



International

Virtual

Observatory

Alliance

IVOA Data Access Layer Interface Version 1.0

IVOA Internal Working Draft 2012-10-14

Interest/Working Group:

<http://www.ivoa.net/cgi-bin/twiki/bin/view/IVOA/IvoaDAL>

This version:

WD-DALI-1.0-20121014

Latest version:

Not yet issued

Previous version(s):

WD-DALI-1.0-20120724

WD-DALI-1.0-10120202

Editors:

Patrick Dowler

Authors:

Patrick Dowler, Markus Demleitner, Mark Taylor, Doug Tody

Data Access Layer Interface

Abstract

This document describes the Data Access Layer Interface (DALI). DALI defines the base web service interface common to all Data Access Layer (DAL) services. This standard defines the behaviour of common resources, the meaning and use of common parameters, success and error responses, and DAL service registration. The goal of this specification is to define the common elements that are shared across DAL services in order to foster consistency across concrete DAL service specifications and to enable standard re-usable client and service implementations and libraries to be written and widely adopted.

Status of This Document

This is a working draft internal to the DAL-WG.

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](http://www.ivoa.net/Documents/) can be found at <http://www.ivoa.net/Documents/>.

Acknowledgements

TODO

Contents

1	Introduction.....	5
1.1	General Principles.....	5
1.2	The Role in the IVOA Architecture.....	5
1.3	Example Usage of DALI Specification.....	6
2	Resources.....	7
2.1	Asynchronous Execution: DALI-async.....	7
2.2	Synchronous Execution: DALI-sync.....	9
2.3	Examples: DALI-examples.....	9
2.4	Availability: VOSI-availability.....	10
2.5	Capabilities: VOSI-capabilities.....	10
2.6	Content: VOSI-tables.....	11
3	Parameters.....	13
3.1	General Rules.....	13
3.1.1	Case Sensitivity.....	13
3.1.2	Literal Values: Numbers, Boolean, Date, and Time.....	13
3.1.3	Multiple Values.....	14
3.2	Standard Parameters.....	14
3.2.1	REQUEST.....	14
3.2.2	VERSION.....	14
3.2.3	RESPONSEFORMAT.....	15
3.2.4	MAXREC.....	16

Data Access Layer Interface

3.2.5	UPLOAD.....	17
3.2.6	RUNID.....	18
4	Responses.....	19
4.1	Successful Requests.....	19
4.2	Errors.....	19
4.3	Redirection.....	20
4.4	Use of VOTable.....	21
4.4.1	Overflow.....	21
4.4.2	Errors.....	22
4.4.3	Additional Information.....	23
5	References.....	24

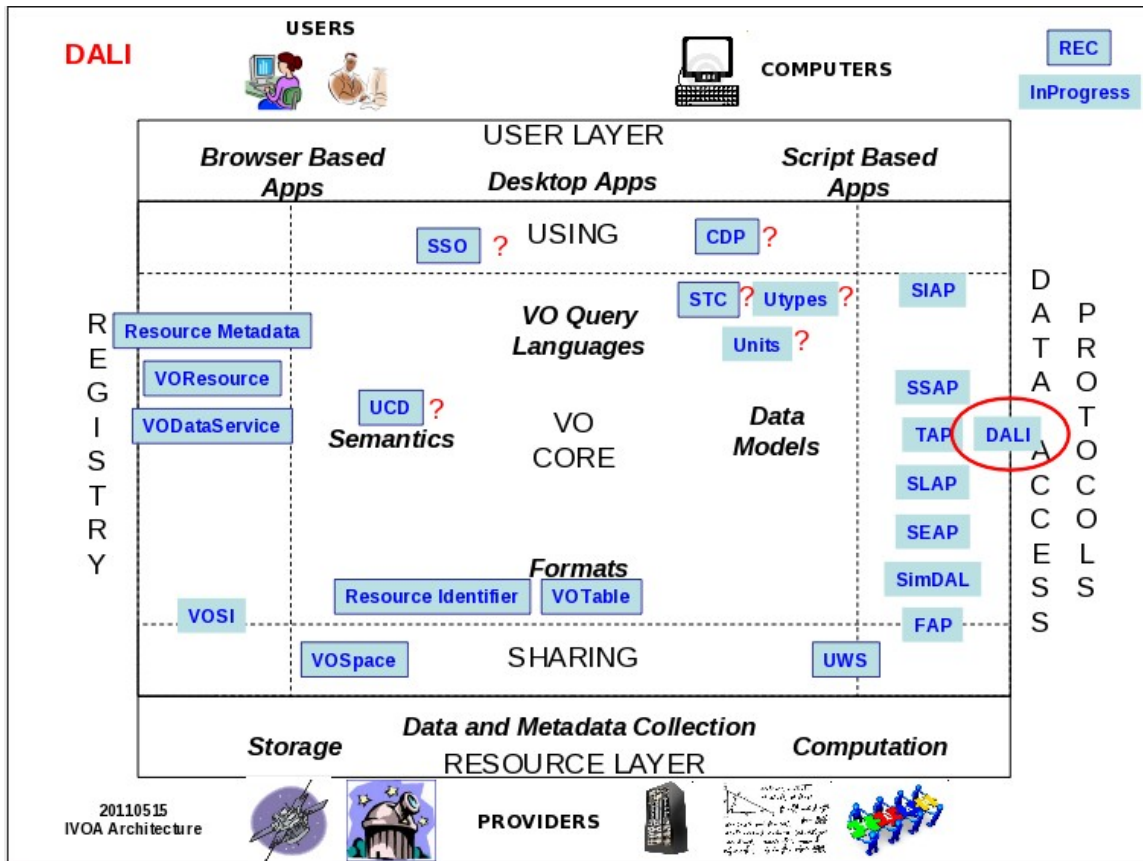
1 Introduction

The Data Access Layer Interface (DALI) defines resources, parameters, and responses common to all DAL services so that concrete DAL service specifications need not repeat these common elements.

1.1 General Principles

1.2 The Role in the IVOA Architecture

DALI defines how DAL service specifications use other IVOA standards as well as standard internet designs and protocols.



DAL services use the Universal Worker Service (UWS) pattern for asynchronous requests. All DAL services include Virtual Observatory Support Interfaces (VOSI) resources. DAL services generally use VOTable as the default output format for successful queries and always use VOTable to return error documents. Astronomical coordinate values accepted and returned by DAL services use a string representation of the Space-Time Coordinates (STC) data model. Other

Data Access Layer Interface

data models maybe describe the content of a specific DAL service; in such cases, the Utypes in the response used that describe the content will be those specified by the data model.

1.3 Example Usage of DALI Specification

2 Resources

DAL services are implemented as HTTP REST [18] web services. The primary resource in a DAL service is a job. A DAL job is defined by parameters (see 3) and can be executed either synchronously or asynchronously. A concrete service specification defines the job parameters and the manner of execution is defined by separate resources below.

In addition to job list resources, DAL services also implement several Virtual Observatory Support Interface (VOSI) resources to describe service availability, capabilities, and content.

A concrete DAL service must define at least one DALI-async or DALI-sync resource. It may define both with the same job semantics (e.g. TAP-1.0) or it may define one with one kind of job and the other with a separate kind of job (a service that does some things synchronously and others asynchronously).

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

A simple query-only DAL service like ConeSearch can be easily described as having a single DALI-sync resource where the job is a query and the response is the result of the query.

2.1 Asynchronous Execution: DALI-async

Asynchronous resources are resources that represent a list of asynchronous jobs as defined by the Universal Worker Service (UWS) pattern [10]. Requests can create, modify, and delete jobs in the job list. Special requests to modify the phase of the job cause the job to execute or abort.

As specified in UWS, a job is created by using the HTTP POST method to modify the job list. The response will always be an HTTP redirect (status code 303) and the Location (HTTP header) will contain the URL to the job (a child resource of the job list).

```
POST http://example.com/base/async-jobs
```

Data Access Layer Interface

```
< HTTP/1.1 303 See Other
< Location: http://example.com/base/async-jobs/123
```

The job description (an XML document defined by the UWS schema) can always be retrieved by accessing the job URL with the HTTP GET method:

```
GET http://example.com/base/async-jobs/123

<?xml version="1.0" encoding="UTF-8"?>
<uws:job
xmlns:uws="http://www.ivoa.net/xml/UWS/v1.0"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xlink="http://www.w3.org/1999/xlink">
  <uws:jobId>123</uws:jobId>
  <uws:runId>test</uws:runId>
  <uws:ownerId xsi:nil="true" />
  <uws:phase>PENDING</uws:phase>
  <uws:quote>2013-01-01T12:34:56</uws:quote>
  <uws:startTime/>
  <uws:endTime/>
  <uws:executionDuration>600</uws:executionDuration>
  <uws:destruction>2013-02-01T00:00:00</uws:destruction>
  <uws:parameters>
    <uws:parameter id="LANG">ADQL</uws:parameter>
    <uws:parameter id="REQUEST">doQuery</uws:parameter>
    <uws:parameter id="QUERY">select * from tab</uws:parameter>
  </uws:parameters>
  <uws:results/>
</uws:job>
```

In addition to the UWS job metadata, DAL jobs are defined by a set of parameter-value pairs. The client may add new parameters by modifying the current list of parameters via the HTTP POST method:

```
POST F00=bar http://example.com/base/async-jobs/123/parameters
```

The UWS standard allows parameters to be POSTed along with the initial job-creation request, or POSTed to the job URL, or POSTed to the parameter list (the parameters child resource) directly (as in the above example). This is easily implemented by simply applying all UWS-specific parameters to the UWS job itself and putting all remaining parameters into the parameter list. Job parameters may only be POSTed while the job is in the PENDING phase; once execution has been requested and the job is in any other phase, job parameters may not be modified.

A concrete DAL service specification will specify one or more asynchronous job-list resources and whether they are mandatory or optional. It may mandate a specific resource name to support simple client use, or it can allow the resource name to be described in the service metadata (see 2.5).

2.2 Synchronous Execution: DALI-sync

Synchronous resources are resources that accept a request (a DAL job description) and return the response (the result) directly. Synchronous requests can be made using either the HTTP GET or POST method. The parameters used to specify the job are the same for synchronous and asynchronous DAL jobs. A synchronous job is created by a GET or POST request to a synchronous job list, executed automatically, and the result returned in the response. The web service is permitted to split the operation of a synchronous request into multiple HTTP requests as long as it is transparent to standard clients. This means that the service may use HTTP redirects (status code 302 or 303) and the Location header to execute a synchronous job in multiple steps. For example, a service may

- immediately execute and return the result in the response, or
- the response is an HTTP redirect (status code 303) and the Location (HTTP header) will contain a URL; the client accesses this URL with the HTTP GET method to execute the job and get the result

Clients must be prepared to get redirects and follow them (using normal HTTP semantics) in order to complete requests.

A concrete DAL service specification will specify one or more synchronous job-list resources and whether they are mandatory or optional. It may mandate a specific resource name to support simple client use, or it can allow the resource name to be described in the service capability metadata (see 2.5).

2.3 Examples: DALI-examples

The DALI-examples resource returns a document with examples to help users use a service effectively. The document must be a human-readable HTML document with RDFa markup [19] to make it easy for software to extract the examples. The details of the microformat are left to specific service specifications; here we simply define the concept and fixed resource name for the example document...

TBD: define the actual vocab/attributes for specifying examples or example values for parameters? It probably just needs tagging example param values and explanatory text.

2.4 Availability: VOSI-availability

VOSI-availability [9] defines a simple web resource that reports on the current ability of the service to perform. In DAL services, this resource is always accessed as a resource named *availability* that is a child of the base URL for the service.

All DAL services must implement the */availability* resource, which responds with a dynamically generated document describing the current state of the service operation, e.g.:

```
<?xml version="1.0" encoding="UTF-8"?>
<vosi:availability
  xmlns:vosi="http://www.ivoa.net/xml/VOSIAvailability/v1.0">
  <vosi:available>true</vosi:available>
  <vosi:note>service is accepting queries</vosi:note>
</vosi:availability>
```

2.5 Capabilities: VOSI-capabilities

VOSI-capabilities [9] defines a simple web resource that returns an XML document describing the service. In DAL services, this resource is always accessed as a resource named *capabilities* that is a child of the base URL for the service. The VOSI-capabilities should describe all the resources exposed by the service, including which standards each resource implements.

All DAL services must implement the */capabilities* resource. The following capabilities document shows the three VOSI resources and a TAP base resource:

```
<?xml version="1.0" encoding="UTF-8"?>
<vosi: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/tap/capabilities
      </accessURL>
    </interface>
```

Data Access Layer Interface

```
</capability>
<capability standardID="ivo://ivoa.net/std/VOSI#availability">
  <interface xsi:type="vod:ParamHTTP" version="1.0">
    <accessURL use="full">
      http://example.com/tap/availability
    </accessURL>
  </interface>
</capability>
<capability standardID="ivo://ivoa.net/std/VOSI#tables">
  <interface xsi:type="vod:ParamHTTP" version="1.0">
    <accessURL use="full">
      http://example.com/tap/tables
    </accessURL>
  </interface>
</capability>
<capability xmlns:tr="http://www.ivoa.net/xml/TAPRegExt/v1.0"
  standardID="ivo://ivoa.net/std/TAP" xsi:type="tr:TableAccess">
  <interface xsi:type="vod:ParamHTTP" role="std" version="1.0">
    <accessURL use="full">
      http://example.com/tap/
    </accessURL>
  </interface>
  <!-- service details from TAPRegExt go here -->
</capability>
</vosi:capabilities>
```

Note that while this example shows the use of a registry extension schema (the inline `xmlns:tr="http://www.ivoa.net/xml/TAPRegExt/v1.0"` in the last capability element) this is not required; services may be registered and described without such an extension.

2.6 Content: VOSI-tables

VOSI-tables [9] defines a simple web resource that returns an XML document describing the content of the service. In DAL services which include it, this resource is always accessed as a resource named *tables* that is a child of the base URL for the service. The document format is defined by the VODataService-1.1 [12] standard and allows the service to describe their content as a tableset: schemas, tables, and columns.

Data Access Layer Interface

A concrete DAL service specification will specify if the */tables* resource is mandatory or optional.

```
<?xml version="1.0" encoding="UTF-8"?>
<vos:tableset
  xmlns:vosi="http://www.ivoa.net/xml/VOSITables/v1.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:vod="http://www.ivoa.net/xml/VODDataService/v1.1">
  <schema>
    <table type="output">
      <name>ivoa.ObsCore</name>
      <column>
        <name>dataprodukt_type</name>
        <description>type of product</description>
        <ucd>meta.id;class</ucd>
        <utype>obscure:Obs.dataProductType</utype>
        <dataType xsi:type="vod:TAPType" size="128">
          VARCHAR</dataType>
        </column>
        <column>
          <name>calib_level</name>
          <description>calibration level (0,1,2,3)</description>
          <ucd>meta.id;class</ucd>
          <utype>obscure:Obs.calibLevel</utype>
          <dataType xsi:type="vod:TAPType">
            INTEGER</dataType>
          </column>
          ...
        </table>
      </schema>
```

Note that the ellipsis (...) in the above example is not valid XML; it simply implies more column elements would be included there.

3 Parameters

A DAL job is defined by a set of parameter-value pairs. Some of these parameters have a standard meaning and are defined here, but most are defined by the service specification or another related standard.

3.1 General Rules

3.1.1 Case Sensitivity

Parameter names are not case sensitive; a DAL service must treat upper-, lower-, and mixed-case parameter names as equal. Parameter values are case sensitive unless a concrete DAL service specification explicitly states that the values of a specific parameter are to be treated as case-insensitive. For example, the following are equivalent:

```
F00=bar
Foo=bar
foo=bar
```

Unless explicitly stated by the service specification, these are not equivalent:

```
F00=bar
F00=Bar
F00=BAR
```

In this document, parameter names are typically shown in uppercase for typographical clarity, not as a requirement.

3.1.2 Literal Values: Numbers, Boolean, Date, and Time

Integer numbers **must** be represented in a manner consistent with the specification for integers in *XML Schema Datatypes* [10].

Real numbers **must** be represented in a manner consistent with the specification for double-precision numbers in *XML Schema Datatypes* [10]. This representation allows for integer, decimal and exponential notations.

Boolean values must be represented in a manner consistent with the specification for Boolean in *XML Schema Datatypes* [10]. The values *0* and *false* are equivalent. The values *1* and *true* are equivalent.

```
F00=1
F00=true

BAR=0
BAR=false
```

Data Access Layer Interface

Date and time values must be represented as ISO 8601 formatted strings with a T character separating the date and time components. The time component is optional, in which case the T separator is not used. Fractions of a second are permitted but not required. For example:

```
2000-01-02T15:20:30.456
2001-02-03T04:05:06
2002-03-04
```

are all legal date or date plus time values. Values never include a timezone indicator and are always interpreted as UTC [17]. In cases where values may be expressed using Julian Date (JD) or Modified Julian Date (MJD), these follow the rules for double precision numbers above and are not restricted to UTC.

3.1.3 Multiple Values

Parameters may be assigned multiple values with multiple parameter=value pairs using the same parameter name. Whether or not multiple values are permitted and the meaning of multiple values is specified for each parameter by the specification that defines the parameter. For example, the UPLOAD parameter (3.2.5) permits multiple occurrences of the specified pair (table,uri), e.g.:

```
UPLOAD=foo,http://example.com/foo
UPLOAD=bar,http://example.com/bar
```

Services must respond with an error if the request includes multiple values for parameters defined to be single-valued.

3.2 Standard Parameters

3.2.1 REQUEST

The REQUEST parameter specifies the type of the DAL job at the highest level. In many cases, a service will have only one possible value. This parameter is still used in such cases so that future versions or non-standard (site-specific) features may support additional values.

A service must respond with an error if the REQUEST parameter is missing or the value is not recognised.

The REQUEST parameter is always single-valued.

3.2.2 VERSION

The VERSION parameter is used so the client can specify which version of the service standard they are using to make the request. This allows implementers to support multiple versions of a standard in a single web service and with a single

Data Access Layer Interface

resource for the DAL job list. If the client does not specify a value for the VERSION, the service must interpret the request using the rules and semantics of the latest recommended version supported by the service. For example, if a TAP service supports version 1.0 and 1.1 then the client may make requests with VERSION=1.0 or VERSION=1.1; if the VERSION is not specified in the client request, the behaviour must be the same as VERSION=1.1). If version 1.2 of the standard exists and is supported by the service but is not a recommendation, the default is still VERSION=1.1 but the caller can include VERSION=1.2 to have the request interpreted according to the draft specification.

The value of the version advertised by the service and requested by the client must follow the IVOA version number scheme [14]. According to [14], the version number includes a publication date. For the purposes of this specification, these date tags are not supported, i.e. Services only accept VERSION without dates. Clients can thus not request the behaviour of a specific WD or PR. This is intentional since implementation against unstable standards are not supposed to be stable.

A service must respond with an error if the caller requests an unsupported version using the VERSION parameter.

The VERSION parameter is always single-valued.

3.2.3 RESPONSEFORMAT

The RESPONSEFORMAT parameter is used so the client can specify the format of the response (e.g. the output of the job). For DALI-sync requests, this is the content-type of the response. For DALI-async requests, this is the content-type of the result resource(s) the client can retrieve from the UWS result list resource; if a DALI-async job creates multiple results, the RESPONSEFORMAT should control the primary result type, but details can be specific to individual service specifications. While the list of supported values are specific to a concrete service specification, the general usage is to support values that are content-types (mimetypes) for known formats as well as shortcut symbolic values.

table type	MIME type(s)	short form
VOTable	application/x-votable+xml text/xml	votable
comma separated values	text/csv	csv
tab separated values	text/tab-separated-values	tsv
FITS file	application/fits	fits
pretty-printed text	text/plain	text
pretty-printed Web page	text/html	html

Data Access Layer Interface

A DAL service **must** accept a *RESPONSEFORMAT* parameter indicating a format that the service supports and **should** fail (xref to error handling section) where the *RESPONSEFORMAT* parameter specifies a format not supported by the service implementation.

A concrete DAL service specification will specify any mandatory or optional formats as well as new formats not listed above; it may also place limitations on the structure for formats that are flexible. For example, a resource that responds with tabular output only may impose a limitation that FITS files only contain FITS tables, possibly only of specific types (ascii or binary).

If a client requests a format by specifying the mimetype (as opposed to one of the short forms), the response that delivers that content must set that mimetype in the Content-Type header. This is only an issue when a format has multiple acceptable mimetypes (e.g. VOTable).

Individual DAL services (not just specifications) are free to support custom formats by accepting non-standard values for the *RESPONSEFORMAT* parameter.

The *RESPONSEFORMAT* parameter should not be confused with the *FORMAT* parameter used in many DAL services. The latter is generally used as a query parameter to search for data in the specified format; only in the case of TAP-1.0 are *FORMAT* and *RESPONSEFORMAT* equivalent.

The *RESPONSEFORMAT* parameter is always single-valued.

3.2.4 MAXREC

For resources performing discovery (querying for an arbitrary number of records), the resource **must** accept a *MAXREC* parameter specifying the maximum number of records to be returned. If *MAXREC* is not specified in a request, the service **may** apply a default value or **may** set no limit. If the size of the result exceeds this value, the service **must** only return the requested number of rows. If the result set is truncated in this fashion, it must include an overflow indicator as specified in 4.4.1 .

The service **must** support the special value of *MAXREC=0*. This value indicates that, in the event of an otherwise valid request, a valid response be returned containing metadata, no results, and an overflow indicator as above. The service is not required to execute the request and the overflow indicator does not necessarily mean that there is at least one record satisfying the query. The service **may** perform validation and may try to execute the request, in which case a *MAXREC=0* request can fail.

The *MAXREC* parameter is always single-valued.

3.2.5 UPLOAD

The *UPLOAD* parameter is used to reference read-only external tables via their URI, to be uploaded for use as input tables to the query. The value of the *UPLOAD* parameter is a table name-URI pair. For example:

```
UPLOAD=table1,http://example.com/t1
```

would define an input table *table1* at the given URIs. Table names must be simple strings made up of alphabetic, numeric, and the underscore characters only and must start with an alphabetic character.

Services that implement *UPLOAD* **must** support *http* as a URI scheme (e.g. must support treating an *http* URI as a URL). A VOSpace URI (*vos:<something>*) is a more generic example of a URI that requires more service-side functionality; support for the *vos* scheme is optional.

To upload a table inline, the caller must specify the *UPLOAD* parameter (as above) using a special URI scheme “param”. This scheme indicates that the value after the colon will be the name of the inline content. The content type used is *multipart/form-data*, using a “file” type input element. The “name” attribute must match that used in the *UPLOAD* parameter.

For example, in the POST data we would have this parameter:

```
UPLOAD=table3,param:t3
```

and this content:

```
Content-Type: multipart/form-data; boundary=AaB03
[...]
--AaB03x
Content-disposition: form-data; name="t3"; filename="t3.xml"
Content-type: application/x-votable+xml
[...]
--AaB03x
[...]
```

If inline table upload is used by a client, the client **must** POST both the *UPLOAD* parameter and the associated inline content in the same request. Services that implement table upload **must** support the *param* scheme for inline uploads.

In principle, any number of tables can be uploaded using the *UPLOAD* parameter (see for details) and any combination of URI schemes supported by the service as long as they are assigned unique table names in the request. For example:

```
UPLOAD=table1,http://example.com/t1.xml
```

Data Access Layer Interface

```
UPLOAD=table2,vos://example.authority!tempSpace/t2.xml
```

```
UPLOAD=table3,param:t3
```

Services may limit the size and number of uploaded tables; if the service refuses to accept the entire table it **must** respond with an error as described in 4.2 . Specific service specifications must specify how uploaded tables are referenced in other request parameters (for example, in a query), and interpreted.

3.2.6 RUNID

The service **should** implement the *RUNID* parameter, used to tag service requests with the identifier of a larger job of which the request may be part. The *RUNID* value is a string with a maximum length of 64 characters.

TBD: do we need to place a limit on the length? is 64 enough?

For example, if a cross match portal issues multiple requests to remote services to carry out a cross-match operation, all would receive the same *RUNID*, and the service logs could later be analyzed to reconstruct the service operations initiated in response to the job. The service **should** ensure that *RUNID* is preserved in any service logs and **should** pass on the *RUNID* value in calls to other services made while processing the request.

The *RUNID* parameter is always single-valued.

4 Responses

All DAL service requests eventually result in one of three kinds of responses: successful HTTP status code (200) and a service- and resource-specific representation of the results, an HTTP status code (??) and an unspecified error document, or a redirect HTTP status code (302 or 303) with a URL in the HTTP header.

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 (see 3.2.3) or in the default format for the service. The service should set the following HTTP headers to the correct values where possible.

Content-Type	mimetype of the response
Content-Encoding	encoding/compression of the response
Content-Length	size of the response in bytes (generally not known for dynamically generated and streamed response)
Last-Modified	timestamp when the resource was last changed (not applicable to dynamically generated response)

For jobs executed using a DALI-async resource, the result(s) **must** be made available as child resources of the result list and directly accessible there. For jobs that inherently create a fixed result, service specifications may specify the name of the result explicitly. For example, TAP-1.0 has a single result and it **must** be named *result* in the result list and be directly accessible by that name, e.g.:

GET <http://example.com/base/joblist/123/results/result>

For concrete DAL service specifications where multiple result files may be produced, the specification may dictate the names or it may leave it up to implementations to choose suitable names.

4.2 Errors

If the service detects an exceptional condition, it **must** return an error document with an appropriate HTTP-status code. DAL services distinguish three classes of errors:

- Errors in the use of the HTTP protocol

Data Access Layer Interface

- Errors in the use of the specific DAL protocol, including an invalid request
- Errors caused by a failure of the service to complete a valid request

Error documents for HTTP-level errors are not specified since responses to these errors may be generated by service containers and cannot be controlled by service implementations. There are several cases where a DAL service could return an HTTP error. First, a DALI-async resource return a 404 (not found) error if the client accesses a job within the UWS joblist that does not exist, or accesses a child resource of the job that does not exist (e.g. the error resource of a job that has not run and failed, or a specific result resource in the result list that does not exist). Second, access to a resource could result in an HTTP 401 (not authorized) response if authentication is required or an HTTP 403 (forbidden) error if the client is not allowed to access the requested resource.

Error documents describing errors in use of the DAL service protocol **must** be VOTable documents as described in [11]; any result-format specified in the request is ignored. In all cases, these are errors that occur when the job is executed and do not override any error behaviour for a UWS resource which specifies the behaviour and errors associated with interacting with the job itself.

If the invalid job is being executed using a DALI-async resource, the error document must be accessible from the <DALI-async>/<jobid>/error resource (specified by UWS) and when accessed via that resource it must be returned with an HTTP status code 200, e.g.:

```
GET http://example.com/base/joblist/123/error
```

If the error document is being returned directly after a DALI-sync request, the service should use a suitable HTTP status code (4xx or 5xx) to describe the failure and include the error document in the body of the response. The Content-Type header will tell the client that there is a VOTable error document in the body of the response.

Error documents describing the failure of the service to execute a valid job are returned as above, but the VOTable document must ??? to indicate that the request was valid but failed due to (i) an internal limitation or (ii) a transient failure and may succeed in the future.

TBD: How do we generically divide errors into these 3 types: invalid, internal limitation/refusal, or transient failure?

4.3 Redirection

A concrete DAL service specification may require that HTTP redirects (302 or 303) be used to communicate the location of an alternate resource which should be accessed by the client via the HTTP GET method. For example, the UWS

pattern used for DALI-async (2.1) requires this behaviour. Even when not required, concrete DAL service specifications must allow implementors to use redirects and clients should follow these redirects using normal HTTP semantics [5].

4.4 Use of VOTable

VOTable is a general format. In DAL services we require that it be used in a particular way. The result VOTable document **must** comply with VOTable v1.2 or greater [11]. For columns containing coordinate values, the coordinate system metadata should be provided as described in [15]. For columns containing photometric values (including fluxes), the system should be described as in [16].

For resources where the job is a query, the VOTable **must** contain a *RESOURCE* element identified with the attribute *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, before the *TABLE* element, an *INFO* element with attribute *name = "QUERY_STATUS"*. The *value* attribute **must** contain one of the following values:

OK	the query executed successfully and the result is included in the resource
ERROR	an error was detected at the level of the protocol, the job failed to execute, or an error occurred while writing the table data
OVERFLOW	the query executed successfully, the result is included in the resource, and the result was truncated at MAXREC rows

The content of the *INFO* element conveying the status **should** be a message suitable for display to the user describing the status.

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

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

```
<INFO name="QUERY_STATUS" value="ERROR">
value out of range in POS=45,91
</INFO>
```

4.4.1 Overflow

If an overflow occurs (result exceeds MAXREC), the service must include an *INFO* element in the *RESOURCE* with *name="QUERY_STATUS"* and the

Data Access Layer Interface

value="OVERFLOW". If the initial info element (above) specified the overflow, no further elements are needed, e.g.:

```
<RESOURCE type="results">
<INFO name="QUERY_STATUS" value="OVERFLOW"/>
...
<TABLE>...</TABLE>
</RESOURCE>
```

If the initial info element specified a status of OK then the service must append an INFO element for the overflow after the table, e.g.:

```
<RESOURCE type="results">
<INFO name="QUERY_STATUS" value="OK"/>
...
<TABLE>...</TABLE>
<INFO name="QUERY_STATUS" value="OVERFLOW"/>
</RESOURCE>
```

In the above examples, the TABLE should have exactly MAXREC rows.

4.4.2 Errors

If an error occurs, the service must include an INFO element with *name="QUERY_STATUS"* and the *value="ERROR"*. If the initial info element (above) specified the error, no further elements are needed, e.g.:

```
<RESOURCE type="results">
<INFO name="QUERY_STATUS" value="ERROR"/>
...
<TABLE>...</TABLE>
</RESOURCE>
```

If the initial info element specified a status of OK then the service must append an INFO element for the overflow after the table, e.g.:

```
<RESOURCE type="results">
<INFO name="QUERY_STATUS" value="OK"/>
...
<TABLE>...</TABLE>
<INFO name="QUERY_STATUS" value="ERROR">
unexpected IO error while converting something
</INFO>
</RESOURCE>
```

Data Access Layer Interface

The use of trailing INFO element allows a service to stream output and still report overflows or errors to the client. The content of these trailing INFO elements is optional and intended for users; client software **should not** depend on it.

4.4.3 Additional Information

Additional *INFO* elements **may** be provided, e.g., to echo the input parameters back to the client in the query response (a useful feature for debugging or to self-document the query response), but clients **should not** depend on these. For example:

```
<RESOURCE type="results">
<INFO name="QUERY_STATUS" value="ERROR">
missing REQUEST parameter
</INFO>
<INFO name="standardID" value="TAP"/>
<INFO name="VERSION" value="1.0"/>
...
</RESOURCE>
```

The following names for INFO elements should be used if applicable, but this list is not definitive.

name	meaning	example value
standardID	IVOA standardID for the service specification	ivo://ivoa.net/std/TAP
VERSION	Version number used to interpret and process request (3.2.2)	1.0
CITATION	Reference to a publication that can/should be referenced if the result is used	Hubble, E. P., 1915, ApJ 42, 283

TBD: It would be nice to use an INFO param with name attribute CITATION, a second attribute saying what kind of vauue there was, and then put the value in as text content of the INFO element. Possible? It could in thery permit bibcodes, ADS uris/urns, DOIs, or free-form references as above, maybe even in a usable enough fashion.

5 References

- [1] D. Tody, F. Bonnarel, M. Dolensky, J. Salgado, DAL-WG, *IVOA Data Access Layer Service Architecture and Standard Profile*, IVOA Note 5 October 2008. http://www.ivoa.net/internal/IVOA/SiaInterface/DAL2_Architecture.pdf
- [2] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*, IETF RFC 2119. <http://www.ietf.org/rfc/rfc2119.txt>
- [3] T. Berner-Lee, R. Fielding L. Masinter, *Uniform Resource Identifiers (URI): Generic Syntax*, IETF RFC 2396. <http://www.ietf.org/rfc/rfc2396.txt>
- [4] P. Biron & A. Malhotra, *XML Schema Part 2: Datatypes Second Edition*, W3C Recommendation 28 October 2004. <http://www.w3.org/TR/xmlschema-2/>
- [5] R. Fielding, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, *Hypertext Transfer Protocol – HTTP/1.1*, IETF RFC 2616. <http://www.rfc-editor.org/rfc/rfc2616.txt>
- [6] N. Freed & N. Borenstein, *Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies*, IETF RFC 2045. <http://www.ietf.org/rfc/rfc2045.txt>
- [7] Y. Shafranovich, *Common Format and MIME Type for Comma-Separated Values (CSV) Files*, IETF RFC 4180. <http://www.ietf.org/rfc/rfc4180.txt>
- [8] IANA, *MIME Media Types*, <http://www.iana.org/assignments/media-types/text/tab-separated-values>
- [9] M. Graham & G. Rixon (ed.), GWS-WG, *IVOA Support Interfaces Version 1.0*, IVOA Recommendation, 31 May 2011. <http://www.ivoa.net/Documents/VOSI/>
- [10] P. Harrison & G. Rixon, *Universal Worker Service Version 1.0*, IVOA Recommendation, 10 October 2010. <http://www.ivoa.net/Documents/UWS/>
- [11] F. Ochsenbein (ed.), R. Williams, *VOTable Format Definition Version 1.2*, IVOA Recommendation 30 November 2009. <http://www.ivoa.net/Documents/VOTable/1.2>
- [12] R. Plante (ed.), A. Stébé, K. Benson, P. Dowler, M. Graham, G. Greene, P. Harrison, G. Lemson, A. Linde, G. Rixon & IVOA Registry-WG, *VODataService: a VOResource Schema Extension for Describing Collections and Services Version 1.1*. IVOA Recommendation, 02 December 2010 <http://www.ivoa.net/Documents/VODataService/1.1>
- [13] P. Dowler, G. Rixon, D. Tody, DAL-WG, *Table Access Protocol Version 1.0*, IVOA Recommendation 27 March 2010. <http://www.ivoa.net/Documents/TAP/1.0>
- [14] R. J. Hanisch, C. Arviset, F. Genova, B. Rhino, *IVOA Document Standards*, IVOA Recommendation 13 April 2010. <http://www.ivoa.net/Documents/DocStd/>

Data Access Layer Interface

- [15] M. Demleitner, F. Ochsenbein, J. McDowell, A. Rots, Space-Time Coordinate Metadata in VOTables, IVOA Note 18 June 2010
- [16] J. Salgado (ed.), C. Rodrigo, p. Osuna (ed.), M. Allen, M. Louys, J. McDowell, D. Baines, J. Maiz, E. Hatzimnaoglou, S. Derriere, *IVOA Photometry Data Model Version 1.0*, IVOA Proposed Recommendation 17 May 2012
- [17] A. H. Rots (ed.), *Space-Time Coordinate Metadata for the Virtual Observatory Version 1.33*, IVOA Recommendation 30 October 2007
<http://www.ivoa.net/Documents/STC/1.33>
- [18] [R.Fielding, *Architectural Styles and the Design of Network-based Software Architectures*, PhD thesis, University of California Irvine, 2000](#)
- [19] M. Sporny (ed.) *RDFa Lite 1.1*, W3C Recommendation 07 June 2012,
<http://www.w3.org/TR/2012/REC-rdfa-lite-20120607/>