



*International
Virtual
Observatory
Alliance*

IVOA Server-side Operations for Data Access

Version 1.0

IVOA Working Draft 2015-12-21

Working group

DAL

This version

<http://www.ivoa.net/documents/SODA/20151221>

Latest version

<http://www.ivoa.net/documents/SODA>

Previous versions

WD-AccessData-1.0-20151021

WD-AccessData-1.0-20140730

WD-AccessData-1.0-20140312

Author(s)

François Bonnarel, Markus Demleitner, Patrick Dowler, Douglas

Tody

Editor(s)

François Bonnarel

Abstract

This document describes the SODA web service capability. SODA is a low-level data access capability or server side data processing that can act upon the data files, performing various kinds of operations: filtering/subsection, transformations, pixel operations, and applying functions to the data.

Status of This Document

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”.

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

Contents

1	Introduction	3
1.1	The Role in the IVOA Architecture	3
1.2	Motivating Use Cases	4
1.2.1	Retrieve Subsection of a Datacube	4
1.2.2	Retrieve subsection of a 2D Image	4
1.2.3	Retrieve subsection of a Spectrum	4
1.3	Provide the data in different formats	4
1.3.1	Flatten a Datacube into a 2D Image	4
1.3.2	Flatten a Datacube into a 1D Spectrum	4
1.3.3	Rebin Data by a Fixed Factor	5
1.3.4	Reproject Data onto a Specified Grid	5
1.3.5	Compute Aggregate Functions over the Data	5
1.3.6	Apply Standard Function to Data Values	5
1.3.7	Apply Arbitrary User-Specified Function to Data Values	5
1.3.8	Run Arbitrary User-Supplied Code on the Data	5
2	Resources	5
2.1	{sync} resource	6
2.2	{async} resource	6
2.3	Examples: DALI-examples	7
2.4	Availability: VOSI-availability	7
2.5	Capabilities: VOSI-capabilities	7
2.6	Parameter Description and Three-Factor Semantics	8
2.6.1	Three-factor Semantics	8
2.6.2	Discovery of Supported Parameter. Implementation strategies	9
2.6.3	SODA Service Descriptor	10
2.6.4	Client Handling of Discovered Parameters	10

3	Parameters for {sync} and {async}	11
3.1	Common Parameters	11
3.1.1	ID	11
3.2	Filtering Parameters	11
3.2.1	POS	12
3.2.2	BAND	13
3.2.3	TIME	14
3.2.4	POL	14
4	{sync} Responses	15
4.1	Successful Requests	15
4.2	Errors	15
5	{async} Responses	16
A	Changes from Previous Versions	16
A.1	WD-SODA-1.0-20151120	16
A.2	WD-AccessData-1.0-20151021	16
A.3	WD-AccessData-1.0-20140730	16
A.4	WD-AccessData-1.0-20140312	16

Acknowledgments

The authors would like to thank all the participants in DAL- WG discussions for their ideas, critical reviews, and contributions to this document.

1 Introduction

The SODA web service interface defines a RESTful web service for performing server-side operations and processing on data before transfer.

1.1 The Role in the IVOA Architecture

TODO: new diagram from TCG

SODA services conform to the Data Access layer Interface (Dowler, Demleitner, Taylor and Tody, 2013) specification, including the Virtual Observatory Support Interfaces (Grid and Web Services Working Group, 2011) resources.

1.2 Motivating Use Cases

Below are some of the more common use cases that have motivated the development of the SODA specification. While this is not complete, it helps to understand the problem area covered by this specification.

1.2.1 Retrieve Subsection of a Datacube

Cutout a subsection using coordinate axis values. The input to the cutout operation will include one or more of the following:

- a region on the sky
- an energy value or range
- a time value or range
- one or more polarization states

The region on the sky should be something simple: a circle, a range of coordinate values, or maybe a polygon.

1.2.2 Retrieve subsection of a 2D Image

This is a special case of 1.2.1, where the cutout is only in the spatial axes.

1.2.3 Retrieve subsection of a Spectrum

This is a special case of 1.2.1, where the cutout is only in the spectral axis.

1.3 Provide the data in different formats

Examples are images in PNG, or JPEG instead of FITS and spectra in csv, FITS or VOTable.

1.3.1 Flatten a Datacube into a 2D Image

This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.2 Flatten a Datacube into a 1D Spectrum

This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.3 Rebin Data by a Fixed Factor

This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.4 Reproject Data onto a Specified Grid

This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.5 Compute Aggregate Functions over the Data

This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.6 Apply Standard Function to Data Values

It could be “denoising” with standard methods or “on the fly” recalibration. This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.7 Apply Arbitrary User-Specified Function to Data Values

This use case will be developed and supported in the SODA-1.1 (or later) specification.

1.3.8 Run Arbitrary User-Supplied Code on the Data

This use case will be developed and supported in the SODA-1.1 (or later) specification.

2 Resources

SODA services are implemented as HTTP REST (Richardson and Ruby, 2007) web services with a {sync} resource that conforms to the DALI- sync resource description.

A stand-alone SODA service may have one or both of the {sync} and {async} resources. For either type, it could have multiple resources (e.g. to support alternate authentication schemes). The SODA service may also include other custom or supporting resources.

Either the {sync} or {async} SODA capability may be included as part of other web services. For example, a single web service could contain the SIA-2.0 {query} capability, the DataLink-1.0 {links} capability, and the SODA {sync} capability. Such a service must also have the VOSI-availability and

resource type	resource name	required
{sync}	service specific	
{async}	service specific	
DALI-examples	/examples	no
VOSI-availability	/availability	yes
VOSI-capabilities	/capabilities	yes

Table 1: Endpoints for AccessData services

VOSI-capabilities resources to report on and describe all the implemented capabilities.

2.1 {sync} resource

The {sync} resource is a synchronous web service resource that conforms to the DALI-sync description. Implementors are free to name this resource however they like, except that the name must consist of one URI segment only (i.e., contain no slash). This is to allow clients, given the access URL, can reliably find out the URL of the capabilities endpoint. Clients, in turn, can find the resource path using the VOSI-capabilities resource, but will in general be provided the access URLs through a previous data discovery query or through direct user input.

The {sync} resource performs the data access as specified by the input parameters and returns the data directly in the output stream. Synchronous data access is suitable when the operations can be quickly performed and the data stream can be setup and written to (by the service) in a short period of time (e.g. before any timeouts).

2.2 {async} resource

The {async} resource is an asynchronous web service resource that conforms to the DALI-async description. The considerations on naming the resource given in sect. 2.1 apply for it.

The {async} resource performs the data access as specified by the input parameters and either (i) stores the results for later transfer or (ii) pushes the results to a specified destination (e.g. to a VOSpace location). Asynchronous data access usually introduces resource constraints on the service (which may be limited) and usually imposes a higher latency before any results can be seen because the location of results does not have to be valid until the data access job is complete. Asynchronous data access is intended for (but not limited to) use when the operations take considerable time and results must be staged (e.g. some multi-pass algorithms or operations that result in multiple outputs).

2.3 Examples: DALI-examples

SODA services should provide a DALI-examples resource with one example invocation that shows the variety operations the service can perform. Example operations using the {sync} resource and that output a small data stream are preferred, as the examples may be used by automatic validators doing relatively frequent (of order daily) queries.

Parameters to be passed to the service must be given using the DALI `generic-parameter` term.

2.4 Availability: VOSI-availability

A SODA web service must have a VOSI-availability resource (Grid and Web Services Working Group, 2011) as described in DALI (Dowler, Demleitner, Taylor and Tody, 2013).

2.5 Capabilities: VOSI-capabilities

A web service that includes SODA capabilities must have a VOSI-capabilities resource (Grid and Web Services Working Group, 2011) as described in DALI (Dowler, Demleitner, Taylor and Tody, 2013). The standardID for the {sync} resource is

```
ivo://ivoa.net/std/SODA#sync-1.0.
```

The standardID for the {async} resource is

```
ivo://ivoa.net/std/SODA#async-1.0.
```

All DAL services must implement the `/capabilities` resource. The following capabilities document shows the minimal metadata for a stand-alone SODA service and does not require a registry extension schema:

```
<?xml version="1.0"?>
<capabilities
  xmlns:vosi="http://www.ivoa.net/xml/VOSICapabilities/v1.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:vod="http://www.ivoa.net/xml/VODDataService/v1.1">
  <capability standardID="ivo://ivoa.net/std/VOSI#capabilities">
    <interface xsi:type="vod:ParamHTTP" version="1.0">
      <accessURL use="full">http://example.com/data/capabilities</accessURL>
    </interface>
  </capability>
  <capability standardID="ivo://ivoa.net/std/VOSI#availability">
    <interface xsi:type="vod:ParamHTTP" version="1.0">
      <accessURL use="full">
```

Note

As SODA builds upon several concepts of DataLink (Dowler, Bonnarel, Michel, Donaldson and Languignon, 2013), that document should be read before trying to understand the following material.

```
        http://example.com/data/availability
    </accessURL>
</interface>
</capability>
<capability standardid="ivo://ivoa.net/std/SODA#sync-1.0">
  <interface xsi:type="vod:ParamHTTP" role="std" version="1.0">
    <accessurl use="full">
      http://example.com/data/sync
    </accessurl>
  </interface>
  <!-- service details from extension schema could go here -->
</capability>
<capability standardid="ivo://ivoa.net/std/SODA#async-1.0">
  <interface xsi:type="vod:ParamHTTP" role="std" version="1.0">
    <accessurl use="full">
      http://example.com/data/async
    </accessurl>
  </interface>
  <!-- service details from extension schema could go here -->
</capability>
</capabilities>
```

Note that the {sync} and {async} resources do not have to be named as shown in the accessURL(s) above. Multiple capability elements for the {sync} and the {async} resources may be included; this is typically used if they differ in protocol (http vs. https) and/or authentication requirements.

2.6 Parameter Description and Three-Factor Semantics

2.6.1 Three-factor Semantics

Parameters in SODA are defined by triples of name, UCD, and unit (the “SODA triple”). Data services are free to support as many such parameters as is appropriate for their datasets, in addition to supporting standard parameters. With the three factors, it is unlikely that two operators will by accident use the same three factors for parameters of differing semantics. In this way, when a new standard parameter is adopted, no existing service should require changes to remain compliant.

With standard parameters as defined in this document, clients can rely on certain semantics and exploit that knowledge in the provision of special UIs or APIs. Standard parameters defined so far are given in table 2. Instructions for how to propose additional standard parameters are given on the landing page of the IVOA DAL working group¹.

Name	UCD	Unit	Semantics
ID	meta.ref.url;meta.curation		cf. sect. 3.1.1
POS	pos		cf. sect. 3.2.1
BAND	em	m	cf. sect. 3.2.2
TIME	time	d	cf. sect. 3.2.3
POL	pol		cf. sect. 3.2.4

Table 2: SODA triples for the standard parameters defined here.

Both standard and non-standard parameters should follow DALI conventions if at all possible. Roughly, float-valued parameters should be mapped to interval-valued parameters (i.e., do not split up minimum and maximum into separate parameters). Depending on their semantics, integer parameters should either be intervals or enumerated parameters (which typically can be repeated). String-valued parameters should always be enumerated.

2.6.2 Discovery of Supported Parameter. Implementation strategies

There are two principal ways a client will be led to a SODA service: From a service descriptor that is part of a DAL discovery response, or by direct access mediated or not via a DataLink service. In the latter case an empty query to the service will answer by giving the service descriptor. . In both cases, clients must obtain the set of supported parameters from the applicable service descriptor.

In contrast to previous IVOA DAL protocols, SODA is not a data discovery protocol but instead operates on concrete datasets. In a typical case, some combinations of parameters (e.g., a positional or spectral cutout) may yield no output at all, as the coverage of an individual dataset is very limited. To provide meaningful user interfaces, clients would therefore need detailed information on the service parameters, in particular as regards their domains (i.e., set of values that are likely to yield non-empty results). This, in regard, requires efficient and rich service APIs, not unlikely involving service-specific parameters.

But the double ways towards SODA service execution makes the providing of this information difficult to describe and is not done in this version.

¹At the time of writing, this is <http://wiki.ivoa.net/twiki/bin/view/IVOA/IvoaResReg>.

As a consequence of this, when the SODA service is driven from a data discovery phase, parameters domains can be inferred by the end-user from the discovery response. In the other case, it may happen that these parameters have to be furnished blindly by the end-user.

2.6.3 SODA Service Descriptor

The DataLink (Dowler, Bonnarel, Michel, Donaldson and Languignon, 2013) specification describes a mechanism for describing a service within a VOTable resource and recommends that services can describe themselves with a special resource with name="this". SODA responses for empty sync queries should include a descriptor describing both standard and custom query parameters (if applicable). The descriptor for a service with standard parameters (see sect. 3) would be:

```
<RESOURCE type="meta" utype="ad hoc:service" name="this">
  <PARAM name="standardID" datatype="char" arraysize="*"
    value="ivo://ivoa.net/std/SODA#sync-1.0" />
  <PARAM name="accessURL" datatype="char" arraysize="*"
    value="http://example.com/SODA/sync" />
  <GROUP name="inputParams">
    <PARAM name="ID" ucd="meta.id" datatype="char"
      arraysize="*" xtype="ivoident" />
    <PARAM name="POS" ucd="pos" unit="deg" datatype="char"
      arraysize="*" xtype="circle" />
    <PARAM name="POS" ucd="pos" unit="deg" datatype="char"
      arraysize="*" xtype="range" />
    <PARAM name="POS" ucd="pos" unit="deg" datatype="char"
      arraysize="*" xtype="polygon" />
    <PARAM name="BAND" ucd="em" unit="m" datatype="double"
      arraysize="*" xtype="interval" />
    <PARAM name="TIME" ucd="time" unit="d" datatype="double"
      arraysize="*" xtype="interval" />
    <PARAM name="POL" ucd="pol" datatype="char" arraysize="*"
      xtype="Stokes" />
  </GROUP>
</RESOURCE>
```

This VOTable resource should be output for empty sync queries; Thus all inputs and outputs would be fully described.

2.6.4 Client Handling of Discovered Parameters

To keep SODA clients useful even when advanced data products are being accessed, the following procedure should be followed by clients:

1. Obtain the parameter triples domains as described in sect. 2.6.2.

2. Identify standard parameters understood by the client and provide suitable input widgets (e.g., a rubberband on top of a sky rendering, or a `spatial_cutout` method possibly allowing library-specific coverage objects) for them.
3. For the remaining parameters, provide generic UI or API elements (e.g., sliders or text boxes with domains annotated for intervals, pop-up menus for enumerated values).

3 Parameters for `{sync}` and `{async}`

The `{sync}` and `{async}` resources accept the same set of parameters.

3.1 Common Parameters

3.1.1 ID

The ID parameter is used to specify the dataset or file to be accessed. The values for the ID parameter are generally discovered from data discovery or DataLink requests. The values must be treated as opaque identifiers that are used as-is. The DataLink specification (Dowler, Bonnarel, Michel, Donaldson and Languignon, 2013) describes mechanisms for conveying opaque parameters and values in service descriptor resources that can be used by clients to set the ID parameter.

The ID parameter is single-valued in `{sync}` requests, so `{sync}` soda requests access a single dataset or file. Multiple ID parameters may be submitted in `{async}` requests on order to bundle access to multiple datasets or files in a single job.

The ID ucd is “meta.id”, and its unit is blank. In addition its xtype is “ivoident” and its datatype “char”, with arraysize “*”.

3.2 Filtering Parameters

Filtering parameters are used to extract subsets of larger datasets or data files. In general, filtering parameters are single-valued in `{sync}` requests and multi-valued in `{async}` requests (exceptions noted below). When multiple values of filtering parameters are used in an `{async}` job, each combination of values produces zero or one result. For example, if an `{async}` job included two POS and two BAND values, there could be as many as four results (or fewer if some combinations do not produce a result because the filter does not intersect the bounds of the data).

3.2.1 POS

The POS parameter defines the positional region(s) to be extracted from the data. The value is made up of a shape keyword followed by coordinate values. The allowed shapes are:

Shape	Coordinate values
CIRCLE	<longitude> <latitude> <radius>
RANGE	<longitude1> <longitude2> <latitude1> <latitude2>
POLYGON	<longitude1> <latitude1> ... (at least 3 pairs)

Table 3: POS Values in Spherical Coordinates

As in DALI, open intervals use -Inf or +Inf as one limit.

Examples for POS values:

- A circle at (12,34) with radius 0.5:

POS=CIRCLE 12 34 0.5

- A range of [12,14] in longitude and [34,36] in latitude:

POS=RANGE 12 14 34 36

- A polygon from (12,34) to (14,34) to (14,36) to (12,36) and (implicitly) back to (12,34):

POS=POLYGON 12 34 14 34 14 36 12 36

The inside is always assumed to be the smaller of the region to the left and the region to the right so only polygons smaller than half the sphere can be specified.

- A band around the equator:

POS=RANGE 0 360 -2 2

- The north pole:

POS=RANGE 0 360 89 +Inf

This syntax is in the same style as STC-S, but with no reference positions, coordinate systems, units, or geometric operators like union, intersection, not, etc.

All longitude and latitude values (plus the radius of the CIRCLE) are expressed in degrees in the ICRS. A future version of this specification may

allow the use of other reference systems (specifically the native system of the data).

The POS parameter is single-valued for {sync} requests and multi-valued for {async} jobs.

The unit of POS is “deg” and the ucd is “pos”. However the datatype of the POS parameter is “char”, and the xtype can take one of the three values “circle”, “range” and “polygon” as defined in DALI.

3.2.2 BAND

The BAND parameter defines the energy interval(s) to be extracted from the data. The value is an open or closed numeric interval of values in the native spectral axis coordinate system and units of the data. The intervals always include the bounding values. As in DALI, open intervals use -Inf or +Inf as one limit.

If there is one single value the interval is assumed to be infinitely small (a scalar value).

- The closed interval [500,550]:

BAND=500 550

- The open interval (-inf,300]:

BAND=-Inf 300

- The open interval [750,inf):

BAND=750 +Inf

- The scalar value 550, equivalent to [550,550]:

BAND=550

Extracting using a scalar value should normally extract a single pixel along the energy axis of the data; extracting using an interval should extract one or more pixels.

All energy values are expressed as barycentric wavelength in meters. A future version of this specification may allow the use of other reference systems (specifically the native system of ther data).

The BAND parameter is single-valued for {sync} requests and multi-valued for {async} jobs.

The ucd of the BAND parameter is “em”, the unit is “m”. Its datatype is double with an arraysize of 2, and the xtype is “interval” as defined in DALI.

3.2.3 TIME

The TIME parameter defines the time interval(s) to be extracted from the data. The value is an open or closed interval with either numeric values (interpreted as Modified Julian Dates). As in DALI, open intervals use -Inf or +Inf as one limit.

If there is one single value the numeric interval is assumed to be infinitely small (a scalar value).

- An open interval from the MJD 55100.0 and all later times:

TIME= 55100.0 +Inf

- A range of MJD values:

TIME=55123.456 55123.466

- An instant in time using Modified Julian Date:

TIME=55678.123456

Time values are always UTC. The TIME parameter is single-valued for {sync} requests and multi-valued for {async} jobs.

The ucd of the TIME parameter is “time” and the unit is “d”. The datatype is “double” with an arraysize of 2, and the xtype is, again, “interval” as defined in DALI

3.2.4 POL

The POL parameter defines the polarization state(s) (Stokes) to be extracted from the data.

- Extract the unpolarized intensity:

POL=I

- Extract the standard circular polarization:

POL=V

- The POL parameter is multi-valued; multiple values can be included in a single request and all will be extracted. Extract only the IQU components:

POL=I

POL=Q

POL=U

The POL is multi-valued for both {sync} and {async} requests. Unlike general filtering parameters, all values of POL are combined into a single filter; for example, if the request includes the three values above, the job would generate one result with some or all of these polarization states (per combination of ID and other filtering parameters).

The ucd of the POL PARAMETER is “pol” and the unit is blank. The datatype is “char” with arraysize “*”, and the xtype is “stokes”.

4 {sync} Responses

All responses from the {sync} resource follow the rules for DALI-sync resources, except that the {sync} response allows for error messages for individual input identifier values.

4.1 Successful Requests

Successfully executed requests should result in a response with HTTP status code 200 (OK) and a response in the format requested by the client or in the default format for the service.

If the values specified for cutout parameters do not include any pixels from the target dataset/file, the service must respond with HTTP status code 204 (No Content) and no response body.

The service should set the following HTTP headers to the correct values where possible.

Content-Type	mime-type of the response
Content-Encoding	encoding/compression of the response (if applicable)

Since the response is usually dynamically generated, the Content-Length and Last-Modified headers cannot usually be set.

4.2 Errors

The error handling specified for DALI-sync resources applies to service failure. Error documents should be text using the text/plain content-type and the text must begin with one of the following strings:

Error	General error (not covered below)
AuthenticationError	Not authenticated
AuthorizationError	Not authorized to access the resource
ServiceUnavailable	Transient error (could succeed with retry)
UsageError	Permanent error (retry pointless)

Table 4: error messages with their meaning

5 {async} Responses

The {async} resource conforms to the DALI-async resource description, which means the job is a UWS job with all the job control features available. All result files are to be listed as children of the UWS results resource. The service provider is free to name each result.

A Changes from Previous Versions

A.1 WD-SODA-1.0-20151120

Change the name of the protocol. Suppression of SELECT and COORD. xtype description are in DALI. Reference to this has been added.

A.2 WD-AccessData-1.0-20151021

Added general introduction on PARAMETER description to section 3. Modified SELECT and COORD sections in order to detach them from SimDal. Added Appendix on xtype description with BNF syntax.

A.3 WD-AccessData-1.0-20140730

- Removed REQUEST parameter since the DAL-WG decision to not include it when there is only one value.
- Clarified that ID and filtering parameters are single valued for {sync} and multi-valued for {async}, with POL being multi-valued but still being treated as a single filter.

A.4 WD-AccessData-1.0-20140312

This is the initial document version.

References

Dowler, P., Bonnarel, F., Michel, L., Donaldson, T. and Languignon, D. (2013), 'Datalink', IVOA Working Draft.

URL: <http://www.ivoa.net/documents/DataLink/>

Dowler, P., Demleitner, M., Taylor, M. and Tody, D. (2013), 'Data access layer interface, version 1.0', IVOA Recommendation.

URL: <http://www.ivoa.net/documents/DALI>

Grid and Web Services Working Group (2011), 'IVOA support interfaces version 1.0'.

URL: <http://www.ivoa.net/documents/VOSI/index.html>

Richardson, L. and Ruby, S. (2007), *RESTful Web Services*, O'Reilly.