



*International
Virtual
Observatory
Alliance*

Simple Spectral Access Protocol

Version 1.00

IVOA Working Draft May 8, 2007

This version (release candidate 2 Rev 1):

<http://www.ivoa.net/Documents/WD/SSA/WD-SSA-yyyymmdd.html>

Latest version:

<http://www.ivoa.net/Documents/WD/Identifiers/WD-SSA.html>

Previous version(s):

Version 0.97, November 2006
Version 0.96, September 2006
Version 0.95 May 2006
Version 0.91 October 2005
Version 0.90 May 2005

Editors:

D.Tody, M. Dolensky

Contributors:

D.Tody, M. Dolensky, J. McDowell, F. Bonnarel, T.Budavari, I.Busko, A. Micol, P.Osuna, F.Valdes,
and the data access layer working group.

Abstract

The Simple Spectral Access (SSA) Protocol (SSAP) defines a uniform interface to remotely discover and access one dimensional spectra. SSA is a member of an integrated family of data access interfaces altogether comprising the Data Access Layer (DAL) of the IVOA. SSA is based on a more general data model capable of describing most tabular spectrophotometric data, including time series and spectral energy

distributions (SEDs) as well as 1-D spectra; however the scope of the SSA interface as specified in this document is limited to simple 1-D spectra, including simple aggregations of 1-D spectra.

The form of the SSA interface is similar to that of the older Simple Image Access (SIA) interface for accessing 2-D image data, and the cone search interface for accessing astronomical catalogs. Clients first query the global resource **registry** to find services of interest. Clients then issue a **data discovery query** to selected services to determine what relevant data is available from each service; the candidate datasets available are described uniformly in a **VOTable** format document which is returned in response to the query. Finally, the client may **retrieve selected datasets** for analysis.

Spectrum datasets returned by an SSA spectrum service may be either precomputed, archival datasets, or they may be **virtual data** which is computed on the fly to respond to a client request. Spectrum datasets may conform to a standard **data model** defined by SSA, or may be **native** project spectra, and may be returned in any of a number of standard **data formats**. Spectral data is generally stored externally in a format specific to each spectral data collection; currently there is no standard way to represent astronomical spectra, and virtually every project does it differently. Hence spectra may be actively **mediated** to the standard SSA-defined data model at access time by the service, so that client analysis programs do not have to be familiar with the idiosyncratic details of each data collection to be accessed. Services are self describing, and provide a **service metadata query** operation which may be called to determine the capabilities of a specific service instance. Metadata returned by a service metadata query may be cached in the registry to facilitate registry-based service discovery.

Since SSA is part of a family of interfaces, much of the SSA interface described herein is common with the other DAL interfaces and not specific to SSA. In particular, the HTTP-based basic service profile, the main query parameters, and most of the dataset metadata returned in the query response, are generic and apply equally well to any type of data, and are (or will be, as interfaces are updated) shared by all the DAL interfaces.

Status of This Document

This is a Working Draft. The first release of this document was in April 2005.

This is an IVOA Working Draft for review by IVOA members and other interested parties. It is a draft document and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use IVOA Working Drafts as reference materials or to cite them as other than "work in progress".

Comments on this document can be posted to the mailing list dal@ivoa.net, uploaded to the collaborative web page ivoaDAL, or sent to the authors directly.

It is expected that the Simple Image Access specification, which predates SSA, and the new Table Access Protocol (TAP) will be homogenized with SSA when next revised. Additional changes are expected when support for the Astronomy Data Query Language

(ADQL) is integrated into the DAL interfaces, and when functional enhancements are made to support features such as asynchronous data staging and authentication, and interface enhancements such as support for SOAP.

A list of [current IVOA Recommendations and other technical documents](http://www.ivoa.net/Documents/) can be found at <http://www.ivoa.net/Documents/>.

Acknowledgements

This document has been developed with support from the 5th Framework Programme of the European Community for research, technological development and demonstration activities, contract HPRI-CT-2001-50030, and via a grant from the [National Science Foundation's](#) Information Technology Research program to develop the U.S. National Virtual Observatory.

Many of the ideas in this document originated from others involved in developing Virtual Observatory concepts and standards. In particular, the idea of using association in the query response to group similar datasets was originally proposed by Roy Williams. Arnold Rots originated the idea of ranking query results via a *score* heuristic, and helped put the coordinate systems used in SSA on firm theoretical foundation via the development of STC. The technique for using dimensional analysis to consistently scale the physical axes of a spectrum came from Pedro Osuna, Jesus Salgado, and others at ESAC. Francois Bonnarel, Mireille Louys, Alberto Micol, and others made notable contributions to the representation of astronomical metadata and in particular the Characterization data model. Laszlo Dobos contributed early implementations of the access protocol using the spectral archive at JHU. Randy Thompson and Inga Kamp were especially helpful in providing feedback on test implementations.

Many thanks to all who contributed to the DAL survey among spectral data providers and consumers (Dolensky/Tody 2004): Ivo Busko, Mike Fitzpatrick, Satoshi Honda, Stephen Kent, Tom McGlynn, Pedro Osuna Alcalaya, Benoît Pirenne, Raymond Plante, Phillippe Prugniel, Enrique Solano, Alex Szalay, Francisco Valdes and Andreas Wicenec.

Parts of this protocol were adapted from the OpenGIS (Open Geospatial Consortium, Inc.) Web-Mapping Service (WMS) specification. In particular, the basic service elements and certain details of the use of the HTTP protocol to formulate requests and responses is patterned after the OpenGIS WMS service. Parts of the text of this specification were adapted directly from the WMS service specification.

Contents

1	Introduction	6	
1.1	Architecture	6	
1.2	Basic Usage	7	
1.3	Basic Service Elements	8	
1.3.1	Request Format		8
1.3.2	Parameters		8
1.3.3	Parameter Values		9
1.3.4	Error Response		9
1.4	Requirements for Compliance	9	
1.4.1	Levels of Compliance		9
2	SSA/DAL Concepts	10	
2.1	Data Model	10	
2.2	Data Representation	11	
2.3	Virtual Data	11	
2.4	Data Derivation	12	
2.4.1	Data Source		12
2.4.2	Creation Type		12
2.5	Service Type	13	
2.6	Provenance	14	
2.7	Data Association	14	
3	SSA Operations	15	
3.1	Introduction	15	
3.2	Methods & Protocols	15	
3.3	QueryData Operation (required)	15	
3.3.1	Input Parameters		16
3.3.2	Mandatory Query Parameters		16
3.3.2.1	POS		17
3.3.2.2	SIZE		17
3.3.2.3	BAND		17
3.3.2.4	TIME		18
3.3.2.5	FORMAT		19
3.3.3	Recommended and Optional Query Parameters		19
3.3.3.1	APERTURE		20
3.3.3.2	SPECRP		21
3.3.3.3	SPATRES		21
3.3.3.4	TIMERES		21
3.3.3.5	SNR		21
3.3.3.6	REDSHIFT		21
3.3.3.7	VARAMPL		22
3.3.3.8	TARGETNAME		22
3.3.3.9	TARGETCLASS		22
3.3.3.10	FLUXCALIB		22
3.3.3.11	PUBDID		22
3.3.3.12	CREATORID		23
3.3.3.13	COLLECTION		23
3.3.3.14	TOP		23
3.3.3.15	MAXREC		23

3.3.3.16	MTIME	24
3.3.3.17	COMPRESS	24
3.3.3.18	RUNID	24
3.3.4	Service-Defined Parameters	24
3.3.5	Query Response	24
3.3.6	Query Response Metadata	25
3.3.7	Types of Metadata	26
3.3.8	Query Metadata	27
3.3.8.1	Query.Score	27
3.3.8.2	Query.Token	28
3.3.9	Association Metadata	28
3.3.9.1	MultiFormat Association	28
3.3.9.2	Association.Type	29
3.3.9.3	Association.ID	29
3.3.9.4	Association.Key	29
3.3.10	Access Metadata	29
3.3.10.1	Access Reference	30
3.3.10.2	Output Format	30
3.3.10.3	Dataset Size Estimate	30
3.3.11	Data Model Metadata	31
3.3.11.1	General Dataset Metadata	31
3.3.11.2	Dataset Identification Metadata	31
3.3.11.3	Curation Metadata	32
3.3.11.4	Astronomical Target Metadata	33
3.3.11.5	Coordinate System Metadata	33
3.3.11.6	Dataset Characterization Axis Metadata	33
3.3.11.7	Characterization Coverage Metadata	34
3.3.11.8	Characterization Accuracy and Error Metadata	34
3.3.12	Additional Service-Defined Metadata	35
3.3.13	Metadata Extension Mechanism	35
3.4	GetData (reserved)	35
3.5	GetCapabilities (recommended)	36
3.6	StageData (reserved)	36
3.7	GetAvailability (recommended)	36
4	Data Retrieval	36
4.1	Access Reference URL	36
4.2	Data Format	37
4.3	Data Compression	37
4.4	Error Response	37
5	Basic Service Elements	38
5.1	Introduction	38
5.2	Version numbering and negotiation	38
5.2.1	Version number form and value	38
5.2.2	Version number changes	38
5.2.3	Appearance in requests and in service metadata	38
5.2.4	Version number negotiation	38
5.3	General HTTP request rules	39
5.3.1	Introduction	39
5.3.2	Reserved characters in HTTP GET URLs	39
5.3.3	HTTP GET	39
5.3.4	HTTP POST	40

5.4	General HTTP response rules	40	
5.5	Numeric and Boolean values	41	
5.6	Output formats	42	
5.7	Request parameter rules	42	
5.7.1	Parameter ordering and case		42
5.7.2	Range-List Parameters		42
5.8	Common request parameters	43	
5.8.1	VERSION		43
5.8.2	REQUEST		43
5.8.3	Extended capabilities and operations		43
5.9	Service result	44	
5.10	Error Response and Other Unsuccessful Results	44	
5.10.1	Service Error		45
5.10.2	Overflow		45
5.10.3	Other Errors		45
	Appendix A: Sample Query Response	45	
	References	49	

1 Introduction

The Simple Spectral Access protocol (SSAP, SSA) defines a uniform interface to remotely discover and access simple 1-D spectra. SSA is based on a more general data model capable of describing tabular spectrophotometric data including time series and SEDs as well as 1-D spectra. Basic usage is similar to the Simple Image Access (SIA) protocol (Tody/Plante 2004) and the simple cone search (SCS) protocol for simple access to astronomical catalogs. Unlike these earlier interfaces however, spectral data access via SSA may involve active transformation of data as stored externally into a standard data format and data model defined by SSA, in order to deal with the problem of heterogeneous spectral data formats as stored externally.

1.1 Architecture

All the IVOA data access interfaces share the same basic interface, differing mainly in the type of data being accessed. A query is used for data discovery, and to negotiate with the service the details of the *static* or *virtual* (dynamically created) datasets to be retrieved. Subsequent data access requests can then be made to retrieve individual datasets of interest. SSA differs from some other data access interfaces in that a service may **mediate** not only dataset metadata, but the actual dataset itself, to allow a client to do detailed analysis on a spectrum without having to understand how it is represented externally. Direct access to native project data is also provided.

All of the newer DAL interfaces share the same basic service profile, although services may define additional operations specific to the service. A single service may support multiple **operations** or **methods** to perform various functions. The current DAL interfaces use an HTTP GET-based interface to submit parameterized requests, with responses

being returned as structured documents, e.g., FITS or VOTable. The service operations defined for SSA are the following:

1. A **queryData** operation returns a table (VOTable format) describing candidate datasets which can be retrieved, including standard metadata describing each dataset, and an access reference which can be used to retrieve the data. The queryData operation is used both for data discovery and for negotiation with the service of the details of virtual datasets to be generated.
2. A **getData** operation (currently implemented in SSA merely as a data access URL) is used to access an individual dataset. The accessed data may be generated on-the-fly by the service at access time, e.g., to reformat or subset the data.
3. A **getCapabilities** operation is defined by which an external client (such as the registry or a client application) can retrieve service metadata to determine the capabilities of a given SSA service instance.

A fourth **stageData** operation is planned but it not currently defined. This will be used to perform asynchronous or bulk staging of data. Finally, a simple **getAvailability** operation (provided for all services) will provide a simple mechanism for monitoring the basic function and availability of a service.

A spectrum conforming to the SSA data model may be returned serialized in any of a number of different data formats, including VOTable, FITS binary table, and native XML.

1.2 Basic Usage

Although SSA is a complex interface, the most common usage can be quite simple. A query can be entered in a Web browser, viewing the results as XML in the browser and downloading selected spectra by a copy-paste operation on the given access reference URL. A simple query might search for 1-D spectra by position on the sky – the classic “cone search” type of query. More complex queries are little more complicated, merely adding additional query constraint parameters, e.g., to constrain the waveband or spectral resolution, or to find spectra by redshift.

In a simple case of a positional query the SSA query URL is very similar to that for SIA or SCS. For example,

Example:

<http://www.myvo.org/ssa?REQUEST=queryData&POS=22.438,-17.2&SIZE=0.02>

The query response is a VOTable describing each candidate dataset as defined later in this document.

Dataset retrieval is then a simple matter of examining the query response, selecting the dataset or datasets to be retrieved, if any, and retrieving them by reading the document

pointed to by the *access reference* (a URL) for the dataset. Interpretation of the returned spectrum dataset is the responsibility of the client application.

For a fully compliant SSA service, the data returned by the service will be in one of the SSA defined standard data formats, conformant to the SSA data model. Due to the need to mediate external data or support features such as format conversion or data subsetting, the service may have to compute on demand the dataset to be returned, however this is transparent to the client.

1.3 Basic Service Elements

The basic form of a SSA service (or any other second generation DAL service) is specified in detail in section 5. In the current section we merely summarize the basic elements of a SSA or other DAL service.

1.3.1 Request Format

In general a service may implement multiple *operations*, e.g., `queryData`, `getCapabilities`, and so forth; altogether these define the *interface* to the service. Interfaces may change with time hence are versioned. It is possible for a given service instance to simultaneously expose multiple interfaces or versions of interfaces.

The SSA interface described in this document is based on a distributed computing platform (DCP) comprising Internet hosts that support the Hypertext Transfer Protocol (HTTP). Thus, the online representation of each operation supported by a service is an HTTP Uniform Resource Locator (URL).

A request URL is formed by concatenating a *baseURL* with zero or more operation-defined request *parameters*. The baseURL defines the network address to which request messages are to be sent for a particular operation of a particular service instances on a particular server. Each service operation may have a different baseURL.

1.3.2 Parameters

Parameters may appear in any order. If the same parameter appears multiple times in a request the operation is *undefined* (if alternate values are desired the *range-list* syntax may be used instead). Parameter names are case-insensitive. Parameter values are case-sensitive unless defined otherwise in the description of an individual parameter.

All operations define the following standard parameters:

REQUEST The request or operation name, e.g., “getCapabilities” (required).

VERSION The version number of the interface (optional).

The values of both the REQUEST and VERSION parameters are case-insensitive.

The VERSION parameter is optional but recommended. A given service instance may support multiple versions of the SSA interface, which includes both the input parameters and the query response with all of its complex metadata, and by default the service assumes the highest standard version which is implemented. Explicit specification of the

interface version assumed by the client is necessary to ensure against a runtime version mismatch, e.g., if the client caches the service endpoint but a newer version of the service is subsequently deployed. The main reason VERSION is optional is to allow the client to disable runtime version checking.

All other request parameters are defined separately for each operation.

1.3.3 Parameter Values

Integer numbers are represented as defined in the specification of integers in XML Schema Datatypes. Real numbers are represented as specified for double precision numbers in XML Schema Datatypes. Sexagesimal formatting is not permitted, either for parameter input or in output metadata, other than in ISO 8601 formatted time strings (sexagesimal format is fine for a user interface but inappropriate for a lower level machine interface, where it only complicates things).

SSA defines a special *range-list* format for specifying numerical ranges or lists of ranges. For example, “1E-6/3E-7;source” could specify a spectral bandpass defined in the rest frame of the source. The syntax supports both open and closed ranges. Ranges or range lists are permitted only when explicitly indicated in the definition of a particular parameter. For a full description of range list syntax refer to section 5.7.2.

1.3.4 Error Response

In the case of an error, service operations will normally return a VOTable (which in this case is merely a very small, essentially fixed format XML document), containing an `INFO` element with name `QUERY_STATUS` and the value set to “ERROR”. More fundamental service or protocol errors may result in an HTTP level error. A null query, that is a queryData which does not find any data, is not considered an error. More information on error responses is given in section 5.10.

1.4 Requirements for Compliance

The keywords “**must**”, “**required**”, “**should**”, and “**may**” as used in this document are to be interpreted as described in the W3C specifications (RFC 2119). Mandatory interface elements are indicated as **must**, recommended interface elements as **should**, and optional interface elements as **may** or simply “may” without the bold face font.

Sometimes the extent to which a given interface element is required depends upon the mode of operation of the service. For example, a service which performs spectral extraction **must** implement the APERTURE query parameter, but it is not used for other subtypes of SSA services.

1.4.1 Levels of Compliance

In order to be *minimally compliant* a service must implement all elements of the SSA protocol identified as “must” in this document. In summary the minimal service includes the following:

1. The SSA query method **must** implement the HTTP GET interface, returning the query response encoded as a VOTable document. At least the `POS`, `SIZE`, `TIME`, `BAND`, and `FORMAT` query parameters **must** be supported by the service. The query response **must** include all metadata fields identified as mandatory in the protocol.
2. The direct URL-based *getData* method **must** be provided capable of returning data in at least one of the SSA-compliant data formats (VOTable is preferred if only one format is supported).
3. The *getCapabilities* operation **should** be provided to return service metadata encoded as defined herein.

If a service cannot return data which is SSA-compliant, it may still be useful to implement a service which has a SSA-compliant query but which returns native or external data. Such a service is said to be **query compliant** if the query operation is at least minimally compliant.

A service is said to be **fully compliant** if, in addition to the functionality required to be minimally compliant, the service implements all the “should” elements of the interface defined herein.

A top of the line service will be fully compliant plus will implement some of the optional (“may” provide) elements of the interface. For example the service may support additional query parameters or may return additional metadata; the service may provide access to native data as well as SSA-compliant data, or may be capable of returning data in any supported standard data format requested by the client.

2 SSA/DAL Concepts

2.1 Data Model

SSA consists of both an access protocol and interface, and an underlying data model describing the data to be accessed. The term **data model** as used here refers to a logical model for the data detailing the decomposition of a complex dataset into simpler elements, including specifying the meaning of each element, the relationships between elements, the metadata used to describe the data elements and the overall dataset, and the concepts upon which the data model is based. In this document we refer to the underlying data model interchangeably as the SSA data model or the **spectrum or spectral data model**. The data model used in SSA is described in McDowell, Tody, et.al, 2007.

Explicitly defining the data model assumed by a data object is important for a variety of reasons. Doing so helps greatly to document the structure and meaning of the data. Data analysis software has to understand data at a fundamental level in order to function correctly.

Data model **mediation** - the process of transforming data from some externally-defined data model to a prescribed data model (the SSA data model in our case) - makes it possible for a client application to deal uniformly with external data without having to understand the idiosyncratic details of each external data collection. SSA does data model mediation on the fly, at data access time, in the service used to publish a data collection to the VO. A data publishing service is written for a specific data collection by the creators or curators of the data who understand the data well, and may thereafter be accessed by any number of independently written client applications; hence mediation to a standard model is best performed by the service. If more detailed knowledge of a specific data collection is required than is possible using a standard model, direct pass-through of the native project data is also possible.

2.2 Data Representation

A data model defines the logical content of data, but says nothing about how the data is represented externally. The same data object may be represented externally in many different ways, e.g., as a FITS file or VOTable, as a direct XML serialization, in a RDBMS, and so forth. So long as the data model does not change, and the data representation is expressive enough, data may be transformed from one representation to another without loss of information. If transformation between different data models is required, some loss of information may occur. This can happen, for example, during mediation of external data to a known data model by a SSA service.

In the most general case SSA uses a container-component approach to represent datasets. In this case a general container such as VOTable or FITS is used to represent a Spectrum object. A similar approach is used for the SSA query response, which is returned as a VOTable. The container is used to aggregate **component data models** which are associated in some fashion to model more complex objects such as a spectrum. The advantage of this approach is flexibility, in that there is no fixed structure for the overall dataset, and extensibility, as it is easy to add custom components to describe the details of a specific data collection while conforming to the standard core model.

Application programs typically manipulate a data object by directly accessing the elements of the data model via some language-specific API. **UTYPE** tags are used to provide a uniform means to identify the elements of a data model in any language or environment. For example, given the component data model "DataID", the UTYPE "DataID.Title" identifies the data model field containing the title string for the dataset; "DataID.Collection" identifies the parent data collection, and so forth.

2.3 Virtual Data

A **virtual dataset** is one which can be described, but which may not physically exist until it is accessed, at which time it is created on the fly by the service. A typical example is a cutout (subset) of an image or spectrum. Most data access in the VO is to virtual data. Physical datasets can also be accessed, but this is a far less powerful technique as physical datasets are often too large to transmit efficiently over the network, particularly when only a small portion of the data is needed, and capabilities such as mediation to a standard model or transformations of various kinds are not possible.

When a query is made to a SSA service which can return virtual data, the service computes the parameters of any virtual datasets it can generate to satisfy the query. What can be generated depends upon what the client has requested, the input data available to the service, and the capabilities of the service. The metadata returned in the query response will describe the virtual dataset and its relationship to any parent dataset or datasets. The access reference is in effect a token to be passed back to the service to generate the virtual dataset. The client can either access the virtual data (in which case it is realized by the service, and returned), or further refine the query to more finely specify the data to be returned by the service.

2.4 Data Derivation

Data can come from a variety of sources, and may go through various types of processing, including by the data access service itself, before being delivered to a client analysis application. It is important for analysis to understand the origin of the data and what processing it has undergone. To address this issue we introduce two new concepts, *data source* and *creation type*.

2.4.1 Data Source

The data source specifies where the data originally came from. The following values are currently defined:

survey	A survey dataset, which typically covers some region of observational parameter space in a uniform fashion, with as complete as possible coverage in the region of parameter space observed.
pointed	A pointed observation of a particular astronomical object or field. Typically these are instrumental observations taken as part of some PI observing program. The data quality and characteristics may be variable, but the observations of a particular object or field may be more extensive than for a survey.
custom	Data which has been custom processed, e.g., as part of a specific research project.
theory	Theory data, or any data generated from a theoretical model, for example a synthetic spectrum.
artificial	Artificial or simulated data. This is similar to theory data but need not be based on a physical model, and is often used for testing purposes.

2.4.2 Creation Type

The creation type describes the process used to produce the current dataset from the data source. Typically this describes *only* the processing performed by the data service, but it could describe some additional earlier processing as well, e.g., if data is partially precomputed. The creation type is especially important for virtual data and for data which is derived from the parent data set by some complex form of processing. The following values are currently defined:

archival	The entire archival or project dataset is returned.
----------	---

	Transformations such as metadata or data model mediation or format conversions may take place, but the content of the dataset is not substantially modified (e.g., all the data is returned and the sample values are not modified).
cutout	The dataset is subsetted in some region of parameter space to produce a subset dataset. Sample values are not modified, e.g., cutouts could be recombined to reconstitute the original dataset.
filtered	The data is filtered in some fashion to exclude portions of the dataset, e.g., passing only data in selected regions along a measurement axis, or processing the data in a way which recomputes the sample values, e.g., due to interpolation or flux transformation. Filtering is often combined with other forms of processing, e.g., projection.
mosaic	Data from multiple non- or partially-overlapping datasets are combined to produce a new dataset.
projection	Data is geometrically warped or dimensionally reduced by projecting through a multidimensional dataset.
spectral extraction	Extraction of a spectrum from another dataset, e.g., extraction of a spectrum from a spectral data cube through a simulated aperture.
catalog extraction	Extraction of a catalog of some form from another dataset, e.g., extraction of a source catalog from an image, or extraction of a line list catalog from a spectrum (not valid for a SSA service).

The full creation type may involve more than one of these operations, for example, both projection and filtered, or both spectral extraction and filtered.

This list is by no means complete in general astronomical data processing terms, but is intended to express only the types of operations which might take place during VO data access, where subsetting, filtering, projection, spectral extraction, etc., are all defined operations. Other values may be added in the future.

2.5 Service Type

Not all SSA services are of the same type: services are further classified by their subtype, indicating how they generate the spectra returned by the service (the subtype of a SSA service is specified in the service metadata, as outlined in section 3.5). The subtype of a SSA service is similar to the dataset creation type as described in section 2.4.2; usually the creation type and the SSA service subtype are the same, but this is not always the case. A simple service which returns only entire archival spectra is an “**archival**” SSA service. A service which can return subregions of larger spectra is a “**cutout**” service. A SSA service which can combine multiple input spectra is a “**mosaic**” service (a mosaic service can also do cutouts if presented with a sufficiently small spectral bandpass). A SSA service which dynamically generates spectra from more fundamental data, e.g., a spectral data cube or event list, is a “**spectral extraction**” service.

2.6 Provenance

The combination of a data source with a creation type provides us with a primitive capability for describing the *provenance* of a dataset, i.e., where it came from, and how it was produced. This is important because SSA and other DAL services can generate virtual data products where complex processing may be performed at access time.

To be able to describe the provenance of a virtual data product we need one additional concept, the *dataset identifier* of the parent dataset, as assigned by the entity which created the dataset (typically a survey project, observatory, modeling program, etc.).

Given a virtual data product we can then say how the data product was derived from the parent dataset or datasets (the creation type), identify the parent dataset (the creator-assigned dataset ID), and the origin and type of data from which the virtual data product was derived (data source, collection, and so forth). In the more complex cases such as a mosaic a virtual data product may have multiple parent datasets.

If a process which produces data products is complex enough, with many inputs, ultimately the result is a new data collection, but in most runtime data access scenarios the simple provenance model presented here should be enough to identify a virtual data product or other dataset and how it was produced.

2.7 Data Association

There are many cases where it is desirable to be able to associate multiple datasets, for example to model a multi-spectral observation such as an Echelle, or to group datasets that represent the same data made available in several different data formats. Spectra of the member galaxies in a cluster might be a completely different type of association. In the case of images, a multi-band observation could be viewed as an association of several independent images, each in a single spectral band and with some shared observational metadata.

The approach taken in SSA to address this problem of *complex data* is to keep the basic data objects as simple as possible but use *association* to describe more complex entities. Hence, an Echelle observation could be viewed as a collection of independently accessible 1-D spectra which are logically associated. The spectra would include the individual Echelle orders and possibly an overall combined high resolution spectrum. Some extension metadata might also be provided to provide additional information describing the overall association. The individual spectra would be usefully accessible without requiring that a client application understand the complex instrument (an Echelle spectrograph) which produced the data, however the more complex view would optionally be accessible as well.

Associations are described in the SSA query response since this has the ability to relate multiple datasets. How this is done will be described further in the specification of the SSA query response, but the main technique is to define a new query response field Association.ID for which all members of an association share the same value. An association key may also be provided for each member of the association to uniquely identify their role within the association (e.g., the Echelle order in our example above). Finally, an association Type field or param tells what type of association this is. The ID

may be used to link to extension metadata providing further information describing the specific extension.

3 SSA Operations

3.1 Introduction

The operations currently defined by the SSA protocol are **queryData**, **getData**, **stageData**, **getCapabilities**, and **getAvailability**. Of these, currently only queryData, getCapabilities and getAvailability are defined as explicit parameterized operations. GetData is currently not implemented as a service operation, rather an explicit access reference URL is used to retrieve datasets.

The specification herein of whether support for a parameter is required, recommended, or optional refers to the service, not to the query submitted by the client. Unless otherwise specified by the operation, all parameters except `REQUEST` are optional for the client (depending upon the operation, invoking an operation with *no* parameters may however result in an invalid operation).

3.2 Methods & Protocols

As with all the DAL interfaces, a SSA service may eventually define interfaces for multiple “distributed computing platforms” or transports, e.g., HTTP GET/POST and SOAP. At this time only a HTTP GET interface is defined.

If the SSA query is transmitted as an HTTP GET request then the URL to express a data query is formed like this:

```
<Service.BaseURL>?VERSION=1.0&REQUEST=queryData<&param=value...>
```

Example:

```
http://www.myvo.org/ssa?VERSION=1.0&REQUEST=queryData&POS=22.438,-17.2&SIZE=0.02
```

The `Service.BaseURL` is stored in the IVOA resource registry (Hanisch et al. 2005).

3.3 QueryData Operation (required)

The purpose of the SSA query is to determine the availability and characterization of data satisfying the given search constraints. The result is returned encoded as a VOTable document wherein each row of the table describes one candidate dataset.

3.3.1 Input Parameters

A simple query is defined in terms of a 4-dimensional physical parameter space:

- spatial region (for SSA an aperture on the sky defined by `POS`, `SIZE`)
- temporal region (`TIME`)
- spectral region (`BAND`)

A minimally-compliant SSA service **must** support at least these four parameters, plus the `FORMAT` parameter, which specifies the format in which data (spectra) are to be returned. Various other parameters specific to spectrophotometric data are also defined, and **should** also be supported by the service if possible, to be fully compliant.

Unless otherwise specified, if the service does not support a query parameter defined by the protocol it **must** permit the parameter to be present in the query without error, even if the parameter is not actually used as a query constraint by the service. Most parameters are used to constrain the query; if a given parameter is not specified or is not supported by the service, a logical value of “all” is generally assumed. This allows a query to succeed even if it includes parameters which the service does not support, so that the same query can be submitted to multiple service instances. Since queries can be imprecise it is up to the client to analyze the returned query metadata to further refine the query.

If a service is required to support certain input parameters, that means that the service must be prepared to use such a parameter to constrain a query. If this is not done and the service merely ignores an input parameter which is required to be supported, then it may be easy for the client to pose a query which results in an overflow of the query response or some other error condition. For example, if a client queries for data based only on the spectral bandpass and the service does not support the `BAND` parameter, the query may overflow or be declared invalid even though valid data is available.

Specific parameters may or may not have meaningful values for a given data collection. For example, for theory data, anything having to do with time or position on the sky may be undefined. For solar or planetary data, time is defined but the spatial position on the celestial sphere may be undefined or not meaningful. In such a case, where a specific value is specified for an attribute which is undefined for a given data collection, the service should respond by finding no matching data. For data collections where all physical measurement parameters are meaningful, for example spectra of galactic or extragalactic astronomical targets, all parameters should be supported and used to constrain the query, even if only imprecise values of the parameters are known for a given dataset.

3.3.2 Mandatory Query Parameters

The following parameters **must** be implemented by a compliant service:

<i>Parameter</i>	<i>Sample value</i>	<i>Physical unit</i>	<i>Datatype</i>
POS	52, -27.8	degrees; defaults to ICRS	string
SIZE	0.05	degrees	double
BAND	2.7E-7/0.13	meters	string

TIME	1998-05-21/1999	ISO 8601 UTC	string
FORMAT	votable	-	string

While a compliant service must implement these parameters, a valid query can be composed from any combination of parameters; in other words, a SSA query does not have to be positional. All services must support queries by at least these five parameters, representing position in the fundamental physical measurement axes, and the desired client data formats.

3.3.2.1 POS

The center of the region of interest. The coordinate values are specified in list format (comma separated) with no embedded white space, as defined in section 5.7.2.

Example: `POS=52,-27.8`

POS defaults to right-ascension and declination in decimal degrees in the ICRS coordinate system. A coordinate system reference frame **may** optionally be specified to specify a coordinate system other than ICRS. The reference frame is specified as a list format modifier, with the acceptable values as defined in the spectral data model.

Example: `POS=52,-27.8;GALACTIC_I`

Coordinates requiring more than two values are possible merely by having more than two comma-delimited values before the qualifier.

Whether or not a service supports coordinate systems other than ICRS for POS is optional and is a service-defined capability. It is an error if a coordinate reference frame is specified which the service does not support.

3.3.2.2 SIZE

The diameter of the search region specified in decimal degrees.

Example: `SIZE=0.05`

A valid query does not have to specify a `SIZE` parameter. If `SIZE` is omitted in a positional query, the service should supply a default value intended to find nearby objects which are candidates for a match to the given object position. At most TOP candidate objects near the given position should be returned (see the discussion of the `TOP` parameter below).

3.3.2.3 BAND

The spectral bandpass is given in range-list form as defined in section 5.7.2, with each list element specified either numerically as a wavelength value or range, or as a spectral

bandpass identifier, e.g., a filter name or instrumental bandpass. The service **must** support at least one bandpass list element, which may be either a single value or (for numerical ranges) an open or closed range. If a single numerical value is specified it matches any spectrum for which the spectral coverage includes the specified value. If two values are given, a spectrum matches if any portion of it overlaps the given spectral region.

For a numerical bandpass the units are wavelength in vacuum in units of meters [RSM, Hanisch *et.al*, 2005]. The spectral rest frame **may** optionally be qualified as either “source” or “observer”, specified as a range-list qualifier.

Example: `BAND=1E-7/3E-6;source`

For most queries the precision with which the spectral bandpass is specified *in the query* probably does not matter. A rough bandpass broad enough to find all the interesting data will suffice; the more precise spectral bandpass specified in the query response for each spectrum can then be used to refine the query. In some cases, for example a cutout service operating upon high resolution spectra, support at the service level for specifying the spectral rest frame could be important. If the service does not support specification of the spectral frame the syntax should be permitted but ignored.

If a bandpass is specified as a string it is assumed to be a bandpass identifier such as a filter name or instrumental bandpass, as specified in the resource metadata [RSM, Hanisch *et.al*, 2005] for `Coverage.Spectral.Bandpass`. The service **may** choose to support bandpass names in the BAND parameter. Since there is no standard list of filter names or instrumental bandpasses, and these can overlap, there is no apriori way to know what to call a bandpass in a query; however it is possible to learn these values in a prior query to the same service and then use them to refine the query. If the service does not recognize a bandpass name, or does not support query by bandpass name, it should not match anything.

Example: `BAND=J`

Bandpass names are often not useful for spectra (they are probably more useful for image or time series data) but there are cases where they are useful for spectra, for example for a velocity spectrum of a specific emission line.

3.3.2.4 TIME

The time range (epoch) of the search given in range-list form as defined in section 5.7.2, specified in ISO 8601 format. If the time frame is not specified UTC is assumed. The value specified may be a single value or an open or closed range. If a single value is specified it matches any spectrum for which the time coverage includes the specified value. If a range is specified it matches any spectrum which contains any data in the specified range.

3.3.2.5 FORMAT

The `FORMAT` parameter defines the data formats the client is interested in retrieving via a subsequent `getData` call. The value is a comma-delimited list where each element can be any recognized MIME-type such as

```
application/x-votable+xml, application/fits, application/xml,
text/plain, text/html, image/jpeg
```

and so forth. If finer discrimination is necessary, MIME type parameterization can be used to more finely specify any allowable format specified by a service, for example

```
FORMAT=application/fits; convention=STScI-STIS
```

might specify the project specific format defined by the HST STIS instrument (this is an example; the specification of such MIME types is not defined by SSA). Normally this should not be required as `FORMAT=native` may be used to specify the native format for a data collection, if available. The `getCapabilities` operation of a service, or the registry cache of this information, may be used to determine the standard or custom formats supported by a given service instance. If not otherwise specified, the format-specific default serializations defined by SSA are assumed. In addition, the following special shorthand values are defined:

FORMAT	Meaning
<code>ALL</code>	All formats supported by the service
<code>COMPLIANT</code>	Any SSA data model compliant format
<code>NATIVE</code>	The native project specific format for a spectrum
<code>GRAPHIC</code>	Any of the graphics formats: JPEG, PNG, GIF
<code>votable</code>	Shorthand for <code>application/x-votable+xml</code> , the SSA VOTable format
<code>fits</code>	Shorthand for <code>application/fits</code> , the SSA FITS format
<code>xml</code>	Shorthand for <code>application/xml</code> , the SSA native XML serialization

These shorthand values all assume an SSA-compliant serialization. If `FORMAT` is omitted, `FORMAT=ALL` should be assumed.

The `FORMAT` parameter describes the desired format of returned data. If no data is available in the specified format, a null query response should be returned. If data is dynamically generated the service may generate the data in the requested format on the fly. Note `FORMAT` applies only to the data; the query response itself is always a VOTable document.

3.3.3 Recommended and Optional Query Parameters

The following additional parameters **should** or **may** be implemented by a service; all the recommended parameters are required for a fully compliant service. In the table below

and those following, mandatory parameters are indicated by MAN, recommended parameters by REQ, and optional parameters by OPT.

Parameter	Sample value	Unit	Req	Datatype
APERTURE	0.00028 (=1")	degrees	OPT	double
SPECRP	2000	$\lambda/d\lambda$	REC	double
SPATRES	0.05	degrees	REC	double
TIMERES	86760 (=1yr)	seconds	OPT	double
SNR	5.0	dimensionless	OPT	double
REDSHIFT	1.3/3.0	dimensionless	OPT	double
VARAMPL	0.77	Dimensionless	OPT	double
TARGETNAME	mars		OPT	string
TARGETCLASS	star		OPT	string
FLUXCALIB	absolute		OPT	string
PUBDID	ADS/col#R5983		REC	string
CREATORID	ivo://auth/col#R1234		REC	string
COLLECTION	SDSS-DR4		REC	string
TOP	20	dimensionless	REC	int
MAXREC	5000		REC	string
MTIME	2005-01-01/2006-01-01	ISO 8601	REC	string
COMPRESS	true		REC	boolean
RUNID	<string>		REC	string

The spectral resolution (specified here as the spectral resolving power) is an important characteristic of any spectrum. The spatial resolution is important as it is usually comparable to the diameter of the spectral aperture on the sky (larger values of the observational aperture are possible but can be determined from the query response metadata). The creator and publisher data identifiers and data collection name are important to identify the data. The remaining recommended parameters are important for full functionality of the protocol. All parameters are explained in more detail below.

3.3.3.1 APERTURE

The aperture parameter is used only for spectral extraction, i.e., computation of spectra derived from more fundamental data such as a spectral data cube or event list, using a synthetic aperture. The aperture is specified as a diameter in decimal degrees. The aperture parameter is only used for spectral extraction; a spectral extraction SSA service **must** support this parameter.

If no aperture is specified by the client the service should supply a default value appropriate to the data in question, for example, a circular aperture large enough to capture 98% of the signal from a point source in the aperture, knowing the spatial resolution of the data in the desired spectral band. The size of the aperture used to generate the simulated data should be returned in the description of the data in the query response table. The service should not normally use SIZE as or the default value of the aperture for spectral extraction, as this will generally represent an upper limit on the

maximum separation of the target position and observed position when the SSA query is used to search for data potentially matching a given target object, not knowing whether a given service is searching pre-existing spectral collections or computing extracted spectra.

Note that SSA makes it possible to find data for a specific point source object on the sky merely by specifying the estimated object position. For catalog spectra, SIZE defaults to whatever is appropriate for a possible match for an object in the catalog. For extracted spectra, the measurement aperture should be a value judged to be appropriate for the spatial resolution of the data.

If these simple heuristics are not adequate, the client data analysis application should explicitly specify the diameter of the synthetic aperture to be used.

3.3.3.2 SPECRP

The minimum spectral resolution, specified as the spectral resolving power $d\lambda/\lambda$ in dimensionless units.

3.3.3.3 SPATRES

The minimum spatial resolution (corresponding to the PSF of the observed signal) specified in decimal degrees.

3.3.3.4 TIMERES

The minimum time resolution, specified in seconds. For a typical spectrum the time resolution corresponds to the bounds of the time coverage of the exposure.

3.3.3.5 SNR

The minimum signal-to-noise ratio of a candidate dataset, for example specified as the ratio of the mean signal to the RMS noise of the background (see the SSA data model document for more detailed recommendations on how to compute the SNR).

3.3.3.6 REDSHIFT

A photometric (observed) redshift range specified as a single element open or closed range-list as defined in section 5.7.2. A negative redshift indicates a “blueshift”, e.g., an object in the local neighborhood with a proper motion towards the Earth (a negative redshift is not proper terminology but this is thought to be simpler than other alternatives such as defining new terminology or adding additional parameters). An open range may be used to specify a minimum or maximum value. The optical redshift convention should be used ($d\lambda/\lambda$).

Example: 1.2/3

3.3.3.7 VARAMPL

The desired range of variability amplitude, specified as a range, with values in the range 0.0 to 1.0.

3.3.3.8 TARGETNAME

The target name, suitable for input to a name resolver. In general target name resolution should take place on the client side, where greater sophistication is possible without raising the minimal requirements for a service. The reason that `TARGETNAME` is included here is to make it possible to find spectra of objects that do not have a known position, for example, spectra of solar system planets or asteroids. For a service which can supply spectra for moving objects, `TARGETNAME` is a required parameter; for other SSA services it is not normally required and name resolution is assumed to take place on the client side.

3.3.3.9 TARGETCLASS

A comma delimited list of strings denoting the types of astronomical objects to be searched for. At the moment there is no standard classification for astronomical objects but it is suggested to use the “condensed” names from the list at <http://simbad.u-strasbg.fr/guide/chF.htx>.

Examples: `star, galaxy, pulsar, PN, AGN, QSO, GRB`

3.3.3.10 FLUXCALIB

Specifies whether only flux calibrated data is to be found. Possible values are “true” and “false”. Data which has any kind of flux calibration (either relative or absolute) qualifies as flux calibrated here.

3.3.3.11 PUBDID

The IVOA publisher’s dataset identifier, assigned by the publisher of a dataset. The same dataset published in different places may have a different PUBDID assigned by each publisher, however, unlike `CREATORID`, where data creators may often not assign IVOA identifiers; it is guaranteed that a publisher can always assign a unique PUBDID when a dataset is published to the VO. ADS dataset identifiers are an example of a PUBDID, but in general any publisher may assign their own unique publisher dataset identifier. Publisher dataset identifiers may be determined by a prior query or some external means, such as another form of archive query.

Note:

A special case of a publisher's dataset identifier is the **ADS dataset identifier**, used to reference published IVOA datasets in journal articles.

3.3.3.12 CREATORDID

The IVOA dataset identifier, assigned at creation time by the creator of the parent data collection (survey project, observatory, etc.). Datasets can have a unique CreatorDID assigned prior to publication of the data to the VO, for example when the data is generated by a processing pipeline, or ingested into the master archive for the data collection. This is possible since the Creator entity for a data collection (e.g., an observatory or survey project) controls its own namespace. When a CreatorDID has been assigned this is the most universal way to refer to a dataset; all replicated versions will share the same CreatorDID. Creator dataset identifiers may be determined by a prior query or by some other means, such as another form of archive query.

Example: `ivo://nrao.edu/vla#1998s2/4992a`

3.3.3.13 COLLECTION

The IVOA identifier or “shortName” of a data collection as defined by the service, for example `SDSS-DR2`, or `NRAO-VLA`. By data collection we refer to an organized, uniform collection of datasets from a single source, for example a single data release from a survey, or an instrumental data collection from an observatory. The service should treat the search term as a case insensitive, minimum match string. For instance, “dss” would match either `dss1` or `ESO-DSS2`. Allowable data collection references are specified in the service capabilities.

3.3.3.14 TOP

`TOP` limits the number of returned records in the query response table to the specified number of top ranking ones. Records are ranked according to a “score” heuristic [see for example Dolensky 2006]. The details of the actual heuristic used are up to the service, but the general idea is that the better a candidate dataset matches the query, the higher the score it receives. Metrics such as distance from the specified position, or the degree of overlap with a specified bandpass or time interval, determine the score. If two datasets would otherwise have the same score, the service may use other unspecified dataset characteristics, such as some intrinsic data quality metric, to further rank candidate datasets. If the service implements a ranking heuristic the query response table should normally be returned sorted in order of decreasing score. `TOP` can also be used by the client to limit the size of the query response table in cases where the query might find a very large number of candidate objects.

3.3.3.15 MAXREC

The maximum number of records to be returned. This may be used to increase the built-in default limit set in the service, up to some maximum service-specified default (this is provided in an attempt to permit larger queries without having to page through the query response, which requires saving state on the server). A service should typically have a

fairly small default MAXREC, provided to improve the query response time, and a large upper limit on MAXREC, provided to enable large queries.

3.3.3.16 MTIME

Find only datasets modified or created in the given range of dates, specified in range-list format, as an open or closed range, with the dates specified in ISO 8601 format. Note this is not the same thing as TIME, which refers to time of observation. MTIME may be used to periodically query services for new or updated data.

3.3.3.17 COMPRESS

If this flag is present, datasets returned via the *getData* method may optionally be returned to the client in compressed form. Valid values are “true” and “false”; if no value is given, “true” is assumed. By default compressed data is not permitted.

Compression is performed by applying a whole-file compression algorithm such as *gzip*, and updating the HTTP content type of the returned document accordingly. Dataset-level compression should be distinguished from protocol-level compression, which is performed at the level of the HTTP protocol, on the entire data stream, and is transparent to the client.

3.3.3.18 RUNID

The RUNID is an opaque string used to associate multiple service invocations in service logs, e.g., to identify them as all belonging to the same job or application. RUNID is not used by SSA in any way, except in cases where SSA may call another VO service, in which case the RUNID parameter should be passed on to the called service.

3.3.4 Service-Defined Parameters

The service **may** support additional service-defined parameters. Parameter names must not match any of the reserved parameter names defined herein, independent of case.

Service defined parameters should be defined by the service in the service metadata returned by the *getCapabilities* operation (3.5).

3.3.5 Query Response

The output returned by a query is an XML document compliant with VOTable V1.1 or higher (VOTable 2004) and should be returned with a MIME-type of `application/x-votable+xml`.

Note:

The `FORMAT` parameter has no influence on the query response. `FORMAT` applies only to the returned datasets, not to the query response. The query response is always returned as a VOTable.

The VOTable **must** contain a RESOURCE element, identified with the tag `type = "results"`, containing a single TABLE element with the results of the query. Additional RESOURCE elements may be present, but the usage of any such elements is not defined here.

The RESOURCE element **must** contain an INFO with `name="QUERY_STATUS"`. Its value attribute should be set to "OK" if the query executed successfully, regardless of whether any matching data were found. All other possible values for the value attribute are described in section 5.10.

Examples:

```
<INFO name="QUERY_STATUS" value="OK"/>
<INFO name="QUERY_STATUS" value="OK">Successful Search</INFO>
```

In the response table each row represents a different dataset which is potentially available to the client. VOTable GROUP constructs are used to associate related groups of fields.

Hint:

Put constant values in PARAM elements instead of repeating them in each table row.

3.3.6 Query Response Metadata

Names of fields and parameters are left to the service provider. UTYPEs of standard fields are required for identification of interface elements and **must** be given and **must** comply with the SSA protocol (this document) and the SSA data model (McDowell, Tody, et. al. 2005). UCDs **should** also be given but are not used to identify interface or data model elements. Values for the UCDs of standard interface and data model elements are given in the SSA and Spectrum data model documents.

Note:

UTYPE values **must** be provided to identify interface or data model elements.
UCD values for standard data model elements **should** be provided as well.
Omit the leading "spectrum." in the UTYPE for Spectrum data model attributes.

The SSA query response consists of a number of fields, identified by UTYPE, grouped into component data models of the form "`<component-name>'.<field-name>`". Some components of the query response are defined directly by the SSA protocol (this document), while others are taken directly from the Spectrum data model. Unless otherwise specified, the leading "spectrum." in the UTYPE values specified in the Spectrum data model is omitted in the SSA query response since this metadata is not specific to spectra and we use the same metadata for other types of data objects. Hence most of the query response metadata consists of generic component data models. For example, if the Spectrum data model specifies "Spectrum.Target.Name", this appears in

the SSA query response as “Target.Name”. Applications can refer to “Target.Name” regardless of whether the data to be accessed is a Spectrum or some other data object such as an Image.

In the following, query response parameters which are mandatory, recommended, or optional are indicated as such in the tables or specified more precisely in the text. Additional attributes from the Spectrum data model not shown here may appear in the query response table. The SSA query response does not include any actual data values, and elements of the Spectrum data model used to represent data values are not included here (the client needs to download the full dataset to get the data).

When a generic data model is applied in a specific context, the requirements for what is required, what is optional, and flexibility in what is permitted will vary depending upon how the data model is being used. Hence when data model attributes are indicated as **mandatory or recommended** in this document, this overrides any similar requirements specified in the Spectrum data model document. The SSA query response is also more restrictive than the underlying model; in particular the **allowable units** are more restrictive than what is permitted in the model, in order to be more consistent with other elements of SSA, and to provide more uniformity to make multiband data discovery by the client easier. Hence within SSA, characterization restricts the allowable units for spatial coordinates to decimal degrees, for spectral coordinates to wavelength in meters, and for time measures to seconds, except where MJD is used (all represented in floating point).

It is difficult to specify every detail of every metadata element in this document without burdening the text with too much detail; furthermore, many optional metadata values are omitted from the summary tables shown here. Full details are given in the Spectrum data model document, and in a convenient summary form in a spreadsheet which lists all metadata elements with full details for each. All this information can be found on the SSA TWiki page at <http://www.ivoa.net/twiki/bin/view/IVOA/SsaInterface>.

Query metadata may be mapped to VOTable fields in any order, so long as fields which are part of the GROUP construct (all the component data models are GROUP elements) appear in consecutive table fields.

3.3.7 Types of Metadata

Metadata in the query response is grouped into a number of component data models as summarized in the table below, and explained in more detail in the sections which follow.

Service Metadata	
Query	Describes the query itself
Association	Logical associations
Access	Dataset access-related metadata
Data Model Metadata	
Dataset	General dataset metadata

DataID	Dataset identification (creation)
Curation	Publisher metadata
Target	Observed target, if any
Derived	Derived quantities
CoordSys	Coordinate system frames
Char	Dataset characterization
Characterization Metadata	
Char.FluxAxis	Observable, normally a flux measurement
Char.SpectralAxis	Spectral measurement axis, e.g., wavelength
Char.TimeAxis	Temporal measurement axis
Char.SpatialAxis	Spatial measurement axis
Char.*.Coverage	Coverage in any axis
Char.*.Accuracy	Resolution and error in any axis

Service metadata is specific to the functioning of the service itself, for example to step through large queries or retrieve selected datasets. **Data model metadata** describes each dataset, and is common between the SSA protocol and the Spectrum data model. **Characterization metadata** physically characterizes the dataset in terms of the spatial, spectral, and temporal measurement axes and the observable. Characterization is part of the data model but is broken out separately in the table above to show the major elements of the characterization model. Most of the metadata returned by SSA is generic dataset metadata, which means it is not actually specific to spectra and may be used in other DAL interfaces to describe other types of dataset, for example an image or catalog. For data model metadata, please refer to the Spectrum data model for details such as the UCD and units, unless specified otherwise in this document.

Each of these types of query response metadata is discussed in more detail in the sections which follow.

3.3.8 Query Metadata

Query metadata describes the query itself.

UTYPE	UCD	Description	Req
Query.Score		Degree of match to query params	REC
Query.Token		Continuation reference for large queries	REC

3.3.8.1 Query.Score

A record with a higher score more closely matches the query parameters. The score is expressed as a floating point number with an arbitrary scale (different queries may return results with different scale factors and cannot be inter-compared). If scoring is used, the query response table should be returned sorted in order of decreasing values of score, with the top-scoring items at the top of the list. The details of the heuristic used to

compute the score are left to the service. See the discussion of the TOP parameter in section 3.3.3.14.

3.3.8.2 Query.Token

[This has been removed in the most recent version of the interface – see MAXREC for the new strategy.] If the query response is too large to return in a single query the service **should** return a string token to be used by the client to obtain the next page of the query response (otherwise it may be omitted, or an error status returned upon query overflow). To obtain the next page of the query response the client invokes the *queryData* operation again with the `TOKEN` parameter set to the given value. Since `TOKEN` refers to the entire query response it should be returned as a `PARAM` in the output `VOTable`.

The content of `TOKEN` is up to the service. For example, if the service maintains state for the query the value could be a database key used to indicate the next page of query output results. If the service does not want to maintain state and re-execution of the query will return the same results, token could be an index into the query response, with the query re-executed in each call while updating the index until the query response is exhausted.

3.3.9 Association Metadata

Association metadata is used to describe logical associations relating datasets described in the query response, as described in section 2.7. Logical associations between query response records may refer to the data access operation itself, e.g., where the same data object is available in multiple output formats, or to logical associations relating the physical data, e.g., where multiple primary datasets are part of the same observation. The same dataset may belong to multiple associations.

UTYPE	UCD	Description	Req
Association.Type		type of association	OPT
Association.ID		unique ID for entire association	OPT
Association.Key		unique ID for element of association	OPT

3.3.9.1 MultiFormat Association

A pre-defined case is the **MultiFormat** association, where several records refer to the same dataset which is available in several different data formats. In this case `Association.Type` should be set to “MultiFormat”, `Association.ID` can be anything (an object name suitable for use as a file or record name would be a good choice), and `Association.Key` should be set to “@Access.Format” to indicate that the key which differentiates the elements of the association is the value of the `Access.Format` field of the record. If several query response records are of this type the association **should** be specified to indicate the association. In all other cases (currently undefined by the protocol) the association **may** be specified.

3.3.9.2 Association.Type

A service-defined type used to indicate what type of association is being referred to. The value should be unique within the scope of the query response. There can be many types of logical associations. Associations provide a means of describing complex data aggregations relating multiple datasets (spectra in the case of SSA). Association is a type of extension mechanism, and the definition of associations is beyond the scope of SSA; SSA merely provides the means to define and manipulate associations. Examples of possible associations might be an Echelle observation consisting of 100 orders, each of which appears in the query response as an individual 1-D spectrum, or a group of query response records which all refer to the same dataset but differ only in the output format.

Since the association type may be shared by many table records, it may be best specified as a PARAM in the output VOTable, using an ID-REF to link it to the association it refers to. An association type **should** be provided for each association in the table.

3.3.9.3 Association.ID

The association ID is a string, unique within the scope of a given VOTable, identifying the records belonging to a given association. All members of the association have the same ID. The association ID **must** be provided for any association. The content of the string is up to the service. All records in the association share the same association ID. Multiple association IDs may be provided if a record belongs to more than one association.

Extension metadata may optionally be provided to describe an association in more detail. Extension metadata appears in the output VOTable as optional additional RESOURCE elements. The ID-REF mechanism may be used to link such an extension record to the association in the main table. The contents of an association metadata extension record are externally defined and beyond the scope of SSA.

3.3.9.4 Association.Key

The association key **should** be used to identify what is “different” for each member of an association. The value is a string and may be either an arbitrary value defined by the association, or a reference to one or more table fields which form the association key. If a table field is referenced the ‘@’ character should be prefixed to the field UTYPE to indicate the indirection, otherwise the literal string is used as the key. A key may contain multiple elements delimited by commas.

3.3.10 Access Metadata

Access metadata is required to tell a client how to access the datasets described in the SSA query response.

UTYPE	UCD	Description	Req
Access.Reference		dataset access URL	MAN
Access.Format		MIME type of dataset	MAN

Access.Size		rough dataset size	REC
-------------	--	--------------------	-----

3.3.10.1 Access Reference

The access reference is a URI (typically a URL) which can be used to synchronously retrieve the specific dataset described in a row of the query table response. If the dataset pointed to by the access reference does not exist at query time, it will be computed on the fly when accessed.

Since the datasets supported by SSA are typically small (compared to images), SSA does not currently support data staging and asynchronous data access. Support for this may need to be added in the future, e.g., to support generation of simulated or synthetic data.

3.3.10.2 Output Format

The file format of a candidate dataset is specified by its MIME type. Both uncompressed and compressed data can be indicated in this fashion.

The file format says nothing about the data model used by whatever data object is stored in the file; this is specified by the `DataModel` attribute discussed in section 3.3.11.1.

A single data object may be available in multiple file formats. In such a case an association **should** be defined to indicate that the entries all refer to the same data object.

3.3.10.3 Dataset Size Estimate

The estimated size of the dataset in bytes **should** be given to help the client estimate the download time required to retrieve the data. Only an approximate, order of magnitude value is required.

3.3.11 Data Model Metadata

The following metadata components are in common with the Spectrum data model

3.3.11.1 General Dataset Metadata

General dataset metadata describes the overall dataset.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
Dataset.DataModel	Datamodel name and version	MAN	Spectrum-1.0
Dataset.Type	Type of dataset	OPT	Spectrum
Dataset.Length	Number of points in spectrum	MAN	
Dataset.TimeSI	SI factor and dimensions	REC	
Dataset.SpectralSI	SI factor and dimensions	REC	
Dataset.FluxSI	SI factor and dimensions	REC	

`DataModel` is a string identifying the data model type and version used in the retrievable dataset being described. For SSA-compliant data this should be a value such as "Spectrum-1.0", as specified in the Spectrum data model document for the version of the data model being used. For pass-through of foreign or native project data some other value should be used which identifies the specific project data model used, e.g., "HST-STIS-1.0".

For SSA, `Type` is always "Spectrum" and can normally be omitted. `Length` is mandatory and specifies the length of the spectrum, i.e., the number of data points or samples. The SI parameters provide a simplified approach to defining the units of the spectral coordinate and flux density quantities, e.g., for overplotting spectra from different sources.

3.3.11.2 Dataset Identification Metadata

Dataset identification metadata is used to describe the fundamental identify of a dataset, including where it came from and how it was created.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
DataID.Title	Dataset title	MAN	
DataID.Creator	Creator name (string)	REC	
DataID.Collection	Collection shortname (string)	REC	
DataID.DatasetID	IVOA Dataset ID	OPT	
DataID.CreatorDID	Creator assigned dataset identifier	REC	
DataID.Date	Data processing/creation date	OPT	
DataID.Version	Version of creator-produced dataset	OPT	
DataID.Instrument	Instrument name	OPT	
DataID.Bandpass	Bandpass name, e.g., filter	OPT	
DataID.DataSource	Original source of data	REC	survey
DataID.CreationType	Dataset creation type	REC	archival

`Title` is a short, human-readable description of a dataset, and should be less than one line of text. Information such as the instrument or survey, filter, target name, etc., is typically included in a condensed form. The exact contents of `Title` are up to the data provider. `Creator` identifies the entity which created the dataset, and should be a short string consistent with the RSM specification, e.g., "SDSS". `Collection` is the shortName of the data collection to which the dataset belongs, e.g., "SDSS-DR5".

The `CreatorDID` is the IVOA dataset identifier assigned by the entity which created the data content, e.g., an observatory or survey project. If the dataset referred to is virtual data, `CreatorDID` refers to the parent dataset from which the virtual data will be created (see 2.4.2 for further details). `Date`, specified in ISO time format, specifies the date when the dataset was created or last modified by the `Creator` entity. If a dataset is modified or replaced without changing its `CreatorDID`, `Date` and `Version` should be updated accordingly.

`Instrument` is a short string identifying the instrument used to create the data (instrument may be an actual instrument or something else, e.g., a program in the case of theory data). `Bandpass` is a short string specifying the bandpass name if any, e.g., a filter name or an instrumental bandpass such as I, J, K, Q, HI, and so forth. Values specified with `Bandpass` may be used as input to the `BAND` parameter (3.3.2.3) to refine a query.

`DataSource` and `CreationType` describe the original source of the data, and how the dataset returned by the service was or will be created, as defined in section 2.4.

3.3.11.3 Curation Metadata

Curation metadata describes who curates the dataset and how it is published to the VO.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
Curation.Publisher	Publisher	MAN	
Curation.Reference	URL or Bibcode for documentation	REC	
Curation.PublisherDID	Publisher's ID for the dataset	REC	
Curation.Date	Data curated dataset last modified	OPT	
Curation.Version	Version of curated dataset	OPT	
Curation.Rights	Restrictions: public, proprietary, etc	OPT	public

`Publisher` is a short string identifying the publisher of the data, e.g., a data archive or data center, or an indexing service such as the ADS. The `PublisherDID` is the IVOA dataset identifier (URI) assigned by the publisher to identify the dataset within its holdings. `Reference` is a forward link to publications which reference the dataset; multiple instances are permitted. `Date` and `Version` refer to the dataset *as curated by the publisher*, hence can differ from the same values given in `DataID`, which refer to the *content* of the dataset as generated by the dataset `Creator`. `Rights` specifies whether the dataset is "public" or "proprietary". Proprietary data requires authentication and authorization by the data provider to access, and once downloaded should be protected from subsequent access on the client side.

Note:

If the same dataset is replicated at several locations with multiple publishers, it is possible to set up an association group to indicate this fact.

3.3.11.4 Astronomical Target Metadata

Target metadata describes the astronomical target observed, if any.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
Target.Name	Target name	OPT	
Target.Class	Target or object class	OPT	
Target.Redshift	Target redshift	OPT	
Target.VarAmpl	Target variability amplitude, typical	OPT	
Derived.SNR	Signal-to-noise for spectrum	OPT	

The target `Name` is a short string identifying the observed astronomical object, suitable for input to a name resolver. `Class` is the object class if known, e.g., Star, Galaxy, AGN, QSO, and so forth (see section 3.3.3.9). `Redshift`, `VarAmpl`, and `SNR` are as defined in the data model. Either standard target values, or derived quantities, may be used in the query response.

3.3.11.5 Coordinate System Metadata

Coordinate system metadata describes the coordinate system reference frames used in the SSA query response.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
CoordSys.SpaceFrame.Name	Spatial coordinate frame	REC	ICRS
CoordSys.SpaceFrame.Equinox	Equinox	OPT	2000.0
CoordSys.TimeFrame.Name	Timescale	OPT	TT
CoordSys.TimeFrame.Zero	Zero point of timescale in MJD	OPT	0.0

These reference frames apply to all spatial (sky), spectral, and time coordinates used in the SSA query response (including Characterization) unless otherwise specified. Note that spatial coordinates are not limited to the celestial sphere; any spatial coordinate frame specified in the data model may be specified, including solar and planetary coordinate systems, although the default is ICRS.

3.3.11.6 Dataset Characterization Axis Metadata

The Characterization axis metadata specifies the type of physical quantity on each physical measurement axis as well as the observable.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
Char.FluxAxis.Ucd	ucd for flux	REC	
Char.SpectralAxis.Ucd	ucd for spectral coord	REC	

Values are specified as UCDs, as defined in the data model. For example, to specify that the flux axis is flux density per unit wavelength, the value "phot.fluDens;em.wl" would be given.

3.3.11.7 Characterization Coverage Metadata

The `Coverage` component of the Characterization data model describes the coverage of the dataset in each of the three primary measurement axes.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
Char.SpatialAxis.Coverage.Location.Value	Observed position, e.g., RA DEC	MAN	
Char.SpatialAxis.Coverage.Bounds.Extent	Aperture angular diameter, deg	MAN	
Char.SpatialAxis.Coverage.Support.Area	Aperture region	OPT	
Char.SpatialAxis.Coverage.Support.Fill	Sampling filling factor	OPT	1.0
Char.TimeAxis.Coverage.Location.Value	Midpoint of exposure (MJD)	MAN	
Char.TimeAxis.Coverage.Bounds.Extent	Total elapsed time	REC	
Char.TimeAxis.Coverage.Bounds.Start	Start time	OPT	
Char.TimeAxis.Coverage.Bounds.Stop	Stop time	OPT	
Char.TimeAxis.Coverage.Support.Extent	Effective exposure time	OPT	
Char.TimeAxis.Coverage.Support.Fill	Sampling filling factor	OPT	1.0
Char.SpectralAxis.Coverage.Location.Value	Midpoint of Spectral coord range	MAN	
Char.SpectralAxis.Coverage.Bounds.Extent	Width of spectrum in A or other unit	MAN	
Char.SpectralAxis.Coverage.Bounds.Start	Start in spectral coordinate	REC	
Char.SpectralAxis.Coverage.Bounds.Stop	Stop in spectral coordinate (see text)	REC	
Char.SpectralAxis.Coverage.Support.Fill	Sampling filling factor	OPT	1.0

Within Char, `Coverage` specifies the *location* (central or characteristic value), *bounds* (measurement limits), *support* (region covered within the bounds), and *filling factor* (fraction of total area covered) for each measurement axis. The coordinate system reference frames specified in `Coordsys` apply here. Spatial coordinates are specified in units of decimal degrees, spectral coordinates in units of meters, and time coordinates in units of days.

3.3.11.8 Characterization Accuracy and Error Metadata

The `Accuracy` component of Characterization specifies the sampling, resolution, and error estimates for the dataset.

<i>UTYPE</i>	<i>Description</i>	<i>Req</i>	<i>Default</i>
Char.FluxAxis.Accuracy.StatError	Statistical error	OPT	
Char.FluxAxis.Accuracy.SysError	Systematic error	OPT	
Char.FluxAxis.Accuracy.Calibration	Type of flux calibration	REC	Calibrated
Char.SpectralAxis.Accuracy.BinSize	Wavelength bin size	OPT	
Char.SpectralAxis.Accuracy.StatError	Spectral coord measurement error	OPT	
Char.SpectralAxis.Accuracy.SysError	Spectral coord measurement error	OPT	
Char.SpectralAxis.Accuracy.Calibration	Type of coord calibration	REC	Calibrated

Char.SpectralAxis.Resolution	Spectral resolution FWHM	REC	BinSize
Char.TimeAxis.Accuracy.BinSize	Time bin size	OPT	
Char.TimeAxis.Accuracy.StatError	Time coord statistical error	OPT	
Char.TimeAxis.Accuracy.SysError	Time coord systematic error	OPT	
Char.TimeAxis.Accuracy.Calibration	Type of coord calibration	OPT	Calibrated
Char.TimeAxis.Resolution	Temporal resolution FWHM	OPT	BinSize
Char.SpatialAxis.Accuracy.StatError	Astrometric statistical error	REC	
Char.SpatialAxis.Accuracy.SysError	Systematic error	OPT	
Char.SpatialAxis.Accuracy.Calibration	Type of coord calibration	REC	Calibrated
Char.SpatialAxis.Resolution	Spatial resolution of data	REC	

Both overall statistical and systematic error estimates may be specified. The calibration status of all three primary measurement axes as well as the observable **should** be given, otherwise "calibrated" is assumed. The spatial and spectral resolution **should** be specified. Note that, for consistency within Char, the spectral resolution is specified here in spectral coordinate units (FWHM in meters), unlike the SPECRES query parameter, which is specified as $\lambda/d\lambda$.

3.3.12 Additional Service-Defined Metadata

A given service **may** return additional query response metadata not defined by the SSA protocol. This additional metadata may take the form of additional table columns, or additional `RESOURCE` elements in the query response VOTable.

Service-defined output metadata **should** use service-defined UTYPEs and UCDs as long as they do not clash - and can be easily distinguished - from mandatory and reserved SSA output columns.

3.3.13 Metadata Extension Mechanism

[To be added]. The metadata extension mechanism allows a data provider to add additional custom metadata to the query response to describe collection-specific details of the data. Extension metadata appears in the query response as additional `RESOURCE` elements in the VOTable. The format and contents of these `RESOURCE` elements is up to the data provider. The ID-REF mechanism of VOTable is used to link to associated fields of the main query response VOTable.

3.4 GetData (reserved)

The current SSA protocol does not include an explicit specification for a *getData* operation; an access reference URL is used instead, to provide maximum flexibility in how the *queryData* operation refers datasets back to the service for access. A more explicit *getData* operation remains an option in the future for accessing referenced datasets. This would still allow a URL to be used as at present, but could provide finer control over access-related options such as the output data format, or the ability to refer to a virtual dataset externally by a publisher assigned dataset identifier.

3.5 *GetCapabilities (recommended)*

The *getCapabilities* operation is used to query the service for its capabilities and interface. This replaces the old `FORMAT=METADATA` mechanism used in earlier DAL interfaces. The *getCapabilities* operation requires an agreement between multiple working groups and is still in the process of being specified. The most recent proposal is that *getCapabilities* should return the “Capability” element of a *VOResource*, within a simple XML container defined by SSA.

3.6 *StageData (reserved)*

The *stageData* operation is used in DAL interfaces to request asynchronous generation of one or more, possibly virtual datasets, as identified in a prior call to *queryData*. Greater flexibility in staging data, specifying where generated data is to be delivered, including support for third-party data delivery, will also be possible in a *stageData* request. Providing this capability is not a priority for SSA V1.0, but it may be added in a later version of the interface.

3.7 *GetAvailability (recommended)*

[To be added]. The *getAvailability* operation is used by an external client (normally the VO or grid infrastructure) to monitor the availability and health of a service. This is planned to be included in the final V1.0 interface, but has not yet been fully specified.

4 Data Retrieval

The data retrieval request allows a client to retrieve a single spectrum given an access reference `acref` as returned by a prior spectrum query. Since a spectrum query is required to obtain an `acref`, no further constraints are placed on the form the `acref` takes. This has the effect of hiding the details of the `acref` URL from the client, making it easy to layer an implementation of a getter web method on top of an existing retrieval service, and making it easier to hide changes to the implementation of existing services.

4.1 *Access Reference URL*

The access reference or `acref` is a simple URL (RFC 1738). In principle the URL may reference transports other than HTTP but at the present time this is not recommended. If the client issues a HTTP GET request using this URL, and the request is successful, the client will receive a document of the type given in query response column with the `UTYPE Access.Reference`. Since a prior query to the service is required to obtain an `acref`, no requirements are placed on the form the `acref` takes; this is completely up to a given service implementation. This has the effect of hiding the details of the `acref` URL from the client, making it easy to layer an implementation of the GET web method on top of an existing data retrieval service, and making it easier to hide changes to the implementation of existing services.

4.2 Data Format

The response to a data retrieval request is a single Spectrum instance. Both the data mode of the returned spectrum, and the file data format, may vary, but must agree with what is specified in the query response. The data model of a spectrum will be either some version of the Spectrum data model, or an externally-defined data model such as the native project data model. The available file formats will in general depend upon the data model, but for Spectrum they include at least the following:

- `application/x-votable+xml`
- `application/fits`
- `application/xml`
- `text/csv` – comma separated values
- `image/jpeg` - graphics preview
- `text/html`

The graphics formats and `text/html`, if available, provide a directly viewable, rendered version of the spectrum. All the other formats return science data.

4.3 Data Compression

If the query parameter `COMPRESS` is present then the service **may** return a compressed dataset, using some standard compression technique such as *gzip*, in place of a normal dataset, without indicating this in the query response. Basically the client is indicating that it is prepared to receive either compressed or uncompressed datasets and does not care which is delivered (the service should pick whichever is more efficient). This should be distinguished from protocol-level compression, which is transparent to the client, and may occur at the level of the HTTP protocol if both client and server support HTTP protocol compression.

In case of an HTTP GET the keyword `Content-Encoding` informs the receiver about the encoding of the output document, like for instance `x-gzip`. Care needs to be taken to treat the encoding distinct from the MIME-type of the encoded data stream.

4.4 Error Response

If possible, unsuccessful data retrieval should return a standard VOTable-format service error response, as outline in section 5.10. Depending upon the nature of the error or how data retrieval is implemented, this may or may not be possible, and a HTTP error may result instead. The client should be prepared to handle either form of error. In particular, if the operation is successful at the HTTP level, the client must check for a VOTable error response to be sure that an error has not occurred.

5 Basic Service Elements

5.1 Introduction

This clause specifies aspects of SSA service behaviour that are independent of particular operations or are common to several operations.

5.2 Version numbering and negotiation

5.2.1 Version number form and value

The SSA protocol defines a protocol version number. The version number applies to the XML schema and the request encodings defined in this document. The version number follows IVOA document conventions and contains two non-negative integers, separated by decimal points, in the form “x.y”, for example, “1.0”, or “1.13”.

5.2.2 Version number changes

The protocol version number shall be changed with each revision of this document. The number shall increase monotonically and shall comprise no more than two integers separated by decimal points, with the first integer being the most significant. There may be gaps in the numerical sequence. Some numbers may denote draft versions. Servers and their clients need not support all defined versions, but shall obey the negotiation rules below.

A version number change at the first level (1.0 – 2.0) indicates a major change. A version number change at the second level indicates a minor change which is not necessarily backwards compatible. A version number change at the third level is considered backwards compatible, and should not affect the pre-existing functionality of the interface.

5.2.3 Appearance in requests and in service metadata

The version number may appear in at least two places: in the service metadata and in the parameter list of client requests to a server. The version number used in a client’s request of a particular server shall be equal to a version number which that server has declared it supports (except during negotiation, as described below). A server may support several versions, whose values clients may discover according to the negotiation rules.

5.2.4 Version number negotiation

If a SSA client does not specify the version number in a request, the server assumes the highest standard version supported by the service, and no explicit version checking takes place. If the client specifies an explicit version number, and this does not match a version available from the service at level two, the service returns a version number mismatch error. The client can determine what versions of the protocol the service supports by a prior call to getCapabilities. All service metadata shall include a protocol version number and shall comply with the XML Schema defined for that version.

5.3 General HTTP request rules

5.3.1 Introduction

This document defines the implementation of the SSA service on a distributed computing platform (DCP) comprising Internet hosts that support the Hypertext Transfer Protocol (HTTP) (see IETF RFC 2616). Thus, the Online Resource of each operation supported by a server is an HTTP Uniform Resource Locator (URL). The URL may be different for each operation, or the same, at the discretion of the service provider. Each URL shall conform to the description in IETF RFC 2616 (section 3.2.2 “HTTP URL”) but is otherwise implementation-dependent; only the query portion comprising the service request itself is defined by this document.

While the SSA protocol currently only supports HTTP as the DCP for general parameterized operations, data access references are more general and may use other internet protocols, e.g., FTP, or potentially grid protocols.

HTTP supports two request methods: GET and POST. One or both of these methods may be offered by a server, and the use of the Online Resource URL differs in each case. Support for the GET method is mandatory; support for the POST method is optional.

5.3.2 Reserved characters in HTTP GET URLs

The URL specification (IETF RFC 2396) reserves particular characters as significant and requires that these be escaped when they might conflict with their defined usage. This document explicitly reserves several of those characters for use in the query portion of SSA requests. When the characters “?”, “&”, “=”, and “,” appear in one of the roles defined in Table 1, they shall appear literally in the URL. When those characters appear elsewhere (for example, in the value of a parameter), they shall be encoded as defined in IETF RFC 2396. The server shall be prepared to decode any character escaped in this manner.

Table 1 — Reserved characters in SSA query string

Character	Reserved usage
?	Separator indicating start of query string.
&	Separator between parameters in query string.
=	Separator between name and value of parameter.
,	Separator between individual values in list-oriented parameters (such as POS, BAND and TIME)

5.3.3 HTTP GET

A SSA service shall support the “GET” method of the HTTP protocol (IETF RFC 2616).

An Online Resource URL intended for HTTP GET requests is in fact only a URL prefix to which additional parameters are appended in order to construct a valid Operation request. A URL prefix is defined in accordance with IETF RFC 2396 as a string including, in order,

the scheme (“http” or “https”), Internet Protocol hostname or numeric address, optional port number, path, mandatory question mark “?”, and optional string comprising one or more server-specific parameters ending in an ampersand “&”. The prefix defines the network address to which request messages are to be sent for a particular operation on a particular server. Each operation may have a different prefix. Each prefix is entirely at the discretion of the service provider.

This document defines how to construct a query part that is appended to the URL prefix in order to form a complete request message. Every SSA operation has several mandatory or optional request parameters. Each parameter has a defined name. Each parameter may have one or more legal values, which are either defined by this document or are selected by the client based on service metadata. To formulate the query part of the URL, a client shall append the mandatory request parameters, and any desired optional parameters, as name/value pairs in the form “name=value&” (parameter name, equals sign, parameter value, ampersand). The “&” is a separator between name/value pairs, and is therefore optional after the last pair in the request string.

When the HTTP GET method is used, the client-constructed query part is appended to the URL prefix defined by the server, and the resulting complete URL is invoked as defined by HTTP (IETF RFC 2616).

Table 2 summarizes the components of an operation request URL when HTTP GET is used.

Table 2 — Structure of SSA request using HTTP GET

URL component	Description
http://host[:port]/path[?{name=[value]&}]	Base-URL (prefix) of service operation. [] denotes 0 or 1 occurrence of an optional part; {} denotes 0 or more occurrences.
name=value&	One or more standard request parameter name/value pairs as defined for each operation by this document.

5.3.4 HTTP POST

SSA does not currently support the “POST” method of the HTTP protocol (IETF RFC 2616), but may do so in the future. POST could be used, for example, to permit large range-lists to be specified.

5.4 General HTTP response rules

Upon receiving a valid request, the server shall send a response corresponding exactly to the request as detailed in Clause 7 of this document, or send a service exception if unable to respond correctly. Only in the case of Version Negotiation (see 6.2.4) may the server offer a differing result. Upon receiving an invalid request, the server shall issue a service exception as described in 6.11.

A server may send an HTTP Redirect message (using HTTP response codes as defined in IETF RFC 261 6) to an absolute URL that is different from the valid request URL that was sent by the client. HTTP Redirect causes the client to issue a new HTTP request for the new URL. Several redirects could in theory occur. Practically speaking, the redirect sequence ends when the server responds with a SSA response. The final response shall be a SSA response that corresponds exactly to the original request (or a service exception).

Response objects shall be accompanied by the appropriate Multipurpose Internet Mail Extensions (MIME) type (IETF RFC 2045) for that object. A list of MIME types in common use on the internet is maintained by the Internet Assigned Numbers Authority (IANA) [2]. Allowable types for operation responses and service exceptions are discussed below. The basic structure of a MIME type is a string of the form “type/subtype”. MIME allows additional parameters in a string of the form “type/subtype; param1=value1; param2=value2”. A server may include parameterized MIME types in its list of supported output formats. In addition to any parameterized variants, the server should offer the basic unparameterized version of the format.

Response objects should be accompanied by other HTTP entity headers as appropriate and to the extent possible. In particular, the Expires and Last-Modified headers provide important information for caching; Content-Length may be used by clients to know when data transmission is complete and to efficiently allocate space for results, and Content-Encoding or Content-Transfer-Encoding may be necessary for proper interpretation of the results.

5.5 Numeric and Boolean values

Integer numbers shall be represented in a manner consistent with the specification for integers in XML Schema Datatypes ([8], section 3.3.13). This document shall explicitly indicate where an integer value is mandatory. Real numbers shall be represented in a manner consistent with the specification for double-precision numbers in XML Schema Datatypes ([8], section 3.2.5). This representation allows for integer, decimal and exponential notations. A real value is allowed in all numeric fields defined by this document unless the value is explicitly restricted to integer.

Sexagesimal formatting is not permitted other than in ISO 8601 formatted time strings unless otherwise specified in this document. In particular, astronomical coordinates should be rendered as real numbers as specified above.

Positive, negative and zero values are allowed unless explicitly restricted.

Boolean values shall be represented in a manner consistent with the specification for Boolean in XML Schema Datatypes ([8], section 3.2.2). The values “0” and “false” are equivalent. The values “1” and “true” are equivalent. Absence of an optional value is equivalent to logical false. This document shall explicitly indicate where a Boolean value is mandatory.

5.6 Output formats

The response to a SSA request is always a computer file that is transferred over the Internet from the server to the client. The file may contain text, or the file may be a graphics or FITS-formatted file. As stated in 6.4, the type of the returned file shall be indicated by a MIME type string.

Text output formats are usually formatted as Extensible Markup Language (XML; MIME type text/xml). Text formats are used to convey service metadata, descriptions of error conditions, or responses to data queries. In particular, the response to a data query is always returned as an XML file in VOTable format.

5.7 Request parameter rules

5.7.1 Parameter ordering and case

Parameter names shall not be case sensitive, but parameter values shall be. In this document, parameter names are typically shown in uppercase for typographical clarity, not as a requirement.

Parameters in a request may be specified in any order.

When request parameters are duplicated with conflicting values, the response from the server may be undefined. This document does not mandate which of the duplicated values sent by the client are to be used by the server.

A SSA service shall be prepared to encounter additional request parameters that are not part of this document without reporting an error. In terms of producing results per this document, a SSA service shall not require such parameters, but may define additional service-defined parameters.

5.7.2 Range-List Parameters

Parameters consisting of lists (for example, POS, BAND and TIME in SSA) shall use the comma (",") as the separator between items in the list. Additional white space shall not be used to delimit list items. If a list item value includes a space or comma, it shall be escaped using the URL encoding rules (see 6.3.2 and IETF RFC 2396).

In some lists, individual entries may be empty, and shall be represented by the empty string (""). Thus, two successive commas indicate an empty item, as does a leading comma or a trailing comma. An empty list ("") shall be interpreted either as a list containing no items or as a list containing a single empty item, depending on context.

Some parameters (for example BAND and TIME) may allow a parameter value to be specified as a range. Such range-valued parameters shall use the forward slash ("/") character as the separator between elements of the range specification. For example, "5E-7/8E-7" would specify a range consisting of all values from 5E-7 to 8E-7, inclusive. If

a third field is specified it is a step size for traversing the indicated range. If a parameter permits a step size the semantics of the step size are defined by the specific parameter.

An open range may be specified by omitting either value. If the first value is omitted the range is open toward lower values. If the second value is omitted the range is open toward higher values. Omitting both values indicates an infinite range which accepts all values. For example, “/5” is an open range which accepts all values less than or equal to 5.

A list may be qualified by appending the character “;” followed by the qualifier string. For example “1E-6/3E-7;source” could specify a spectral bandpass in the rest frame of the source.

List and range syntax may be combined, e.g., to indicate a list of scalar or range-valued parameter values. Range lists may be ordered or unordered, and may contain either numeric or string data. An “ordered” list is one which requires values to be given in sequential order, and is sorted by the service before being used. Items in an “unordered” list will be processed in the order specified by the client. Ranges are limited to numeric values or ISO dates.

5.8 Common request parameters

5.8.1 VERSION

The VERSION parameter specifies the protocol version number. The format of the version number, and version negotiation, are described in 6.2.

5.8.2 REQUEST

The REQUEST parameter indicates which service operation is being invoked. The value shall be the name of one of the operations offered by the server.

5.8.3 Extended capabilities and operations

The SSA service allows for optional extended capabilities and operations. Extensions may be defined within an information community when needed for additional functionality or specialization. A generic client shall not be required or expected to make use of such extensions. Extended capabilities or operations shall be defined by the service metadata. Extended capabilities provide additional metadata about the service, and may or may not enable optional new parameters to be included in operation requests. Extended operations allow additional operations to be defined.

A server shall produce a valid response to the operations defined in this document, even if parameters used by extended capabilities are missing or malformed (*i.e.* the server shall supply a default value for any extended capabilities it defines), or if parameters are supplied that are not known to the server.

Service providers shall choose extension names with care to avoid conflicting with standard metadata fields, parameters and operations.

5.9 Service result

The return value of a valid Service request shall correspond to the output type specified for the operation, or requested in the FORMAT parameter in the case of an operation which can return data in a choice of output formats. In an HTTP environment, the Content-type header of the response shall be exactly the MIME type associated with the valid request.

5.10 Error Response and Other Unsuccessful Results

Upon receiving a request that is invalid according to this document, the server shall issue a service exception report. The service exception report is meant to describe to the client application or its human user the reason(s) that the request is invalid. The allowed service exception formats are defined below.

If a service operation throws an error response and exits, the default action of the service should be to return a VOTable noting that an error has occurred, and describing the error. An INFO element within the "results" RESOURCE element of the VOTable is used to indicate success or failure of the operation. As described in the previous section, the INFO element must have name="QUERY_STATUS"; if the operation is successful (regardless of whether any data is returned) the value attribute is set to "OK". The remainder of this section defines other possible values to indicate that the query was unsuccessful in some way. When the query is unsuccessful, the contents of INFO element (i.e. its PCDATA child node) **should** contain an error message suitable for display.

When the query is unsuccessful (in any of senses described below), the resulting VOTable is not required to contain any other elements as specified for a successful operation; however, it is not an error to do so. For example, additional INFO elements may be returned to echo back the input parameters of the operation which failed, as in the following example.

Example:

```
<VOTABLE ... version="1.1">
  <RESOURCE type="results">
    <INFO name="QUERY_STATUS" value="ERROR">unrecognized operation</INFO>
    <INFO name="REQUEST" value="findData"/>
    <INFO name="baseUrl" value="http://webtest.aoc.nrao.edu/ivoa-dal"/>
    <INFO name="serviceVersion" value="1.0"/>
    <INFO name="serviceName" value="ssap"/>
    <INFO name="ServiceEngine" value="ssap: SSAP 1.0 DALServer version 0.1"/>
  </RESOURCE>
</VOTABLE>
```

The other allowed values for value attribute besides "OK" are as specified below.

5.10.1 Service Error

The server failed to process the operation. Typical reasons include:

- The input query contained a syntax error.
- The way the query was posed was invalid for some reason, e.g., due to an invalid query region specification.
- A constraint parameter value was given an illegal value; e.g. DEC=91.
- The server trapped an internal error (e.g., failed to connect to its database) preventing further processing.

In this case, the inclusion of a descriptive error message **should** be returned.

5.10.2 Overflow

The operation produced results that exceeded the limits of the service in some way. For instance, a data query matched too many candidate datasets. In this case, the service **should** include an error message indicating the nature of the overflow condition (see also section 4.4).

Example:

```
<INFO name="QUERY_STATUS" value="ERROR">DEC out of range: DEC=91</INFO>
<INFO name="QUERY_STATUS" value="OVERFLOW">Number of matching spectra
exceeds default limit of 500</INFO>
```

If overflow occurs the query may be repeated, requesting a higher value than the default for the maximum number of output records, up to the hard limit defined by the service capabilities. Alternatively the query parameters may be adjusted to more carefully constrain the query.

5.10.3 Other Errors

Although the intention is that service should catch all errors and return a uniform error response in the prescribed VOTable format, informing the client of the nature of the error which occurred in service-specific terms, this is not always possible. More fundamental errors may result in a HTTP level error. The client should be prepared to handle either form of error. Which is returned in a given case, may depend upon the operation performed, the nature of the error which occurred, and the details of how a given service is implemented.

Appendix A: Sample Query Response

The output below illustrates the query response from a working SSA service. In the interests of brevity some of the FIELD and GROUP definitions have been omitted, and only data for a single table row is included. Minimal effort has been made to “pretty print” the table.

```

<VOTABLE xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="xmlns:http://www.ivoa.net/xml/VOTable/VOTable-1.1.xsd"
  xmlns:ssa="http://www.ivoa.net/xml/DalSsap/v1.0" version="1.1">

<RESOURCE type="Results">
<DESCRIPTION>DALServer proxy service for JHU spectrum services</DESCRIPTION>
<INFO name="QUERY_STATUS" value="OK"/>
<INFO name="REQUEST" value="queryData"/>
<INFO name="POS" value="180.0,1.0"/>
<INFO name="SIZE" value="0.2"/>
<INFO name="FORMAT" value="all"/>
<INFO name="Collection" value="ivo://jhu/sdss/dr5"/>
<INFO name="ServiceEngine" value="JhuProxySsap: SSAP 1.0 DALServer version 0.1"/>
<INFO name="TableRows" value="42"/>
<TABLE>
<FIELD ID="Score" name="Score" datatype="float" utype="ssa:Query.Score">
  <DESCRIPTION>Degree of match to query parameters</DESCRIPTION>
</FIELD>
<FIELD ID="AssocID" name="AssocID" datatype="char" utype="ssa:Association.ID"
  arraysize="*">
  <DESCRIPTION>Association identifier</DESCRIPTION>
</FIELD>
<FIELD ID="AcRef" name="AcRef" datatype="char" ucd="meta.ref.url"
  utype="ssa:Access.Reference" arraysize="*">
  <DESCRIPTION>URL used to access dataset</DESCRIPTION>
</FIELD>
<FIELD ID="Format" name="Format" datatype="char" utype="ssa:Access.Format"
  arraysize="*">
  <DESCRIPTION>Content or MIME type of dataset</DESCRIPTION>
</FIELD>
<FIELD ID="DataModel" name="DataModel" datatype="char" utype="ssa:Dataset.DataModel"
  arraysize="*">
  <DESCRIPTION>Datamodel name and version</DESCRIPTION>
</FIELD>
<FIELD ID="DataLength" name="DataLength" datatype="long" utype="ssa:Dataset.Length">
  <DESCRIPTION>Number of points</DESCRIPTION>
</FIELD>
<FIELD ID="Title" name="Title" datatype="char" ucd="meta.title;meta.dataset"
  utype="ssa:DataID.Title" arraysize="*">
<DESCRIPTION>Dataset Title</DESCRIPTION>
</FIELD>
<FIELD ID="Creator" name="Creator" datatype="char" utype="ssa:DataID.Creator"
  arraysize="*">
<DESCRIPTION>Dataset creator</DESCRIPTION>
</FIELD>
<FIELD ID="Collection" name="Collection" datatype="char" utype="ssa:DataID.Collection"
  arraysize="*">
  <DESCRIPTION>Data collection to which dataset belongs</DESCRIPTION>
</FIELD>
<FIELD ID="CreatorDID" name="CreatorDID" datatype="char" ucd="meta.id"
  utype="ssa:DataID.CreatorDID" arraysize="*">
  <DESCRIPTION>Creator's ID for the dataset</DESCRIPTION>
</FIELD>
<FIELD ID="CreatorDate" name="CreatorDate" datatype="char" ucd="time;meta.dataset"
  utype="ssa:DataID.Date" arraysize="*">
  <DESCRIPTION>Data processing/creation date</DESCRIPTION>
</FIELD>
  [more FIELDS omitted]

<GROUP ID="Query" name="Query" utype="ssa:Query">
  <DESCRIPTION>Query Metadata</DESCRIPTION>
  <FIELDref ref="Score"/>

```

```

</GROUP>
<GROUP ID="Association" name="Association" utype="ssa:Association">
  <DESCRIPTION>Association Metadata</DESCRIPTION>
  <FIELDref ref="AssocID"/>
  <PARAM ID="AssocType" datatype="char" name="AssocType"
    utype="ssa:Association.Type" value="MultiFormat" arraysize="*">
    <DESCRIPTION>Type of association</DESCRIPTION>
  </PARAM>
  <PARAM ID="AssocKey" datatype="char" name="AssocKey" utype="ssa:Association.Key"
    value="@Format" arraysize="*">
    <DESCRIPTION>Key used to distinguish association elements</DESCRIPTION>
  </PARAM>
</GROUP>
<GROUP ID="Access" name="Access" utype="ssa:Access">
  <DESCRIPTION>Access Metadata</DESCRIPTION>
  <FIELDref ref="AcRef"/>
  <FIELDref ref="Format"/>
  <PARAM ID="DatasetSize" unit="bytes" datatype="long" name="DatasetSize"
    utype="ssa:Access.Size" value="800000">
    <DESCRIPTION>Estimated dataset size</DESCRIPTION>
  </PARAM>
</GROUP>
<GROUP ID="Dataset" name="Dataset" utype="ssa:Dataset">
  <DESCRIPTION>General Dataset Metadata</DESCRIPTION>
  <FIELDref ref="DataModel"/>
  <FIELDref ref="DataLength"/>
  <PARAM ID="DatasetType" datatype="char" name="DatasetType"
    utype="ssa:Dataset.Type" value="Spectrum" arraysize="*">
    <DESCRIPTION>Dataset or segment type</DESCRIPTION>
  </PARAM>
</GROUP>
<GROUP ID="DataID" name="DataID" utype="ssa:DataID">
  <DESCRIPTION>Dataset Identification Metadata</DESCRIPTION>
  <FIELDref ref="Title"/>
  <FIELDref ref="Creator"/>
  <FIELDref ref="Collection"/>
  <FIELDref ref="CreatorDID"/>
  <FIELDref ref="CreatorDate"/>
  <FIELDref ref="CreatorVersion"/>
  <FIELDref ref="Instrument"/>
  <PARAM ID="DataSource" datatype="char" name="DataSource"
    utype="ssa:DataID.DataSource" value="survey" arraysize="*">
    <DESCRIPTION>Original source of the data</DESCRIPTION>
  </PARAM>
  <PARAM ID="CreationType" datatype="char" name="CreationType"
    utype="ssa:DataID.CreationType" value="Archival" arraysize="*">
    <DESCRIPTION>Dataset creation type</DESCRIPTION>
  </PARAM>
</GROUP>
  [More GROUPs omitted]

```

<DATA>

<TABLEDATA>

<TR>

<TD>1.0</TD>

<TD>MultiFormat.12</TD>

<TD><http://webtest.aoc.nrao.edu/ivoa-dal/JhuProxySsap?>

REQUEST=getData&FORMAT=csv&

PubDID=ivo%3A%2F%2Fjhu%2Fsdss%2Fdr5%2380442261170552832

</TD>

<TD>text/csv</TD>

<TD>Spectrum 1.0</TD>

<TD>4000</TD>

<TD>SDSS J115923.80+000000.00 Galaxy 0285-51663-01</TD>
<TD>sdss</TD>
<TD>ivo://sdss/dr5/spec</TD>
<TD>ivo://sdss/dr5/spec#80442261170552832</TD>
<TD>2000-04-29T03:22:00.7900000-04:00</TD>
<TD>3.13.1branch.1</TD>
<TD>SDSS 2.5-M SPEC2 v4_5</TD>
<TD>ivo://jhu/sdss/dr5#80442261170552832</TD>
<TD>SDSS J115923.80+000000.00</TD>
<TD>Galaxy</TD>
<TD>0.451652</TD>
<TD>0</TD>
<TD>FK5</TD>
<TD>2000</TD>
<TD>TAI</TD>
<TD>179.849160 .984768</TD>
<TD>0.000833333333333333333339</TD>
<TD>em.wl</TD>
<TD>6518.40990663334</TD>
<TD>5389.1347401044286</TD>
<TD>3823.8425365811263</TD>
<TD>9212.9772766855549</TD>
<TD>0</TD>
<TD>0</TD>
<TD>Absolute</TD>
<TD>0</TD>
<TD>51663.30695358796</TD>
<TD>3600</TD>
<TD>51663.2713056713</TD>
<TD>51663.319500115744</TD>
<TD>phot.fluDens;em.wl</TD>
<TD>0</TD>
<TD>Absolute</TD>

</TR>

[More table rows omitted]

</TABLEDATA>

</DATA>

</TABLE>

</RESOURCE>

</VOTABLE>

References

[Allen et al. 2003]

Allen, M., Boch, T., 2003, Plotting SEDs in the AVO prototype: Catalog Requirements, Conventions and Conversions, http://www.euro-vo.org/internal/Avo/WorkPackageTwoTwo/SEDs_in_AVO.pdf

[Bonnarel et al. 2004]

Bonnarel, F. et al., 2004, Proposal for an evolution of the SIA protocol V1.00, <http://www.ivoa.net/Documents/latest/SIAPEvol.html>

[Bunclark/Rots 1996]

Bunclark, P., Rots, A. 1996, Precise re-definition of DATE-OBS Keyword encompassing the millennium,

<http://www.cv.nrao.edu/fits/documents/standards/year2000.txt>

[Dolensky/Tody 2004]

Dolensky, M., Tody, D., 2004, Survey among Spectral Data Providers and Consumers, ASP Conf. Ser., 314, 338

[Dolensky 2006]

Dolensky, M, Ranking Query Result Sets,

<http://www.ivoa.net/Documents/latest/Ranking.html>

[Hanisch et al. 2005]

Hanisch, R. 2005, Resource Metadata for the Virtual Observatory V1.1,

<http://www.ivoa.net/Documents/REC/ResMetadata/RM-20051115.html>

[McDowell/Tody et al. 2006]

McDowell, J., Tody, D., et al. 2006, Work in progress: IVOA Spectral Data Model V0.98,

<http://hea-www.harvard.edu/~jcm/vo/docs/spec0.98c.html>

[Osuna/Salgado 2004]

Osuna, P., Salgado, J. 2004, Simple Spectral Access for ISO data V1.0,

<http://www.ivoa.net/Documents/latest/SADimEq.html>

[RFC 1738]

Berners-Lee et al. 1994, Uniform Resource Locators (URL), IETF RFC 1738,

<http://www.ietf.org/rfc/rfc1738.txt>

[RFC 2119]

Bradner et al. 1997, Key words for use in RFCs to Indicate Requirement Levels, IETF

RFC 2119, <http://www.ietf.org/rfc/rfc2119.txt>

[Rots 2005]

Rots, A., Space-Time Coordinate Metadata for the Virtual Observatory V1.10,

<http://www.ivoa.net/Documents/latest/STC.html>

[Tody/Plante 2004]

Tody, D., Plante, R. 2004, Simple Image Access Specification V1.0,

<http://www.ivoa.net/Documents/WD/SIA/sia-20040524.html>

[UCD 2004]

Derriere, S., Preite Martinez, A., Williams, R. 2004, Work in progress: UCD (Unified Content Descriptor) - moving to UCD1+,

<http://www.ivoa.net/Documents/PR/UCD/UCD-20041026.html>

[Valdes 2003]

Valdes, F. 2003, A Virtual Observatory Data Model,

<http://iraf.noao.edu/projects/vo/dal/datamodel.html>

[VOTable 2004]

Ochsenbein, F. et al., 2004, VOTable Format Definition V1.1,

<http://www.ivoa.net/Documents/REC/VOTable/VOTable-20040811.html>