacy of the VAO to be used after the termination of funding at the end of FY2014. A VAO Closeout Repository was designed and populated during this period. This combines several elements of Google's free cloud services (GoogleCode, GoogleSite, YouTube and GoogleDrive) to enable the VAO to preserve all digital artifacts in a fashion easily accessed by users at very low cost. A GoogleSite for the US Virtual Observatory (http://sites.google.com/site/usvirtualobservatory) serves as the entry point in the repository. A GoogleCode code repository is used to mirror the operational VAO repository and the GoogleDrive and YouTube are used to save documentation. During the past quarter an operational script to update the code elements of the closeout repository on a weekly basis has been implemented. This will continue until the termination of VAO activities.

The GoogleSite includes pages for each of the major development efforts as well as summaries of the documentation and other artifacts. Many pages are already completed. There are at least placeholders for each software project.

During the year significant hardware upgrades were made at ST and JHU, the later in response to a catastrophic failure of the servers on which the logging services were hosted.

The upgrade to TAP compatible interfaces to the registry was nearly completed at the end of the first quarter of calendar 2014 and should be completely finished in the second quarter.

While there were several issues that caused interruptions in non-science VAO services, notably the hardware failure that affected the logging services and a recurrent issue with the validation services, there was only one significant problem with the science services during the past year. The data query services hosted at the HEASARC and used in the VAO portal were disabled during the shutdown of the federal government. However even including this major disruption science services averaged over 98% availability.

Late in 2013 a dramatic decrease in the validation statistics for the VO as whole was traced to changes in the Vizier data services. The responsible party was informed and a fix was quickly in place. However VAO monitoring personnel noted that the initial response addressed only some of the services and provided additional support noting other services that were still failing. A full recovery has since been effected.

No significant changes to infrastructure services were released in this period although there were many incremental changes. A test version of the new notification service, which uses the VAO SSO (Single Sign-On) service to authorize users to send notifications of downtimes for VO services has been released at the HEASARC. A public release awaits the finalization of the SSO architecture. This will allow the notification service to be useful generally throughout the VO after the VAO termination.

2.2 Science Operations

Despite the disruption of the federal shutdown, usage of VAO science services remained stable during the past year for most elements. A dramatic surge in Registry service queries caused a high plateau in science requests during late 2013 but these returned to their normal levels in early 2014. Missing data from CfA and MAST slightly depresses the last two data points for the science services.

A gradual increase in VAO infrastructure requests has leveled out since the summer of 2013. The number of VO data requests from VO resources hosted at VAO institutions (MAST, IRSA and the HEASARC) has continued stable.

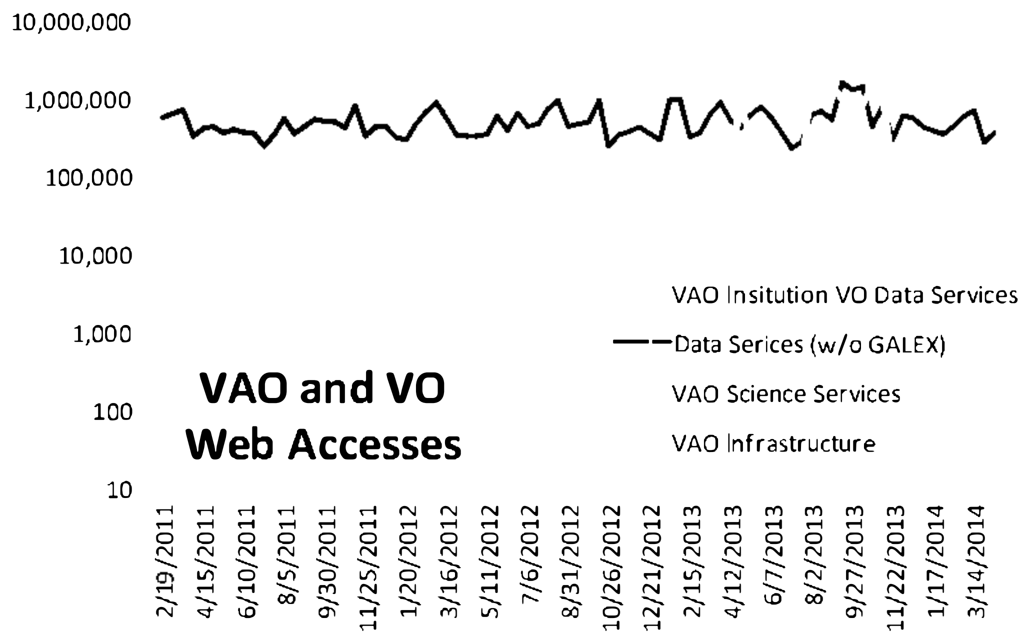

Figure 2-1. *Usage of VAO science services.*

2.3 Service Monitoring and Validation

The only major interruption to the VAO's science services during the past year was due to the federal shutdown, with average availability of science services declining to below 90%. The HEASARC-hosted DataScope and SimpleQuery capabilities, used by the Portal and Registry respectively, were unavailable for several weeks. Otherwise science services had excellent availability. A number of significant downtimes in non-science services occurred due to interruptions at JHU and NCSA, but these have limited impact on the science community.

The average availability for science services over the past year was 98.9%. Figure 2-2, below, gives the uptime for VAO services.

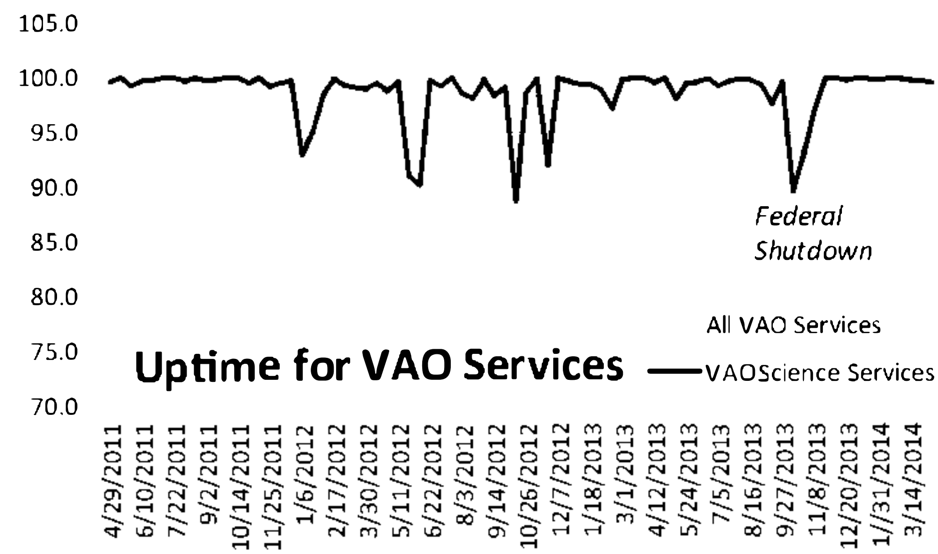

Figure 2-2. *Uptime for VAO services.*

Validation of deployed VO data services also continued. During the last quarter of 2013 the pass-rate for VizieR services plunged towards 0 and we informed the CDS of the issue. They quickly addressed part of the problem, but we noted continue problems and provided links to services that continued failing. With it taking a few weeks for the CDS to completely resolve the problem and a one-month cycle before services were revalidated it took well into 2014 before we achieved a complete recovery of the validation statistics. This incident is reminiscent of a similar problem detected in March 2011 shown at the beginning of the graph, but is even more dramatic.

The gap in the validation numbers for the VAO services (Fig. 2-3) is due to the federal shutdown.

The validation statistics for VAO institution hosted services (at the HEASARC, MAST, CfA, IRSA, and NOAA) were stable during the year.

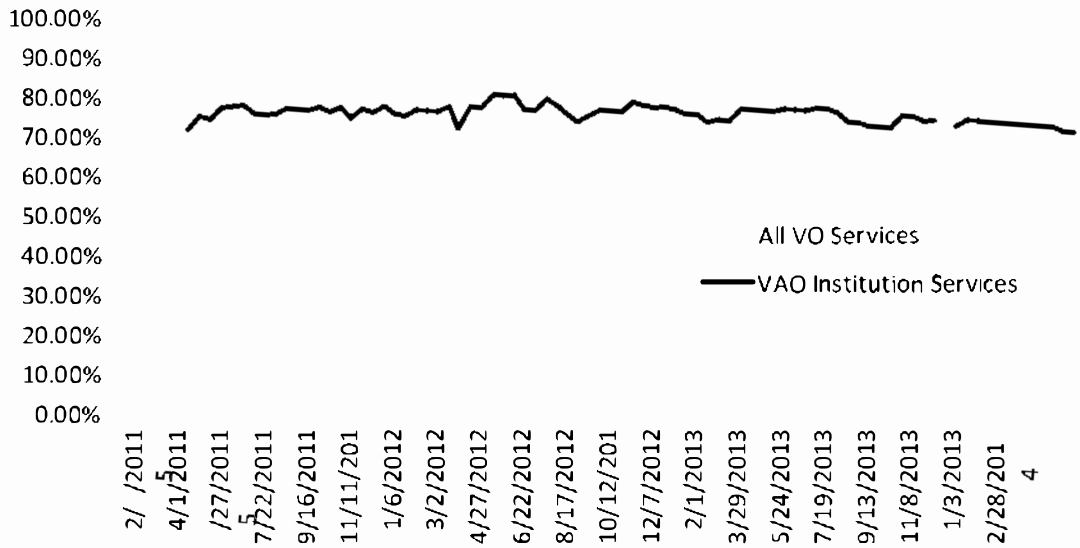

Figure 2-3. *Validation success rate for VO and VAO services.*

2.4 Internal Operations

The internal operations of the VAO are also a critical area for operations. The Web site, Wiki, Trac site, test environment and SVN repository are all maintained and tested.

A weekly telecon has been held for several years to facilitate rapid interaction and resolution of any operational issues that arise. Minutes for these telecons are kept.

Between March 30, 2013 and April 11, 2014, 327 JIRA tickets were opened and 245 were closed. During the past six months most ticket activity has been in service interruptions (labeled Ops in the table below) since the lessened rate of software development means that we had few bug reports or enhancement requests for our software tools.

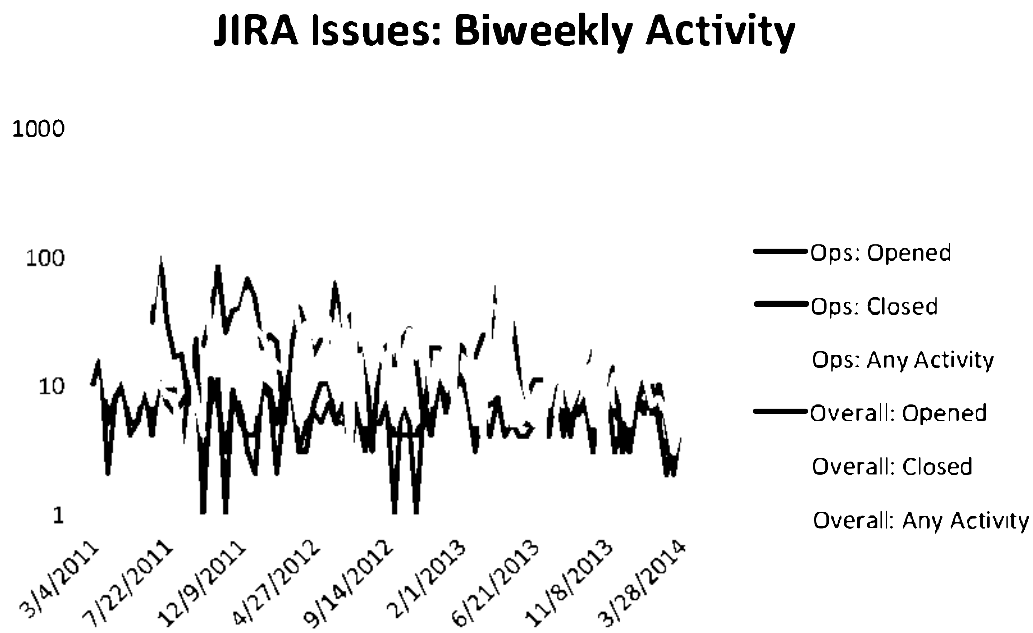

Figure 2-4. *JIRA ticket activity in biweekly periods.*

The SVN repository is maintained to provide a repository for all VAO developed code and data. The pace of commits is down substantially during the current year compared to previous years. This is anticipated as software development activities continue to wrap up.

3 Documentation Services

The VAO Documentation Services Group provides the primary interface between the VAO and its user community. The VAO assembles and curates documentation from across the project, makes the information easy to find through the web site, and notifies the community of updates and changes relevant to the VO infrastructure.

3.1 End-User Documentation

This year's work focuses on the documentation section of the VAO's closeout repository. A documentation area of this repository is populated with tutorials and user guides for VAO tools & services, outreach materials from "VO Days" held in 2011 and 2012, project documents and reports. Also included is a list of technical presentations and research publications by VAO Team members.

As part of the closeout, we will also add final packages of the VAO web sites, wiki, and mailing list archives to this repository. More details about the closeout repository can be found in the Operations section of this report.

3.2 Web Sites

During the period January-April 2014, there were five news items posted to the VAO web site (http://www.usvao.org). Over the course of the past year, there have been a total of 26 news items posted. These news items, summarizing Virtual Observatory news for the astronomical research community, are also pushed to subscribers via RSS, Facebook (164 followers), and Twitter (172 followers). The number of followers on these social media sites continues to grow steadily.

Work continues on the transfer of IVOA communication assets (web site, wiki, and mailing lists) from hosting at a VAO organization (Caltech CACR) to other organizations (VO Italy and VO India). The transfer of these assets is scheduled to be complete on June 30, 2014.

We are currently in the process of scheduling transfer of VAO communication assets (web site and mailing lists) from Caltech CACR to Caltech IPAC. The handover of these assets is scheduled to be completed by the end of May 2014.

3.3 Help Desk and User Forum

Documentation Services staff monitor and respond to VAO Help Desk inquiries and refer them to other team members for resolution/answers. Statistics on help desk tickets can be found in the Operations section of this report.

The VAO Forum (Astrobabel.com) is maintained and backed up by Documentation Services staff. During the past quarter there have been 1,187 page views of forum discussions. The total page views for the year are 6,704. A list of support requirements has been developed for discussions about hosting of the forum post-closeout.

4 Science Applications

The activities in the science applications task area over the past year focused on conclusion of the VAO development efforts and closeout activities related to Science Applications as mandated by the Program Execution Plan (PEP).

The group focused on adding useful capabilities to the applications and addressed high priority bug fixes in the first half of the year. Public releases of the Catalog Cross Comparison, Data Discovery Tool, and Iris were completed. The Science Applications release summary is provided in Table 1.

The group turned attention to closeout activities in the second half of the year. Priorities included updated user documentation, updated user training videos, and science application contributions to the VAO close out repository.

Tool/Application	Version	Release Date	New Features
Catalog Cross Comparison	V1.2	Apr 27	Replaced SDSS Rel 7 (DR7) catalog with DR9
Data Discovery Tool	V1.5	Jun 21	Replaced flash in Astroview, Metadata preservation in table data, scatter plots, footprint accuracy improvement
Iris	V2.0	Jul 01	Co-plotting, redshift/blueshift, interpolation, integrated quantities, smoothing
Iris	2.0.1	Dec 03	2 bug fixes: conversion of magnitude errors, text double printout

Table 4-1. Science Applications release highlights.

4.1 Data Discovery Tool

The Data Discovery tool (DDT) is a web application for discovering all resources that are known to the VAO about an astrophysical object or a region of the sky. Using protocols defined by the International Virtual Observatory Alliance (IVOA), the DDT searches those widely distributed resources and presents the results in a single unified web page.

Many of the most popular US archives and catalog holdings are available for searches in DDT, including HST, MAST, Chandra, HEASARC, SDSS, Spitzer, and 2MASS, to name a few. A powerful filtering mechanism allows the user to quickly narrow the initial results to a short list of likely applicable data. Guidance on choosing appropriate data sets is provided by a variety of integrated displays, including an interactive data table, basic histogram and scatter plots, and an all-sky browser/visualizer with observation and catalog overlays (see Figure 4-1).

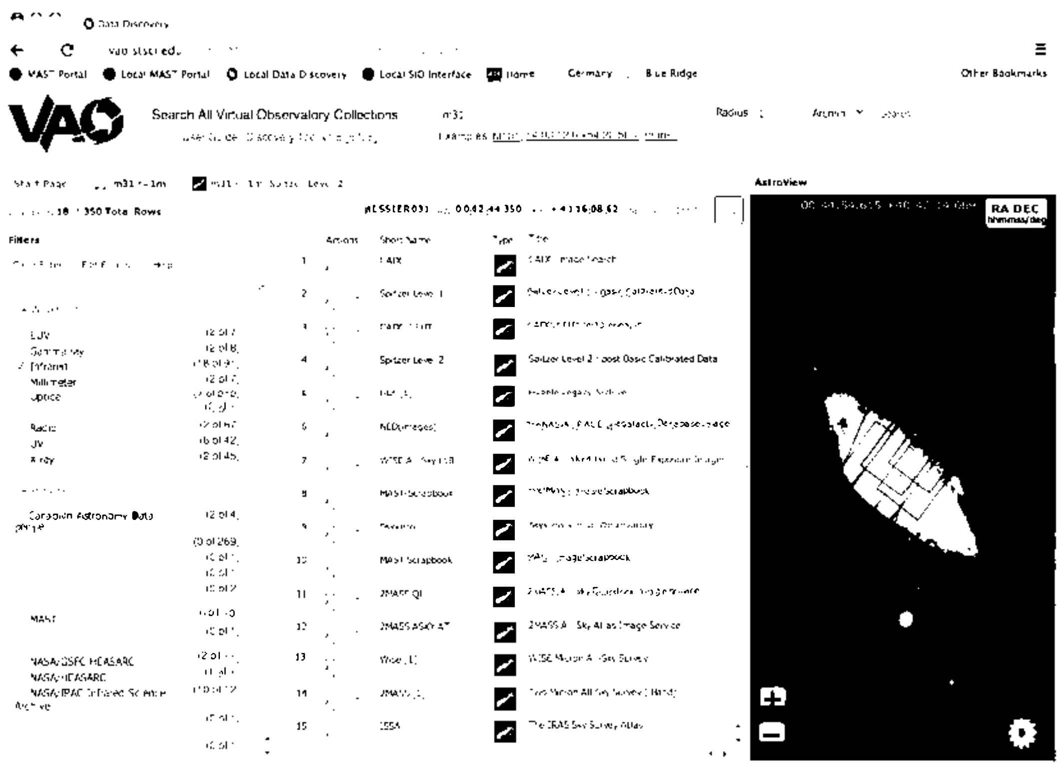

Figure 4-1. *Appearance of the DDT after a search for **M31** within radius = 1 arcmin, showing the filters (left panel) that can be applied to the Search Results (center panel), and the AstroView component with FOV overlays representing the available datasets.*

Version 1.5 was completed and approved for public release on June 21. New features added to the Data Discovery Tool in version 1.5 include:

- The AstroView all-sky display no longer requires Flash or any other browser plug-in.
- All source metadata is preserved when data, such as catalogs or image lists, are exported from the DDT to a VOTable.
- Scatter plots are available for any result tables that have at least two numeric columns.

- More accurate footprint displays for cases where the image data resource provides more than the minimum set of image metadata.

A new DDT video was developed in late October that covers material in the original video and includes highlights of the new features available in the final VAO release of DDT.

The DDT contributions to the close out repository pages were initiated over the summer and completed in the previous quarter.

4.2 Interoperable SED Access and Analysis (*Iris*)

Iris is a downloadable Graphical User Interface application that enables astronomers to build and analyze wide-band Spectral Energy Distributions (SEDs). SED data may be loaded into Iris from a file on the user's local disk, from a remote URL, or imported directly from the NASA Extragalactic Database (NED) for analysis. A plug-in component enables users to extend the science function of Iris. Communication between Specview and Sherpa is managed by a Simple Application Messaging Protocol (SAMP) connection. Data can also be read into Iris and can be written out via the SAMP interface. A separable library for SED I/O (SEDLib) is also included and available independently from Iris. See Fig. 4-2.

Version 2.0 was completed and approved for public release on July 1, 2013. New features added to Iris in version 2.0 include:

- Redshifting/blueshifting capabilities
- Interpolation and calculation of integrated quantities
- Improved visualization for co-plotting
- Upgraded plug-in for including data from the Italian Space Agency (ASI) Science Data Center (ASDC)
- Latest Sherpa fitting package (V4.5)

A patch release of Iris (v2.0.1) was completed on December 3, 2013. The patch fixed two bugs: one related to the conversion of magnitude errors, and the second to saved text file where model parameter results were printed twice. Tooltips were also provided to annotate a component list of fitting models and a documentation fix was included correcting a displayed URL.

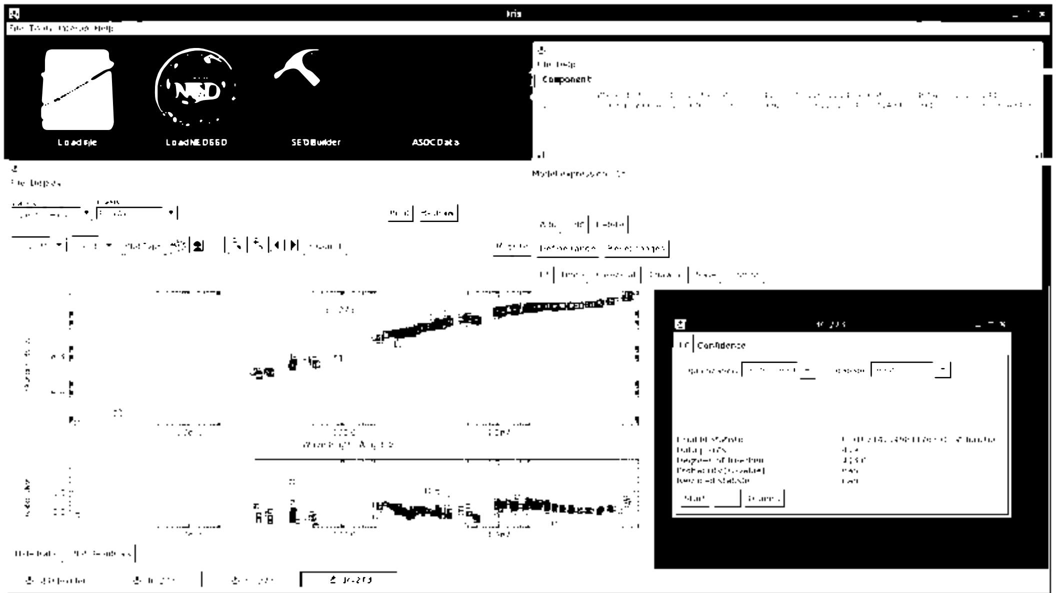

***Figure 4-2.** The SED Access and Analysis tool Iris in operations. The Iris Desktop holds the interactive windows for SED data review and analysis.*

A new Iris video was developed over the summer that covers material in the original video and includes highlights of the new features available in release v2.0.

The group supported a focus demo of Iris at the ADASS XXIII meeting Sept 29 – Oct 3, 2013. The team supported demonstrations to the Astronomy department at UMASS Amherst on December 13, 2013 and to the High Energy Astrophysics Division at SAO on March 5, 2014.

The Iris contributions to the close out repository pages were initiated and completed over the summer.

4.3 Catalog Cross-Comparison Service

The Scalable Cross-Comparison (SCC) Service performs fast positional cross-matches between an input table of up to 1 million sources and common astronomical source catalogs for a user-specified match radius. The service returns a list of cross-identifications to the user. The output is a composite table consisting of records from the first table, joined to all the matching records in the second table, and the angular distance and position angles of the matches. See Fig. 4-3.

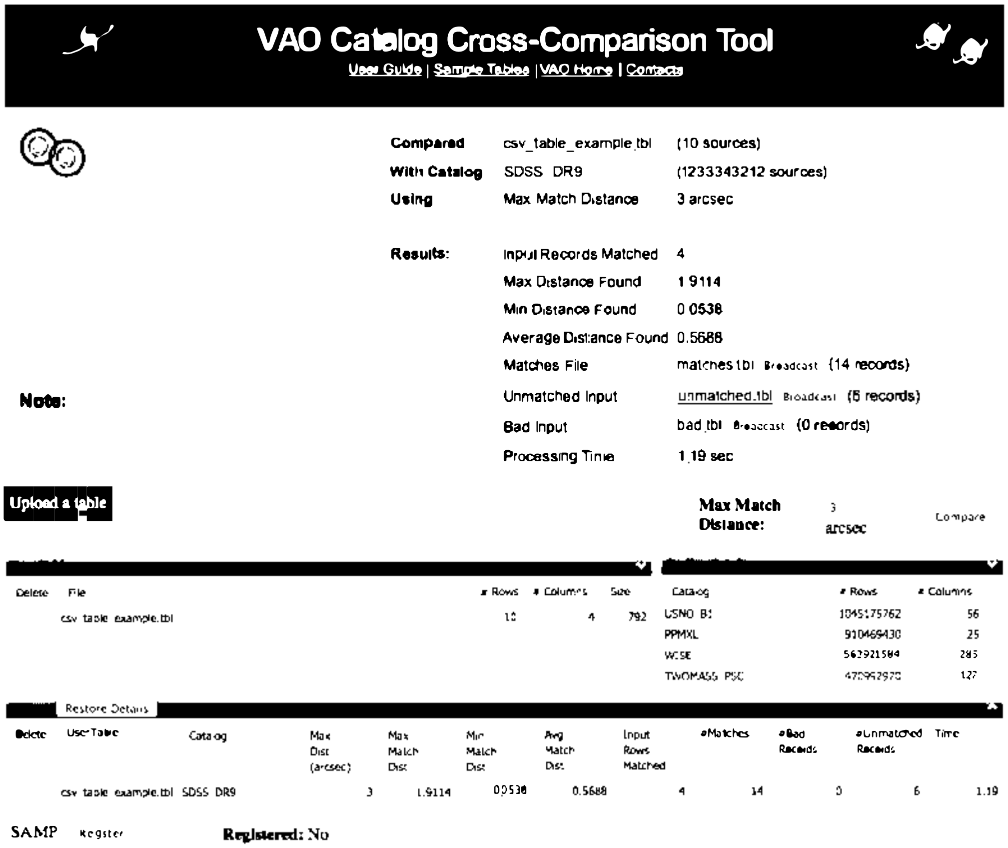

Figure 4-3. *The results view of the Scalable Cross Comparison service. The service shows results of a user table cross-matched with SDSS DR9.*

SCC version 1.2 was released on April 27, 2013. The new feature for SCC included an upgrade of the old Sloan Digital Sky Survey Data Release 7 (SDSS DR7) catalog with the current SDSS DR9.

SCC contributions to the closeout repository pages were developed over the reporting period. Completion of the repository pages and efforts to upgrade the documentation and video are goals for the next few months.

4.4 Community Applications and Assistance

SAO has continued coordination with AAVSO to support their efforts to expose their data to the Virtual Observatory. The SAO archive team successfully assisted AAVSO in getting an SAO-provided TAP server up and running with their APASS database that contains photometric all-sky survey data.

5 Standards and Infrastructure

In the last year, a major focus of our efforts has been to bring the VAO infrastructure components to a stable state in anticipation of the VAO close-out and transfer to NASA data centers. We also have worked to make significant advances in two important, leading-edge topics: access to N-dimensional data cubes and data publishing.

5.1 IVOA Standards Development

5.1.1 Data Cubes

The IVOA identified the support for access to data cubes as one of its key science initiatives, and at the May 2013 IVOA Interoperability Conference (which include consultations with specially invited data providers like NRAO, LSST, and LOFAR), projects were called on to develop prototype services to demonstrate the key features required of a new data-cube supporting standard, *Simple Image Access Version 2* (SIAv2). The VAO set about putting together such a demonstration for presentation at the September 2013 meeting. This required a revision of the relevant draft standards, namely SIAv2 and the Image Data Model. D. Tody (NRAO) and M. Cresitello-Dittmar (SAO) led the revisions of the two documents, respectively, to clarify the key pieces required for our planned prototype. Interim versions were published to be the basis for subsequent development for the summer of 2013 (see next section).

Our successful demonstration of a prototype data cube service at the September 2013 IVOA meeting, alongside the demonstrations from other projects, proved important for identifying the key requirements for a data cube standard. This allowed the working groups to craft a modular architecture that allows easy access to simple static image cubes while also enabling advanced access services for generated products like sub-cubes (cut-outs), reprojections, moments, and other virtual products. This architecture was broken down into four separate standards that work together:

- *The Image Data Model* (IDM)—this defines the semantics used by the service standards to describe the images.

- *Simple Image Access, version 2* (SIAv2)—describes a service provided by an archive that lets a user ask what images are available and to do simple downloads of discovered images.
- *DataLink*—when an archive responds with a list of images match a user's query, this standard allows the archive to point to related images (e.g. calibration data or source data) and services that can transform the image (like a cutout service).
- *AccessData*—describes a service that can transform images in advanced ways such as create cutouts or moments.

Working with our IVOA partners, we have since been actively developing these documents. In particular, Tody and Cresitello-Dittmar have been revising the SIAv2 and IDM documents, respectively to fit the revised architecture. Mature versions of the first three documents are expected to be presented at the May 2014 meeting.

5.1.2 Registry Standards

G. Greene (STScI), as chair of the IVOA Registry Working Group, has been coordinating the development of new standards for discovery data collections and other resources via the registry. In November 2013, the SimpleDALRegExt (*Describing Simple Data Access Services*) was ratified as an IVOA Recommendation. With the R. Plante (NCSA) as the lead author, this is considered the last key metadata standard needed for describing resources in a registry. As a result, attention has since moved to improving the interfaces for searching a registry.

It has long been recognized that the advanced searching features of the original Registry Interfaces standard (v1) is difficult for a typical registry to support and inconvenient for clients to use. This reflected the rapidly evolving standards and technologies when this key standard was put into place. It was based on the SOAP standard and a nascent version of ADQL that have both been deprecated by the IVOA since then. Fortunately, the advent of the general Table Access Protocol (TAP) provided a well-suited and well-supported basis for advanced searching of a registry database. To that end, we have been working on a new TAP-based search standard that spells out a standard registry schema. Last year, we produced a prototype version of our registry database that demonstrated such a schema; this year we have worked to bring that new schema into production. The TAP schema standard, the IVOA Registry Relational Schema (led by M. Demlietner of the GAVO project), is now in RFC.

Meanwhile, Greene is coordinating a second standard for simple text-based searching of a registry. While the TAP interface works well for precise and sophisticated queries to the registry, it is more complicated to use effectively when finding resources related to topics identified by keywords. Thus, the free-text search interface is meant to address these very common, but simple queries with an interface similar to the data access protocols.

Related to these efforts, R. Plante (NCSA) has been discussing a model for best practices for registering resources, intended to be the basis for our planned registry curation drive that is part of our close-out plans.

5.1.3 UTypes

UTypes are labels that can be used to identify metadata and what they represent. A comprehensive standard for formally defining such labels has been a major focus of the IVOA over the past year. Last year, the IVOA established a UTypes "tiger team", of which O. Laurino (SAO/Harvard) has been a leading participant, to better understand requirements and features.

This group presented a structured framework for defining Utype labels that can enable sophisticated metadata processing. One important use case examined was the problem of building spectral energy distributions (SEDs) from diverse spectra and photometry. Working with IVOA partners, Laurino built and presented some prototype tools that demonstrated when source data are labeled according to this structured framework, tools like the VAO's Iris can figure out how to join the data properly into an SED.

Bringing a UType standard together has been a challenge, as it has required balancing support for the simple applications currently in use against the benefits of a unified approach, as well as the needs of the expert users against those of the non-experts.

5.2 Data Sharing and Publishing

5.2.1 Data Cube Support

As mentioned in the previous section, an important activity driving of our standards development for accessing data cubes is our SIAv2 prototyping efforts. Following the call for SIAv2 prototypes at the May 2013 IVOA Interoperability Meeting, we set down an aggressive development plan to complete a prototype by the next IVOA meeting. Our goal was to deploy a service based on a draft version of the SIAv2 document that could deliver the various types of data cubes coming from modern telescopes. For this demonstration, we gathered representative data from a range instruments and projects, including ALMA, JVLA, JCMT, Chandra, OSIRIS (Keck), as well as author-contributed data sets available in NED. This required an inventory of the different FITS metadata conventions used and a mapping of those to the schema provided by the (draft) Image Data Model. Tody led the development of the service code using the framework of our existing product, DALServer. In accordance with our goals and those laid out by the IVOA, we focused on the discovery and downloading of static FITS datasets; however, the design allowed for sub-cube access (e.g. cut-outs) and other image transformations. The completed service software was delivered in early September, and a prototype service for our test collection was deployed at NRAO. This service was then made available for in-house testing and subsequent evaluation by the broader IVOA community. Plante assembled a wiki-based report to the IVOA (http://wiki.ivoa.net/twiki/bin/view/IVOA/SIA2VAOPrototype) which includes test queries with live links to the response. Our subsequent presentation was critical input into the IVOA evaluation of the different techniques for data cube access and the eventual consensus reached for a unified standard.

Building on the demonstration at the Fall IVOA meeting (and reported last quarter), we established a collaboration with a data provider, NRAO. As they are now building their next-generation archive, they want to provide scalable access to the very large datasets coming from JVLA and ALMA. In particular, they want seamless, programmatic access from their CASA-based visualization software. To ensure that the emerging standard can address these needs, we laid out a coordinated development plan (led by D. Tody) that provides NRAO developers access to early versions of our general-purpose data access software. In December 2013, we shared with NRAO developers for review the client API (to be incorporated with our VOClient product) that an application like the NRAO visualizer can use to pull image cube data from the archive. An initial implementation of that software was delivered in January 2014. By March 2014, we delivered an enhanced version of our DALServer product to NRAO, which they used to deploy an initial image cube service that would support further development of the archive. By April 2014, NRAO had successfully integrated the VOClient libraries into the CASA software, allowing it to access and visualize data from VO services. Also in April, we delivered another update to DALServer with improved support for sub-cube access. This project was reported as a poster presentation at the AAS meeting in January 2014.

5.2.2 DALServer: a Toolkit for Serving Data

A key release we are working toward in advance of the close-out is for our DALServer Toolkit. This toolkit aims to help data providers deploy services that are compliant with IVOA data access service standards. We have been using this software internally for a number of years including as a framework for prototyping new standards like SIAv2. This year, we put in a special effort to release a production version that can be easily used by others to share data through standard services. This release, planned for July 2014, is focused on IVOA's first four so-called "simple data access" services: Simple Cone Search, Simple Image Access (SIAv1), Simple Spectral Access, and Simple Line Access. In many cases, a data provider will be able to deploy services for a simple data collection with no programming required. In more complex environments, it will possible to use the underlying DALServer Java Library to connect the service to one's local archive infrastructure.

5.2.3 SciDrive: Platform for Data Sharing and Publishing

Developed initially under the name VOBox, SciDrive is our remote storage application developed at JHU, led by D. Mishin. It provides users with a platform for storing data remotely, where it can be shared with collaborators and VO services. It is based on the OpenStack software, and has been designed to allow multiple services to be deployed throughout the VO. This year, the service was rolled out (http://www.scidrive.org/) into production with 100TB of available storage (with triple redundancy), and we have been exploring different modes of use in aid of data sharing and publishing. Though similar to Dropbox and other widely used remote storage services, this service is distinguished in ways that are important to astronomy. First, it does not have the size limits typical of the free versions of the commercial offerings. Also, with multiple installations across the VO, SciDrive makes it easy to move data between storage platforms. We envision installations at major archives that provide data services for processing, integrating and analyzing data. As it supports VOSpace, it becomes possible to route data from VO data services directly into a user's SciDrive space.

This spring and summer, we plan to get feedback from friendly users to understand how best to leverage this platform for sharing and even publishing data. SciDrive was presented as a poster at the AAS meeting in January.

5.2.4 Single Sign-On Services

To support applications like SciDrive, we developed a federated, web-based single sign-on service based on the community standard, OpenID and the VAO standard for single sign-on authentication. In the last year, that service was put into production (http://sso.usvao.org). This service was designed to make it easier for VO web portals created and administered by different organizations to work together. In particular, it can let users securely pass private data between portals. It also provides users with a single username and password that can be securely used at different web portals. SciDrive was the first production portal to support the VAO SSO service. Since its rollout we have integrated the VAO notification service as well as the new registry publishing service. The NOAO portal is scheduled to convert its current support for NVO's predecessor service to the current VAO service this summer.

To make it easier for a portal to support VAO usernames, we developed the VAO Login package (http://dev.usvao.org/wiki/Products/vaologin). It provides helper software libraries in Java, python, and even command-line tools that a portal developer can use to start supporting VAO logins. This package was used to add VAO login support to the VAO notification service, and it will be used by the NOAO portal.

5.2.5 Registry Upgrade

In coordination with the new emerging registry standards, this year we updated the registry database which is now largely complete (see also previous section on standards support). A key activity this year was to ensure that the revised registry provided backward-compatibility with the applications that already rely on its services. This included key client applications like our Data Discovery Tool and Microsoft's World Wide Telescope. Integration with the VO is now under way as we complete the final testing. This work is considered critical to ensuring a stable registry service prior to its hand-over to the NASA data centers.

5.2.6 Registry Publishing Tool

Another registry improvement critical to the VAO close-out is an improved tool for registering data collections and services into the registry. This new tool brings a wizard-like interface to the process that makes it easier not only to create new records but also to keep them updated.

5.3 VO on the Desktop

5.3.1 VOClient

We completed first production release (v1.0) of the VOClient product this year and advertised in a poster at the AAS meeting. This package allows developers to create scripts and tools that access VO services from a variety of languages. The audience is not only developers that want to integrate VO capability into their desktop applications, but also individual astronomers who wish to create customized scripts that pull data from the VO and into their preferred analysis environment. The focus of this first release of VOClient was the rich set of tools that can be run on the command-line. With these tools, astronomers may either run interactive VO sessions or create simple, re-runnable shell scripts.

Meanwhile, M. Fitzpatrick (NOAO) continues to develop VOClient for a second release planned for spring 2014. This includes an update of the underlying C libraries that will let developers build in access to the VO into C and C++ applications. A key customer for this functionality is NRAO, who is building VO data access directly into the CASA image viewer. This new release will also feature new Python bindings.

5.3.2 PyVO

In parallel with our VOClient effort, we developed a product dedicated to supporting Python developers called PyVO. With its development led by Plante and M. Graham (Caltech), PyVO features a pure-Python implementation built on the community package, Astropy (www.astropy.org). It allows one to find archives with data via queries to the VAO registry, and then allows one to find and retrieve data from those archives, all from a Python script or interactive setting. Like VOClient, the audience includes both tool developers and research astronomers.

We have made a special effort to connect with the astronomy Python developers community through this package. We established its source repository in GitHub (https://github.com/pyvirtobs/pyvo) instead of our dedicated VAO repository in order to engage the community in development, and since them we have accepted contributions from several people outside of the VAO team. Because we built on top of the Astropy package, we successfully applied to become a so-called "affiliate package" of Astropy. Not only does this bring Astropy users to PyVO, is also allows PyVO to serve as a development platform for bringing more VO capabilities into the Astropy core. Two beta releases have been made available in the past year through the Python Package Index, pypi.python.org, where it has been downloaded extensively. Plante has presented papers on the PyVO package as a talk at the SciPy conference in June 2013 and as a poster at the AAS meeting in January 2014.

Appendix A. Presentations and Publications

Conference Presentations

Laurino, O. 2013, "Iris: Constructing and Analyzing Spectral Energy Distributions with the Virtual Observatory," 23rd Astron. Data Analysis Software & Systems (ADASS) Conference, Waikoloa, HI

Mishin, D. 2013, "Data sharing and publication using the SciDrive service," 23rd Astron. Data Analysis Software & Systems (ADASS) Conference, Waikoloa, HI

Plante, R. 2013, "Accessing the VO with Python," 23rd Astron. Data Analysis Software & Systems (ADASS) Conference, Waikoloa, HI

Hanisch, Robert J., Berriman, G. B., Lazio, J., Szalay, A. S., Fabbiano, G., Plante, R. L., McGlynn, T. A., Evans, J., Emery Bunn, S., Claro, M., VAO Project Team 2014, "Virtual Astronomy: The Legacy of the Virtual Astronomical Observatory," Amer. Astron. Soc. 223rd Meeting, National Harbor, MD; paper #255.04

Mishin, D., Medvedev, D., Szalay, A. S., & Plante, R. L. 2014, "Data publication and sharing using the SciDrive service," Amer. Astron. Soc. 223rd Meeting, National Harbor, MD; paper #255.26

Plante, R. L., Fitzpatrick, M. J., Graham, M., Tody, D., US Virtual Astronomical Observatory 2014, "The Virtual Observatory for the Python Programmer," Amer. Astron. Soc. 223rd Meeting, National Harbor, MD; paper #253.04

Tody, D., Plante, R. L., Berriman, G. B., Cresitello-Dittmar, M., Good, J., Graham, M., Greene, G., Hanisch, R. J., Jenness, T., Lazio, J., Norris, P., Pevunova, O., & Rots, A. H. 2014, "Accessing Multi-Dimensional Images and Data Cubes in the Virtual Observatory," Amer. Astron. Soc. 223rd Meeting, National Harbor, MD; paper #255.05

IVOA Standards

The following lists IVOA Standards published since 2013 April 1. Listed are standards and recommended standards, in chronological order, and VAO team members are highlighted.

VOTable Format Definition, Version 1.3, IVOA Recommendation 20 September 2013 (F. Ochsenbein, **R. Williams**, C. Davenhall, M. Demleitner, D. Durand, P. Fernique, D. Giaretta, **R. Hanisch**, **T. McGlynn**, **A. Szalay**, M. Taylor, A. Wicenec)

Describing Simple Data Access Services, Version 1.0, IVOA Recommendation 25 November 2013 (**R. Plante**, J. Delago, P. Harrison, **D. Tody**, and the IVOA Registry Working Group)

Data Access Layer Interface, Version 1.0, IVOA Recommendation 29 November 2013 (P. Dowler, M. Demleitner, M. Taylor, **D. Tody**)

IVOA Registry Relational Schema, Version 1.0, IVOA Proposed Recommendation 27 February 2014 (M. Demleitner, P. Harrison, M. Molinaro, **G. Greene**, **T. Dower**, M. Perdikeas)

MOC - HEALPix Multi-Order Coverage map, Version 1.0, IVOA Proposed Recommendation 10 March 2014 (T. Boch, **T. Donaldson**, D. Durand, P. Fernique, W. O'Mullane, M. Reinecke, M. Taylor)

Appendix B. Participants

Name	Role	Effort (%)	Organization
Berriman, Graham Bruce	Program Manager, IPAC Manager, Chair Program Council	0.50	IPAC
Busko, Ivo	Software Developer	0.05	STScI
Budynkiewicz, Jamie	Software Developer	0.90	SAO
Claro, Maricel	Interim Business Manager	0.75	AUI
Cresitello-Dittmar, Mark	Software Developer	0.32	SAO
D'Abrusco, Raffaele	Post-doctoral research fellow	0.33	SAO
Djorgovski, George	Scientist	0.01	Caltech
Doe, Stephen	Software Developer	0.02	SAO
Donaldson, Tom	Software Engineer	0.20	STScI
Dower, Theresa	Software Developer	0.50	STScI
Drake, Andrew	Scientist	0.05	Caltech
Economou, Frossie	Operations	0.05	NOAO
Emery Bunn, Sarah	User Support – Documentation Manager, Operations	0.50	Caltech
Evans, Janet	Software Development Manager	0.14	SAO
Fabbiano, Guissepina	Scientist, Chair of Science Council	–	SAO
Fitzpatrick, Michael	Software Engineer	0.45	NOAO
Good, John	Software Engineer	0.50	IPAC

Name	Role	Effort (%)	Organization
Graham, Matthew	Scientist, Software Developer; Program Council representative	0.35	Caltech
Greene, Gretchen	System Engineer; Program Council representative	0.10	STScI
Hanisch, Robert	Director	0.60	STScI
Lazio, Joseph	Project Scientist	0.15	IPAC
Laurino, Omar	Software Developer	0.33	SAO
Low, Stephen	VAO/IVOA Hardware	0.04	Caltech
Mahabal, Ashish	Scientist, Software Developer	0.10	Caltech
McGlynn, Tom	Lead, Operations; Program Council representative	0.40	HEASARC
Mishin, Dmitry	Software Developer	0.50	JHU
Muench, Gus	User Support – Portal QA&T Lead, Forum Lead, Documentation	0.05	SAO
Nandrekar-Heinis, Deoyani	Software Developer	0.25	JHU
Norris, Patrick	Test Engineer	0.41	NOAO
Pevunova, Olga	Software Engineer	0.25	IPAC
Plante, Raymond	Lead, Product Development; Program Council representative	0.80	NCSA
Preciado, Michael	Operations	1.00	HEASARC
Rots, Arnold	Program Council representative	0.05	SAO
Stobie, Elizabeth	Program Council representative	0.05	NOAO

Name	Role	Effort (%)	Organization
Szalay, Alex	Technical Advisor, Member of VAO Executive	0.01	JHU
Thakar, Ani	Deputy, Operations; Program Council representative	0.10	JHU
Tody, Doug	Deputy Lead, Standards and Protocols; Program Council representative	1.00	NRAO
Yenkkirala, Venkat	Software Developer	0.06	NCSA

– indicates in-kind contribution

Appendix C. Acronyms and Abbreviations

Term	Full Name (proper nouns are capitalized)	Definition
2MASS	Two Micron All Sky Survey	
AAAS	American Association for the Advancement of Science	
AAS	American Astronomical Society	
ADEC	Astronomy Data Centers Executive Council	
ADQL	Astronomical Data Query Language	
ADS	NASA Astrophysics Data System Virtual Library	
AGN	active galactic nucleus (or nuclei)	A compact region at the center of a galaxy that has a much higher luminosity over at least some portion, and possibly all, of the electromagnetic spectrum.
AIP/SPS	American Institute of Physics/Society of Physics Students	
AIPS	Astronomical Image Processing System	A software package by NRAO used for interactive and batch calibration and editing of radio interferometric data, and for the calibration, construction, display and analysis of astronomical images made from those data using Fourier synthesis methods.

Term	Full Name (proper nouns are capitalized)	Definition
ALMA	Atacama Large Millimeter Array	An array of 66 high-precision antennae located in Chile used for sub-millimeter observing.
APOD	Astronomy Picture of the Day	A popular web site that showcases an astronomical photo each day. The site is a service of NASA's Astrophysics Science Division (ASD) and Michigan Technological University
APS	American Physical Society	
arXiv	Pronounced as "archive," with the X as a hard K sound as in "LaTex"	An automated electronic archive and distribution server for research articles. Maintained and operated by the Cornell University Library.
AUI	Associated Universities, Inc.	
AURA	Associated Universities for Research in Astronomy, Inc.	
Ajax		A group of interrelated web development methods used on the client-side to create interactive web applications. With Ajax, web applications can retrieve data from the server asynchronously in the background without interfering with the display and behavior of the existing page.
ASCL	Astronomy Source Code Library	
ASKAP	Australian Square Kilometre Array Pathfinder	

Term	Full Name (proper nouns are capitalized)	Definition
BuildBot		A system that automates the compile/test cycle required by most software projects to validate code changes. By automatically rebuilding and testing the tree each time something has changed, build problems are pinpointed quickly, before other developers are inconvenienced by the failure.
CACR	Center for Advanced Computing Research	
CANDELS	Cosmic Assembly/Near-infrared Deep Extragalactic Legacy Survey	
CASA	Common Astronomy Software Applications	A software packaged developed by NRAO with the goal of supporting the data post-processing needs of the next generation of radio astronomical telescopes.
CDS	Centre de donneés astronomiques de Strasbourg	
CFHT	Canada France Hawaii Telescope	
CIAO	From "s'sciavo," meaning "I am your servant" in Italian (Venetian dialect)	Software developed to analyze data returned by the Chandra X-ray Observatory.
CIS	continuous integration system	A system that periodically and automatically builds a software tool and (possibly) runs its unit tests.

Term	Full Name (proper nouns are capitalized)	Definition
CiTO	Citation Typing Ontology	An ontology for describing the nature of reference citations in scientific research articles and other scholarly works.
CM	configuration management, or Configuration Manager	
COROT	COnvection ROtation and planetary Transits	A European space mission that searches for extrasolar planets with short orbital periods, and performs asteroseismology by measure solar-like oscillations in stars.
COSMA	Cosmic Sky Machine for Astrophysics	
CRTS	Catalina Real Time Survey (CRTS)	
CUDA	Compute Unified Device Architecture	A proprietary parallel computing architecture that enables dramatic increases in computing performance by harnessing the power of the GPU (graphics processing unit).
CXC	Chandra X-Ray Observatory	
DALServer	Data Access Layer Server	
DM	data mining	
EPO	Education and Public Outreach	A variety of activities conducted by research institutes, universities, and institutions such as science museum that are aimed at promoting public awareness and understanding of science.

Term	Full Name (proper nouns are capitalized)	Definition
EVLA	Expanded Very Large Array	An array of 27 antennas with 25-meter diameters that is located in New Mexico.
GAVO	German Astrophysical Virtual Observatory	
GALEX	Galaxy Evolution Explorer	An orbiting, ultraviolet space telescope launched in April 2003 to determine the distances of several hundred thousand galaxies, as well as the rate of star formation in each galaxy.
GPU	graphical processing units	
GSC2	Guide Star Catalogue II	Also known as the Hubble Space Telescope, Guide Catalog (HSTGC). It is a star catalog compiled to support HST with targeting off-axis stars.
GSFC	Goddard Space Flight Center	
GWS	Grid and Web Services	
HCO	Harvard College Observatory	
HEASARC	High Energy Astrophysics Science Archive Research Center	
Hudson		An open-source build server.
HST	Hubble Space Telescope	
IPAC	Infrared Processing and Analysis Center	A multi-mission center of expertise for long-wavelength astrophysics that carries out data-intensive processing tasks for NASA's infrared and sub-millimeter astronomy programs.

Term	Full Name (proper nouns are capitalized)	Definition
IRAF	Image Reduction and Analysis Facility	A general purpose software system for the reduction and analysis of astronomical data.
IRSA	Infrared Science Archive	An online archive of data and science products for NASA's infrared and sub-millimeter missions.
IVOA Interoperability Meeting	International Virtual Observatory Alliance Interoperability meeting	
IVOA TCG	International Virtual Observatory Alliance Technical Coordination Group	A group composed of the Chairs and Vice Chairs of IVOA working group and interest groups that reviews and proposed IVOA recommendations.
Java		A compiled programming language that allows application developers to write code that can run on any platform.
JIRA	Truncation of "Gojira," the Japanese name for Godzilla	A proprietary issue-tracking product the VAO will use as its help desk issue/bug tracking system.
JPL	Jet Propulsion Laboratory	
LLC	limited liability company	A flexible form of business enterprise that blends elements of partnership and corporate structures. It is a legal form of a company that provides limited liability to its owners within the vast majority of United States jurisdictions.
LSST	Large Synoptic Survey Telescope	A facility that will produce a wide-field astronomical survey of the Universe using an 8.4-meter ground-based telescope.

Term	Full Name (proper nouns are capitalized)	Definition
MAST	Multimission Archive at Space Telescope	
MoU	Memorandum of Understanding	
NASA	National Aeronautics Space Administration	
NCSA	National Center for Supercomputing Applications	
NED	NASA/IPAC Extragalactic Database	An online, knowledge-based database and archive for extragalactic astronomy.
NOAO	National Optical Astronomical Observatory	
NRAO	National Radio Astronomy Observatory	
NSF	National Science Foundation	
NStED	NASA/IPAC/NExScl Star and Exoplanet Database	A general purpose stellar and exoplanet archive to support NASA's planet finding and characterization activities.
NVO	National Virtual Observatory	
MERLIN	Multi-Element Radio Linked Interferometer Network	
PEP	Project Execution Plan	
PI	Principal Investigator	
PQL	Program Query Language	A programming language for expressing patterns of events on objects.

Term	Full Name (proper nouns are capitalized)	Definition
ObsCore	Observation Core	An IVOA data model that describes the necessary metadata for multi-wavelength data discovery queries.
ObsTAP	Observation Table Access Protocol	
Python		An open source, interpreted, object-oriented programming language.
PQL	Program Query Language	A language for expressing patterns of events on objects. It provides a front end to static and dynamic program analyses to go find those sequences on the program as it runs.
OPO	Office of Public Outreach	
REST	REpresentational State Transfer	A style of software architecture for distributed hypermedia systems, such as the World Wide Web.
RLG	Research Libraries Group	A U.S.-based library consortium that developed standards and tools for computerized library systems. It merged with OCLC in 2006.
SAMP	Simple Application Messaging Protocol	A messaging protocol that enables astronomy software tools to interoperate and communicate.
SAO	Smithsonian Astrophysical Observatory	
SEDs	Spectral Energy Distributions	
SemDB	Data-Literature Semantic Database	

Term	Full Name (proper nouns are capitalized)	Definition
SDSS	Sloan Digital Sky Survey	
SHA	Spitzer Heritage Archive	The interface to all data gathered by the Spitzer Space Telescope.
Sherpa		A modeling and fitting software tool that is part of the CIAO software packaged developed by Chandra X-ray Observatory
SIAPV2	Simple Image Access Protocol 2	The second generation of an IVOA standard that defines a protocol for retrieving image data from a variety of astronomical image repositories through a uniform interface.
SkyAlert		A clearinghouse of VOEvent notices that individuals may subscribe to and/or monitor from various event streams and feeds.
SSO	single sign-on	Technology that allows users to gain access to multiple systems using a single, central authentication method or interface.
SIMBAD	Set of Identifications, Measurements, and Bibliography for Astronomical Data	An astronomical database that provides basic data, cross-identifications and a bibliography for astronomical objects outside the solar system.
SMC	Small Magellanic Cloud	
Solr		An open-source search server based on the Lucene Java search library.
Specview		A tool developed at STScI for 1-D spectral visualization and analysis of astronomical spectrograms.

Term	Full Name (proper nouns are capitalized)	Definition
STScI	Space Telescope Science Institute	
TA	Technology Assessment	One of the seven work areas of the VAO project that advises the VAO on technology choices that best support VAO goals.
TAP	Table Access Protocol	A VO standard that defines a service protocol for accessing general table data, including astronomical catalog and general database tables.
TSC	Time Series Center	A project at Harvard University dedicated to creating the world's largest data center for time series, and developing algorithms to understand various aspects of those time series.
triplestore (also Triple-Store)		A type of database that is very efficient at storing character information. Queries consist of short, three-part statements (subject-predicate-object).
USNO-B	United States Naval Observatory B1.0 catalog	An all-sky catalog of digitized photographic plates from observations covering the years 1949-2002.
UWS	Universal Worker Service	An IVOA specification that defines how to manage asynchronous execution of jobs on a service.
VAO	Virtual Astronomical Observatory	
VAO, LLC	Virtual Astronomical Observatory, Limited Liability Corporation	

Term	Full Name (proper nouns are capitalized)	Definition
VO-CLI	Virtual Observatory Command Line Interface	A set of tools developed by the NVO that use a command-line interface.
VO	Virtual Observatory	
VAOBase	Virtual Astronomical Observatory Base ontology	
VAOBib	Virtual Astronomical Observatory Bibliographic ontology	
VAOObserv	Virtual Astronomical Observatory Observational Ontology	
VOClient	Virtual Observatory Client	A software package that provides a high-level, programmable interface between desktop applications and the distributed VO framework, providing access to remote VO data and services, reference implementations for VO data-providers and end-user applications.
VOG	VAO Oversight Group	A group convened by NSF Division of Astronomical Sciences (AST) to provide oversight of the VAO and act as a point-of-contact for NSF and NASA and the Project Director.
VOSI	Virtual Observatory Support Interfaces	
VOSpace	Virtual Observatory Space	The IVOA interface to distributed storage.
VOView	Virtual Observatory View	A web utility used to view large data tables within a web browser.

Term	Full Name (proper nouns are capitalized)	Definition
WBS	Work Breakdown Structure	A definition and description of a project's discrete work elements and tasks that helps organize and define the total work scope of the project.
WSBP	Web Services Basic Profile	
WISE	Wide-Field Infrared Survey Explorer	A space telescope that is scanning the entire sky in infrared light to photograph distant objects.

VAO

Final Report

For the Period:

15 May 2010
through
30 April 2015

Associated Universities Inc.

The Virtual Astronomical Observatory (VAO) is managed by the VAO, LLC, a non-profit company chartered in the District of Columbia and established as a partnership of the Associated Universities, Inc. and the Association of Universities for Research in Astronomy, Inc. The VAO is sponsored by the National Science Foundation and the National Aeronautics and Space Administration.

Submitted to the National Science Foundation
Pursuant to Cooperative Agreement
AST-0834235

Contents

Executive Summary

The Virtual Astronomical Observatory (VAO) was funded from May 15, 2010 to September 30, 2014. All software and documentation developed during the lifetime of the project is publicly accessible through a repository at https://sites.google.com/site/usvirtualobservatory/. The software repository provides all the information third parties need to use and extend each software component. This information includes the code itself, requirement, test results, test data sets, build instructions, design documents, usage information, dependencies, license information, API specifications and tutorials. The documentation repository includes all project reports, posters, presentations, papers and exhibit material. Beginning October 1, 2014, the essential components of the infrastructure built by the VAO will be sustained by the NASA-funded archives at MAST (STScI), HEASARC (GSFC) and NED/IRSA (IPAC). These components are the registry, operational monitoring, repository, and the web page. The VAO staff worked closely with staff at the NASA archives to ensure a smooth transition, and at the request of the archives, VAO management participated in their planning process. On behalf of the International Virtual Observatory Alliance (IVOA), the VAO managed the IVOA webpage, document repository, and mailing lists. These assets have been transferred to the IVOA partners in Italy and India; the VAO Program Manager coordinated this transfer at the request of the IVOA.

A close-out review attended by the VAO team, and representatives of NASA, NSF, the VAO, LLC Board of Directors, and the NASA archives, was held in Pasadena, CA, on July 10 and 11, 2014. The review focused on project achievements and its legacy, financial status and corporate close-out, lessons learned, and identified tasks (now completed) to fulfill the terms of the close-out plan by September 30, 2014. A report on the review was distributed to the agencies and posted on the VAO web site.

The VAO adopted a financial and business model that included management costs of only 10-15% of the total budget, met all audit requirements, and ensured sufficient funds were available throughout the project lifetime to ensure all contractual obligations could be met. The VO concept and management approach has been endorsed by other disciplines, including a panel of neuroscientists convened by the Kavli Foundation and General Electric as a means for improving access and interoperability to the vast data sets being collected in the European Brain Project and U.S. Brain Initiative.

Throughout the project, the VAO has been active in leadership roles within the IVOA. The VAO staff have been major contributors to 23 IVOA standards. These contributions include authorship of standards, including those recommended by the IVOA Executive Committee and those submitted to the Executive Committee for recommendation, and in the development of reference implementations. The VAO led the IVOA in the development of service validators, which check whether a service is compliant with VO standards. Moreover, staff members of the VAO and the NVO before it held 29 IVOA leadership positions, including chair of the Executive Committee, and chairs and vice-chairs of the technical Working Groups.

The key technical project for the VAO in the close-out period was to advance access to multidimensional image sets ("image-cubes"), now routinely produced by major telescope such as ALMA. The VAO provided leadership within the IVOA for the development of requirements, engagement of data providers, preparation of a white paper, and development of a prototype service that drove the specifications of the protocols needed to support access to data cubes, and the development of a reference implementation. The protocols are the Image Data Model, which define the semantic labels that describe image cubes; the Simple Image Access Protocol V2 (SIAP V2) uses these labels to annotate image search results; and Access Data, which defines how to request operations and transformations. Throughout this work, the VAO worked with the NRAO to support access to cutouts from archived data cubes in their CASA Viewer software package.

The registry, a database containing descriptions of VO resources, is the heart of the VAO data access and discovery services. The VAO extended the registry developed by NVO into an operations-quality product. It was overhauled to include the latest standard for searching registries, based on the IVOA Table Access Protocol (TAP). The usability of the registry enhanced by a new web-based resource publishing service, and a web-based directory service for keyword-based searches data collections and services.

The VAO delivered a data sharing service, SciDrive, aimed at storing, archiving and serving the "long-tail" of astronomical data; that is, those data produced by small teams that are not readily accessible from project or thematic archives. SciDrive operates in the same fashion as drag-and-drop services such as Dropbox, but with the addition of powerful science capabilities such as seamless movement of files among SciDrive instances via the VOSpace 2.0 interface, the capability to publish data to the VO, and support for interfacing with existing services. The SDSS CasJobs is one example of this.

The VAO has delivered scripting toolkits that allow scientists to incorporate VO data directly into widely-used reduction and analysis environments. VO services have been incorporated into the IRAF reduction package. VOClient offers command-line access to VO services in a Unix/Linux shell environment. A subsequent release deploys a set of libraries that allow VO capabilities to be plugged into existing C and C++ application—the NRAO CASA Viewer is one example—and is the basis of bindings to other languages. A second product, PyVO, represents a native Python instantiation of VO capabilities, requested by the Python community.

Three high-level science applications for data discovery, integration, and analysis have been deployed and maintained. The Data Discovery Tool enables discovery and visualization of, and access to, data archived worldwide. Iris allows astronomers to build and analyze broad-band spectral energy distributions. Finally, the Scalable Cross-Comparison Service cross-matches user-supplied tables of sources with large online source catalogs.

The tools and services developed by the VAO are underpinned by a robust operational environment, where distributed operational VO services worldwide are routinely checked hourly for aliveness and compliance with IVOA protocols. The operations team developed a process for fast response to service interruptions and downtime, and worked with service providers to rectify defects and problems. Services that were abandoned were deprecated in the VAO Registry and removed from responses to user queries. In large part because of the monitoring and active responses, the percentage of compliant services increased steadily, and the fraction of services that were operational returned close to 100% after service interruptions were identified and corrected.

Take-up of VO standards and infrastructure is now seen within essentially every major data center and survey project in the United States, with ~1M VO-based data requests per month and ~2,000 unique users.

A vigorous program of outreach brought the VAO to the attention of the community. Modes of outreach included VAO Community Days, exhibits at AAS meetings, and demos at topical conferences. VAO staff actively promoted the project with presentations at meetings, including invited and keynote presentations, and with papers in journals and in conference proceedings. Altogether, the staff gave 56 presentations or posters, and published 18 papers. VAO participants are now playing leading roles in the international Research Data Alliance and the newly formed U.S. National Data Service.

1 Management

1.1 Reporting and Communications

An internal reporting schedule was maintained throughout the project's lifetime. VAO Director R. Hanisch led a weekly executive telecon attended by the VAO Program Manager, B. Berriman, the Project Scientist, J. Lazio, the Business Manager, M. Claro, and the Task Area leads. Other members of the team, such as the Chair of the Science Council, were invited to join the meeting as needed. Berriman led a weekly team meeting throughout the project, primarily to discuss schedule and technical issues. Each work area held weekly telecon or WebEx meetings.

Hanisch and Berriman represented the VAO at the telecons of the Astronomy Data Centers Executive Committee (ADEC), which coordinates activities among the NASA-funded data centers. In the context of the VAO, these telecons provided an opportunity to discuss both the benefits and the concerns of the data centers associated with some of the recent and more complex IVOA standards. In order to assure the NASA data centers were aware of and had an opportunity to comment on the impact of new and modified VO standards, the VAO routinely forwarded all notices of IVOA RFCs (Requests for Comment) to the ADEC mailing list.

From spring 2013 onwards, VAO Team Members participated in a biweekly telecon with members of the Atacama Large Millimeter/submillimeter Array (ALMA) software development and the archive development team at the National Radio Astronomy Observatory (NRAO) to review progress in the development of an operational service that supports discovery of, and access to, multidimensional data sets from the Jansky Very Large Array (JVLA) and ALMA within their Common Astronomy Software Applications (CASA) software.

Requirements and design reviews for project deliverables were generally carried out as needed using the WebEx e-conferencing software. These reviews were led by the development leads and were attended by representatives of management, User Support, and Operations, as well as by the development team itself. The Configuration Change Board (CCB) met biweekly by telecon or WebEx as needed to review requests for changes to released software and documentation. The CCB was chaired by the Program Manager, and the CCB consisted of the task leads of the work areas. All members of the VAO team were welcome to attend the CCB meetings. The Board's proceedings were recorded on the wiki.

Quarterly reports were submitted to NSF and NASA on schedule throughout the life of the project. The agencies waived the July 2014 Quarterly Report in lieu of the Summary Report from the VAO Close-out review, delivered August 15, 2014. The agencies also waived the October 2014 Quarterly Report in lieu of the Final Project Report (this document), given that technical work ended on September 30, 2014.

1.2 Meetings

The project held team meetings, generally every six months, each usually hosted by a participating organization. The meetings reviewed schedule and technical status, and included working sessions where team members collaborated on studying technical issues. The meetings were held in Charlottesville (hosted by NRAO; October 2010), Tucson (NOAO; April 2011), Cambridge (SAO; September 2011), Pasadena (CACR/IPAC; December 2011), Washington DC (AUI; April 2012), Seattle (Microsoft; September 2012), Annapolis (AUI; June 2013), Tucson (NOAO; January 2014), and Pasadena (IPAC/CACR; July 2014).

Berriman convened the Program Council in August 2010 (Chicago), July 2011 (Cambridge, MA), and January 2013 (Atlanta) to negotiate work packages and deliverables for participating organizations. The meetings were attended by the Director, Program Manager, Project Scientist, and Program Council representatives of the participating organizations.

The Science Council, chaired by Dr. Giuseppina Fabbiano (SAO), met in March 2010 and July 2011. The Science Council advised the VAO on science priorities and presented their recommendations in summary reports. Meetings were attended by Council members and VAO team members. The Science Council became inactive after the second meeting, when the VAO was de-scoped to emphasize infrastructure over science applications.

The VAO Board of Directors met regularly, either in person at the AUI/VAO Corporate office in Washington, DC or via WebEx. Meetings were held in July 2010, August 2011, October 2011, November 2011, October 2012, and April 2013. The meetings concerned business management, fiduciary issues, and technical progress and budget matters. Re-scoping the project in the light of budget restrictions was discussed in detail. The meetings were attended by the VAO Board of Directors, Secretary, Treasurer, the VAO Director, Program Manager, Project Scientist, Business Manager, and Technology Advisor (A. Szalay), and the Science Council Chair. In March 2011, a subset of the VAO team met with the data management team of the LSST (Large Synoptic Survey Telescope) at NOAO to discuss formalizing the collaboration between VAO and LSST. LSST intends to use VO protocols for data access (images, catalogs) and for disseminating transient event notices. As a result of a request made by LSST this meeting, J. Good (IPAC) developed a prototype VO-compliant image access service that will allow the LSST to serve image data sets used to support project development. The service was built in less than two days.

VAO staff were active participants in presenting papers, including invited reviews, at AAS and topical conferences and AAS meetings, organizing and participating in workshops, and serving as members of review boards.

1.2.1 AAS Meetings

The VAO was active in the Winter AAS meetings in January 2012, 2013, and 2014, with presentations, exhibits, and participation in workshops. The project had exhibits at the 219th meeting in Austin, TX in January 2012 and at the 221st meeting in Long Beach, CA in January 2012. The 2012 meeting included a workshop entitled "Science Tools for Data Intensive Astronomy," and the 2013 exhibit featured a collection of 12 tutorials on how to use VO/VAO tools and services were prepared.

VAO team members presented posters at the 2013 and 2014 meetings. In 2013, VAO staff and affiliated researchers contributed 12 posters in the "Computation, Data Handling, and Image Analysis" poster session. The posters described Virtual Observatory tools and illustrated their use, and highlighted research results that made use of them. In 2014, members of the VAO team presented four posters at the 223rd AAS meeting: "Virtual Astronomy: The Legacy of the Virtual Astronomical Observatory," "Accessing Multi-Dimensional Images and Data Cubes in the Virtual Observatory," "Data Publication and Sharing Using the SciDrive Service," and "The Virtual Observatory for the Python Programmer." At the same meeting, Szalay gave an invited presentation entitled "From AISR to the Virtual Observatory" at the session "Building the Astronomical Information Sciences: From NASA's AISR Program to the New AAS Working Group on Astroinformatics and Astrostatistics."

Also at the 223rd meeting, Hanisch and Berriman attended the workshop on "Metadata Semantics and Data Formats," organized by the AAS. They emphasized how VO services can support publishing of data through peer-reviewed journals. In their capacity as members of the Advisory Board of the Astronomy Source Code Library, Hanisch and Berriman were co-organizers of, and presenters at, a special session "Astrophysics Code Sharing II: The Sequel." This was a follow-on to a similar special session they helped organize one year earlier.

1.2.2 Topical Meetings, Workshops, and Review Boards.

Several team members presented papers and demonstrations of VAO applications at the annual Astronomical Data Analysis Software and Systems (ADASS) conferences in Champaign, IL, November 2012, and in Waikoloa, HI, in October 2013. Again in their capacity as ASCL Advisory Board members, Hanisch and Berriman were co-organizers of a Birds Of A Feather session at ADASS XXIII entitled "Ideas for Advancing Code Sharing."

The VAO had a strong presence at the SPIE Telescopes & Instrumentation Conferences held in Amsterdam, The Netherlands, July 2012. Five oral presentations were given by VAO staff (see Appendices A and B for details).

Hanisch and Berriman both gave invited presentations at the Innovations in Data-Intensive Astronomy meeting held in Green Bank, WV, May 2011. Hanisch gave a presentation entitled "Data discovery, access, and management with the Virtual Observatory," while Berriman spoke on "Approaches to investigating technology solutions in data-intensive astronomy."

Berriman gave a keynote presentation entitled "Astronomy in the Cloud" at the .Astronomy 4 conference in Heidelberg, Germany, held July 9-11, 2012. VAO services were demonstrated as part of the conference's extensive hands-on sessions.

Berriman was a member of the program committee for the Workshop on Maintainable Software Practices in e-Science, October 8-9, 2012, part of the e-Science conference 2012, Chicago, IL. He gave an invited presentation at this workshop entitled "Adoption of Software by a User Community: The Montage Image Mosaic Engine Example." The Montage engine was built specifically for the VO.

Hanisch presented status reports to the AUI Board (February 15, 2013) and AURA Board (February 22, 2013). Lazio gave a presentation at the annual AURA Member Representatives meeting in Tucson, May 17-19, 2013. In addition, team members gave interactive presentations on VAO services to attendees.

Hanisch was member of the review panel for the Common Astronomy Software Applications (CASA) review in Socorro, NM in March 2013.

Berriman, G. Djorgovski, C. Donalek, and M. Graham (Caltech), and A. Szalay and T. Budavari (JHU) gave presentations at the AstroInformatics conference in Sydney, Australia, December 9-12, 2013.

Hanisch and Lazio attended Radio Astronomy in the Era of LSST, a meeting held at NRAO-Charlottesville, May 6-8, 2013, where Hanisch gave a presentation "Virtual Astronomy in the Era of LSST."

Hanisch and Szalay were invited to a workshop, Sustaining Domain Repositories, sponsored by the Sloan Foundation and held at the University of Michigan, Ann Arbor, June 24-25, 2013.

G. Djorgovski and A. Mahabal gave presentations at the conference Interface 2013 held at Chapman University, Orange, CA, April 4-6, 2013.

R. Plante gave a presentation at the SciPy 2013 conference in Austin, TX, June 26, 2013.
Hanisch and Plante gave presentations to the NanoGrav consortium on November 6, 2013 describing the VO and the Python interfaces.

Hanisch described the VO infrastructure in a presentation to a small group of theoretical physicists and neuroscientists organized by the Kavli Foundation (Santa Monica, CA, February 19, 2014), and spoke at an additional workshop of neuroscientists organized by Kavli and the General Electric Corporation, May 8-10, 2014, in England.

1.3 International Collaboration

Members of the VAO attended the biannual International Virtual Observatory Alliance (IVOA) Interoperability Meetings ("Interops"). They were held in Nara, Japan (December 2010); Naples, Italy (May 2011); Pune, India (October 2011); Urbana, IL (May 2012); Sao Paulo, Brazil (October 2012); Heidelberg, Germany (May 2013); Waikoloa, HI (October 2013; sponsored by the VAO); and Madrid, Spain (May 2014). The Interoperability meetings are the focal point for discussion of the international VO standards and increasingly a venue for collaboration and consultation on VO-enabled research tools.

VAO team members held leadership roles in the IVOA as follows:

- IVOA leadership roles
  - Executive chair: Hanisch (2002-3), De Young (2007-8)
  - Executive secretary: Evans (2013-)
  - Executive members: Fabbiano, Lazio
  - Technical Coordination Group: Williams chair (2002-8), Graham deputy chair (2012-15)
  - IVOA Document Coordinator: Emery Bunn (2010-14)
- IVOA Working Groups and Interest Groups
  - Applications chair: McGlynn (2008-11)
  - Data Access layer chair: Tody (2003-7) deputy: Fitzpatrick (2010-13)
  - Data Models chair: McDowell (2003-7) deputy: Laurino (2011-15)
  - Grid and Web Services chair: Graham (2007-11) deputy: Graham (2006-7)
  - Registry chair: Plante (2006-10), Greene (2011-14) deputy: Greene (2009-11)
  - Standards and Processes chair: Hanisch (2003-5)
  - UCD chair: Williams (2003-5)
  - VOEvent chair: Williams (2005-8), Seaman (2008-11), Graham (2011-12) deputy: Seaman (2006-8), Williams (2010-11)
  - VOQL chair: Nieto-Santisteban (2005-6)
  - Data Curation and Preservation chair: Hanisch (2007-10), Accomazzi (2010-14)
- Knowledge Discovery in Databases chair: Djorgovski (2012-15)

Altogether, VAO staff contributed to 23 IVOA standards documents as editors or authors. Appendix C lists the standards recommended by the IVOA Executive Committee and those submitted to the Executive Committee for recommendation.

The VAO hosted the IVOA web site, document repository, wiki, and e-mail distribution lists, and was responsible for editing the IVOA Newsletter.

Just prior to the Sao Paulo meeting, VAO team members assisted with a VO School hosted by the Brazilian VO project, BRAVO, on November 18-19, 2012. Hanisch (STScI), M. Graham

(Caltech), and M. Fitzpatrick (NOAO) were also invited speakers at the National Astronomy Meeting of Brazil, held in Aguas de Lindoia, November 16-17, 2012.

Virtual Observatory-related events featured strongly at the International Astronomical Union General Assembly in Beijing, China in August 2012. Special Session 15: Data-Intensive Astronomy spanned four days and featured a number of presentations in which VO facilities play or will play a prominent role. Attendees included many people who were not already involved with VO programs. Hanisch served as a member of the Science Organizing Committee and gave the closing summary presentation.

IAU Commission 5 also held business meetings for the main Commission and its working groups, including Virtual Observatories, Data Centers, and Networks (also chaired by R. Hanisch). Commission 5 agreed to wind down the Working Group on Astronomical Data, instead folding its activities into the Commission itself, and added a new Working Group on Astroinformatics and Astrostatistics that is chaired by E. Feigelson (Penn State) and co-chaired by P. Shastri (Indian Institute of Astrophysics, Bangalore). The Libraries Working Group also held a two-day meeting that included discussion of a new astronomy thesaurus to be constructed in collaboration with the IVOA.

1.4 Close-Out Activities

Formal approval of the VAO Project Management Plan (PEP) for the period May 1, 2013 through September 30, 2014, was received on June 8, 2013. A formal close-out review and final team meeting was held on July 10-11, 2014 at IPAC/Caltech. NSF and NASA representatives, members of the VAO Board, and members of the NASA and VAO team members attended. A summary report of the close-out review was delivered to the agencies on August 15, 2014, in lieu of the quarterly report due on July 30, 2014.

The goal of the VAO close-out review was to assess the status of all technical deliverables of the project, understand the curation of project assets and their transfer to the NASA data centers and other organizations, and to review the overall budgetary status and fiduciary responsibilities of the VAO, LLC. In particular:

- What deliverables have been made by the VAO and what tasks have been completed?
- What is the plan for completing all remaining deliverables and tasks?
- What tasks are at a high risk of non-completion and how will the risk be mitigated?
- What are the plans for sustaining the infrastructure developed by the VAO?

In addition, we examined two broader issues:

- What is the legacy of the VAO? What is its impact on the astronomical and broader communities?

- What are the lessons learned for follow-on projects and the user community?

The review contained presentations on the following topics:

- A Brief History of the U.S. Virtual Observatory Efforts (Hanisch)
- Fiscal and Contractual Matters (Claro)
- Science Applications (J. Evans)
- Operations (T. Mc Glynn)
- Standards and Infrastructure (R. Plante)
- Community Engagement and User Support (Lazio)
- Long-Term Curation of VAO Assets (Berriman)
- Remaining Tasks and Schedule (Berriman)
- The VAO Legacy (Hanisch)

The review identified a set of tasks for completion before project close on September 30, 2014:

Management and Contractual Tasks

- Close-out Review Report
- Quarterly Report
- Final Annual Report
- Contractual and Fiscal Close-Out
- SMC Cross-match catalog to repository

Transfer of Assets to the IVOA

- Complete web site/documentation
- Mailing lists
- Reports to IVOA Executive

Standards and Infrastructure Deliverables

- SciDrive
- VAO Log-in
- Registry Publishing Interface
- PyVO
- VOClient
- DAL Server

Documentation Services

- Add mailing lists to repository
- Add web page content to repository
- Archive of JIRA tickets
- Final review of repository contents
- Final updates to contents of repository

Transfer to NASA Data Centers and Archives

- Transfer VAO web page and mailing lists to IPAC
- Agreement on how VAO services are branded post-VO
- Initiation of bi-weekly VO telecons (led by Project Scientist)

NASA directed the archives at STScI (MAST), GSFC (HEASARC), and IPAC (IRSA and NED) to develop a proposal to support the essential components of the VAO-developed infrastructure as part of their core ("in-guide") activities, beginning in October 2014. This proposal was submitted to NASA on November 8, 2013, and a technical and managerial response to the review of the proposal was submitted on April 4, 2014. The partners in the proposal consulted with Hanisch and Berriman on the essential VAO components that can be supported within their core budgets, and have developed a transition plan to transfer the necessary VAO assets to them. The transition plan was straightforward: the VAO web page would be transferred from CACR to IPAC, whereas the Registry and the Operational Services would be sustained *in situ* at STScI and HEASARC by staff who currently support the VAO. JHU agreed to support SciDrive.

1.5 Milestones Met

1.5.1 Year One (May 15, 2010– April 30, 2011) Milestones

- Held Science Council Meeting and delivered technical response to it.
- Delivered PEP and updated it.
- Released a new project web page.
- Published guidelines for testing of VAO services.
- Developed a VAO exhibit for the AAS meeting in Seattle.
- Represented the VAO at the IVOA Meeting in Nara, Japan, in December 2010. The contributions of the VAO team members are described in subsequent sections of this report.
- Published the VAO Documentation Plan.
- Deployed a Configuration Management Plan.
- Deployed a User Support and issue tracking system.

1.5.2 Year Two (May 1, 2011 – April 30, 2012) Milestones

MANAGEMENT

- Science Council Meeting and Response
- 1st Quarterly Report
- Program Execution Plan
- 2nd Quarterly Report
- Response to NSF Portfolio Review
- 3rd Quarterly Report

- Annual Report

STANDARDS AND INFRASTRUCTURE

- Spring IVOA Meeting
- Winter IVOA Meeting

SCIENCE APPLICATIONS

- Data Discovery Portal
- Cross-Match Engine Beta1
- Cross-Match Engine Beta2
- Time Series Service Beta1
- Desktop Integration Tools
- Data Mining Report
- Time Domain (transients) Report

OPERATIONS

- Monthly Operations Report
- Configuration Management Plan
- Monitoring Service and Operations Metrics

USER SUPPORT

- User Forum
- VO Community Day – Cambridge, MA
- VO Community Day – Los Angeles, CA
- Exhibit at 219th AAS Meeting – Austin, TX
- "Tools for Data Intensive Astronomy" Workshop, 219th AAS Meeting – Austin, TX
- Community Day – Tucson, AZ

1.5.3 Year Three (May 1, 2012 – April 30, 2013) Milestones

MANAGEMENT

- Program Execution Plan
- 1st Quarterly Report
- Response to PEP
- 2nd Quarterly Report
- 3rd Quarterly Report
- Program Council Meeting
- Program Execution Plan

- Annual Report (this document)
- Response to PEP Review

STANDARDS AND INFRASTRUCTURE

- Spring IVOA Meeting Winter IVOA Meeting Python API

SCIENCE APPLICATIONS

- Data Discovery Portal v1.3 Data Discovery Portal v1.4 Data Discovery Portal v1.5 Iris (SED Builder) v1.1
- Iris (SED Builder) v.1.2
- Iris (SED Builder) v 2.0 Scalable Cross-Comparison (SCC) v1.1 Scalable Cross-Comparison (SCC) v1.2

OPERATIONS

- Monthly Operations Report VAO Assets Inventory

Professional Engagement

- Summer School for Statistics Sagan Workshop, LSST All-Hands Meeting, VO Community Day – Michigan
- VO Community Day – Baltimore, MD
- Exhibit at 221st AAS Meeting – Austin, TX
- "Tools for Data Intensive Astronomy" Workshop, 219th AAS Meeting – Austin, TX
- Community Day– Tucson, AZ
- VAO Documentation Inventory

1.5.4 Year Four (May 1, 2013 – April 30, 2014) Milestones

MANAGEMENT

- 1st Quarterly Report
- 2nd Quarterly Report
- 3rd Quarterly Report
- Annual Report (this document)

STANDARDS AND INFRASTRUCTURE

- Spring IVOA Meeting
- Winter IVOA Meeting
- Python API

- Upgrade Registry
- Sci Drive released to astronomers at JHU
- Command-Line VO Services
- PyVO Release
- VO Client Implementations in C Test Release October 2013
- NRAO Review of Client API for Access To Data Cubes
- White Paper: Mapping Use Cases On To Proposed Standards
- Demonstration data access and discovery service for data cube
- VAO Single-Sign-On Services
- Deliver Registry Update to Test
- Deliver Publishing Interface To Test
- VO Client C Delivered to test

SCIENCE APPLICATIONS

- Archive software in repository
- Data Discovery Tool v 1.5
- SED Builder (Iris) v.1.2

OPERATIONS

- Monthly Operations Report
- Deploy Close-out Code Repository
- Archive SVN Repository

DOCUMENTATION SUPPORT

- Updated videos for the Science Applications
- Identified IVOA Partner to Curate Assets
- Include initial set of documents in repository
- Transfer Material from Presentations, Community Days, etc. to Repository
- Deploy Close-Out Documentation Repository

YEAR FIVE (MAY 1, 2014 – SEPTEMBER 30, 2014)

- 1st Quarterly Report
- 2nd Quarterly Report
- Close-Out Review
- Close-Out Review Report
- Close-Out deliverables as identified in Section 1.4.

1.6 Personnel

There were no major changes in VAO key personnel during the final year of the project. The key personnel are:

Director	Robert Hanisch (STScI)
Program Manager	Bruce Berriman (IPAC)
Project Scientist	Joseph Lazio (JPL)
Technology Advisor	Alexander Szalay (JHU)
Business Manager	Maricel Claro (AUI)
Operations Lead	Thomas McGlynn (HEASARC)
Documentation (formerly User Support) Lead	Sarah Emery Bunn (CACR)
Science Applications Lead	Janet Evans (SAO)
Standards and Infrastructure Lead	Ray Plante (NSCA)

Lazio assumed the position of Project Scientist following the death of Dr. David DeYoung in December 2011. In October 2012, Maricel Claro assumed the position of Business Manager following the resignation of Marie Huffman, and Sarah Emery Bunn assumed the role of Documentation Lead from Elizabeth Stobie (NOAO). The positions of E/PO lead and Technology Assessment Leads were eliminated in 2011 as part of project de-scope following budget reductions.

Hanisch left the project at the end of July 2014 to assume the position of Director of the Office of Data and Informatics, Materials Measurement Laboratory, National Institute of Standards and Technology in Gaithersburg, Maryland. He arranged with NIST and cleared with federal ethics lawyers, the ability to carry out fiduciary responsibilities of the VAO, LLC as needed. E. Schreier (AUI) is the formal representative of the VAO project to NSF effective July 27, 2014.

Over the life of the project, approximately 69 FTE were spent on VAO work (NSF and NASA combined); we estimate another 2 FTE in in-kind effort. 70% of the effort was from men and 30% of the effort was from women.

1.7 Financial Status

1.7.1 Award Status

The cooperative agreement award expiration date is April 30, 2015, and since VAO will have continued access to the funds for an additional 120 days, VAO management believes the LLC will not need to file a no-cost extension.

The total funds authorized by NSF from May 15, 2010 to September 30, 2014 were $11M. Actual expenses as of June 2015 were $10.665M or 97% of the authorized funding, leaving available funds of $0.335 M or 3 %.

Table 1-1 illustrates the breakdown of the $10.665M spending.

NSF Total Obligated Budget to VAO:	**$ 11,000,000.00**	
	Expenses to date	
Board Meetings, program council, Science counc	$ 57,252	1%
Auditing/Legal Costs	$ 194,227	2%
Office Expenses	$ 243,889	2%
Travel Expenses/Other meeting expenses	$ 362,495	3%
Sub-award Expenses	$ 9,807,502	92%
	$ 10,665,364.84	96.96%
Remaining Budget as of July 2015	$ **334,635.16**	3.04%

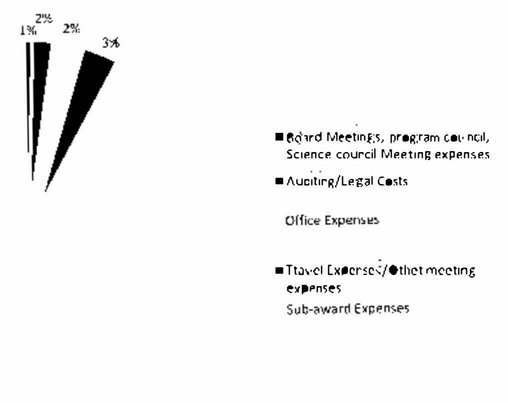

Table 1-1: The bulk of the funding was spent on the sub-awards; $9.808M or 92% of the total $10.665M, roughly 8% was spent on various overhead costs such as Board, Program Council, and Science Council meetings, auditing and legal costs, miscellaneous office expenses such as shared office and utility expenses with AUI, and travel expenses.

Total Sub-award Budget : $ 10,011,406.00

Sub awardees		YTD Expenses	
AUI	$	428,992	4%
CALTECH (excl. IPAC)	$	1,332,518	14%
JHU	$	1,046,768	11%
NOAO	$	1,294,573	13%
NRAO	$	1,358,940	14%
SAO	$	1,744,111	18%
STScI	$	1,539,942	16%
UIUC	$	1,061,656	11%
Total	$	9,807,501	100%

Actual Exp Vs budget 98%

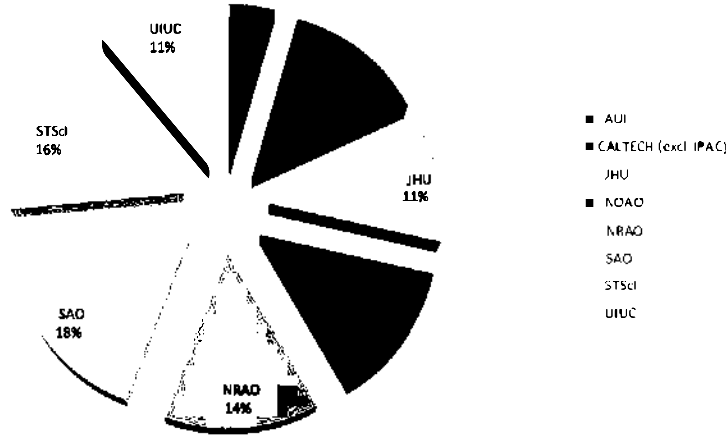

Table 1-2: The $9.808M sub-awardees' spending was allocated among different organizations: all sub-awardees' rates range from 11-18% percentile, except for AUI—4% for Business Management; and SAO and STScI range a little over 16% since they are the larger organizations with more tasks, plus the VAO Director, Dr. Robert Hanisch, was being paid by STScI.

	YTD Expenses		NSF Budget		Remaining	
AUI	$	428,992	$	458,364	$	29,372
UIUC	$	1,061,656	$	1,076,188	$	14,532
JHU	$	1,046,768	$	1,046,768	$	(0)
NRAO	$	1,358,940	$	1,361,272	$	2,332
NOAO	$	1,294,573	$	1,296,547	$	1,974
CALTECH	$	1,332,518	$	1,332,518	$	0
STScI	$	1,539,942	$	1,689,412	$	149,470
SAO	$	1,744,111	$	1,750,337	$	6,226
VAO, LLC	$	857,863	$	978,146	$	120,283
Contingency	$	-	$	10,448	$	10,448
Total	**$**	**10,665,364**	**$**	**11,000,000**	**$**	**334,636**

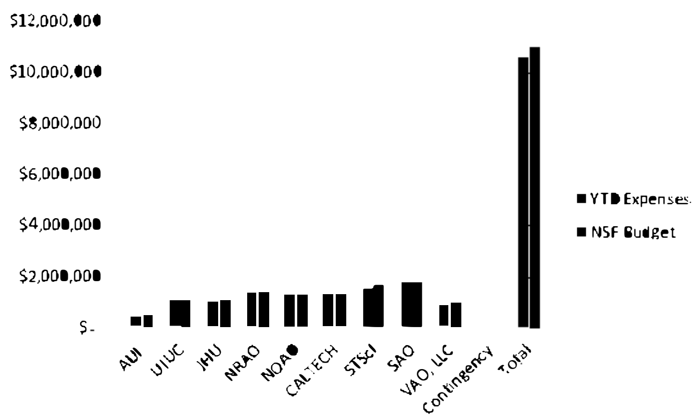

Table 1-3: Actual expenses vs. budget. Every organization managed to stay within budget.

1.8 Accomplishments

The accomplishments of the VAO are described in the following sections, but may be summarized as follows:

- Major contributor to IVOA standards: Appendix C contains a list of IVOA standards to which VAO staff has contributed. The list includes standards recommended by the IVOA Executive Committee and those submitted to the Executive Committee for recommendation.
- Leadership of IVOA Working Groups and Interest Groups. Section 1.3 includes a compilation of the leadership positions.
- High-level science applications for data discovery, integration, analysis, and catalog cross-comparison.
- Scripting toolkits that allow scientists to incorporate VO data directly into their reduction and analysis environments.
- A robust operational environment in which worldwide VO services are routinely checked for aliveness and compliance with IVOA standards.
- Community engagement through AAS meetings, summer schools (NVO), and community days (VAO).
- A comprehensive web site with online tutorials, announcements, links to both U.S. and internationally developed tools and services.
- Take-up of VO standards and infrastructure within essentially every major data center and survey project in the United States, with ~1M VO-based data requests per month and ~2,000 unique users.

1.9 The VAO Legacy

The impact of the U.S. VO programs on the virtual observatory overall can be seen in a number of ways:

- Significant contributions to at least 23 IVOA standards and documents, summarized in Appendix C, from the first basic standards and services (VOTable, Simple Cone Search) to sophisticated data models and advanced data access protocols (Table Access Protocol, ObsCore, Simple Image Access Version 2, etc.).
- Leadership of numerous IVOA Working Groups and Interest Groups, as well as leadership at the IVOA Executive level.
- A rich infrastructure for data discovery and access, with wide deployment and implementation at major data centers in the U.S.
- A robust operational environment in which distributed services are routinely validated against IVOA standards.
- A system of resource registries that enables discover of data and data services through the world.
- Exemplar science applications for data discovery, spectral energy distribution construction and analysis, and catalog cross-comparison.
- Desktop scripting tools, including a native Python implementation.
- Cloud-based data storage for collaborative research and simple data sharing with the research community.
- Creation of a "data scientist" position at the American Astronomical Society whose responsibilities include "to help process and manage the increasing volume of digital data and to integrate it within the Virtual Observatory."
- A repository of all VAO products: software, documentation, tutorials, videos, newsletters, etc.

However, it is not easy to measure impact quantitatively. Since the VAO is mostly about the deployment of software tools and infrastructure services, it can be challenging to attribute data accesses to the VAO as opposed to the underlying data services. Web applications are primarily entry points to VO services; scripting environments are needed for bulk processing. In the astronomy community at least, and probably in many other disciplines, new software can take many years to penetrate the community, and even then, there is not a strong culture of software citation. For example, we find that although some 22,000 peer-reviewed papers mention the VLA radio telescope, only 68 formally acknowledge the use of AIPS and only 59 acknowledge use of CASA, the two dominant reduction and analysis packages for radio interferometry data. Remarkably (or perhaps not, given the situation for software citation) of over 13,000 peer-reviewed publication in astronomy and astrophysics published in 2013, only 4% acknowledge use of the ADS, and the ADS is probably the most widely used software system in the field. Thus, counting acknowledgments to VAO or VO tools is unlikely to reflect accurately on community take-up.

On the other hand, VAO usage logs indicate close to one million VO-based data accesses per month at U.S. data providers, and with ~100 organizations who have published some 10,000 VO-compliant data services worldwide. VAO usage logs also show ~2,000 distinct users of VAO services in the three month period April-June 2014. The ADS lists some 2,500 papers (about half of these peer-reviewed) citing "virtual observatory" in some context, and these papers are read as often and cited as often as other types of papers. Of course, without reading each and every paper one cannot be sure of the level of contamination in this sample (a paper saying "our observatory has photometry measurements of virtually thousands of stars" would count as a hit). We have tried to track research publications that are truly based on VO-tools and resources, but we do not have the time to do this with any level of completeness. The ~100 papers we are aware of are listed at http://www.usvao.org/support-community/vo-related-publications/.

The journal *Astronomy and Computing* (Elsevier) published two issues dedicated to VO technology, and VAO team members are contributing a number of papers. The first issue has 12 articles and the second issue has 14. The VAO team wrote an article for the second issue describing the accomplishments of the VAO project.

In order to maintain a U.S. national presence at the IVOA, and to help ensure ongoing cooperation between ground-based and space-based data centers in astronomy, a U.S. Virtual Observatory Alliance has been established under the auspices of the AAS Working Group on Astronomical Software.

The VO concept has been adopted in numerous other fields, particular in space science (with seven VxOs within NASA), plus the Virtual Solar Observatory (NASA, NSO), Planetary Science Virtual Observatory (Europe), and the Deep Carbon Virtual Observatory (Rensselaer). The VO concept was recently endorsed by a panel of neuroscientists convened by the Kavli Foundation and General Electric as a means for improving access and interoperability to the vast data sets being collected in the European Brain Project and U.S. Brain Initiative. VAO and IVOA participants are now playing leading roles in the international Research Data Alliance and the U.S. National Data Services Consortium.

1.10 Lessons Learned

In looking back over the VAO project and its NVO predecessor, a number of lessons learned are apparent. (A discussion of the lesson learned specific to the science applications development is given in that section.)

- Successful infrastructure is largely invisible and unappreciated.
- Deployment of a distributed infrastructure takes more time than you think. Community consensus and buy-in requires early and ongoing participation.
- It is important to do marketing to the research/user community, but also to manage expectations. Promising too much is as bad or worse than delivering too little.
- Toolkits such as PyVO and VOClient probably have more value to researchers than high-end applications, as they provide more flexibility. VAO should have developed these toolkits earlier.
- Managing a distributed project has pros and cons.
    - Pros: Access to a diversity of skills and different environments for validating technical approaches and implementations
    - Cons: Coordination of efforts takes time; staff members have competing priorities as most are not working on VAO full-time.
- Setting up an independent management entity such as the VAO, LLC is a non-trivial effort, though in the VAO case it has proven worthwhile and effective.
- Top-down imposition of standards is likely to fail.
- One must balance requirements coming from the research community with those coming from innovation in information technology—science-driven, technology-enabled.
- Coordination at the international level is essential, but takes time and effort.
- Data models are important, even in cases where they seem obvious.
- Metadata collection and curation are essential but ongoing tasks.

2 Operations

2.1 Introduction

The VAO operations effort addressed two primary goals. The first was to enable science use of the virtual observatory, especially through the VAO developed interfaces, and through VO-interfaces generally. Tools needed to work consistently, and when problems arose, they required swift resolution. The second goal was to enable the services needed internally for the activities of the VAO itself. VAO personnel needed reliable access to the tools needed for software design and access, user support, testing, configuration management, and bug tracking.

The VAO provided a number of science services and tools directly to the scientific community: its home web site, a data portal and cross-correlation tool, the Iris SED tool, downloadable VO libraries for use by clients and servers, and cloud storage and secure access protocols. Internal services include the VAO infrastructure: our JIRA ticket system, a Jenkins testing service, SVN code repository, a YouTube channel, a blog, and mailing lists. The VAO managed the IVOA web site, document repository and mailing lists; these services have now been transferred to our international partners. A close-out repository using free Google cloud-based services has been established and populated to ensure VAO-developed resources continue to be available after the termination of funding.

VAO services were supported by member institutions of the VAO with significant resources hosted at each of our sites: CfA, JHU, MAST, HEASARC, NRAO, NOAO, and Caltech (IRSA and NED). Elements were by necessity distributed across the country and across the Internet.

Supporting such a distributed system posed special operational concerns. Especially for our science users, we worked to ensure that elements were seen as a coherent whole: science tools were available from a common location, forms had a consistent look-and-feel, and everything was visible through a consistent web-presence, even when the web sites were on various servers. A common VAO style-sheet was developed and distributed to participating sites. The usvao.org web domain was used for disparate elements of the VAO hosted by Caltech (the primary web site), NOAO/NRAO (the issue tracking system) and MAST (the data portal) so that the elements of the VAO appeared seamless to users. This integration was not pursued to completion after it became clear that lifetime of this project was finite.

All elements were continuously monitored and a responsible party identified for each so that issues were rapidly and decisively addressed. The operations staff met frequently (in short weekly telecons) and operational issues were rapidly escalated using our issue tracking software to whatever level needed to ensure they were resolved.

2.2 Monitoring the VAO

All VAO services were monitored hourly and a database of all tests was continuously updated. Each service was tested to ensure not only that the service is up, but that it responded sensibly to some simple request. When services failed a test, they were retested 15 minutes later. If the second test also fails, a message was automatically sent to the responsible parties and to the VAO operations monitor.

A web site is available giving the current status of all operational services and the VAO home site reflected the operations status of VAO science services so users were immediately informed if there was an issue. Statistics were collected in and reported in biweekly periods.

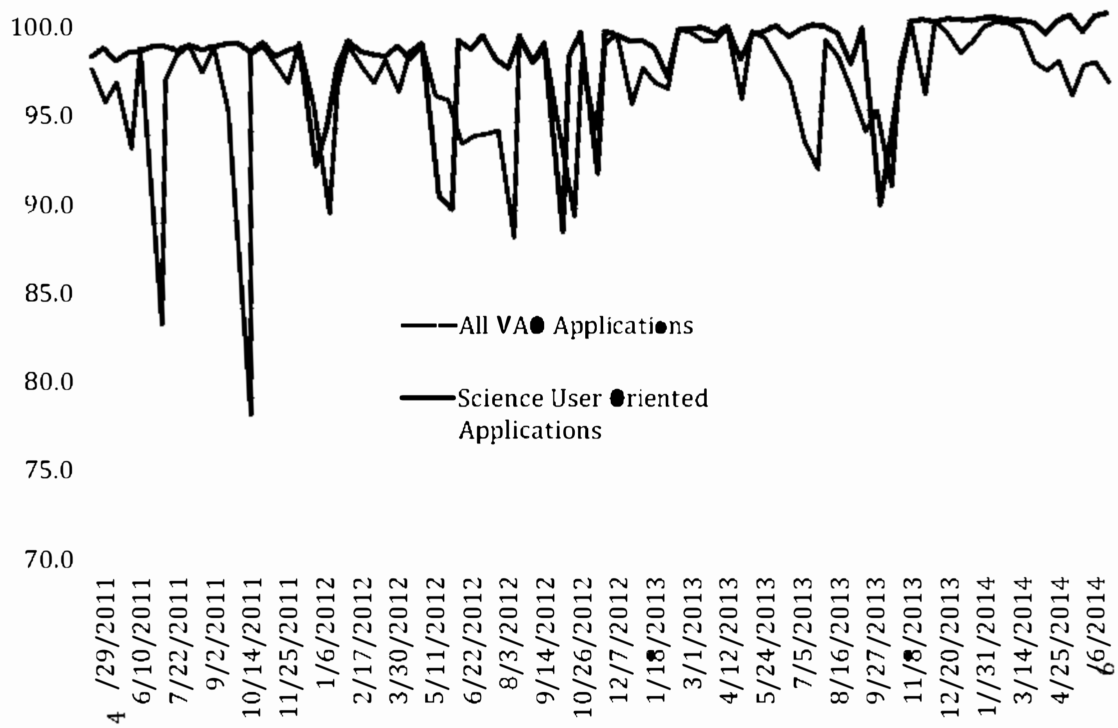

Figure 2-1: This figure gives the operational status for all VAO services since early 2011 in each biweekly period The blue line shows some of the internal VAO services had a number of significant periods with limited availability. This was reflected mostly in our testing and validation tools, which were subject to occasional outages when the validation services hung. This did not affect science users. The more critical red line indicates only one significant lapse, in October 2013, in our science-oriented services since early 2013. This was directly due to the federal shutdown that affected services at NASA sites. Other significant downtimes in our science services included a shutdown of HEASARC services due to a security concern and downtime at MAST due to a hurricane that caused widespread power outages.

2.3 Monitoring and Validation of VO data providers

Since the effective operation of the VAO from our science users' perspectives required that VO data providers' services were available, the VAO tested the aliveness of VAO services and whether other VO data providers were working. Every site that published data through the VO was tested hourly, though not all published services were tested. Instead, a representative service from each of class of services at each site was tested. All tests were recorded and the current status of all VO sites could be accessed on the VO monitoring web site.

When a problem was detected, the VAO Operations monitor contacted the responsible party and noted the problem. In many cases, the VAO assisted sites in rapidly bringing their services back online.

Occasionally, a VO data-providing site goes out of service. When sites were no longer responsive after two months, we deprecated them in the VAO registry so users no longer saw them in typical queries.

Each week, VAO Operations handled 5-10 service interruption issues. In addition to testing for service availability, the VAO also validated every published VAO service using the VO's catalog/table, image, spectral or registry services. Validators for the simple cone search and image search services were developed within the VAO, but we also used validators developed in Europe for the spectral and table access protocols. Each day, approximately 300 services were validated and all validation issues were recorded in a database, so all published services were validated roughly once per month. Periodically, a summary report describing the VAO validation issues was prepared for each site that gave concrete recommendations for resolution of validation issues.

Note that most services that did not pass validation still provide valuable information, but getting more complete agreement with the IVOA standard ensured that tools work more robustly.

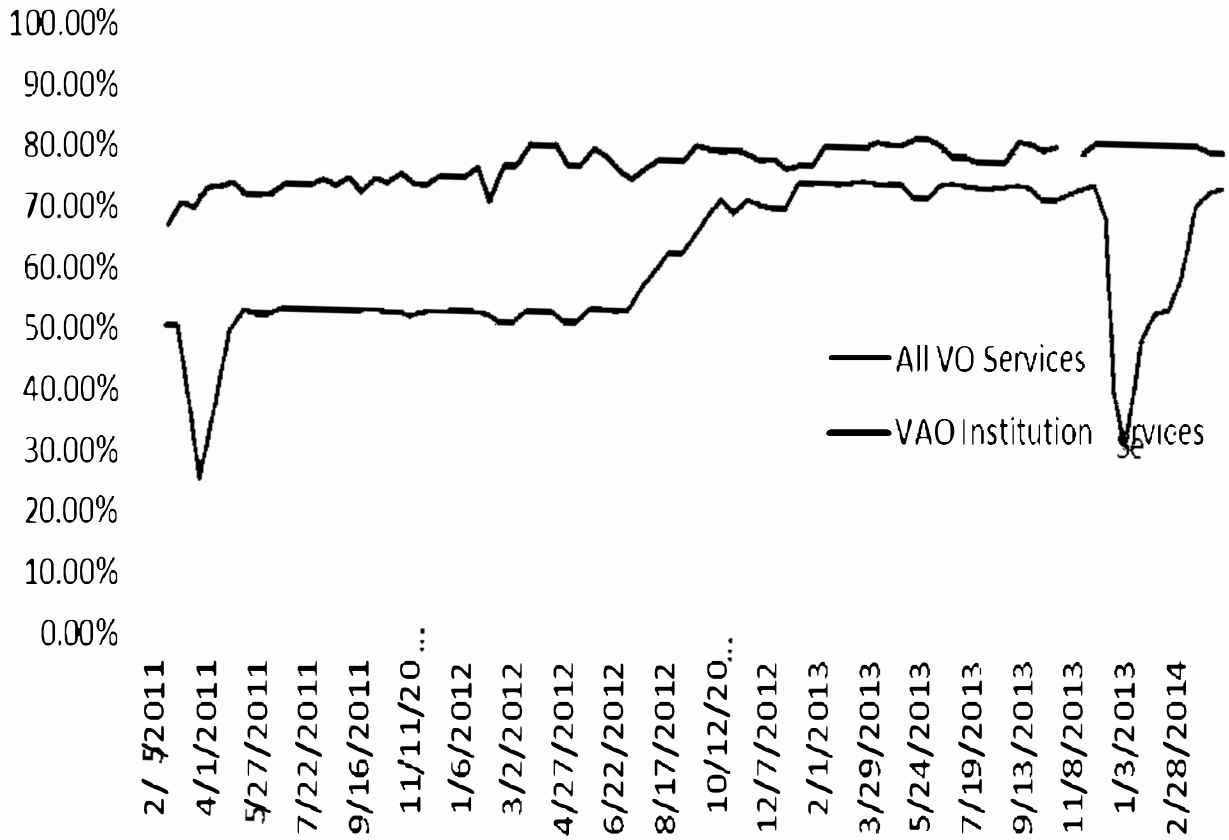

Figure 2-2: The figure above shows the fraction of VO services that completely passed validation. The blue line shows all VO data providers, while the red line gives the services associated with institutions that are part of the VAO (often the VO-data-providing capabilities are distinct from the specifically VAO activities). In both cases, there has been a steady rise in compliance over the past several years. Two major drops in the overall compliance reflect bugs introduced at one of the major VO data providers. Seeing these declines, our operations monitor worked with the provider, identified specific services that were affected after initial bug fixes did not completely rectify the problem, and helped their recovery

2.4 Post-VAO Operations

The VAO web site has been transferred to IPAC and will be maintained by IPAC beginning October 1, 2014 to provide convenient documentation of VAO-developed resources that will, in many cases, continue to be developed or at least maintained by the VAO partners who developed them.

The VAO close-out repository has been fully populated and contains all digital artifacts of the VAO, including, notably, all code developed as part of the project. Each major software project has its own page within the repository. In most cases this includes links to groups who plan to maintain or continue development of the software post-VAO.

All VO monitoring and validation activities will be continued after termination of the VAO project by the HEASARC as part of the NASA's coordinated VO efforts. The HEASARC will use the software and databases developed as part of the VAO and continue to refine them as needed. Monitoring of VAO services will also continue, but at a low level of effort and with no significant upgrades to this capability anticipated. It is likely that equivalent monitoring of the NASA-supported capabilities will be provided as part of the NASA VAO follow-on.

3 Documentation Services, User Support and Professional Engagement

User Support provided the primary interface between the VAO and its user community, assembled and curated documentation from across the project, made the information easy to find through the web site, and notified the community of updates and changes relevant to the VO infrastructure. Toward the end of the project, the name of the work area was changed to "Documentation Services," reflecting the balance of the effort needed to support close-out.

During the course of the VAO, effort was undertaken to ensure products and services delivered were robust and usable by research scientists and to reach out to the broader astronomical community. There were outreach efforts aimed to expose VO products and services to potential users, to assist in the take-up of those products and services, and to gather feedback to assure the maximum utility of the VO for astronomical research.

3.1 Web Sites, Help Desk and User Forum

The VAO web site's intended audience was professional astronomers and software developers. The web site was designed both to serve as an entry portal to the VAO and to provide a means for astronomers to find information about the VO—of the more than 3 million results of a search for "virtual observatory" with Google, the VAO web site is one of the top responses.

From the perspective of the end user, the web site had two key areas. The first was "Science Tools & Services." This web document provided access to the web services and software developed by the VAO. Further, as the project began to mature, community-provided tools and services began to be developed, and links to those tools and services were added.

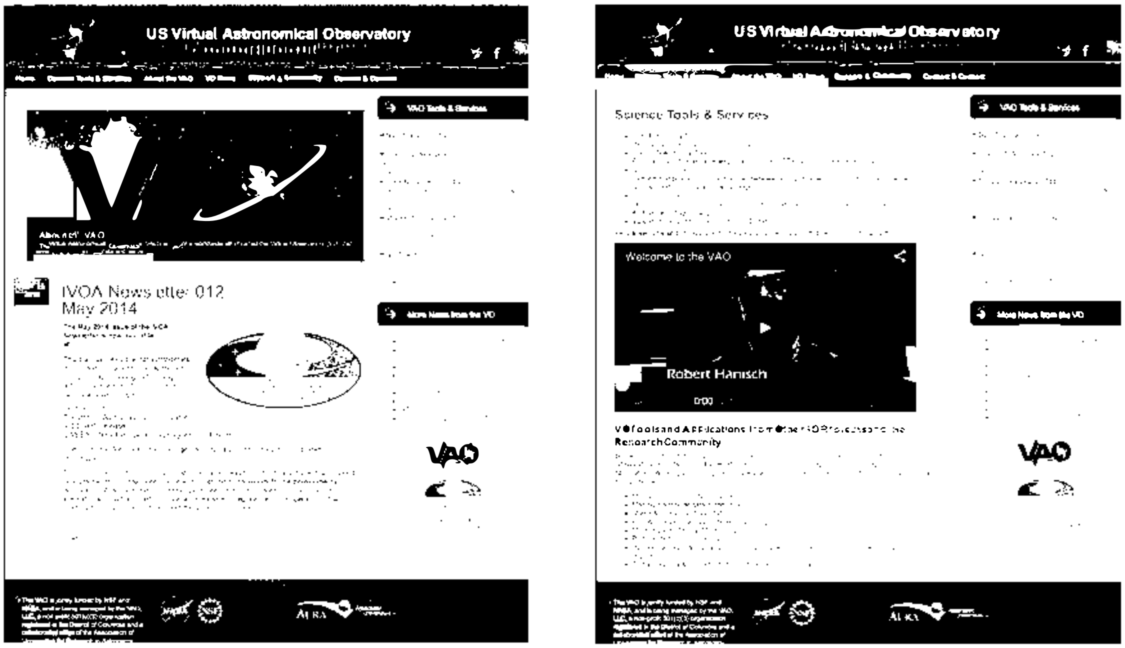

Figure 3-1. (Left) VAO web site home page. (Right) Science Tools & Services area within the VAO web site. Tools and services developed by the VAO appear at the top of the document; VO tools and services provided by the community appear as well

The second area of interest for end users was "Support & Community." Analogous to the "knowledgebase" that might be provided by a commercial software provider, this area was designed to help users find answers to their questions, contact other users, or submit bug reports.

Over the course of the project, a total of 104 news items were posted to the VAO web site. These news items, summarizing Virtual Observatory news for the astronomical research community, are also pushed to subscribers via RSS, Facebook, and Twitter. The social media presence of the VAO will discontinue at the end of the project.

During the final months of the project, the VAO web site contents and ownership were transferred to IPAC.

Documentation Services staff monitored and responded to VAO Help Desk inquiries and referred them to other team members for resolution/answers. Statistics on help desk tickets can be found in the Operations section of this report. The history of help desk tickets will be archived in the VAO's online repository. Issue tracking post-VAO will be managed via a GitHub project.

The VAO Forum (Astrobabel.com) has been maintained and backed up by Documentation Services staff. During the lifetime of the forum, there have been 136 users registered and 42 discussions started, including 108 comments and approximately 25,000 page views of forum discussions. The forum will be supported in its current location for 6 months past the VAO close-out. A list of support requirements has been developed for hosting of the forum post-close-out.

Documentation Services staff have also managed the IVOA web resources, including the web site, mailing lists, and wiki. During the last month of the VAO project, these resources were transferred to VO Italy (wiki, mailing lists, domain ownership) and VO India (web site and document repository).

3.2 Product Testing

At the beginning of the VAO, quality control and testing activities were under the purview of User Support. The motivation for this structure was that User Support could serve as a proxy for the end user and ensure the products and services could be used in a research setting. For most testing activities, the User Support role was to act as the coordinator of the activities and as reviewers. In addition, User Support took the lead for performing User Acceptance Testing (UAT), which was used, along with other tests and quality control reports, to prepare readiness reviews.

3.3 End-User Documentation

Documentation Services staff wrote or completed user documentation to help research scientists have a better understanding of VAO services and applications and how to use them. Documentation packages included deployment instructions, general descriptions, tutorials, cookbooks, and similar documents. Documentation staff and product developers also collaborated to produce video tutorials, which were then made available through a YouTube channel. All software documentation produced is available in the VAO Repository[1] and the video tutorials remain available through YouTube[2].

The close-out work focused on the documentation section of the VAO's close-out repository. A documentation area of this repository is populated with tutorials and user guides for VAO tools and services, outreach materials from "VO Days" held in 2011 and 2012, and project documents and reports. Also included is a list of technical presentations and research publications by VAO

[1] https://sites.google.com/site/usvirtualobservatory/home/documents

[2] https://www.youtube.com/channel/UC08B3AjmWu6bf0sOMH334zA

Team members. The close-out versions of VAO-produced software had their documentation updated as part of their final release. Final archival packages of the VAO web sites, wiki, and mailing list archives are also included in this repository. More details about the close-out repository can be found in the Operations section of this report.

3.4 Scientific Collaborations

During the course of the project, the VAO supported the scientific or technical work of multiple individuals or collaborations. The requests for support resulted both from *ad hoc* proposals to the VAO and from a formal call for proposals that the VAO issued in 2012. The following is a summary of the projects and work supported.

- "NVO Collaborative Proposal for Magellanic Cloud Database and GRAMS Model Grids" (PI: M. Meixner) [2012 Call for Proposals]

  This proposal requested assistance with publishing a catalog and set of models into the VO. After evaluation by the VAO Science Council, it was recommended these data and models could naturally be hosted at the Infrared Science Archive (IRSA). The VAO provided the appropriate contacts within IRSA for the team. The VAO assisted with the implementation of a TAP server to host these data and models.

- "Real-Time Analysis of Radio Continuum Images and Time Series for ASKAP" (PI: T. Murphy) [2012 Call for Proposals]

  This proposal requested assistance in describing multi-dimensional radio wavelength data and publishing it to the VO. Interaction with this team was used as a key use case in developing the VAO Standards & Infrastructure effort toward multi-dimensional data and in interactions with the IVOA.

- "Network Tools for Science Queries to NED" (PI: J. Schombert) [2012 Call for Proposals]

  The proposal was to support travel by the PI to NED to interact with its staff to develop additional tools. The VAO attempted to schedule the requested travel to NED, but it was not possible to confirm a set of dates that were compatible with the PI's availability.

- "Integration of AAVSO Data Archives into the Virtual Astronomical Observatory" (PI: M. Templeton) [2012 Call for Proposals]

  This proposal requested assistance in publishing data from the American Association of Variable Star Observers into the VO. The VAO provided assistance to the AAVSO, and the data are now available. A poster summarizing the support was presented at the 221st AAS Meeting (2013 January; Long Beach, CA; Kinne et al. 2013).

- "Cosmic Assembly/Near-infrared Deep Extragalactic Legacy Survey (CANDELS)" (PIs: S. Faber & H. Ferguson)

  The VAO supported the CANDELS program by distributing supernovae detections with the VOEvent network and providing access to CANDELS images through standard VO image access protocols. CANDELS supported the VAO program by providing guidance on requirements for SED building and analysis tools. A poster summarizing the support was presented at the 221st AAS Meeting (2013 January; Long Beach, CA; Greene et al. 2013).

- "Brown Dwarf Candidate Identification Through Cross-Matching" (PI: S. Metchev)

  The VAO supported a project that continued a search that had begun during the NVO for extremely red L and T=-type brown dwarfs. It involved cross-comparing the 2MASS and SDSS catalogs to identify candidates that were followed-up with spectroscopy at the Infrared Telescope Facility, Mauna Kea. The project identified the two reddest known L dwarfs and nine probable binaries—six of which were new and eight of which likely harbor T dwarf secondary stars—and derived an estimate of the space density of T dwarfs. (Geissler et al. 2011).

In addition to these scientific collaborations, a scientifically motivated sub-award was issued to produce a cross-matched multi-wavelength catalog of more than 1M objects within a 10° radius of the SMC ("A Catalog of Spectral Energy Distributions of Stars in the Small Magellanic Cloud," PI: B. Madore). The catalog has been placed in the VAO Repository, and it has been incorporated into NED's online services with value-added content.

3.5 Booths and Exhibits at American Astronomical Society Meetings

American Astronomical Society (AAS) meetings, principally those occurring during the winter, are one of the focal points for the U.S. (and international) astronomical community. During the course of the project, the VAO had exhibit booths at AAS meetings (Figure 3-2). The use of an exhibit booth built on experience gained from NASA Archives and National observatories, for which it was found that substantial fractions of the community could be engaged at low cost. As an illustration of the value of an AAS meeting, people stopping at the exhibit were offered the opportunity to sign up for the VAO mailing list. At each AAS meeting, the size of the VAO mailing list increased by approximately 20%.

Figure 3-2. Images from the VAO Booth at the 221st American Astronomical Society Meeting, Long Beach, CA (2013 January).

3.6 VO Community Days

VO Community Days were a series of presentations and hands-on activities designed to take the VAO to the astronomy community, demonstrate capabilities, develop and encourage new users, and obtain feedback on VO tools and services (Figure 3-3). Community Days were typically structured with a morning session led by VAO team members, with the option of an afternoon session for attendees to ask more detailed questions to VAO team members or to bring in their research questions to assess how VO tools and services could assist them. Community Days were aimed initially at locations where there were a large number of astronomers with the goal of making it easy for many to attend. Table 3-1 lists the VO Community Days that were held. In addition, two VO Community Days (at the University of Washington and Cornell University) were being planned when the VAO was directed to discontinue them in preparation for its close-out activities.

Center for Astrophysics, Cambridge, MA	November 2011
California Institute of Technology, Pasadena, CA	December 9, 2011
University of Arizona, Tucson, AZ	March 13, 2012
University of Michigan, Ann Arbor, MI	November 14, 2012
Space Telescope Science Institute, Baltimore, MD	November 27, 2012

Table 3-1. VO Community Day Locations and Dates

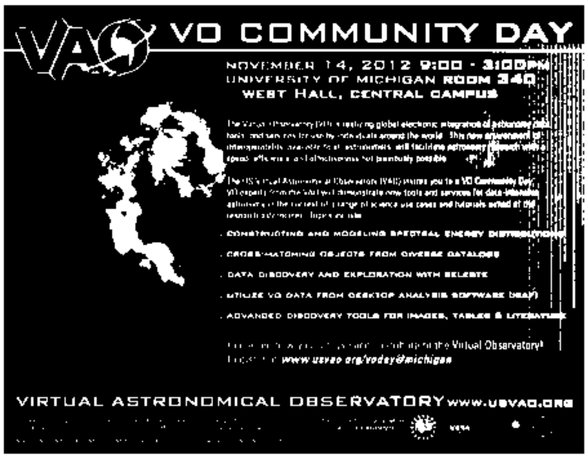

Figure 3-3. (Left) Example of an announcement flyer for a VO Community Day. (Right) Scene from a VO Community Day, in which a VAO team member is demonstrating how a VO tool could be used. Both of these examples are from the Community Day held at the University of Michigan.

3.7 Publications, Presentations, and IVOA Standards

VAO team members were active in giving technical presentations and publishing technical papers. Appendix C contains a list of technical presentations given by VAO staff, and Appendix D lists the technical papers. VAO team members were authors or co-authors of technical papers, conference presentations, and IVOA standards documents. Particularly in the case of IVOA Standards, participation by VAO team members was crucial to achieving consensus and developing standards that could be adopted by a wide community. The VAO Repository[3] contains a full listing of the publications and presentations. A notable aspect of the publication rate is that, in the initial phases of the project, there was a clear indication the number of publications and presentations was increasing with time.

[3] https://sites.google.com/site/usvirtualobservatory/home/documents/publications-presentations

4 Science Applications

The VAO developed three science applications (Data Discovery Tool, Interoperable SED Access and Analysis tool, and the Catalog Cross Comparison Service) and one prototype application (Time Series Search tool). The science applications were developed to provide useful science discovery and analysis tools to a heterogeneous user community. The applications do not serve any one observatory, wavelength, or type of user, but serve astronomers with multi-wavelength data from possibly a variety of telescopes that span the electromagnetic spectrum.

Some of the challenges facing the science applications group were developing to the IVOA standards in order to test and validate the applicable standards, provide feedback to the standards, and be active contributors to resolve any problems identified. The group also worked to deliver useful by-products so other developers could use a library or service developed for the applications. Two examples are the SEDLIB (SED i/o library) and NED/SED service developed for Iris. Lastly, the tools were interoperable with other VO tools and there was an overall goal of paving the way for other VO-enabled tools to be built and contributed by the user community. By way of encouraging contributions, several collaborations (e.g., ASDC archive plug-in for Iris) were fostered during VAO science applications development.

The VAO development framework helped the distributed team focus on the requirements, design, and implementation of complex applications in a shared, lightweight wiki-based environment. This framework enabled the group of developers working on this project, at a distributed set of institutions and working on a part-time basis, to perform their tasks and collaborate efficiently. A science stakeholder was assigned to each application and was key to bringing the view of the user to the development process. The stakeholder provided requirements, developed science use cases, handled technical questions, advised on development priorities, and performed unit tests. The team lead managed priorities, schedule, and communication within the group. Careful management and web-based availability of schedules, meeting notes, and meeting actions helped keep the distributed and part-time science applications team on a productive path.

4.1 Data Discovery Tool

The Data Discovery Tool (DDT) is a web application for discovering all resources known to the VAO about an astrophysical object or a region of the sky. Using protocols defined by the IVOA, the DDT searches those widely distributed resources and presents the results in a single, unified web page.

Many of the most popular U.S. archives and catalog holdings are available for searches in DDT, including HST, MAST, Chandra, HEASARC, SDSS, Spitzer, and 2MASS, to name a few. A powerful filtering mechanism allows the user to narrow the initial results to a short list of likely applicable data quickly. Guidance on choosing appropriate data sets is provided by a variety of integrated displays, including an interactive data table, basic histogram and scatter plots, and an all-sky browser/visualizer with observation and catalog overlays (Figure 4-1).

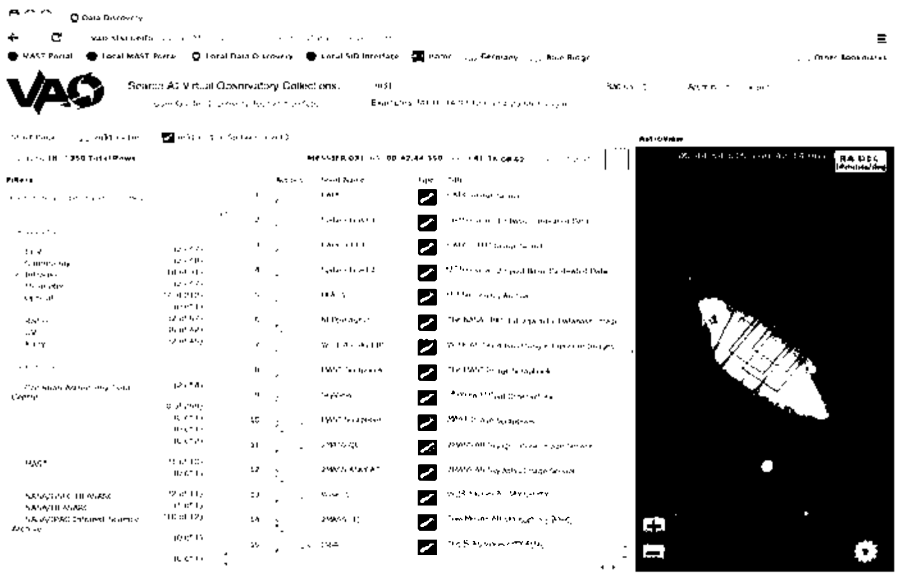

Figure 4-1: Appearance of the DDT after a search for M31 within radius = 1 arcmin, showing the filters (left panel) that can be applied to the Search Results (center panel), and the AstroView component with FOV overlays representing the available data sets.

DDT was developed incrementally with the first release of the application in June 2011. Development continued over the next two years with five incremental releases that added features and addressed any deficiencies. DDT web-based user documentation and training videos were developed and updated for each release (Table 4-1).

The DDT project utilized DataScope (GSFC/NVO) and Astroview (STScI) and shared synergy with the MAST archive development project at STScI. IVOA standards feedback was substantial and included input on registry standards where the DDT project advocated for enhanced metadata, table access protocol improvements, input on data access protocols to ensure support for bulk queries, and co-authorship of MOC (HEALPix Multi-Order Coverage map) standards by the DDT lead developer Tom Donaldson.

DDT Version	Release Date	New features
v1.0	June 30, 2011	Initial release
v1.1	Jan. 05, 2012	Automatic facets, SAMP image broadcasts, bug fixes, performance improvements for large tables, and displays
v1.3	June 25, 2012	Spectral data through SSAP; Server-side paging & filtering; small bug fixes
v1.4	Oct. 08, 2012	Spectral data through SSAP; Server-side paging & filtering; small bug fixes
v1.4.1	Dec. 19, 2012	DDT query length handling, image icon enhancement to display
v1.5	May 30, 2013	Selective source search, preservation of table metadata, remove flash, Jira tickets

Table 4-1: The Data Discovery tool release schedule including version number, release date, and feature highlights. Updated web-based user documentation and user training videos were part of each release.

4.2 Interoperable SED Access and Analysis (Iris)

Iris is a downloadable graphical user interface application that enables astronomers to build and analyze wide-band spectral energy distributions (SEDs). SED data may be loaded into Iris from a file on the user's local disk, from a remote URL, or imported directly from the NASA Extragalactic Database (NED) for analysis via the NED/SED Service. A plug-in component enables users to extend the science function of Iris. Iris utilized Sherpa from CXC/CIAO and Specview from STScI as the components that performed fitting and visualization in the application. Communication between Specview and Sherpa is managed by a Simple Application Messaging Protocol (SAMP) connection. Data can also be read into Iris and can be written out via the SAMP interface. A separable library for SED data input/output (SEDLib) is also included and available independently from Iris (Figure 4-2).

Iris was first released in October 2011. Three incremental releases and one bug fix release followed (Table 4-2). Iris is supported on several flavors of the Mac OS X and Linux. User web-based documentation and user training videos are also provided to the users. Iris was featured on the Astrobetter blog in September 2013.

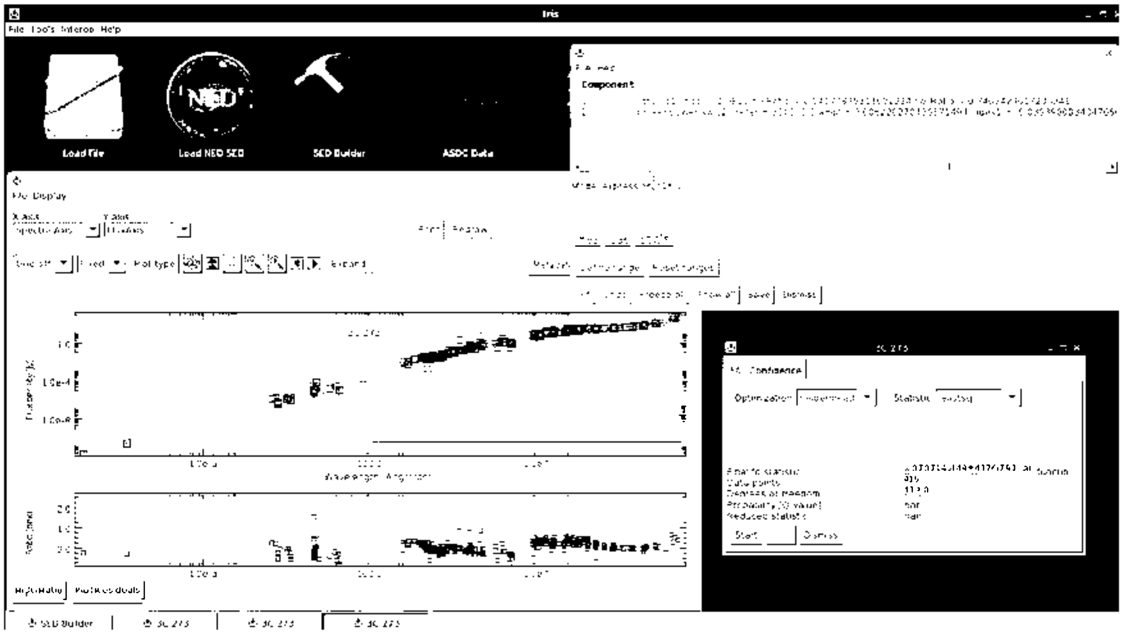

Figure 4-2: The SED access and analysis tool Iris in operations. The Iris desktop holds the interactive windows for SED data review and analysis.

Iris Version	Release Date	New features
v1.0	Oct. 15, 2011	Initial Release
v1.1	Aug. 22, 2012	New Desktop Metadata Browsing, Sherpa 4.4, Advanced Importer, Template & User Model
v1.2	Dec. 19, 2012	Visualization upgrades, Plug-in support, bug fixes
v2.0	July 01, 2013	Co-plot, Redshift/blueshift, integrated quantities, upgraded ASDC plugin, Sherpa 4.5, bug fixes
2.0.1	Nov. 26, 2013	Magnitude err fix; added Tooltips

Table 4-2: The Iris release schedule including version number, release date, and feature highlights. Updated web-based user documentation and user training videos were part of each release.

There were two by-products of the Iris project: the NED/SED service and the SEDLib. There were collaborations with several groups including the ASI Science Data Center (ASDC) and CDS (Strasbourg). The collaborations led to Iris desktop plug-in services to access the respective SED data holdings. The project provided IVOA feedback to the SAMP protocol to enable a full SED into a single file extension, to TOPCAT for better support for SED plots, and inspired work toward a VODML (Virtual Observatory Data Model Language) by lead Iris developer O. Laurino. The group authored several papers (including publications in *SPIE* and *Astronomy and Computing* journal), presented at ADASS and AAS conferences, and participated in outreach presentations including a VO Day in Italy.

Release	Linux	Mac	Total
V1.0	81	84	**165**
V1.1	88	71	**159**
V1.2	43	72	**115**
V2.0	58	121	**203**
V2.0.1	51	63	**114**

Iris averaged 160 unique software downloads per incremental release (Table 4-3). Even though citations for software in published journals are shown to be underrepresented, we can report four citations in ADS to Iris published works. Two of those citations are by non-VAO authors.

4.3 Scalable Cross Comparison Service

The Scalable Cross Comparison (SCC) Service performs fast positional cross-matches between an input table of up to 1 million sources and common astronomical source catalogs for a user-specified match radius. The service returns a list of cross-identifications to the user. The output is a composite table consisting of records from the first table, joined to all the matching records in the second table, and the angular distance and position angles of the matches (Figure 4-3).

The Scalable Cross Comparison Service first released in January of 2012 and was supported with three upgrades over the next ~1.5 years (Table 4-4). The complex indexing schemes that support big data were provided by IPAC for the VAO and later adapted to the WISE and Spitzer projects.

SCC Version	Release Date	New features
V0.9	Jan. 06, 2012	Initial release
v1.0	April 08, 2012	Added caveat for long jobs; bug fix
v1.1	Dec. 19, 2012	Added survey catalogs, added SAMP, few bug fixes
v1.2	April 19, 2013	Added DR9 catalog

Table 4-4: The Scalable Cross Comparison tool release schedule including version number, release date, and feature highlights.

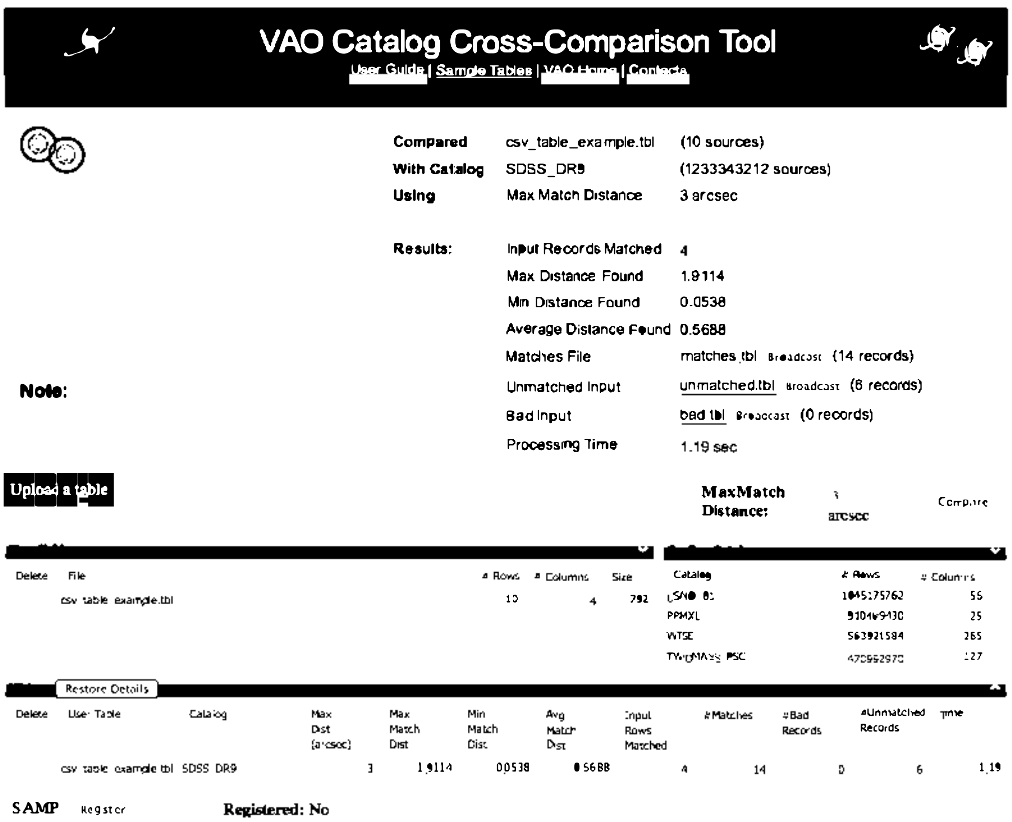

Figure 4-3: The results view of the Scalable Cross Comparison service. The service shows results of a user table cross-matched with SDSS Data Release 9.

4.4 Time Series Search Tool

The Time Series Search tool was used to discover time series data from three major archives and analyze them using the NASA Exoplanet Archive periodogram service. The application was a prototype developed to inform the design of IVOA protocols and data models for time series data sets. The development of the Time Series Search tool ended after the first VAO replan due to budget reductions.

4.5 Lessons Learned

The management strategy for the VAO science applications group was published as an *SPIE* paper ("Managing Distributed Software Development in the Virtual Astronomical Observatory") in 2012. Managing a distributed team with successful results was found to require managing around unknowns. Since managers were unaware of the entire set of external tasks assigned to an individual outside of the VAO efforts, coordinating task assignments and making organizational material and schedules easily available was important. Working within a well-defined lightweight process ensured the project was cohesive and stayed on track. Participation of a stakeholder who represented the user by responding to technical questions and advising on priorities was very important to a science application development effort. Internal product deliveries provided a planned test-and-feedback loop, and incremental releases (rather that one big software release) ensured feedback early in the process that development was progressing as expected. Use cases provided data challenges and provided an opportunity to assess priorities and make course corrections. Frequent communication regarding schedules assured that issues were resolved quickly and the team was working toward a common vision. Easy-to-access project information ensured the team resumed progress quickly when side-tracked due to external project responsibilities. The distance gap of distributed teams required management with diligence. With that, projects stayed on track.

The science applications team evolved into a cohesive team during the first few months of the VAO project. Management efforts were important in building and sustaining the efforts. The team that was assembled worked collaboratively, participating in the projects independent of institutional boundaries. The team was willing to contribute in any way to get the job done successfully.

5 Standards and Infrastructure

A useful way to frame a summary of the infrastructure built over the course of the VAO project is by considering the goals we laid out in the Standards and Infrastructure section of the project close-out plan.

We wanted to ensure that the VAO infrastructure would be on a sure footing for transition to NASA operations. At the time we were developing our close-out plan, it was not clear how much support would available to NASA centers for VAO operations. Thus, we wanted to make sure VAO infrastructure (particularly the Registry) could operate well in a "maintenance mode." If and when funding became available, there would be a firm foundation for new VO development efforts to build on.

We also wanted to make sure data providers have what they need to add new data collections and services to the VO web. We saw two new capabilities that were important to furthering growth of the VO after the VAO project is over, and we set it as a goal to push them forward. The first capability is access to multidimensional image cubes, and the second is sharing and publishing small data collections.

We note that our close-out plan was designed to be aggressive toward finishing the essential work for meeting these goals and making use of all the time and resources in our remaining schedule. Not unexpectedly, our successes diminished somewhat as we approached the project end and key personnel were pulled to work on new projects.

5.1 A VO Architecture Primer

In this chapter of the report, we summarize a variety of services deployed and products released. To understand how all of these components work together, it is worth summarizing the VO Architecture. Figure 5-1 shows the VO architecture as described by the IVOA.

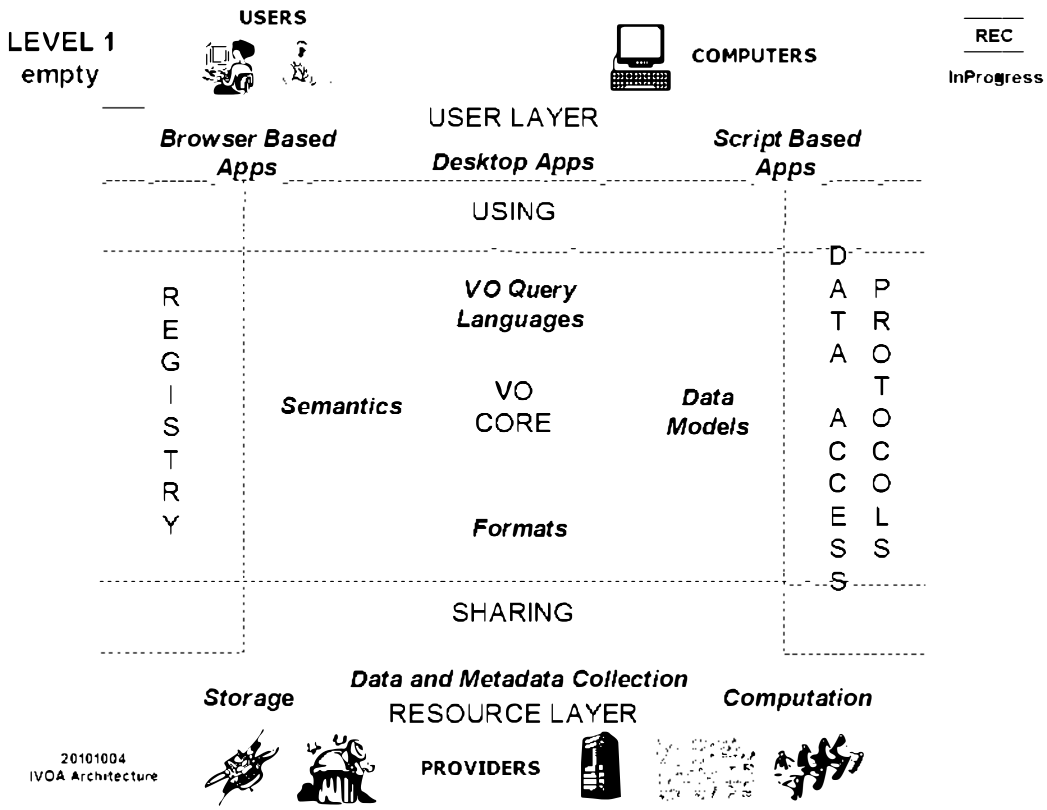

Figure 5-1: The IVOA architecture for VO infrastructure, connecting users to data

In this diagram, the VO infrastructure serves as a bridge between data providers and users, and that bridge is supported by standards. On the provider side, data are connected into the infrastructure through standard services that present that data in terms of standard data models. On the other side, users are connected to the infrastructure via generic tools that understand the VO standards. No longer are tools tied to a single archive, but rather they can talk to any and all archives that speak the common VO language.

A big part of the VO is about data discovery, so it is worth briefly reviewing how discovery is enabled. Figure 5-2 illustrates the discovery framework.

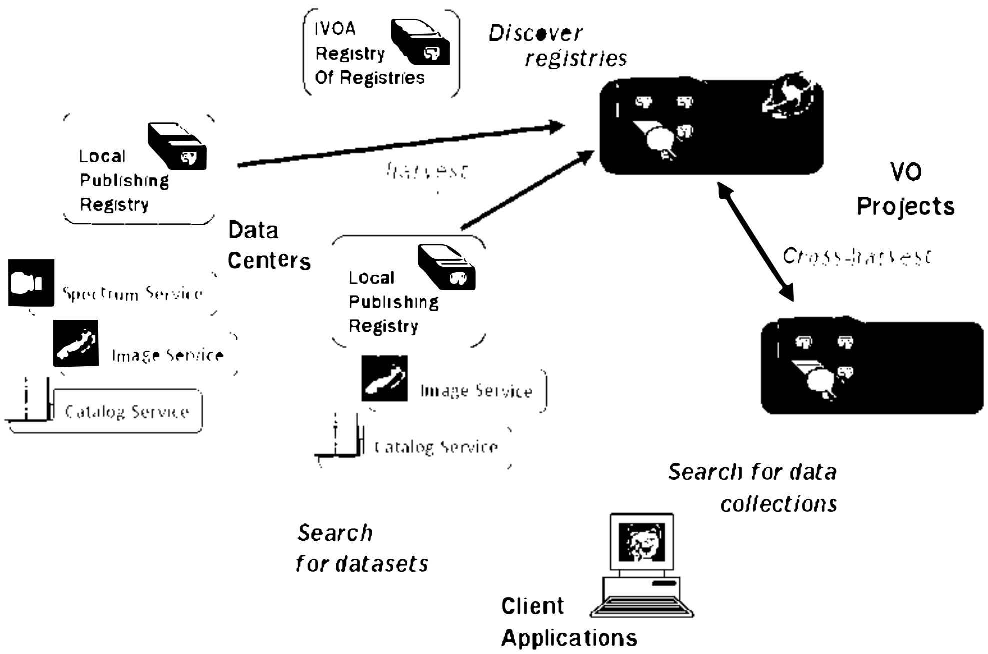

Figure 5-2: Data discovery in the VO.

In this framework, registries represent the first step in data discovery. A *registry* is a database containing descriptions of data collections and services available in the VO. In the VO, there is no master or central registry; however, there are registries called *full searchable registries* that have descriptions of all the data collections, archives, and service providers known to the VO from around the world. This type of registry can populate itself through a process known as *harvesting*; it starts by contacting a special "boot-strapping" registry (run by the VAO) called the *Registry of Registries,* which will return to it all of the other known registries in the VO ecosystem. (When a new registry enters the ecosystem, it registers itself with the Registry of Registries.) Most of the registries it learns about from this query are so-called *publishing registries.* A registry of this type is typically run by a data center, which uses it to advertise the data collections and services it offers to the VO. The full searchable registry then contacts each of the publishing registries and pulls descriptions of all the data collections and services provided by the data center. At this point, the full searchable registry is then populated with descriptions of all of the resources known to the VO. Periodically, it will re-query the other registries to get any new resources added or other changes since the last harvest.

With an up-to-date full searchable registry available to it, a client application can now discover any data known to the VO. It starts by asking the registry for a list of collections and services from each of the data centers that might have data relevant to the user's science question. Most of the services will be standard *data access services* for finding and downloading images, spectra, or catalog information in a particular archive or collection. The application can then send a query to all of the matching services to get back lists of available data sets. By browsing the returned metadata for these data sets, the user can choose which data sets to download.

5.2 IVOA Engagement

Much of the work the VAO does on standards is through engagement with the International Virtual Observatory Alliance (IVOA). The role of the IVOA is two-fold: first, to coordinate the efforts of all of the VO projects around the world, and second, to serve as a standards body for establishing VO interoperability. This activity is organized around a set of working groups and a document standardization process based on that of the World Wide Web Consortium (W3C). From the IVOA's very beginnings, VAO and the NVO before were clear leaders in shaping the VO's global architecture and the standards that enable it. Many NVO/VAO members served as chairs or vice-chairs of key IVOA working groups. The impact of this leadership has also seen in the standard documents; most of the IVOA recommendations across all of the areas of the VO have featured VAO team members, either as first authors, secondary lead authors, editors, or major contributors.

We have also produced many of the key reference implementations—software that demonstrates a standard in action and proves its viability. The NVO project created the first implementations of registries with several different architectures. We have also been instrumental in demonstrating data access services through software packages like DALServer and TAPServer. The NVO/VAO has also led the IVOA in the development of so-called *service validators*. A validator is a service that can check whether another service is compliant with VO standards. It accomplishes this by sending a series of queries to a purported standard service and examines the response and assesses whether is follows all of the rules and recommendations spelled out in the standard. The NVO developed the first validators in the IVOA to assist data providers, allowing them to check their data access services and fix any problems before publishing them to the VO. These quickly became critical pieces of VO infrastructure and other projects joined in to contribute validators for other service standards.

5.3 The Registry

As described above, a registry is database of descriptions of data collections, archives, services, and other resources useful to the VO, and it represents the first step in data discovery. NVO/VAO established itself as an early leader in the area of registries. In addition to creating some of the first registries, the VAO operated the Registry of Registries on behalf of the IVOA, which allows searchable registries to bootstrap their collection of resource descriptions.

While we were developing several different types of registries for the VAO project, we consolidated our support around the production of a full searchable registry service at STScI.

5.3.1 VAO Directory Service

As part of our production registry, we delivered a web browser-based front end called the *Directory Service* (http://vao.stsci.edu/directory). This tool is particularly useful for discovering collections and services related to a topic. By entering keywords into the search input box, the tool returns a list of resources that contain those keywords in their descriptions (Figure 5-3).

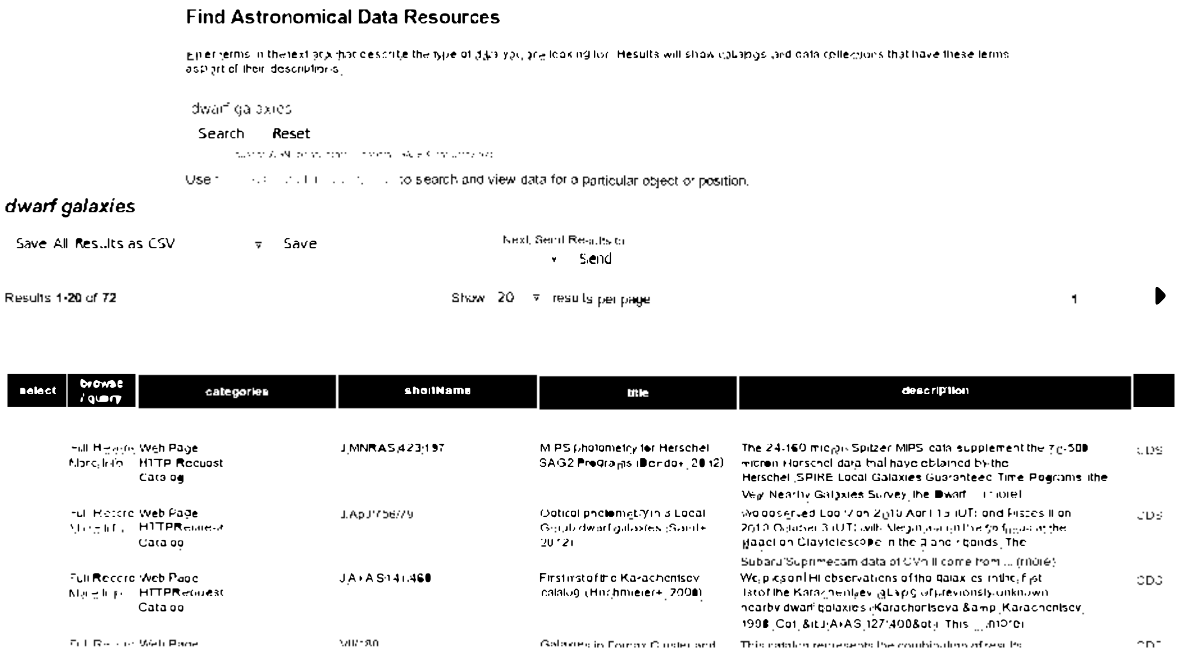

Figure 5-3: Results from a search query submitted through the VAO Directory Service. The Directory Service results allow one to browse the descriptions, filter results, download matching descriptions in VOTable format, and—when the resource is a standard data access service—even send it a sky position-based query.

5.3.2 Registry Upgrades

Over the last two years of the VAO project, we carried out a program to overhaul the underlying registry database. In particular, it was updated to support the latest IVOA registry metadata standards. This overhaul was also necessary to support a newly emerging standard for searching registries. This new standard leverages an existing IVOA standard for querying complex databases generally called the Table Access Protocol (TAP), for which there already existed client software. (The TAP standard did not exist when the first registry search interfaces were standardized.) We considered this upgrade to the database as a critical task in transitioning the registry to its maintenance phase.

In the closing months of the VAO project, we completed some registry curation activities aimed at improving the descriptive content of the registry. In particular, we implemented a specific approach to registering resources intended to make registry searches more effective and their results less confusing. This approach was recently accepted as a best practice by the IVOA Registry Working Group. Our curation work has developed an inventory of existing resources by publisher, and has made recommendations for improving the resource descriptions that bring them in-line with the best practice. The new registry Resource Publishing Tool (described below) communicated these recommendations to the publisher.

5.3.3 ***Publishing Registries and the Resource Publishing Tool***

A publishing registry is the vehicle for making a resource available to the VO. In particular, it can create new descriptions of resources and share them with the rest of the VO through the harvesting process. A data center, which may curate a number of data collections and offer a variety of services to access them, may operate its own publishing registry. Because such a registry does not need to serve end users directly, operating one is much simpler than running a searchable registry. During the NVO project, we developed a product called *VORegistry-in-a-box* that provides a simple but compliant publishing registry implementation through which a data center can maintain its own resource descriptions internally. This product has been used within the VO (including by the Registry of Registries), and so the VAO continued to support it.

We note that a searchable registry can also support the publishing function, and indeed the VAO Registry at STScI does just this. In particular, it maintains resources descriptions on behalf of data providers who only have a few resources to share, relieving them from having to run their own publishing registry. To enable this feature, we created the *Resource Publishing Tool*, a browser-based application that allows a data provider to create and share resource descriptions through the VAO Registry. It features a wizard-like interface that steps a data provider through the process of describing a resource, prompting for metadata along the way. It can check for the validity of values as they are entered, alerting the user of any problems. Draft descriptions can be saved for updating and publishing later, and even published resource entries can be updated with this tool. Various techniques are used to minimize the amount of typing required to create a useful resource description. While the VAO Registry shares records created through this tool, the descriptions are considered "owned" by the user. Thus, to control access, the publishing tool uses the VAO Single Sign-on services (described below). Following a review in August 2014, the documentation was expanded and the tool delivered to end-users. MAST/STScI assumed support maintenance of this service on October 1, 2014.

5.3.4 *Data Access*

Standard services that allow users to find and access to data from an archive are part of the VO architecture known as the *data access layer* (DAL). In the VO architecture, there is a standard service for each type of data set; e.g. the *Simple Image Access* (SIA) protocol enables discovery and downloading of images from an archive, and the *Simple Spectral Access* protocol enables access to spectra. In this section, we describe infrastructure developed to support the data access layer in the VO.

*5.3.5 **DALServer***

To help data providers share their data collections through standard VO services, we created the *DALServer Toolkit*, a Java-based software platform developed primarily by D. Tody (NRAO). When it was first developed as part of the NVO project, it served as a platform for developing reference implementations of standard VO services (like SIA and SSA) that demonstrated features of the standards. This year, we enhanced the toolkit for use directly by data providers; we referred to this effort as "productization," as it focused on making the toolkit easier to use. We focused on a simple class of use cases in which a small data provider had a simple catalog or a simple collection of images or spectra on disk they wished to share. By just editing configuration files and running a few scripts, the provider could deploy highly compliant VO services with no programming required. For more complicated situations, such as for a data center that might already operate custom data access services through its own data management system, it could use the underlying DALServer Library API to adapt the VO services to the local infrastructure.

The first production release of the DALServer was scoped to provide support for the four "simple" standards for data access recommended by the IVOA, namely: Simple Cone Search (for simple position-based querying of object and observation catalogs), Simple Image Access (for finding images), Simple Spectral Access (for finding spectra), and Simple Line Access (for finding rest frequencies for spectral line emissions). All the essential capabilities are available in the code repository and will serve as a basis for post-VAO development by NRAO. In addition to its use within the emerging NRAO archive (see 5.3.8 Image Cube Access), NCSA plans to deploy DALServer to expose the public release of images from the Dark Energy Survey (DES). By doing so, we will give DES users access to a suite of existing software tools that understand the SIA protocol and thus can find and retrieve DES images. This low-cost data access solution will be important to a project lacking the resources to build a traditional archive portal. NOAO is also considering DALServer toolkit integration into their infrastructure.

5.3.6 TAPServer

The *Table Access Protocol* (TAP) is an IVOA standard for querying complex catalogs that may be made up of several tables (such as the 2MASS catalog). When a TAP service is connected to a catalog, users can create complex, SQL-like queries that may join metadata from several tables. Such queries are critical for mining very large catalogs. Not surprisingly, given its power and flexibility, a TAP service is one of the more complex IVOA standards to implement. To make deploying a TAP service easier, we have created the *TAPServer toolkit*.

TAPServer was developed at JHU by D. Nandrekar, and, like DALServer, it is configuration-driven—that is, with no programming required, one can wrap the toolkit around a collection of tables in a database and deploy it as a service accessible to the VO.

Because the VAO close-out schedule limited how much we could complete in the time remaining in the project, TAPServer was not slated for production release. However, the code is included in our close-out repository where it will be available for use. Some targeted deployments are planned. For example, it will be deployed at NCSA to expose the DES Source Catalog. This will allow DES scientists to analyze the catalog using the Seleste TAP client, a tool developed at Harvard/CfA (P. Zografou) that allows users to form complex queries with little or no knowledge of SQL.

5.3.7 ***Service Validators***

During the VAO project, we continued to maintain and extend service validators first developed by R. Plante (NCSA) during the NVO project. These validators have a web browser interface that allows a data center to enter a service access URL and test the service's compliance with the appropriate standards; the result is a listing of errors, warnings, and recommendations for improving the service. These validators share a common Java-based toolkit platform called *DALValidate*. They also support a programmatic interface that allows VAO Operations to automatically test VO services. (Operations also utilizes other validators developed outside of the VAO.) Supported validators include those for Simple Cone Search (SCS), Simple Image Access (SIA), Publishing Registries, and VOResource records. The DALValidate software is available through our close-out repository.

5.3.8 ***Image Cube Access***

As mentioned in the introduction to this section, advancing access to multidimensional image cube data was as a key technical project for the close-out plan. In the IVOA, it was recognized that although the Simple Image Access (SIA) protocol could support image cubes in a limited way, it lacked some of the metadata support and data access mechanisms needed to support the cubes coming out of the leading-edge telescopes such as ALMA, LOFAR, and JWST. For this reason the IVOA identified improved cube support as a key priority for the IVOA as a whole. We saw this project not only as a way to advance a new capability in the VO, but also to bring in and engage the radio and IFU communities in VO activities.

Our work in improved image cube access began in the spring and summer of 2013 when we built an early prototype service that demonstrated a number of the key capabilities needed in a new standard for image cube access. This demonstration was instrumental for mapping out the strategy for a Simple Image Access Protocol Version 2. In particular, the necessary standardization was broken down into three independent documents. First is the Image Data Model, which defines the semantic labels that describe image cubes. These labels are used, then, by the SIAP V2 standard to annotate image search results. The third standard, called Access Data, defines how one can request cutouts or other transformations of image cubes.

While active in the development of the standards, we continued prototyping access to image cube data. To ensure the standards served the needs of real providers of image cubes, we collaborated with NRAO. Our joint goals were first to create a real functional image cube access service based on the emerging SIAP V2 draft serving real data from NRAO instruments, and second, to provide a useful architectural design along with software to support active archive operations. In this collaboration, NRAO provided the VAO project with requirements and use cases. NRAO wanted an image service that can simultaneously provide data to both internal and external clients. One key client is the CASA Viewer that needs to request small, visualizable parts of a larger cube. In return, the VAO project provided NRAO with general purpose software to deliver data over the network to clients. The DALServer product was extended to provide server-side support for SIAP V2, and VOClient (described below) was extended to support the client.

In spring 2014, NRAO, using VAO-provided software, successfully demonstrated a service that provided access to image cube data, including image cutouts. This service allowed their archive and CASA Viewer developers to test against a functional service.

5.4 Data Sharing

The second key project we took on in the close-out plan was more exploratory in its full scope (though it supported an important end-user application). We wanted to understand how to support access to so-called "long-tail" data—the many small collections of data products that ultimately result in published papers. Such data products tend to be highly processed by individual astronomers and, thus, are not typically available from traditional observatory or project archives. In particular, we wanted to understand how these products could be published to the VO in a low-effort way; to enable such access, we surmised, required fitting into the overall science result publishing process that starts even before the first draft of a paper.

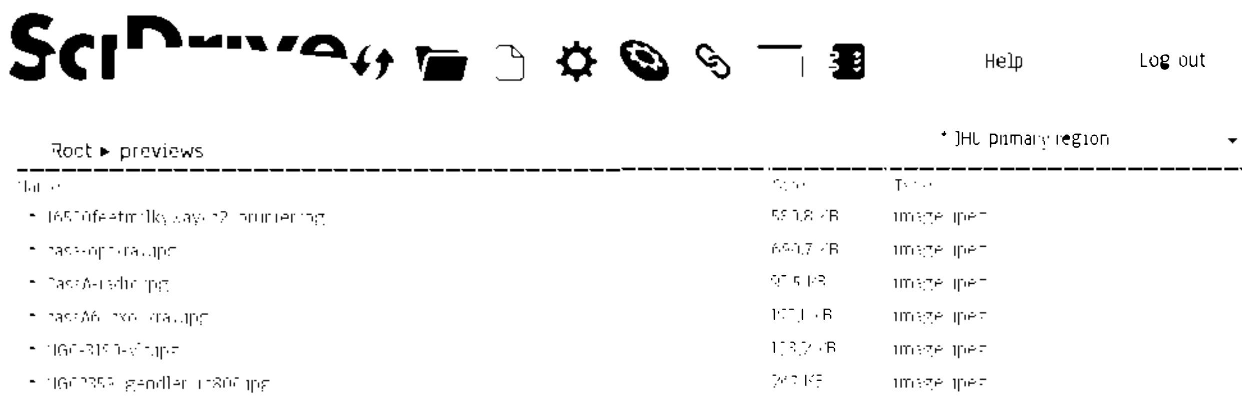

Figure 5-4: The SciDrive file browser, viewed via a web browser

5.4.1 *SciDrive*

SciDrive (scidrive.org) is a Dropbox-like cloud storage application intended for use in scientific research (Figure 5-4). It was developed primarily by D. Mishin (JHU), who based it on the OpenStack software (in particular the OpenStack Swift component for object storage). It can be accessed from a web browser, which presents the user with a view of a personal hierarchical directory space into which a user may save files by drag-and-dropping file icons into the web page interface. Also available is a desktop client that can (like Dropbox) monitor a local directory and automatically upload files that are moved into it. As many researchers already do with Dropbox, SciDrive can be a simple platform for sharing data within a research group; it provides secure means to share read-write access to a collection openly within a restricted group or to send one-off permission (read or read/write) to individuals. One difference from commercial storage providers is SciDrive's ability to scale to larger collections than the typical free versions of storage.

Also, unlike Dropbox, SciDrive supports the VOSpace 2.0 interface, the IVOA standard for managing third-party transfers; this allows a user to seamlessly move files among different SciDrive instances (or other VO-compatible storage systems) located around the network. We see this feature as important for the next phase of the VO in which web-based tools allow users to save tool output to their personal space in the cloud. This output could be reloaded later into the tool for further analysis (e.g. as a "favorite" starting point) or loaded by other tools for synthesis with other data and analysis. A current example of this is the use of SciDrive with the Sloan Digital Sky Survey (SDSS) CasJobs: a SciDrive user can configure a directory to detect uploaded table files automatically and load them into the SkyServer database for correlation with the SDSS catalog.

SciDrive as a Publishing Platform. The CasJobs connection highlights another unique feature of SciDrive: it supports plugins that enable special handling of certain types of data. For this reason, we see this application as a possible platform for publishing data by individual scientists and research groups. The JHU group has experimented with plugins that automatically extract metadata from files needed to expose data to the VO. With such a feature, we can model a data publication process in which a group uses SciDrive to organize a collection of data for publication, engaging plugin tools to extract metadata. When the collection is ready for release, the researcher would press a button to expose the data publically; the metadata would be automatically loaded into a database and the collection would be made available through standard IVOA services (e.g., using DALServer).

Delivering an operational implementation was beyond the scope of the close-out plan; however, we prototyped this scenario through collaboration with NCSA, JHU, and SDSC. Using a development platform at NCSA, we utilized SciDrive's plugin feature to create an interface for adding and editing arbitrary metadata. Further, we added a service that allows that metadata to be exported to other systems, just like the data products themselves. We also experimented with plugging in other federated authentication systems to make it interoperable with external repository and publishing systems.

Further development on SciDrive is planned beyond the VAO project. In particular, operations and development at JHU of the SciDrive platform continues under an NSF-funded DIBBs grant. Further, current VAO partners NCSA and JHU are also collaborating on the emerging, community-driven initiative called the National Data Service (NDS), which aims to address data publishing across all research fields. As the publishing scenario described above is much like one being discussed in the NDS community, we expect to see further progress on SciDrive as a publishing platform.

SINGLE SIGN-ON SERVICES. To restrict access to the user's personal space, SciDrive uses the VAO Login Services for authentication (sso.usvao.org). These services were created so VO users could have a single login to connect any VO-compatible service or portal, even when they are managed by different organizations. More than the simple convenience of a single login, a federated login system allows users to integrate their proprietary data from one archive into the tool of another archive. A participating portal (e.g. an archive) can choose to support VAO logins either as its primary identity or one that augments its local authentication system.

The VAO federated login is built on the OpenID standard that is in broad use across the Internet. Associated with it are all the usual services that help users manage a login: the ability to reset forgotten passwords, edit the user profile, etc. It also leverages an OpenID feature for sharing user information with a portal in a privacy-conscious way; this can make registering users with a portal faster and simpler. One less common feature that is important for VO applications is the ability to transparently deliver X509 certificates to the portal. This allows a portal to access private data at another site on the user's behalf. While the service requires the user's permission to do this, it is worth noting the user never handles the certificates directly.

The VAO Login Services were released and implemented software products. The first is called *VAOSSO*, which provides a user identity server that powers the VAO services. This software can be configured either to run as a mirror of the VAO service (for high availability) or as a completely independent service. A second product, called *VAOLogin*, is a toolkit that helps portal developers add support for VAO logins.

Applications using the VAO Login Services include SciDrive, the VAO Registry's Resource Publishing Tool, and the VAO Notification Service. The NOAO Data Archive (which currently supports its predecessor, the NVO Login Service) is migrating to use of the VAO Login Services to augment their own local authentication system.

5.5 Virtual Astronomy on the Desktop

A key initiative the Standards and Infrastructure program was to make VO capabilities integrated into desktop applications, and to support custom scripts for conducting research. Because of its growing popularity as a scripting language for scientific research, Python was a major focus of our scripting support. Further, we aimed to enable all VO-enhanced applications and scripts running on the desktop to work together through the Simple Application Messaging Protocol (SAMP), the IVOA standard that allows desktop and web applications to talk to each other.

VO-ENHANCED IRAF. Our first product supporting VO on the desktop was a VO-enhanced version of IRAF (developed by M. Fitzpatrick, NOAO). This included some general IRAF infrastructure enhancements, including the ability to load data from arbitrary URLs as well as support for loading data in VOTable format. With these two capabilities, we added a suite of tasks that take advantage of VO services; these include an object name resolver, the ability to search the registry to find archives and services, the ability to search individual archives or catalogs, and the ability to download discovered data products. SAMP support was also added so IRAF could send data to other non-IRAF tools running on the desktop; for example, images could be sent to Aladin and catalogs to TOPCAT for visualization.

VOCLIENT. This downloadable product gives users direct access to VO services outside of a web browser. The first **VOClient** release featured a suite of command-line tools that enables interactive use from UNIX/Linux shell; they can also be used to create customized shell scripts. The capabilities provided by these tools include discovering archives and catalogs via the VAO registry, searching individual archives for images and spectra, downloading discovered data across multiple archives, searching catalogs by position, resolving object names to sky positions, and sending data to other desktop tools (via SAMP).

We continued to develop the VOClient package toward a second release that focuses more on the underlying set of core C libraries. These libraries can be used directly to add VO capabilities to C and C++ applications (as is being done for the NRAO CASA Viewer). These libraries are intended to be the basis for bindings to other languages, such as Python and Perl. The Python bindings, in particular, are another focus of the second release (which includes a common API with PyVO, described below). Code development for the second release was completed; however, full user testing was not. The code was placed in the close-out repository and NOAO plans to complete the release and continue further development.

PyVO. This downloadable product represents a parallel effort to support Python with a slightly different focus. Through our community engagement, we found that many Python users would prefer to use a pure-Python implementation of a VO library, and *PyVO* provides this. Like for VOClient, the audience is two-fold, the first being developers who want to integrate VO capabilities into their own Python applications. As an example, Figure 5-5 shows the Ginga image browser, developed by the Subaru Telescope (by E. Jeschke) to preview observatory images; the right-most panel represents a plugin that allows users to download images and catalogs from the VO for display and overlay in the viewer. Jeschke added this capability using the PyVO Python module.

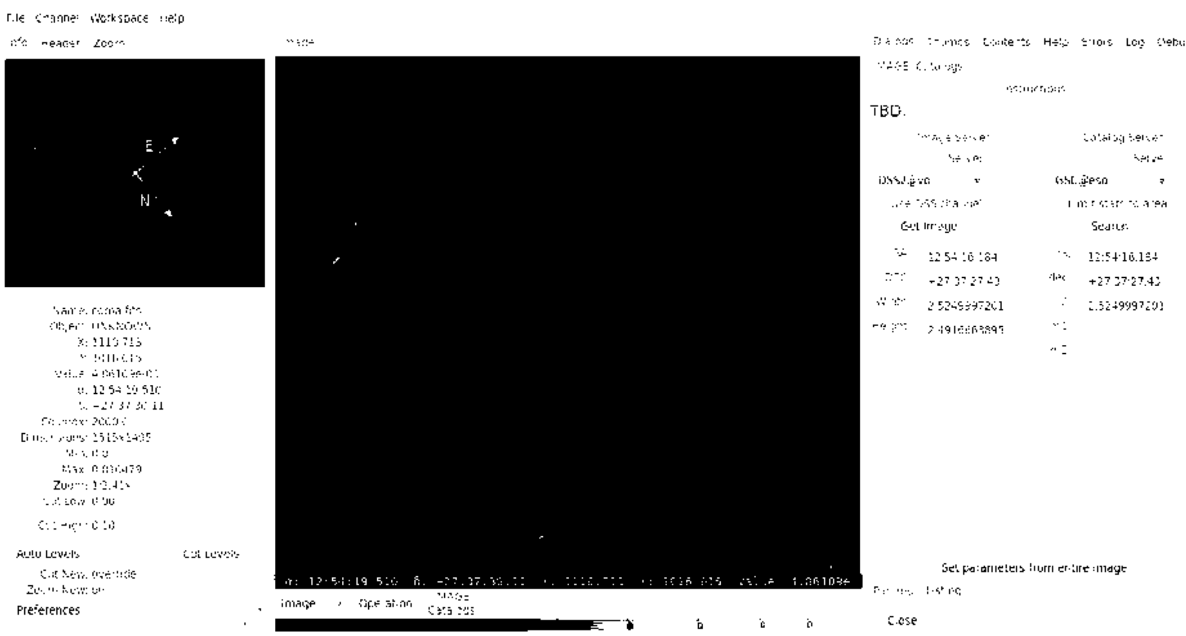

Figure 5-5: The Ginga image browser from the Subaru Telescope showing a VO plugin powered by PyVO.

PyVO is also aimed at the growing community of research astronomers using Python to create custom scripts to carry out their research and analysis. In fact, PyVO is built on top of the widely used Astropy[4] package, an integrated set of astronomically-oriented modules. This allows users to discover and download data and process and analyze it with the robust capabilities of Astropy. This combination is an important key to doing VO science at a large scale, as it becomes

[4] http://www.astropy.org/

very easy to apply a common process to vast arrays of data, either from a single survey or collection or across many. It also becomes possible to continuously monitor the evolving holdings of an archive or the VO in general as new data sets are added.

The first evaluation version of PyVO was released in 2013. As this was within sight of the close of the VAO project, we wanted to ensure further use and development of PyVO beyond the project's end. To that end, we explicitly employed a strategy to build a community around this package. First was to make use of GitHub, the web-based code repository that encourages community contributions. This has enabled important contributions from users outside of the VAO project, including (as of this writing) 22 issue submissions from 7 external users and 7 code submissions from 4 external users. The other part of our strategy was to establish a strong association with the Astropy community, which is quite large and active. (In fact, this association is responsible for much of the external participation via GitHub.) To this end, we were given status as an Astropy "affiliate package." This connection also allows PyVO to become a proving ground for migrating addition VO capabilities into Astropy.

6 Data Curation and Preservation, Technology Assessment and Education and Public Outreach

After Year One, the Data Curation and Preservation, Technology Assessment, and Education and Public Outreach were tasks were discontinued when the project was re-scoped due to budget cuts. This section presents summary reports of the work performed in these three task areas.

6.1 Data Curation and Assessment

The Data Curation and Preservation (DCP) activities concentrated on the design, planning and feasibility of two projects, but budget reductions did not permit deployment.

- Data Capture Project, at IPAC: Its objective was to enable user-friendly capture of data through a web interface as part of the manuscript submission process.
- Semantic Linking Project, at SAO/ADS. Its objectives were to
  - Develop and adapt ontologies to describe astronomical resources and their relationships,
  - Identify, collect, and expose metadata for data sets,
  - Create typed links between data products, bibliographic, and object databases,
  - In the process, create a knowledgebase about instruments and services, and
  - Incorporate all of the above in a "metadata store," exposing resources using semantic web standards.

6.2 Technology Assessment

This task area deployed a wiki page for technology-related material, a blog (*Mitigating the Bleeding Edge*) for more general technology postings, and a Twitter feed for technology announcements. It produced three reports on the topics of collaborative tools, instant messaging, and scheduling tools.

6.3 Education and Public Outreach

A number of collaborative efforts were initiated, and an EPO-focused web site was deployed (www.virtualobservatory. org). Microsoft's WorldWide Telescope (WWT) was adopted as an EPO vehicle, with a special session focused on use of WWT at the January 2012 AAS meeting in Austin, TX in which the VAO had support from Microsoft and the WWT Ambassadors Program. A series of tutorials were developed in which the features of WWT were demonstrated, with detailed write-ups allowing users to explore WWT's capabilities at their own pace (www.virtualobservatory.org/wwt).

Appendix A. Technical Presentations by VAO Team Members

2010

- B. Berriman: "The Virtual Astronomical Observatory and IPAC's Role In It." IPAC Seminar, October 15, 2010.
- R. Plante: "Building Archives in the Virtual Observatory Era", SPIE Astronomical Instrumentation, San Diego, CA/USA. June 2010.

2011

- B. Berriman, "Approaches to Technology Evaluation in the Era of Data Intensive Astronomy,"Innovations in Data Intensive Astronomy Workshop, NRAO Green Bank, May 3-5, 2011.
- B. Berriman, E. Deelman, G. Juve, M. Rynge, J.-S. Vockler, "High-performance compute infrastructure in astronomy: 2020 is only months away." ADASS XXI, Paris, November 2011.
- R. d'Abrusco, G. Fabbiano, G. Longo, O. Laurino, "Knowledge discovery workflows in the exploration of complex astronomical data sets." ADASS XXI, Paris, November 2011 (invited).
- L. Dobos, A. Szalay, J. Blakely, T. Budavari, "An array library for Microsoft SQL Server with astrophysical applications." ADASS XXI, Paris, November 2011 (poster).
- S. Doe et al., "Iris: The VAO SED application." ADASS XXI, Paris, November 2011 (focus demonstration).
- T. Donaldson, A. Rogers, "TRAP—the reusable astronomy portal." ADASS XXI, Paris, November 2011 (poster).
- R. Hanisch, "Science tools developed by the U.S. Virtual Astronomical ry." ADASS XXI, Paris, November 2011.
- R. Hanisch, "Data discovery, access, and management with the Virtual ry." Innovations in Data-Intensive Astronomy, Green Bank, WV, May 2011.
- R. Hanisch, "Data discovery, access, and management with the Virtual ry." Pucon Symposium 2011: Advanced Mathematical Tools for Frontier Astronomy and Other Massive Data-Driven Sciences, Pucon, Chile, August 2011.
- R. Hanisch: Data Discovery and Access for the Next Decade. "Building on New Worlds, New Horizons: New Science from Sub-Millimeter to Meter Wavelengths," held in Santa Fe, NM, March 7-10, 2011. Invited keynote talk.
- R. Plante, V. Yekkirala, B. Baker, "Enabling OpenID authentication for VO-integrated portals." ADASS XXI, Paris, November 2011 (presentation file not available).
- A. Szalay, "Extreme data-intensive computing in astrophysics." Innovations in Data-Intensive Astronomy, Green Bank, WV, May 2011.

2012

- B. Berriman, "A Tale of 160 Scientists, Three Applications, A Workshop, and a Cloud," ADASS XXII, University of Illinois, November 7, 2012.
- B. Berriman, "Adoption of Software by a User Community: The Montage Image Mosaic Engine Example," Workshop on Maintainable Software Practices in e-Science, October 9, 2012.
- B. Berriman, "Astronomy in the Cloud," .Astronomy 4 conference, Heidelberg, Germany, July 9-11.
- B. Berriman, "Adoption of Software by a User Community: The Montage Image Mosaic Engine Example," IEEE e-Science Meeting, Chicago, October 8-12, 2012.
- G. B. Berriman and E. Deelman. "How To Use Cloud Computing To Do Astronomy." Caltech Seminar Series, May 2012.
- D. Fan, T. Budavari, "Efficient Catalog Matching with Dropout Identification," ADASS XXII, University of Illinois, November 5, 2012.
- M. Fitzpatrick, "An Introduction to VO-IRAF," Sociedade Astronômica Brasileira XXXVII Reunião Anual, Aguas de Lindoia, October 17, 2012.
- M. Fitzpatrick, "VO Desktop Tools: IRAF, Command-Line Tasks, and Python," ADASS XXII, University of Illinois, November 5, 2012. (presentation file not available)
- M. Graham, "The Transient Sky and the Virtual Observatory," Sociedade Astronômica Brasileira XXXVII Reunião Anual, Aguas de Lindoia, October 17, 2012.
- R. Hanisch, "The Virtual Astronomical Observatory." NASA Advisory Council, Subcommittee on Information Technology Infrastructure, March 2012.
- R. Hanisch, "The Research Tools of the Virtual Astronomical Observatory," (also in Portuguese) Sociedade Astronômica Brasileira XXXVII Reunião Anual, Aguas de Lindoia, October 17, 2012.
- R. J. Hanisch, "Report on the Virtual Astronomical Observatory." AURA Board of Directors and Member Representatives Meeting, Washington, DC, April 2012.
- R. J. Hanisch, "The Virtual Astronomical Observatory." NASA Headquarters, April 2012.
- R. Hanisch, "Virtual Observatory-Enabled Research," *The Evolving Universe* conference at Catholic University, Washington, DC, July 17, 2012.
- R. J. Hanisch, "Astronomical Data: Virtual and Real, Shared and Open." Open Access Symposium, University of North Texas, Denton. May 2012.
- O. Laurino, I. Busko, M. Cresitillo-Dittmar, R. D'Abrusco, S. Doe, J. Evans, O. Pevunova, "Extending Iris, the VAO SED Analysis Tool," ADASS XXII, University of Illinois, November 5, 2012.
- J. Lazio, "Science Priorities: Radio Astronomy Perspective," IVOA Interoperability Meeting, Sao Paulo, Brazil; 2012 October 22-26
- J. Lazio, "The Virtual Observatory and the Virtual Astronomical Observatory," Sagan Summer School. Pasadena, CA: 2012 July 25
- R. Plante, "Using VO TAP Services to Publish Source Catalogs." Dark Energy Survey Team Meeting, Munich, May 2012.

- A. Szalay, "Extreme Data-Intensive Computing in Astrophysics," ADASS XXII, University of Illinois, November 5, 2012.
- A. Szalay, "Extreme data-intensive computing in astrophysics." AAS Meeting 219, 113.03. (In special session, "Cyber-Discovery for the Decade.") January, 2012

2013

- B. Berriman, R. Hanisch, J. Lazio, "The Role of the Virtual Astronomical Observatory in the Era of Big Data," AAS Meeting 221, 240.27, January 8, 2013.
- T. Donaldson, D. Hinshaw, A. Rogers, G. Wallace, "Discovering Data in the Virtual Observatory," AAS Meeting 221, 240.23, January 8, 2013.
- M. Fitzpatrick, D. Tody, "Desktop Tools for the Virtual Observatory," AAS Meeting 221, 240.24, January 8, 2013.
- M. Graham, G. Djorgovski, A. Mahabal. C. Donalek, A. Drake, "Automatic Discovery of Relationships in Astronomy," AAS Meeting 221, 240.05, January 8, 2013.
- G. Greene, J. Donley, S. Rodney, J. Lazio, A. Koekemoer, I. Busko, R. Hanisch, "VAO Tools Enhance CANDELS Research Productivity," AAS Meeting 221, 240.35, January 8, 2013.
- R. Hanisch, "Virtual Astronomy in the Era of LSST," Radio Astronomy in the LSST Era, Charlottesville, VA; 2013 May 6--8
- R. Hanisch, B. Berriman, J. Lazio, "The Research Tools of the Virtual Astronomical Observatory," AAS Meeting 221, 240.21, January 8, 2013.
- R. Hanisch, "Virtual Astronomy with the Virtual Astronomical Observatory," NANOGrav Science Seminar, WebEx, 2013 November 4.
- R. Hanisch, "(Big) Data Practices in Astronomy," Kavli Foundation, Santa Monica, CA; 2014 February 19.
- R. Kinne, M. Templeton, A. Henden, P. Zografou, P. Harbo, J. Evans, A. Rots, J. Lazio, "Distributing Variable Star Data to the Virtual Observatory," AAS Meeting 221, 240.37, January 8, 2013.
- O. Laurino, I. Busko, M. Cresitillo-Dittmar, R. D'Abrusco, S. Doe, J. Evans, O. Pevunova, P. Norris, "Constructing and Analyzing Spectral Energy Distributions with the Virtual Observatory," AAS Meeting 221, 240.38, January 8, 2013.
- J. Lazio, "Virtual Astronomical Observatory," Greater IPAC Technical Symposium, Pasadena, CA; 2013 July 25
- A. Muench, S. Emery Bunn, "The Virtual Astronomical Observatory User Forum," AAS Meeting 221, 240.22, January 8, 2013.
- R. Plante, D. Mishin, J. Lazio, A. Muench, "Data Sharing and Publishing Using the Virtual Astronomical Observatory," AAS Meeting 221, 240.36, January 8, 2013.
- D. Tody, M. Fitzpatrick, M. Graham, W. Young, "Scripting the Virtual Observatory in Python " AAS Meeting 221, 240.34, January 8, 2013.
- D. Van Stone, P. Harbo, M. Tibbetts, P. Zografou, "Data Discovery and Exploration with Seleste," AAS Meeting 221, 240.25, January 8, 2013.

2014

- R. Hanisch, B. Berriman, J. Lazio, A. Szalay, G. Fabbiano, R. Plante, T. McGlynn, J. Evans, S. Emery Bunn, M. Claro, VAO Project Team, "Virtual Astronomy: The Legacy of the Virtual Astronomical Observatory," AAS Meeting 223, 255.04, January 2014.
- D. Mishin, D. Medvedev, A. Szalay, & R. Plante, "Data publication and sharing using the SciDrive service," AAS Meeting 223, 255.26, January 2014.
- R. Plante, M. Fitzpatrick, M. Graham, D. Tody, "The Virtual Observatory for the Python Programmer," AAS Meeting 223, 253.04, January 2014.
- D. Tody, R. Plante, B. Berriman, M. Cresitello-Dittmar, J. Good, M. Graham, G. Greene, R. Hanisch, T. Jenness, J. Lazio, P. Norris, O. Pevunova, & A. H. Rots, "Accessing Multi-Dimensional Images and Data Cubes in the Virtual Observatory," AAS Meeting 223, 255.05, January 2014.

Appendix B. Technical and Research Publications Authored by VAO Team Members

2010

Stobie, E.; Seaman, R.; Mighell, K.; Bunn, S. E.; Williams, R., "User Support in the Virtual Astronomical Observatory." Astronomical Data Analysis Software and Systems XIX. Proceedings of a conference held October 4-8, 2009 in Sapporo, Japan. Edited by Yoshihiko Mizumoto, Koh-Ichiro Morita, and Masatoshi Ohishi. ASP Conference Series, Vol. 434. San Francisco: Astronomical Society of the Pacific, 2010, p.414. 2010ASPC..434..414S

Seaman, R.; Williams, R., "Transient Response Astronomy: How and Why." Astronomical Data Analysis Software and Systems XIX. Proceedings of a conference held October 4-8, 2009 in Sapporo, Japan. Edited by Yoshihiko Mizumoto, Koh-Ichiro Morita, and Masatoshi Ohishi. ASP Conference Series, Vol. 434. San Francisco: Astronomical Society of the Pacific, 2010., p.112. 2010ASPC..434..112S

Rots, A. H. "When Time Is of the Essence." Astronomical Data Analysis Software and Systems XIX. Proceedings of a conference held October 4-8, 2009 in Sapporo, Japan. Edited by Yoshihiko Mizumoto, Koh-Ichiro Morita, and Masatoshi Ohishi. ASP Conference Series, Vol. 434. San Francisco: Astronomical Society of the Pacific, 2010., p.107. 2010ASPC..434..107R

Hanisch, R. J. "The Virtual Observatory: Retrospective and Prospectus." Astronomical Data Analysis Software and Systems XIX. Proceedings of a conference held October 4-8, 2009 in Sapporo, Japan. Edited by Yoshihiko Mizumoto, Koh-Ichiro Morita, and Masatoshi Ohishi. ASP Conference Series, Vol. 434. San Francisco: Astronomical Society of the Pacific, 2010., p.65. 2010ASPC..434...65H

Budavári, Tamás; Szalay, Alexander S.; Fekete, György, "Searchable Sky Coverage of Astronomical Observations: Footprints and Exposures." Publications of the Astronomical Society of the Pacific, Volume 122, issue 897, pp.1375-1388 (PASP Homepage). 2010PASP..122.1375B

Plante, Raymond L.; Greene, Gretchen; Hanisch, Robert J.; McGlynn, Thomas A.; Miller, Christopher J.; Tody, Doug; White, Richard, "Building archives in the virtual observatory era." Software and Cyberinfrastructure for Astronomy. Edited by Radziwill, Nicole M.; Bridger, Alan. Proceedings of the SPIE, Volume 7740, pp. 77400K-77400K-12 (2010). (SPIE Homepage). 2010SPIE.7740E..17P

Williams, Roy; Drake, A. J.; Djorgovski, S. G.; Donalek, C.; Graham, M. J.; Mahabal, A., "SkyAlert: a Platform for Event Understanding and Dissemination." American Astronomical So-

ciety, AAS Meeting #215, #477.01; Bulletin of the American Astronomical Society, Vol. 42, p.563. 2010AAS...21547701W

2011

R. Seaman, et. al., "Using the VO to study the time domain." New Horizons in TIme Domain Astronomy Proceedings IAU Symposium No. 285, 2011 (http://arxiv.org/pdf/1201.0577.pdf)
G. Fabbiano et al., "Recommendations of the Virtual Astronomical Observatory (VAO) Science Council for the VAO second year activity."
(http://adsabs.harvard.edu/abs/2011arXiv1108.4348F)

2012

G. B. Berriman et al., "The organization and management of the Virtual Astronomical Observatory." Submitted to SPIE Conference 8449: Modeling, Systems Engineering, and Project Management for Astronomy V, June 2012 (http://arxiv.org/abs/1206.4079v1).

G. B. Berriman R. J. Hanisch, T. J. W. Lazio, "The role of the Virtual Astronomical Observatory in the era of massive data sets." SPIE Conference 8448: Observatory Operations: Strategies, Processes, and Systems IV, June 2012 (http://arxiv.org/abs/1206. 4076v1)

J. Evans et al., "Managing distributed software development in the Virtual Astronomical Observatory." SPIE Conference 8449: Modeling, Systems Engineering, and Project Management for Astronomy V, June 2012. (http://arxiv.org/abs/1206.6161)

M. J. Graham et al., "Connecting the time domain community with the Virtual Astronomical Observatory." SPIE Conference 8448: Observatory Operations: Strategies, Processes, and Systems IV, June 2012. (http://arxiv.org/abs/1206.4035)

T. A. McGlynn et al., "Running a distributed virtual observatory: U.S. Virtual Astronomical Observatory operations." SPIE Conference 8448: Observatory Operations: Strategies, Processes, and Systems IV, June 2012. (http://arxiv.org/abs/1206.4493v1)

R. J. Williams et al. 2012. "Responding to the Event Deluge." Submitted to SPIE Conference 8448: Observatory Operations: Strategies, Processes, and Systems IV (http://arxiv.org/abs/1206.0236v1)

B. Lawton et. al., "Education and outreach with the Virtual Astronomical Observatory." AAS Meeting 219, #347.02.

2013

Laurino, O. 2013, "Iris: Constructing and Analyzing Spectral Energy Distributions with the Virtual Observatory," 23rd Astron. Data Analysis Software & Systems (ADASS) Conference, Waikoloa, HI

Mishin, D. 2013, "Data sharing and publication using the SciDrive service," 23rd As- tron. Data Analysis Software & Systems (ADASS) Conference, Waikoloa, HI

Plante, R. 2013, "Accessing the VO with Python," 23rd Astron. Data Analysis Software & Systems (ADASS) Conference, Waikoloa, HI

Appendix C. IVOA Standards Where VAO Team Members Were Authors or Editors

- VOTable Format Definition, Version 1.3, IVOA Recommendation 20 September 2013 (F. Ochsenbein, R. Williams, C. Davenhall, M. Demleitner, D. Durand, P. Fernique, D. Giaretta, R. Hanisch, T. McGlynn, A. Szalay, M. Taylor, A. Wicenec)
- Data Access Layer Interface, Version 1.0, IVOA Recommendation 29 November 2013 (P. Dowler, M. Demleitner, M. Taylor, D. Tody)
- IVOA Registry Relational Schema, Version 1.0, IVOA Proposed Recommendation 27 February 2014 (M. Demleitner, P. Harrison, M. Molinaro, G. Greene, T. Dower, M. Perdikeas)
- MOC - HEALPix Multi-Order Coverage map, Version 1.0, IVOA Proposed Recommendation 10 March 2014 (T. Boch, T. Donaldson, D. Durand, P. Fernique, W. O'Mullane, M. Reinecke, M. Taylor)
- Simple Application Messaging Protocol, Version 1.3 IVOA Recommendation 11 April 2012. (M. Taylor, T. Boch, M. Fitzpatrick, A. Allan, J. Fay, L. Paioro, J. Taylor, D. Tody)
- Simple Line Access Protocol, Version 1.0, IVOA Recommendation 9 December 2010. (Jesus Salgado, Pedro Osuna, Matteo Guainazzi, Isa Barbarisi, Marie-Lise Dubernet, Doug Tody)
- Simple Spectral Access Protocol, Version 1.1, IVOA Recommendation 10 February 2012. (Doug Tody, Markus Dolensky, Jonathan McDowell, Francois Bonnarel, Tamas Budavari, Ivo Busko, Alberto Micol, Pedro Osuna, Jesus Salgado, Petr Skoda, Randy Thompson, Frank Valdes, and the Data Access Layer working group)
- Table Access Protocol, Version 1.0, IVOA Recommendation 27 March 2010. (Patrick Dowler, Guy Rixon, Doug Tody)
- TAPRegExt: a VOResource Schema Extension for Describing TAP Services, Version 1.0, IVOA Recommendation 27 August 2012. (Markus Demleitner, Patrick Dowler, Ray Plante, Guy Rixon, Mark Taylor)
- IVOA Spectral Data Model, Version 2.0, IVOA Proposed Recommendation 09 March 2014. (Jonathan McDowell, Doug Tody, Tamas Budavari, Markus Dolensky, Inga Kamp, Kelly McCusker, Pavlos Protopapas, Arnold Rots, Randy Thompson, Frank Valdes, Petr Skoda, Bruno Rino, Sebastien Derriere, Jesus Salgado, Omar Laurino and the IVOA Data Access Layer and Data Model Working Groups.)
- Observation Data Model Core Components and its Implementation in the Table Access Protocol, Version 1.0, IVOA Recommendation 28 October 2011. (Mireille Louys, Francois Bonnarel, David Schade, Patrick Dowler, Alberto Micol, Daniel Durand, Doug Tody, Laurent Michel, Jesus Salgado, Igor Chilingarian, Bruno Rino, Juan de Dios Santander, Petr Skoda)

- VOSpace specification, Version 2.0, IVOA Recommendation 29 March 2013. (Matthew Graham, Dave Morris, Guy Rixon, Pat Dowler, Andre Schaaff, Doug Tody)
- IVOA Credential Delegation Protocol, Version 1.0, IVOA Recommendation 18 February 2010. (Matthew Graham, Raymond Plante, Guy Rixon, Giuliano Taffoni)
- Web Services Basic Profile, Version 1.0, IVOA Recommendation 16 December 2010. (Andre Schaaff, Matthew Graham)
- StandardsRegExt: a VOResource Schema Extension for Describing IVOA Standards, Version 1.0, IVOA Recommendation 08 May 2012. (Paul Harrison,Douglas Burke,Ray Plante,Guy Rixon,Dave Morris, and the IVOA Registry Working Group)
- Describing Simple Data Access Services, Version 1.0, IVOA Recommendation 25 November 2013. (Raymond Plante, Jesus Delago, Paul Harrison, Doug Tody, and the IVOA Registry Working Group)
- VODataService: a VOResource Schema Extension for Describing Collections and Services, Version 1.1, IVOA Recommendation 02 December 2010. (Raymond Plante, Aurélien Stébé, Kevin Benson, Patrick Dowler, Matthew Graham, Gretchen Greene, Paul Harrison, Gerard Lemson, Tony Linde, Guy Rixon)
- IVOA Registry Relational Schema, Version 1.0 IVOA Proposed Recommendation 27 February 2014. (Markus Demleitner Paul Harrison Marco Molinaro Gretchen Greene Theresa Dower Menelaos Perdikeas)
- IVOA Document Standards, Version 1.2, IVOA Recommendation 13 April 2010. (R.J. Hanisch, C. Arviset, F. Genova, B. Rino)
- Sky Event Reporting Metadata, Version 2.0, IVOA Recommendation 11 July 2011. (Rob Seaman, Roy Williams, Alasdair Allan, Scott Barthelmy, Joshua Bloom, John Brewer, Robert Denny, Mike Fitzpatrick, Matthew Graham, Norman Gray, Frederic Hessman, Szabolcs Marka, Arnold Rots, Tom Vestrand, Przemyslaw Wozniak)
- IVOA Support Interfaces Version 1.0, IVOA Recommendation 31 May 2011 (Grid and Web Services Working Group, Matthew Graham, Guy Rixon)
- Web Services Basic Profile Version 1.0, IVOA Recommendation 16 December 2010. (Andre Schaaff, Matthew Graham)

Appendix D. Participants

Name	Role	Staff Months	Organization
Accomazzi, Alberto	Lead, Semantic Linking	...	SAO
An, Rick	Software Developer	2	NRAO
Arora, Jitin	Software Developer	3	NRAO
Baker, William	Software Developer	10	NCSA
Barg, Irene	User Support, Operations	3	NOAO
Benson, John	Software Developer	4	NRAO
Berriman, Graham Bruce	Program Manager, IPAC Manager, Chair Program Council	27	IPAC
Budavari, Tamas	Co-Lead, Scalable Cross-Match	3	JHU
Budynkiewicz, Jamie	Software Engineer	17	SAO
Burke, Douglas	Scientist, Ontology Development	2	SAO
Busko, Ivo	Software Developer	3	STScI
Claro, Maricel	Business Manager	15	AUI
Cresitello-Dittmar, Mark	Software Developer	14	SAO
D'Abrusco, Raffaele	Post-doctoral research fellow	22	SAO
Dave, Rahul	Computational scientist	16	SAO
DeYoung, David	Project Scientist	3	NOAO
Djorgovski, George	Scientist	1	Caltech
Dobis, Laszlo	Software Developer, Graduate Student	2	
Doe, Stephen	Software Developer	4	SAO
Donaldson, Tom	Software Engineer	19	STScI
Donalek, Ciro	Software Developer	9	Caltech
Dower, Theresa	Software Developer	25	STScI
Drake, Andrew	Scientist	5	Caltech
Ebert, Ricardo	Software Engineer	1	IPAC
Economou, Frossoe	Software Engineer	1	NOAO
Emery Bunn, Sarah	User Support – Documentation Manager, Operations	22	Caltech
Evans, Janet	Software Development Manager	8	SAO
Fabbiano, Guissepina	Scientist, Chair of Science Council	...	SAO
Fekete, George	Software Engineer	4	JHU
Fitzpatrick, Michael	Software Engineer	20	NOAO
Fox, Ryan	Software Developer	2	NRAO

Good, John	Software Engineer; VAO Cross-Match Service Co-Technical Lead, Time - Series Service Technical Lead	27	IPAC
Graham, Matthew	Scientist, Software Developer; Program Council representative	29	Caltech
Greene, Gretchen	System Engineer; Program Council representative	8	STScI
Hanisch, Robert	Director	30	STScI
Harbo, Peter	Software Developer	4	SAO
Harbut, Marcy	Technical Writer and Documentation Support	5	IPAC
Hinshaw, Dean	Software Developer	20	HEASARC
Huffman, Marie	Business Manager	15	AUI
Johnson, Avalon	System Administrator	1	Caltech
Kent, Brian	Scientist, Software Developer	1	NRAO
Lau, Christina	Software Engineer, Time Series Service	3	IPAC
Lazio, Joseph	Project Scientist	10	IPAC
Laurino, Omar	Software Developer	30	SAO
Lawton, Brandon	Lead, Education and Public Outreach	30	STScI
Lively, Rick	Software Developer	4	NRAO
Low, Stephen	Systems Engineer	3	Caltech
Lyons, Daniel	Software Developer	2	NRAO
Madore, Barry	Scientific Collaborator, Science Council Team Consultant	1	IPAC
Mahabal, Ashish	Scientist, Software Developer	23	Caltech
Masters, Joe	Software Developer	2	NRAO
McDowell, Jonathan	Scientist, SED Science Support	...	SAO
McGlynn, Tom	Lead, Operations; Program Council representative	25	HEASARC
Mighell, Kenneth	Software Engineer	1	NOAO
Miller, Joseph	Software Developer	5	SAO
Mishin, Dmitry	Software Developer	24	JHU
Muench, Gus	User Support – Portal QA&T Lead, Forum Lead, Documentation	4	SAO
Nandrekar-Heinis, Deoyani	Software Developer	24	JHU
Norris, Patrick	Test Engineer	15	NOAO
Pevunova, Olga	Software Engineer	15	IPAC

Plank, Jennifer	Software Developer	1	NRAO
Plante, Raymond	Lead, Product Development; Program Council representative	40	NCSA
Preciado, Michael	Operations	45	HEASARC
Protopapas, Pavlos	Scientist, Time Series Center	...	SAO
Raddick. Jordan	Outreach Specialist	1	JHU
Refsdal, Brian	Software Developer	4	SAO
Rots, Arnold	Lead, Data Curation & Preservation; Program Council representative	2	SAO
Schiebel, Darrel	Software Developer	2	NRAO
Schmitz, Marion	Software Engineer	2	IPAC
Seaman, Rob	Standards & Protocols - VOEvent 2.0	2	NOAO
Shaw, Richard	Testing, Documentation	2	NOAO
Stobie, Elizabeth	Lead, User Support; Program Council representative	8	NOAO
Spolaor, Max	Software Developer	2	NRAO
Szalay, Alex	Technical Advisor, Member of VAO Executive	1	JHU
Thakar, Ani	Deputy, Operations; Program Council representative	6	JHU
Thomas, Brian	Test Engineer	3	NOAO
Thompson, Randy	Software Engineer	3	STScI
Tody, Doug	Deputy Lead, Standards and Protocols; Program Council representative	40	NRAO
Van Stone, David	Software Developer	1	SAO
Yenkkirala, Venkat	Software Developer	13	NCSA
Young, Wes	Software Engineer	15	NRAO
Zografou, Yulie	Software Engineer	1	SAO

Appendix E. Acronyms and Abbreviations

Term	Full Name (proper nouns are capitalized)	Definition
2MASS	Two Micron All Sky Survey	
AAS	American Astronomical Society	
AAVSO	American Association of Variable Star Observers	A potential VAO collaborating institution that may expose their data holdings to the virtual observatory.
ADEC	Astronomical Data Centers Executive Council	A VAO collaborator that coordinates the work of NASA-funded data centers and archives.
AID	AAVSO International Database	A database of all-sky survey data that may be available through the virtual observatory.
Ajax		A group of interrelated web development methods used on the client-side to create interactive web applications. With Ajax, web applications can retrieve data from the server asynchronously in the background without interfering with the display and behavior of the existing page.
ALMA	Atacama Large Millimeter Array	An array of 66 high-precision antennae located in Chile used for sub-millimeter observing.
APASS	AAVSO Photometric All-Sky Survey	An AAVSO database that may be linked to the virtual observatory.
arXiv	Pronounced as "archive," with the X as a hard K sound as in "LaTex"	An automated electronic archive and distribution server for research articles. Maintained and operated by the Cornell University Library.

Term	Full Name (proper nouns are capitalized)	Definition
ASCL	Astrophysics Source Code Library	
ASDC	ASI Data Center	The ASI Science Data Center (ASDC), a facility of the Italian Space Agency (ASI) is a multi-mission science operations, data processing and data archiving center that provides support to several scientific space missions. It is located at the ESA site of ESRIN in Frascati, near Rome (Italy).
Astropy Project		A collection of software packages written in the Python programming language and designed for use in astronomy.
Astroview		Software that provides a visual context for search results, displaying background images of the sky with selectable graphics overlaid to indicate the location and coverage of image and catalog search results. This visual feedback greatly simplifies the task of identifying which results are most relevant and useful.
AUI	Associated Universities, Inc.	
AURA	Associated Universities for Research in Astronomy, Inc.	
BRAVO	BRAzilian Virtual Observatory	
CACR	Center for Advanced Computing Research	
CANDELS	Cosmic Assembly/Near-infrared Deep Extragalactic Legacy Survey	

Term	Full Name (proper nouns are capitalized)	Definition
CASA	Common Astronomy Software Applications	A software packaged developed by NRAO with the goal of supporting the data post-processing needs of the next generation of radio astronomical telescopes.
CCB	Configuration Change Board	
CDS	Centre de donneés astronomiques de Strasbourg	
CIAO	From "s'sciavo," meaning "I am your servant" in Italian (Venetian dialect)	Software developed to analyze data returned by the Chandra X-ray Observatory.
CM	configuration management, or Configuration Manager	
DDT	Data Discovery Tool	
DENIS3	Deep Near Infrared Survey of the Southern Sky.	Third catalog release, with 355 million sources.
DES	Dark Energy Survey	A survey designed to probe the origin of the accelerating universe and help uncover the nature of dark energy by measuring the 14-billion-year history of cosmic expansion with high precision. Starting in September 2012 and continuing for five years, DES will survey a large swath of the southern sky out to vast distances in order to provide new clues to this most fundamental of questions.
E/PO		Education/Public Outreach
GALEX	Galaxy Evolution Explorer	An orbiting, ultraviolet space telescope launched in April 2003 to determine the distances of several hundred thousand galaxies, as well as the rate of star formation in each galaxy.

Term	Full Name (proper nouns are capitalized)	Definition
GSFC	Goddard Space Flight Center	
HEASARC	High Energy Astrophysics Science Archive Research Center	
HST	Hubble Space Telescope	
IFU	Integral Field Unit	
IPAC	Infrared Processing and Analysis Center	A multi-mission center of expertise for long-wavelength astrophysics that carries out data-intensive processing tasks for NASA's infrared and sub-millimeter astronomy programs.
IRAF (or VO-IRAF)	Image Reduction and Analysis Facility	A general purpose software system for the reduction and analysis of astronomical data.
Iris		A VO tool that enables astronomers to build and analyze wide-band spectral energy distributions (SEDs).
IRSA	Infrared Science Archive	An online archive of data and science products for NASA's infrared and sub-millimeter missions.
IVOA Interoperability Meeting ("Interoperability meeting")	International Virtual Observatory Alliance Interoperability meeting	
JIRA	Truncation of "Gojira," the Japanese name for Godzilla	A proprietary issue-tracking product the VAO will use as its help desk issue/bug tracking system.
JPL	Jet Propulsion Laboratory	

Term	Full Name (proper nouns are capitalized)	Definition
LLC	limited liability company	A flexible form of business enterprise that blends elements of partnership and corporate structures. It is a legal form of a company that provides limited liability to its owners within the vast majority of United States jurisdictions.
LSST	Large Synoptic Survey Telescope	A facility that will produce a wide-field astronomical survey of the Universe using an 8.4-meter ground-based telescope.
MAST	Multimission Archive at Space Telescope	
NASA	National Aeronautics Space Administration	
NCSA	National Center for Supercomputing Applications	
NED	NASA/IPAC Extragalactic Database	An online, knowledge-based database and archive for extra-galactic astronomy.
NOAO	National Optical Astronomical Observatory	
NRAO	National Radio Astronomy Observatory	
NSF	National Science Foundation	
NVO	National Virtual Observatory	

Term	Full Name (proper nouns are capitalized)	Definition
ODI	One Degree Imager	As the flagship of the Wisconsin Indiana Yale NOAO Consortium's new instrument initiatives, ODI is a unique and very competitive instrument that is sensitive to visible light and features a 1,000 megapixel camera that will cover a one-square-degree field of view, allowing ODI to capture vast areas of sky (e.g. greater than four times the area of the full moon) in a single image.
PEP	Project Execution Plan	
PI	Principal Investigator	
Python		An open source, interpreted, object-oriented programming language.
SAMP	Simple Application Messaging Protocol	A messaging protocol that enables astronomy software tools to interoperate and communicate.
SAO	Smithsonian Astrophysical Observatory	
SCC	Scalable Cross Comparison	
SciDrive		A Dropbox-like cloud storage application intended for use in scientific research.
SDSC	San Diego Supercomputer Center	A VAO collaborating institution that will help improve VO-based access to theoretical simulations and share expertise in management of large data collections, including digital libraries.
SDSS	Sloan Digital Sky Survey	
SEDs	Spectral Energy Distributions	

Term	Full Name (proper nouns are capitalized)	Definition
Seleste		A powerful and versatile user interface for querying and performing modest-scale joins between catalogs that have interfaces compliant with the VO Table Access Protocol (TAP).
SIAP	Simple Image Access Protocol	The firsy generation of an IVOA standard that defines a protocol for retrieving image data from a variety of astronomical image repositories through a uniform interface
SIAP V2	Simple Image Access Protocol 2	The second generation of an IVOA standard that defines a protocol for retrieving image data from a variety of astronomical image repositories through a uniform interface.
SSO	single sign-on	Technology that allows users to gain access to multiple systems using a single, central authentication method or interface.
SIMBAD	Set of Identifications, Measurements, and Bibliography for Astronomical Data	An astronomical database that provides basic data, cross-identifications and a bibliography for astronomical objects outside the solar system.
SMC	Small Magellanic Cloud	
Spitzer	Spitzer Space Telescope	An infrared space observatory launched in 2003. It is the fourth and final of the NASA Great Observatories program.
STScI	Space Telescope Science Institute	
TAP	Table Access Protocol	A VO standard that defines a service protocol for accessing general table data, including astronomical catalog and general database tables.

Term	Full Name (proper nouns are capitalized)	Definition
TAPRegExt		A registry extension for describing TAP database services.
VAO	Virtual Astronomical Observatory	
VAO, LLC	Virtual Astronomical Observatory, Limited Liability Corporation	
VO	Virtual Observatory	
VOClient	Virtual Observatory Client	A software package that provides a high-level, programmable interface between desktop applications and the distributed VO framework, providing access to remote VO data and services, reference implementations for VO data-providers and end-user applications.
VOSpace	Virtual Observatory Space	The IVOA interface to distributed storage.
WISE	Wide-field Infrared Survey Explorer	