

**Report on IVOA Action item FM-19A
Assessment of the Implementation Status of VO protocols**

**David Schade & Christophe Arviset
May 09, 2007**

1 Introduction

At the Victoria InterOp meeting in May 2006 the IVOA Executive committee approved an action item (FM19-8A) to evaluate implementation of VO protocols to see how widespread and successful these implementations have been. David Schade and Christophe Arviset were asked to lead this effort and it was suggested that the national projects should be involved. The goal of the report was to produce an assessment of how much progress has been made toward delivering real VO capabilities to users within the framework of IVOA protocols.

A preliminary evaluation was done by CADDC and ESAC by compiling a table with existing VO services. The progress of the project was discussed in Tucson at ADASS on 16/10/2006 where the national projects were engaged in discussions of the procedures and results. At the IVOA telecon December 5, 2006 the revised results were shared with the IVOA Executive and national projects and the Executive was asked to provide inputs/feedback on this table.

This is the final report. No formal input was received from the national project offices or others.

2 Assessment Procedure

We have concentrated on protocols which already have sufficient reference implementations: SIAP, SSAP, Registry, SLAP, and Conesearch. We've extracted from the registry the list of these implementations and tested their compliance and compiled the results in a spreadsheet.

In particular, it is important to note that the table is to be considered as an input (representing sufficient list of VO services implemented) so a report can be drawn from this table. It is recognized that this table is very dynamic and could change every day/week, but the idea is to have a

snapshot of the existing services so conclusions can be drawn from it. The table is not the final goal of the action, this is just a support for the conclusion report. Once the action is completed, it is not planned to maintain/update this table.

For some of the services (SIAP, ConeSearch), results on the compliancy tests have been compiled by CADC on web pages so services providers can see them directly and in detail.

This table / web pages has been distributed to the National VO projects for feedback and inputs, but none was received.

3 SIAP and ConeSearch

Procedures for Simple Image Access and Conesearch

These services were evaluated by CADC staff. We identified VO services using existing Registries. For the SIA services we typed "SIA" into the search box at <http://nvo.stsci.edu/VORegistry/index.aspx> and saved the results. To assess compliance we used the definition at <http://www.ivoa.net/Documents/WD/SIA/sia-20040524.html>. Similarly, we typed "Conesearch" into the registry page above and we used the definition at <http://us-vo.org/pubs/files/conesearch.html> to assess compliance. To probe the services Perl scripts were written to generate queries for the cases where data should be returned, where metadata should be returned and where an error response was tested (giving bad RA input).

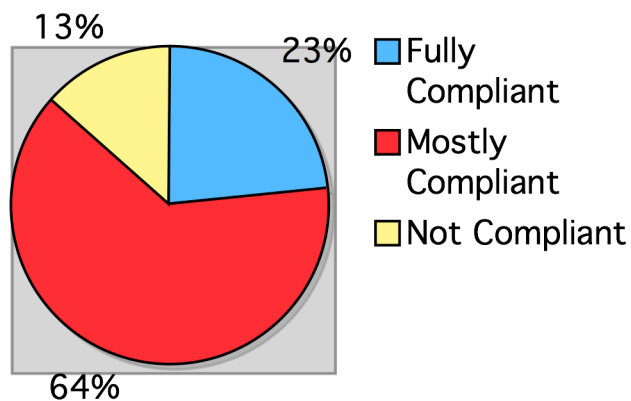
Details of the procedures used at CADC for SIA and Conesearch appear at <http://www.hia-ihp.nrc-cnrc.gc.ca/cadcdev/CVO/conesia/index.html>.

The data that was returned by the services was evaluated and differences between the VO-standard specifications and the actual service performance were classified as either a "Warning" or an "Error". Our intent was that "Warnings" indicated deviations that would not prevent a reasonably flexible program from successfully retrieving results from the service. "Errors" are intended to indicate cases where most programs would fail to retrieve data from the service. We exercised our best judgment to make the decisions where each case of deviance fell.

Results for Simple Image Access and Conesearch

We learned a number of things from this work. We have discovered that some of the CADC and ESAC services were not, in fact, fully compliant in spite of our best intentions. Clearly there is latitude in understanding exactly how to implement the standards so this type of exercise is useful to all of us. We have given the National Projects the opportunity to check our assessments of their services and to provide feedback on our work.

Simple Image Access services



Simple Image Access is important because it represents the primary service that is used to locate to access data. Above we show a summary of the detailed results available at <http://www.hia-ihc.nrc.cnrc.gc.ca/cadcdev/CVO/conesia/index.html>. The majority of the 128 SIA services that we tested are either mostly or fully compliant. Only 13% are judged to be not compliant. Only 23% are judged to be fully compliant.

In the case of Conesearch we evaluated 441 services and found that 356 of them (81%) were fully compliant with the standard and 19% were not compliant. A more detailed discussion of these results appears on the web site above.

4 SSAP

For SSAP, we decided to test the compliancy of the services with the main VO Spectral tools available (mainly VOSpec, SPLAT, SpecView), which at the time were implementing slightly modified SSAP v0.92. Most of the SSAP services could be accessed with this tools, but it is difficult to call them all “fully compliant”.

Indeed, since that date, SSAP has been going through significant changes and still has not reached the IVOA recommended status.

SSAP services providers have been quite eager to be able to publish their spectra through SSAP for them to be accessible through the VO Spectral tool, but they have not upgraded their services each time the SSAP specifications were evolving.

That raises an important question about what will happen to all these existing SSAP services published with an old version of SSAP when the final SSAP recommended version will be accepted. All these SSAP services will become “non compliant” and will have to be upgraded to the latest SSAP specifications.

To adapt the VO spectral tools to the latest SSAP specifications should be reasonably easy and developers will be quite interested in having their applications compatible with the latest SSAP specifications. Nonetheless, these tools may loose their scientific interest if no more SSAP resources are available.

IVOA could argue that data providers should not really publish their resources through “not yet recommended” IVOA protocols or assume the risk that they will have to upgrade all their services. Nonetheless, the IVOA also benefits from these early implementations of draft versions of the protocols to demonstrate the usefulness of these protocols and to get feedback about them. Furthermore, if the development of protocols is too slow then this guarantees there will be many services and applications implemented using incomplete and still-evolving protocols. Therefore, more rapid protocol development is a part of the solution to the problem.

Still, the same problem will apply later on, when the protocols evolve further from v1.0 to v2.0 and the IVOA may have a weaker argument to ask all data providers to upgrade again to the latest version of the protocols specifications.

Therefore, in the context of the relevant IVOA WG, it would be a good idea to think about backward compatibility and/or to offer tools for the data providers to easily upgrade their services to the latest versions of

the specifications. Another option would be to develop a transformation layer that can transform query/results from v1.0 to v2.0. This should be possible with some protocols and services and less feasible with others. Nevertheless, it would be beneficial if the WG's consider this problem. Failure to do so could result in unnecessary loss of availability of services and tools.

5 SLAP

There are few SLAP services (4 coming from 3 different services providers) and they all had been built as prototype implementations of the SLAP v0.5 protocols, so they are all fully compliant with that version of the standard.

6 IVOA Registries

The main compliance tests were to check if the registry OAI harvesting interface was compliant with the IVOA Registry specifications. That includes testing the six OAI verbs with a few different parameters for the verbs that accept them:

- ListRecords
- Identify
- ListIdentifiers
- GetMetadataFormats
- GetRecord
- ListSets

Those verbs were tested “manually” for the main useful functions and the results were analyzed and summarized in the spreadsheet. A more full automatic testing tool could be found (eg OAI interface checker) but that was not felt necessary as the list of IVOA registries is not that long.

Apart from a few details, all the IVOA registries were found to be mostly compliant with the harvesting interface and we can confirm that the harvesting across IVOA registries is basically working fine.

The slight non compliance problems that were encountered were probably due to the fact that some of the registries were built prior to the Registry v0.10 specifications and registry builders have not updated their registry to these specifications, because they were waiting to do so when the formal Registry v1.0 recommendations were accepted. To our knowledge, most of the listed registries are now in the process of upgrading to these v1.0 specifications.

Another test was the Registry Search interface from the Registries User Interface web pages. Keywords search was working fine for all the Searchable Registries (note that the publishing registries do not support the search interface).

The differences were found in the ADQL search interfaces as these were not really defined yet in the IVOA Registry specifications (eg ADQL versions changed, the registry use of the XPath changed, and the Registry search interface was also modified).

Again, all these differences should soon disappear when IVOA registries are upgraded to the recommended Registry v1.0 specifications (Registry Interface and Registry Resource Metadata).

7 "Mostly compliant" vs "Fully compliant"

The IVOA clearly has an interest in encouraging the proliferation of services that use the standards that it develops. Broad adoption is necessary for the VO to be successful. The purpose of this exercise is to see where we stand with services and to encourage full compliance so that the services are more powerful. At the same time we can get input from those that create implementations that will help us to design clear standards that can be implemented with a reasonable level of effort.

In order to interpret these results it is necessary to decide what the relative values are of services that are "mostly compliant" compared to "fully compliant". In other words, does "mostly compliant" represent success or failure?

It is our view that a rating of "mostly compliant" indicates that a little more work needs to be done to make a useful service into a much more useful service that will have long-lasting value to the community.

It is our view that we should not accept "mostly compliant" services as successes.

There are lessons to be learned from the history of HTML. The standards for HTML are rather loose and authors/tools produce many "mostly compliant" web pages. Web browsers try to interpret these as well as they can, making guesses as to what the author meant or sometimes just deciding how to render the page in some suitable fashion. This state of affairs has gradually evolved into web authors creating content for the browsers rather than for the web (standards), so now all browsers have to implement all the work-arounds of all the other browsers. This is very

costly and indicates the acceptance of a situation where there is some uncertainty in interpretation. Such uncertainty is not acceptable in a scientific environment.

XML was developed later and includes many lessons the community learned from HTML. In contrast to HTML, XML parsers (the code that reads and interprets the XML document) are required to return an error if the document is not well-formed. The lesson here is that "mostly compliant" was recognized in retrospect to be a disaster. In fact, in the XML world, clients are not allowed to be smart enough to handle non-compliance.

We believe that we should not be satisfied with "mostly compliant" services. If we do we will never realize the full potential of the VO standards that we are creating.

8 Conclusions and Recommendations:

Making a list of the VO services and testing their compliance against the IVOA standards has allowed us to draw some conclusions and to propose some recommendations to the IVOA Executive and National VO projects.

"Fully compliant" VO services can be used easily by VO applications which "consume" them, but "partially compliant" or "non-compliant", or "test" implementations may generate wrong results or too many useless results at the client side. There should be a way in the IVOA Registries to identify/query only "fully compliant" and "useful" VO services.

Various VO services are implemented during the lifetime of the IVOA standards (eg on version prior to the IVOA final standard recommendation). IVOA WGs should think about backward compatibility especially after the initial standard recommendation (eg from v1.0 to v2.0 of the standard).

The recommendations resulting from this work are the following:

1. Full and complete compliance should be the goal of those that implement VO services. Anything less diminishes the power of the particular service and the contribution it is able to make to the overall VO system.

2. IVOA needs to release, as soon as possible, stable versions (e.g. v1.0) of IVOA protocols so that service providers may confidently implement stable and useful VO services.
3. Protocols must be as clear and unambiguous as possible to minimize compliance failures due to misunderstanding by those that implement services using the protocols.
4. Automated curation tools for the VO services in the IVOA Registries should be developed to check the full compliance of these services with IVOA standards. This should be put in the Roadmap for the Registry WG or other WGs.
5. Services that are identified as being non-compliant with the standard should be differentiated by some means in the registry in a way that will enable registry users to avoid using them.
6. Backward compatibility and/or a supported “portability” path need to be seriously considered by IVOA Working Groups in the process of upgrading existing VO standards with the goal of minimizing the need for service providers to rebuild their services and minimizing the need for VO applications developers to re-write their tools.
7. We recommend that the IVOA develop a strategy for ensuring that an ongoing effort is made to assess the progress toward wide implementation of fully-compliant IVOA standards.