Appendix 2
A Summary of Open GeoSpatial Consortium Standards

Alan Grainger
School of Geography, University of Leeds

1. Introduction

This document complements the summary of IVOA Standards by showing how an astronomical Virtual Observatory can be imitated by a geographical Virtual Observatory based on Open GeoSpatial Consortium (OGC) Standards. It provides an overview compilation of standards for the Registry (*Catalogue Service*), exchanging spatial data in formats of increasing complexity (*Web Map Tile Service, Web Map Service, Web Feature Service* and *Web Coverage Service*), and tools for analysing data (*Web Coverage Processing Service*). Various software packages already comply with these standards.

2. Overview

2.1 Catalogue Service

The Catalogue Service is the geospatial equivalent of the IVOA Registry, or "yellow pages". It enables groups to publish the metadata that summarises the key features of the host so that other groups may search for information with the features which they need and then retrieve this information from the host server where it is stored.

2.2 Web Map Tile Service

The Web Map Tile Service is a standard way for exchanging map tiles, which are spatially referenced units of maps whose content, extent and resolution is predefined by those who produce the maps. Map tiles are also produced as images rather than data which can be analysed. The utility of map tiles has been demonstrated by such services as Google maps, but what they gain in utility they lose in flexibility to meet user requirements.

2.3 Web Map Service

The Web Map Service is a standard way for sharing and combining maps as images. It offers users more flexibility in how they define the area for which they require information and has the additional advantage that map images may be requested by using a simple HTTP interface. The map image files are supplied in standard formats, e.g. JPEG, PNG, etc. which are easy to that display in a browser application.

2.4 Web Feature Service

The Web Feature Service is a more advanced standard because it enables users to obtain data at the feature level which they can manipulate themselves, instead of the image files which are supplied by the Web Map Tile Service and Web Map Service.

2.5 Web Coverage Service

The Web Coverage Service is even more sophisticated than the Web Feature Service, since it encompasses the exchange of more complex forms of spatial information and spatio-temporal information, and offers greater flexibility for how information is coded, stored and exchanged.

2.6 Web Coverage Processing Service

The Web Coverage Processing Service defines a generic language for extracting and analysing multidimensional spatial data in various formats.

3. Catalogue Service

3.1 Introduction

The Catalogue Service is the geospatial equivalent of the IVOA Registry. The following text is directly quoted from the OGC Standard Document OpenGIS Catalogue Service Implementation Specification 2.0.2, which can be found at http://www.opengeospatial.org/standards/cat.

3.2 Overview

"Catalogue services support the ability to publish and search collections of descriptive information (metadata) for data, services, and related information objects. Metadata in catalogues represent resource characteristics that can be queried and presented for evaluation and furtherprocessing by both humans and software. Catalogue services support the use of one of several identified query languages to find and return results using well-known content models (metadata schemas) and encodings.

3.3 Catalogue abstract information model

The abstract information model specifies a grammar for a minimal query language, a set of core queryable attributes (names, definitions, conceptual datatypes), and a common record format that defines the minimal set of elements that should be returned in the brief and summary element sets.

The geospatial community is a very broad-based community that works in many different operational environments.On one extreme there are tightly coupled systems dedicated to well defined functions in a tightly controlled environment. At the other extreme are Web based services that know nothing about the client. This specification is applicable to the full range of catalogue operating environments

The query capabilities of the OGC General Catalogue Model provide a minimum set of data types and query operations. In addition, these Query Capabilities provide a high degree of flexibility enabling alternate styles of query, result presentation, and the potential support any geo-enabled query language. This flexibility is provided by the query operation that contains the parameters needed to select the query result presentation style and to provide a query expression which includes the actual query with an identification of the query language used.

The minimal query language assists the consumer in the discovery of datasets of interest at all sites supporting the OpenGIS Catalogue Services. The ability to specify alternative query languages allows for evolution and higher levels of interoperability among more tightly coupled communities of Catalogue Service Providers and Consumers.

Among assumptions made during the development of OGC_Common Query Language are that the query will have a syntax similar to the SQL "Where Clause."

This specification makes no assumptions about how taxonomies are maintained in a catalogue, or how records are classified according to those taxonomies. Instead, this specification supports classification queries based on taxonomies.

3.4 Core catalogue schema

Metadata structures, relationships, and definitions - known as conceptual schemas - exist for multiple information communities. For the purposes of interchange of information within an information community, a metadata schema may be defined that provides a common vocabulary which supports search, retrieval, display, and association between the description and the object being described.

Common queryable elements include:

Subject:	*The topic of the content of the resource*
Title:	*A name given to the resource*
Abstract:	*A summary of the content of the resource*
AnyText:	*A target for full-text search of character data types in a catalogue*
Format:	*The physical or digital manifestation of the resource*
Identifier:	*A unique reference to the record within the catalogue*
Identifier	
Modified:	*Date on which the record was created or updated within the catalogue*
Type:	*The nature or genre of the content of the resource*
BoundingBox:	*bounding box for identifying a geographic area of interest*
CRS:	*Geographic Coordinate Reference System for the BoundingBox*
Identifier	
Association:	*Complete statement of a one-to-one relationship*

The composition of compound element "BoundingBox" includes:

WestBoundLongitude: Western-most limit
SouthBoundLatitude: Southern-most limit
EastBoundLongitude : Eastern-most limit
NorthBoundLatitude: Northern-most limit

A set of core properties returned from a metadata search is encouraged to permit the minimal implementation of a catalogue service independent of a companion application profile, and to permit the use of metadata returned from different systems and protocol bindings. The core metadata is returned as a request for the Common Element Set, which is a new group of public metadata elements, expressed using the nomenclature and syntax of Dublin Core Metadata, ISO 15836. The elements include:

Title
Creator
Subject
Description (Abstract)
Publisher
Contributor
Date (created or modified)
Type
Format
Identifier (unique reference to the record within the catalogue)
Source
Language
Relation (Association)
Coverage (BoundingBox)
Rights

3.5 General catalogue interface model

The getCapabilities operation allows clients to retrieve service metadata describing Catalogue Service instance.

The Discovery class allows clients to discover resources registered in a catalogue, by providing four operations named query, present, describeRecordType, and getDomain.

The Manager class allows a client to insert, update and/or delete catalogue content. This class has an optional association from the CatalogueService class; it is not required that a catalogue service implement publishing functionality. Two operations are provided: transaction and harvestResource.

The transaction operation allows a client to formulate a transaction, and send it to the catalogue to be processed. The transaction may contain metadata records and elements of the information model that the catalogue understands. To use the transaction operation, the client must know something about the information model that the catalogue implements.

The harvestResource operation, on the other hand, directs the catalogue to retrieve an accessible metadata record and processes it for inclusion in the catalogue, perhaps periodically re-fetching the metadata records to refresh the information in the catalogue. The client does not need to be aware of the information model of the catalogue when using the "harvestResource" operation, since the catalogue itself is doing the work required to process the information. The client is simply pointing to where the metadata resource to be harvested is.

The harvestResource operation facilitates the retrieval of remote resources from a designated location and provides for optional transactions on the local catalogue. Allows a user to request that a catalogue service attempt to retrieve a resource from a specified location, and to optionally create one or more entries for that resource.

The Brokered Access class allows clients to place an order for an identified registered resource, for use when that resource is a data product that is not directly accessible to clients. This class has an optional association from the Catalogue Service class, in which case this interface is implemented by the Catalogue Service implementation. Not all resources can be accessed directly. Brokered access provides for accessing resources that are controlled. Controlled resources might include those for which one or more of the following applies:

a. A fee is charged
b. Have security limitations
c. Require additional processing
d. Are not available electronically

3.6 Hypertext Transfer Protocol (HTTP) protocol

The purpose of the HTTP protocol binding is to describe the request and response messages that are common to all web-based catalogue services. Only the GET and POST methods are employed in the HTTP binding. The mandatory GetCapabilities operation allows CSW clients to retrieve service metadata from a server. The response to a GetCapabilities request shall be an XML document containing service metadata about the server.

3.7 Metadata Extension Package of CS-WebRIM Profile 1.0

3.7.1 Introduction

This standard enables cataloguing of ISO 19115, ISO19115-2 and ISO 19119 compliant metadata. The terms 'catalogue' and 'registry' are often used interchangeably, but the following distinction is made here: a registry is a specialized catalogue that exemplifies a formal registration process such as those described in ISO 19135 or ISO 11179-6. A registry is typically maintained by an authorized registration authority who assumes responsibility for complying with a set of policies and procedures for accessing and managing registry content. This profile does not stipulate any particular registration policies that must be enforced by a conforming implementation.

3.7.2 Publishing metadata

Two distinct and independent functionalities of the system are clearly identified. The first one deals with the management of metadata and metadata records of the resources produced by an organization. The metadata author describes geo-resources by applying ISO 19139. A geo-resource may be a service, a geodataset (single or collection) or an application. The repository publisher then publishes these records to a repository; they can then be managed by the organization. The metadata author and repository publisher act for the producing organization.

The second functionality concerns the cataloguing of metadata records that enables their discovery by users. To achieve this, the registry publisher publishes descriptions of the metadata records to a registry through the catalogue service. In order to give the user access to the full metadata records stored in the repository, the registry publisher links the descriptions of the catalogue to the appropriate items of the repository. By doing so, he gives a requestor access to descriptions of metadata records in the registry and to the full ISO 19139 metadata records in the repository.

3.7.3 Discovering metadata

A requestor discovers metadata entries and then accesses specific ISO 19139 metadata records one at a time through the operations provided by the catalogue service. If a service is discovered that fits his search terms, he can bind to this service in accordance with the information in the result sets of the catalogue service. Depending on the context of the catalogue service implementation, the requestor can be a member of the responsible organization or an external individual.

3.7.4 Harvesting metadata

A catalogue service may harvest metadata records from a given repository that contains ISO 19139 metadata records. If the catalogue successfully retrieves the resource and successfully processes it, then one or more corresponding registry objects are created or updated. Brief representations of all modified records are returned to the client when processing is complete.

3.7.5 Keyword and Thesaurus information

Keywords related to the resource are handled through classifications of the Resource ExtrinsicObjects by keywordNodes which either belong to one of the default Keyword Classification schemes like KeywordSchemeUntyped (which contains all the keywords not defined in a specific thesaurus) or classifications by keywordNodes which belong to a classification Scheme associated to a thesaurus.

Spatial representations include:

vector:	*vector data is used to represent geographic data*
grid:	*grid data is used to represent geographic data*
textTable:	*textual or tabular data is used to represent geographic data*

	in triangulated irregular network
stereoModel	*three-dimensional view formed by the intersecting homologous rays of an overlapping pair of images*
video scene	*from a video recording*

3.8 ISO Metadata Application Profile

Catalogue services are the key technology for locating, managing and maintaining distributed geo-resources (i.e. geospatial data, applications and services). With OGC catalogue services, client applications are capable of searching for geo-resources in a standardized way (i.e. through standardized interfaces and operations) and, ideally, they are based on a well-known information model, which includes spatial references and further descriptive (thematic) information that enables client applications to search for geo-resources in very efficient ways.

Whereas interfaces and operations of OGC catalogue services are well defined, it is left up to the developer of the system to define a specific information model which a catalogue service instance provides. This includes, but is not limited to, the information which can be inserted in the catalog, supported query languages, available search terms, response/result sets, etc. This point is of major importance with respect to interoperability between different catalogue service instances.

In Europe, running catalogue instances result from work being done within different SDI initiatives (e.g. SDI NRW Initiative1, Germany/Netherlands cross-border initiative, JRC EU Portal, EUROSTAT, Inspire, German SDI initiative). Members of these initiatives have developed an ISO-based application profile for ISO19115 metadata for geodata/geospatial applications and ISO19119-based metadata for tightly and loosely-coupled geospatial services."

4. Web Map Tile Service

4.1 Introduction

The Web Map Tile Service is the simplest form of data exchange since it is highly aggregated spatially and works with image files. This section directly quotes from text in the OpenGIS® Web Map Tile Service Implementation Standard Vs. 1.0.0. which can be found at http://www.opengeospatial.org/standards/wmts.

4.2 Overview

"The Web Map Tile Service (WMTS) described in this standard builds on earlier efforts to develop scalable, high performance services for web based distribution of cartographic maps.

This Web Map Tile Service (WMTS) Implementation Standard provides a standard based solution to serve digital maps using predefined image tiles. The service advertises the tiles it has available through a standardized declaration in the ServiceMetadata document common to all OGC web services. This declaration defines the tiles available in each layer (i.e. each type of content), in each graphical representation style, in each format, in each coordinate reference system, at each

scale, and over each geographic fragment of the total covered area. The ServiceMetadata document also declares the communication protocols and encodings through which clients can interact with the server. Clients can interpret the ServiceMetadata document to request specific tiles.

The WMTS standard complements the existing Web Map Service standard of the OGC. The WMS standard focuses on flexibility in the client request enabling clients to obtain exactly the final image they want. A WMS client can request that the server creates a map by overlaying an arbitrary number of the map layers offered by the server, over an arbitrary geographic bound, with an arbitrary background color at an arbitrary scale, in any supported coordinate reference system. The client may also request that the map layers be rendered using a specific server advertised style.

However, all this flexibility comes at a price: server image processing must scale with the number of connected clients and there is only limited potential to cache images between the server and client since most images are different. As web service clients have become more powerful, it has become possible to consider an alternative strategy which forces the clients to perform image overlays themselves and which limits the clients to requesting map images which are not at exactly the right position thereby forcing the clients to mosaic the tiles obtained from the server and clip the set of tiles into a final image. This restriction of image requests to a fixed, predefined set allows for servers to scale based on communication processing abilities rather than image processing abilities because servers can pre-render some or all of their images and can use image caching strategies. The fixed set of images also enables network providers to cache images between the client and the server, reducing latency and bandwidth use.

Popular, non-standardized, commercial implementations of this approach, such as Google Maps, Microsoft Virtual Earth and Yahoo! Maps have already shown that there are clear performance benefits to adopting this methodology.

4.3 Definitions

Feature info: information related to a particular pixel of a map that refers to the geographic data portrayed on that area.

Layer: basic unit of geographic information that may be requested as a map from a server.

Tile: a rectangular pictorial representation of geographic data, often part of a set of such elements, covering a spatially contiguous extent and sharing similar information content and graphical styling, which can be uniquely defined by a pair of indices for the column and row along with an identifier for the tile matrix.

Tile matrix: a collection of tiles for a fixed scale.

Tile matrix set: a collection of tile matrices defined at different scales.

4.4 Major components

The goal of providing a WMTS enabled service is to be performance oriented and scalable. Therefore, servers must be able to return tiles quickly. A good way to achieve that is to use locally stored pre-rendered tiles that will not require any image manipulation or geo-processing. Server developers will decide if pre-rendered tiles will be generated in a previous tile-preparation process or generated on the fly utilizing a caching mechanism. With tile-based mapping it is important that the server will be able to handle asynchronous access to tiles as most clients will simultaneously query for multiple tiles to fill a single view.

The purpose of a WMTS service is to serve maps divided in individual tiles. The WMTS interface allows a client to receive three types of resources either in response to a resource request in the resource oriented architectural style or in response to an operation in the procedure oriented architectural style. Those resources and operations are:

a. A ServiceMetadata resource (in response to a GetCapabilities operation for the procedure oriented architectural style) (required implementation by servers). It describes the abilities and information holdings of the specific server implementation. In procedure oriented architectural style this operation also supports negotiation of the standard version being used for client-server interactions.

b. A tile resource (in response to a GetTile operation for the procedure oriented architectural style) (required implementation by servers). It shows a fragment of a map representation of a layer.

c. A FeatureInfo resource (in response to a GetFeatureInfo operation for the procedure oriented architectural style) (optional implementation by servers). It provides information about the features located at a particular pixel of a tile map, in a similar way to the WMS GetFeatureInfo operation, by providing, for example, the thematicattribute name and value pairs in textual form.

These operations have many similarities to the Web Map Service (WMS), Web Feature Service (WFS), and Web Coverage Service (WCS). Many of the aspects of this WMTS interface that are shared in common with other OWSs are specified in the OpenGIS® Web Services Common Implementation Specification [OGC 06-121r3].

The WMTS serves a single tile of a single layer of a map. Unlike WMS, there is no specified way to request a server to combine and return a map tile with information coming from more than one layer in a single fetching process. WMTS clients that want to show a combination of layers must make independent requests for the layer tiles and then combine or overlay the responses. Also bounding boxes and scales of these WMTS tiles are constrained to a discrete set of values.

4.5 Tile matrix set – the geometry of the tiled space

In a tiled map layer, the representation of the space is constrained in a discrete set of parameters. A tile matrix set defines these parameters. Each tile matrix set contains

one or more "tile matrices" defining the tiles that are available for that coordinate reference system. Each tile matrix specifies:

a. The scale of the tiles as a scale denominator. The scale denominator is defined with respect to a "standardized rendering pixel size" of 0.28 mm. Frequently, the true pixel size is unknown and 0.28 mm is a common actual size for current displays.
b. The width and height of each tile in pixels.
c. The top left (minimum x, maximum y) corner of the bounding box of the tile matrix (i.e., the CRS coordinates of the top left corner of the top left pixel of the top left tile).
d. The width and height of the tile matrix in tile units (i.e., number of tiles).

Each tile in a tile matrix is identified by its TileCol and TileRow indices that have their 0,0 origin in the tile next to the top left corner of the tile matrix and that increases towards the right and towards the bottom respectively.

4.6 ServiceMetadata document

The ServiceMetadata document is the response document of a GetCapabilities request in procedure oriented architectural style or of a standard request to the right endpoint in a resource oriented architectural style. It is the entry point resource that represents the resources available on the service and communication requirements for the service. Metadata include:

ServiceIdentification
Capabilities (Abstract)
OperationsMetadata
ServiceProvider
Contents
Themes

The tile resource is generally a rectangular image containing cartographic data. Alternatively, this resource might be a non-image representation of the tile such as a description of the tile or a link to the actual image.

The GetTile operation in procedure oriented architectural style allows WMTS clients to request a particular tile of a particular tile matrix set in a predefined format. This operation has some parameters in common with WMS GetMap but it has been deliberately simplified. For instance, only one layer can be retrieved at a time."

5. Web Map Service

5.1 Introduction

The Web Map Service is an international standard for sharing and combining multiple maps. However, it is limited in its functionality because it only treats maps as pictures and not as information. This section directly quotes from text in the OpenGIS® Web Map Server Implementation Specification Vs. 1.3.0, which can be found at http://www.opengeospatial.org/standards/wms.

5.2 Overview

"A Web Map Service (WMS) produces maps of spatially referenced data dynamically from geographic information. This International Standard defines a "map" to be a portrayal of geographic information as a digital image file suitable for display on a computer screen. A map is not the data itself. WMS-produced maps are generally rendered in a pictorial format such as PNG, GIF or JPEG, or occasionally as vector-based graphical elements in Scalable Vector Graphics (SVG) or Web Computer Graphics Metafile (WebCGM) formats.

This International Standard defines three operations: one returns service-level metadata; another returns a map whose geographic and dimensional parameters are well-defined; and an optional third operation returns information about particular features shown on a map. Web Map Service operations can be invoked using a standard web browser by submitting requests in the form of Uniform Resource Locators (URLs). The content of such URLs depends on which operation is requested. In particular, when requesting a map the URL indicates what information is to be shown on the map, what portion of the Earth is to be mapped, the desired coordinate reference system, and the output image width and height.

When two or more maps are produced with the same geographic parameters and output size, the results can be accurately overlaid to produce a composite map. The use of image formats that support transparent backgrounds (e.g. GIF or PNG) allows underlying maps to be visible. Furthermore, individual maps can be requested from different servers. The Web Map Service thus enables the creation of a network of distributed map servers from which clients can build customized maps.

This International Standard applies to a Web Map Service instance that publishes its ability to produce maps rather than its ability to access specific data holdings. A basic WMS classifies its geographic information holdings into "Layers" and offers a finite number of predefined "Styles" in which to display those layers. This International Standard supports only named Layers and Styles, and does not include a mechanism for user-defined symbolization of feature data. This International Standard is applicable to pictorial renderings of maps in a graphical format; it is not applicable to retrieval of actual feature data or coverage data values

5.3 General HTTP request rules

This International Standard defines the implementation of the WMS on a distributed computing platform (DCP) comprising Internet hosts that support the Hypertext Transfer Protocol (HTTP) 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 is implementation-dependent; only the query portion comprising the service request itself is defined by this International Standard.

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

A WMS shall support the GET method of the HTTP protocol. 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. This International Standard defines how to construct a query part that is appended to the URL prefix in order to form a complete request message. Every WMS operation has several mandatory or optional request parameters.

a. http://host[:port]/path[?{name[=value]&}] : URL prefix of service operation. [ ] denotes 0 or 1 occurrence of an optional part; {} denotes 0 or more occurrences.

b. name=value& : One or more standard request parameter name/value.

Each parameter has a defined name. Each parameter may have one or more legal values, which are either defined by this International Standard 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.

5.4 Output formats

The response to a Web Map Service 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 represent a map image. 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 queries for information about features shown on a map.

Allowed map formats are either "picture" formats or "graphic element" formats. Picture formats constitute a rectangular pixel array of fixed size. Picture formats include file types such as Graphics Interchange Format (GIF; MIME type "image/gif"), Portable Network Graphics (PNG; MIME type "image/png"), Joint Photographics Expert Group (JPEG; MIME type "image/jpeg"), all of which can be displayed by common Web browsers, and file types such as Tagged Image File Format (TIFF: MIME type "image/tiff") that may require additional software (beyond a basic Web browser) for display.

Graphic element formats constitute a scale-independent description of the graphic elements to be displayed (including points, lines, curves, text and images), such that the size of the display may be modified while preserving the relative arrangement of the graphic elements. Graphic element formats include Scalable Vector Graphics (SVG; MIME type "image/svg+xml") or Web Computer Graphics Metafile (WebCGM; MIME type "image/cgm;Version=4;ProfileId=WebCGM") formats.

5.5 Coordinate Systems

5.5.1 Map CS

This International Standard uses two principal classes of Coordinate Systems: a Map CS applicable to the map portrayal generated by the WMS, and a Layer CRS for a Bounding Box applied to the source data. During a portrayal operation, a WMS converts or transforms geographic information from a Layer CRS into a Map CS. In addition, a Layer may have an associated vertical, temporal or other coordinate system.

A Map CS is a coordinate reference system for a map produced by a WMS. A WMS map is a rectangular grid of pixels displayed on a computer screen (or a digital file that could be so displayed). The Map CS has a horizontal axis denoted i, and a vertical axis denoted j. i and j shall have only nonnegative integer values. The origin (i,j) = (0,0) is the pixel in the upper left corner of the map; i increases to the right and j increases downward.

The Map CS is identified by the label "CRS:1". The usual orientation of the Map CS shall be such that the i axis is parallel to the East-to-West axis of the Layer CRS and increases Eastward, and the j axis is parallel to the North-to-South axis of the Layer CRS and increases Southward. This orientation will not be possible in some cases, as (for example) in an orthographic projection over the South Pole. The convention to be followed is that, wherever possible, East shall be to the right edge and North shall be toward the upper edge of the Map CS.

The WIDTH and HEIGHT parameters used in the GetMap request and by inclusion in the GetFeatureInfo request correspond to i and j as follows:

a. WIDTH denotes the size of the map image in pixels along the i axis (that is, WIDTH-1 is the maximum value of i).

b. HEIGHT denotes the size of the map image in pixels along the j axis (that is, HEIGHT-1 is the maximum value of j).

The I and J parameters used in the GetFeatureInfo request denote integer values along the i and j axes, respectively, of the Map CS.

5.5.2 Layer CRS

A Layer CRS is a horizontal coordinate reference system for the geographic information that serves as the source for a map. Many Layer CRSs are possible. A Layer CRS appears in the following entities relevant to the WMS:

a. The <BoundingBox> element in the service metadata;

b. The CRS parameter in the GetMap request;

c. The CRS parameter in the map request part of the GetFeatureInfo request.

A WMS must support at least one CRS, and maps from multiple servers may be overlaid only if all the selected servers have at least one CRS in common. This International Standard does not mandate support for any particular Layer CRS(s). Instead, it only defines how CRSs are identified and discusses several optional Layer CRSs. Map providers may support the CRSs that are most useful and appropriate to their geographic locale or information community. To maximize interoperability among servers, providers should also support geographic coordinates by geocentric coordinate systems such as "CRS:84", "EPSG:4326" or other ITRF-based systems.

Every Layer CRS has an identifier that is a character string. Two types of Layer CRS identifiers are permitted: "label" and "URL" identifiers:

a. Label: The identifier includes a namespace prefix, a colon, a numeric or string code, and in some instances a comma followed by additional parameters. This International Standard defines three namespaces: CRS, EPSG and AUTO2.

b. URL: The identifier is a fully-qualified URL that references a publicly-accessible file containing a definition of the CRS that is compliant with ISO 19111.

The Layer CRS has two axes, denoted x and y. The x axis is the first axis in the CRS definition, the y axis is the second axis. Depending on the particular CRS, the x axis may or may not be oriented West-to-East, and the y axis may or may not be oriented South-to-North. The WMS portrayal operation shall account for axis order, origin and direction in the Layer CRS when projecting geographic information from a Layer CRS to the Map CS.

Coordinates shall be listed in the order defined by the CRS and shall be mapped appropriately to the Map CS i and j axes, swapping axis order as needed during the projection operation. Many projected coordinate reference systems have an axis and coordinate order other than easting, northing. Most coordinate reference systems are orientated with one axis positive east (easting) and the other axis positive north (northing). These map conveniently to the bounding box i and j axes, respectively. However, some CRSs have coordinates incrementing in other directions. Tests for valid bounding box areas shall recognise and account for the positive orientation of the CRS axes.

In a geographic CRS, latitude shall have values within the range [-90, 90] and longitude shall have values within the range [-180, 180] degrees or equivalent if the CRS definition is in other units. See 7.3.5 regarding the projection of Layer CRS that is a geographic CRS into the Map CS. When the CRS code specifies a geographic 2D coordinate reference system with axes in units other than degrees or in a degree representation other than decimal degrees the representation shall be converted to decimal degrees.

5.5.3 Bounding boxes

Bounding box values specify the portion of the Earth to be mapped through two pairs of coordinates in a specified Layer CRS. The first pair specifies the minimum coordinate values in the Layer CRS, the second specifies maximum coordinate values. Although for most CRSs with axes incrementing to the east and north this would be

the lower left and upper right corners of the area of interest, the minimum and maximum values might be at other points in some instances.

5.5.4 Vertical CRS

Some geographic information may be available at multiple elevations (for example, ozone concentrations at different heights in the atmosphere). A WMS may announce available elevations in its service metadata, and the GetMap operation includes an optional parameter for requesting a particular elevation. A single elevation or depth value is a number whose units, and the direction in which ordinates increment, are declared through a onedimensional vertical CRS. Depending on the context, elevation values may appear as a single value, a list of values, or an interval.

5.5.5 Temporal CS

Some geographic information may be available at multiple times (for example, an hourly weather map). A WMS may announce available times in its service metadata, and the GetMap operation includes a parameter for requesting a particular time. Depending on the context, time values may appear as a single value, a list of values, or an interval.

5.5.6 Other coordinate systems

Some geographic information may be available at other dimensions (for example, satellite images in different wavelength bands). The dimensions other than the four space-time dimensions are referred to as "sample dimensions". A WMS may announce available sample dimensions in its service metadata, and the GetMap operation includes a mechanism for requesting dimensional values. Each sample dimension has a Name and one or more valid values.

5.6. Web Map Service operations

5.6.1 Introduction

The three operations defined for a WMS are GetCapabilities, GetMap, and GetFeatureInfo. GetFeatureInfo is optional. This clause specifies the implementation and use of these WMS operations in the Hypertext Transfer Protocol (HTTP) Distributed Computing Platform (DCP).

5.6.2 GetCapabilities

The purpose of the mandatory GetCapabilities operation is to obtain service metadata, which is a machinereadable (and human-readable) description of the server's information content and acceptable request parameter values.

The general form of a WMS request is as defined for GET above. When making the GetCapabilities request of a WMS server, which may offer other service types as well, it is necessary to indicate that the client seeks information about the Web Map Service in particular. Thus, the SERVICE parameter of the request shall have the value "WMS".

5.6.3 MetadataURL

A server should use one or more <MetadataURL> elements to offer detailed, standardized metadata about the data corresponding to a particular layer. The "type" attribute indicates the standard to which the metadata complies. Two "type" attribute values are defined by this International Standard: the value "ISO 19115:2003" refers to ISO 19115:2003; the value "FGDC:1998" refers to FGDC-STD-001-1998 [1]. An information community may define meanings for other "type" attribute values. The enclosed <Format> element indicates the file format MIME type of the metadata record.

5.6.4 Layer attributes

A <Layer> may have zero or more of the following XML attributes: queryable, cascaded, opaque, noSubsets, fixedWidth, fixedHeight. All of these attributes are optional and default to 0. Each of these attributes can be inherited or replaced by subsidiary layers.

The mandatory BBOX parameter allows a client to request a particular Bounding Box. The value of the BBOX parameter in a GetMap request is a list of comma-separated real numbers in the form "minx,miny,maxx,maxy". These values specify the minimum X, minimum Y, maximum X, and maximum Y values of a region in the Layer CRS of the request. The units, ordering and direction of increment of the x and y axes are as defined by the Layer CRS. The four bounding box values indicate the outside limits of the region. The relation of the Bounding Box to the map pixel matrix is that the bounding box goes around the "outside" of the pixels of the map rather than through the centres of the map's border pixels. In this context, individual pixels represent an area on the ground.

The mandatory FORMAT parameter states the desired format of the map. Supported values for a GetMap request on a WMS server are listed in one or more <Request><GetMap><Format> elements of its service metadata. The entire MIME type string in <Format> is used as the value of the FORMAT parameter. There is no default format. In an HTTP environment, the MIME type shall be set on the returned object using the Content-type entity header.

The mandatory WIDTH and HEIGHT parameters specify the size in integer pixels of the map to be produced. The Map CS applies to the map. WIDTH-1 specifies the maximum value of the i axis in the Map CS, and HEIGHT-1 specifies the maximum value of the j axis in the Map CS. If the request is for a picture format, the returned picture, regardless of its MIME type, shall have exactly the specified width and height in pixels. In the case where the aspect ratio of the BBOX and the ratio width/height are different, the WMS shall stretch the returned map so that the resulting pixels could themselves be rendered in the aspect ratio of the BBOX. In other words, it shall be possible using this definition to request a map for a device whose output pixels are themselves non-square, or to stretch a map into an image area of a different aspect ratio."

5.7 OpenGIS® Web Map Services - Profile for EO Products

5.7.1 Introduction

This is an application profile for remote sensing data, and is directly quoted fromWeb Map Services - Application Profile for EO Products (0.3.3), which can be found at http://www.opengeospatial.org/standards/wms.

5.7.2 Overview

"The WMS configuration proposed in this profile is intended to support the interactive visualization and evaluation of Earth Observation (EO) data products. The profile describes a consistent Web Map Server (WMS) configuration that can be supported by any content providers (satellite operators, data distributors …), most of whom have existing (and relatively complex) facilities for the management of these content.

This application profile document describes how:

a. WMS layers, sample dimensions and nested layers can be used to manage the hierarchy of EO product information from the collection level, down through individual products to the sample and quality bitmask dimensions;

b. The WMS GetMap operation can be used to support interactive browse and the full evaluation of sample coverage and quality information.

The acquisition of EO data is subject to a range of factors which effect the degree to which a product is fit for purpose. For optical EO these factors include cloud cover, snow cover, atmospheric aerosols, low illumination angles, sun glint off the ocean or ice surface and suspended sediment in the water column. For imaging radar the usefulness of products can be affected by wind speed, presence of surfactants and soil moisture.

A number of WMS features make them a good choice for EO data providers wanting to open up their product inventories for evaluation:

a. Open Web service specifications are widely and simply supported using common place Web browsers;

b. Web maps provide a geocoded view of data frequently acquired as ungeocoded scan lines;

c. The WMS interface supports roam and zoom functionality not available with a simple image thumbnail;

d. Use of high levels of image compression reduces bandwidth requirements relative to WCS whilst allowing effective visual inspection;

e. The serving of products via highly compressed images ensures that the original data is protected from unauthorised use.

Unfortunately, practical application of WMS technology is hindered by two issues:

a. Inconsistent interpretation of the WMS standard in the context of EO data;

b. Server instances are typically configured to return just a single representation using a sub-set of the dataset. Some or even most, of the dataset is not made available for evaluation.

To illustrate the first point, consider the use of WMS layers. A layer is defined as the basic unit of geographic information that may be requested as a map from a server. A literal interpretation would be that each individual EO data granule (dataset) should be handled as a WMS layer. With large inventories of EO products (>10,000) this results in bloated responses to GetCapability requests. An alternative interpretation is that a WMS layer should represent a collection of datasets sharing the same product specification (i.e. a dataset series). In this context, individual datasets can be retrieved and presented using their acquisition time.

Whilst both interpretations can be justified, the resulting WMS server configurations are incompatible, seriously hindering the client to server and community wide interoperability. Meeting the twin challenges of true interoperability and the full evaluation of the spatial elements of EO products (interactive browse), requires the definition of a consistent, constrained, interpretation of the WMS standard.

The definition of a consistent, constrained, interpretation of the WMS standard has two aspects:

a. A WMS metadata model of the spatial data and metadata structure of EO datasets;

b. A defined set of WMS server responses to GetMap requests.

5.7.3 Metadata Model

The metadata model provides the basic framework, the skeleton, of a WMS instance. Defining a WMS metadata model for EO products is therefore a prerequisite for an application profile that is unambiguous whilst allowing interactive viewing of all of the spatial information within EO products.

5.7.4 GetMap Request/Response

Where the service metadata model provides the skeleton for this WMS profile, the definition of the GetMap request/response adds the functional muscle that ultimately delivers the rendered image maps back to the client application of a user."

6. Web Feature Service

6.1 Introduction

The Web Feature Service allows greater access to geographic information in spatial data files than is permitted by the Web Map Tile Service and Web Map Service. This section directly quotes from text in the OpenGIS Web Feature Service 2.0 Interface Standard, found at http://www.opengeospatial.org/standards/wfs.

6.2 Overview

"The Web Feature Service (WFS) represents a change in the way geographic information is created, modified and exchanged on the Internet. Rather than sharing geographic information at the file level using File Transfer Protocol (FTP), for example, the WFS offers direct fine-grained access to geographic information at the feature and feature property level. Web feature services allow clients to only retrieve or modify the data they are seeking, rather than retrieving a file that contains the data they are seeking and possibly much more. That data can then be used for a wide variety of purposes, including purposes other than their producers' intended ones.

In the taxonomy of services defined in ISO 19119, the WFS is primarily a feature access service but also includes elements of a feature type service, a coordinate conversion/transformation service and geographic format conversion service.

6.3 Operations

This International Standard specifies the behaviour of a service that provides transactions on and access to geographic features in a manner independent of the underlying data store. It specifies discovery operations, query operations, locking operations, transaction operations and operations to manage stored parameterized query expressions.

a. Discovery operations allow the service to be interrogated to determine its capabilities and to retrieve the application schema that defines the feature types that the service offers.

b. Query operations allow features or values of feature properties to be retrieved from the underlying data store based upon constraints, defined by the client, on feature properties.

c. Locking operations allow exclusive access to features for the purpose of modifying or deleting features.

d. Transaction operations allow features to be created, changed, replaced and deleted from the underlying data store.

e. Stored query operations allow clients to create, drop, list and described parameterized query expressions that are stored by the server and can be repeatedly invoked using different parameter values.

This International Standard does not address the access control issues.

6.4 Elements

Elements that are used to describe each feature type listed within the wfs:FeatureTypeList element include:

Name
Title
Abstract
Keywords
DefaultCRS
OtherCRS
NoCRS
OutputFormats
WGS84BoundingBox
MetadataURL
ExtendedDescription

6.5 GetFeature operation

The GetFeature operation returns a selection of features from a data store. A WFS processes a GetFeature request and returns a response document to the client that contains zero or more feature instances that satisfy the query expressions specified in the request. In GML a feature is represented as an XML element. The content of a feature element is a set of elements that describes the feature in terms of a set of properties".

7 Web Coverage Service

7.1 Introduction

The Web Coverage Service offers the most sophisticated form of geospatial data exchange. This document directly quotes from the Web Coverage Service Interface Standard Core 2.0 standard, found at http://www.opengeospatial.org/standards/wcs.

7.2 Overview

"This document specifies a core set of requirements that a WCS implementation must fulfill. WCS extension standards add further functionality to this core; some of these are required in addition to the core to obtain a complete implementation. This document indicates which extensions, at a minimum, need to be considered in addition to this core to allow for a complete WCS implementation. This core does not prescribe support for any particular coverage encoding format. This also holds for GML as a coverage delivery format: while GML constitutes the canonical format for the definition of WCS, it is not required by this core that a concrete instance of a WCS service implements the GML coverage format. WCS extensions specifying use of data encoding formats in the context of WCS are designed in a way that the GML coverage information contents specified in this core is consistent with the contents of an encoded coverage.

The OGC Web Coverage Service (WCS) supports electronic retrieval of geospatial data as "coverages" – that is, digital geospatial information representing space/time-

varying phenomena. This document specifies the WCS core; every implementation of a WCS shall adhere to this standard. This standard defines only basic requirements. Extensions to the core define extensions to meet additional requirements, such as the response encoding. Indeed, additional extensions are required in order to completely specify a WCS for implementation.

7.3 Comparisons

A WCS provides access to coverage data in forms that are useful for client-side rendering, as input into scientific models, and for other clients. The WCS may be compared to the OGC Web Feature Service (WFS) and the Web Map Service (WMS). As WMS and WFS service instances, a WCS allows clients to choose portions of a server's information holdings based on spatial constraints and other query criteria.

Unlike WMS, which returns spatial data to be portrayed as static maps (rendered as pictures by the server), the Web Coverage Service provides available data together with their detailed descriptions; defines a rich syntax for requests against these data; and returns data with its original semantics (instead of pictures) which may be interpreted, extrapolated, etc., and not just portrayed.

Unlike WFS, which returns discrete geospatial features, the Web Coverage Service returns coverages representing space/time-varying phenomena that relate a spatio-temporal domain to a (possibly multidimensional) range of properties. As such, WCS focuses on coverages as a specialized class of features and, correspondingly, defines streamlined functionality.

7.4 Specifications

This document specifies how a Web Coverage Service (WCS) offers multi-dimensional coverage data for access over the Internet. This document specifies a core set of requirements that a WCS implementation must fulfil. WCS extension standards add further functionality to this core; some of these are required in addition to the core to obtain a complete implementation. This document indicates which extensions, at a minimum, need to be considered in addition to this core to allow for a complete WCS implementation.

This core does not prescribe support for any particular coverage encoding format. This also holds for GML as a coverage delivery format: while GML constitutes the canonical format for the definition of WCS, it is not required by this core that a concrete instance of a WCS service implements the GML coverage format. WCS extensions specifying use of data encoding formats in the context of WCS are designed in a way that the GML coverage information contents specified in this core is consistent with the contents of an encoded coverage.

7.5 Coverage Data Model

For reasons of extensibility and flexibility, many components of the core GML structure CoverageOfferings introduced in this standard are left underspecified (e.g., in terms of multiplicity of the elements or proper semantics and use of a component), some-times even in cases where other standards applying (like OWS Common and

SWE Common) leave such details open. Any item thus underspecified can be handled arbitrarily by implementations – among others, a server is free to deliver optional elements or not, and a client is free to ignore optional elements when present. Note, however, that WCS extensions may regulate further syntax and semantics of such underspecified items.

7.6 ServiceMetadata

WCS ServiceMetadata provide service details plus information about the concrete service capabilities of the WCS service as a whole (as opposed to coverage-specific Service-Parameters). The WCS ServiceMetadata component extension OGC 09-110r4 defines a canonical place for additional information provided, e.g., by WCS extension standards.

7.7 WCS operation types

The WCS interface herein specified supports retrieval of geospatial coverage data – that is, digital geospatial information representing space/time-varying phenomena. To this end, the WCS interface specifies the following operations that may be invoked by a WCS client and performed by a WCS server:

a. GetCapabilities – This operation allows a client to request information about the server's capabilities and coverages offered.

b. DescribeCoverage – This operation allows a client to request detailed metadata on selected coverages offered by a server.

c. GetCoverage – This operation allows a client to request a coverage comprised of selected range properties at a selected set of spatio-temporal locations, expedited in some coverage encoding format.

7.8 WCS operations

7.8.1 GetCoverage operation

A GetCoverage request prompts a WCS service to process a particular coverage selected from the service's offering and return a derived coverage. The WCS Core standard defines the domain subsetting operation which delivers all data from a coverage inside a specified request envelope ("bounding box"), relative to the coverage's envelope – more precisely, the intersection of the request envelope with the coverage envelope.

7.8.2 Domain subsetting

Domain subsetting is subdivided into trimming and slicing. A trim operation identifies a di-mension and a lower and upper bound (which both must lie inside the coverage's domain) and delivers a coverage whose domain, in the dimension specified, is reduced to these new, narrower limits. The result coverage's dimension is identical to that of the input coverage.

A domain slice operation receives a dimension and a position (which must lie inside the coverage's domain) and delivers a coverage which is a slice of the offered coverage obtained at the cutting position provided. The dimension of the result coverage is reduced by one as compared to the original coverage.

Both trimming and slicing can be combined arbitrarily in a request and on as many dimensions as desired. However, per request at most one operation can be applied per dimension.

For trimming a coverage in a particular dimension, the corresponding dimension name is indicated as well as the lower and upper bound of the resulting coverage. Both lower and upper bound are optional. A lower bound omitted shall be substituted in the server by the coverage's lower bound in the dimension on hand, an upper bound omitted shall be substituted in the server by the coverage's upper bound. The result coverage shall contain only those range values of the original coverage which lie within the effective lower and upper bound, obtained as described.

A GetCoverage request may contain several subsetting operations; trimming and slicing operations may be combined in a single request in any sequence."

8. Web Coverage Processing Service

8.1 Introduction

The Web Coverage Processing Service is a tool for extracting, processing and analysing geospatial data. This section directly quotes from the OGC® Web Coverage Processing Service (WCPS) Language Interface Standard Vs. 1.0.0, which can be found at http://www.opengeospatial.org/standards/wcps.

8.2 Overview

"The Web Coverage Processing Service defines a protocol-independent language for the extraction, processing, and analysis of multi-dimensional coverages representing sensor, image, or statistics data.

Services implementing this language provide access to original or derived sets of geospatial coverage information, in forms that are useful for client-side rendering, input into scientific models, and other client applications.

Like WCS, WCPS is currently limited to quadrilateral grid coverages, providing information at the grid points, usually with interpolation between these grid points.

8.3 Coverage model

8.3.1 Coverage

A coverage consists of a set of locations bearing some value. Following the mathematical notion of a function that maps elements of a domain (here: spatio-temporal coordinates) to a range (here: "pixel", "voxel", … values), the set of coverage locations bearing values is called the coverage domain while the set of

possible values, i.e., the coverage value data type, is called the coverage range.

A coverage domain with its set of locations (or coordinates) is aligned along some d dimensional grid where d>0 is called the coverage's dimensionality. The coordinate space, i.e. the set of all possible coordinates, is spanned by d independent dimension axes. A dimension axis (abbreviated also as dimension or as axis) is identified by its name which is unique within the coverage.

8.3.2 Dimensions

Each dimension has an dimension type associated. A coverage can have at most one x, y, z, and t dimension. Dimension types include:

x East-West extent, expressed in the coverage's CRS
y North-South extent, expressed in the coverage's CRS
z Geographical elevation, i.e., height or depth
t Time

Each dimension shall have one or more coordinate reference systems (CRSs) associated, one of them being – according to WCS – either an ImageCRS or a GridCRS (henceforth collectively termed the Image CRS). Any number of further CRSs can be associated with a coverage dimension.

8.3.3 Locations

A location L is a set consisting of dimension names, per dimension the coordinate system used, and a coordinate relative to this dimension and CRS; each of the coverage's dimension name shall appear exactly once in this set. The set DimensionPointValues is a generalization of numeric and string values that allows to express all kind of coordinates, including geographic floating-point coordinates and date/time strings. Along each dimension a coverage is delimited by a lower and upper bound value, these border values being part of the coverage extent. Location addresses always are relative to a particular coverage.

8.3.4 Domain

The set of all locations contained in a coverage forms its domain.

8.4 Probing Functions

A set of so-called probing functions allows to extract the constituents listed above from a given coverage. Coverage characteristics covered by these functions include:

Identifier
Grid point values
Domain dimension list
Domain dimension type
Image CRS
Domain extent of coverage, expressed in Image CRS
Domain extent of coverage along dimension, expressed in Image CRS

Extent of coverage along dimension, expressed in arbitrary CRS
Range data type
Range field type
Range field name set
Null value set
Default interpolation method
Interpolation method set
Interpolation type
Null resistance

8.5 Coverage Processing Language

Expressions include:

a. processCoveragesExpr

The processCoveragesExpr element processes a list of coverages in turn.

b. getMetaDataExpr
The getMetaDataExpr element extracts a coverage description element from a coverage.

The following metadata extraction functions are defined for a coverage C:

identifier(C)
imageCrs(C)
imageCrsDomain(
crsSet(C) crsSet(C)
domain(C,a,c)
nullSet(C)
interpolationDefault(C,r) i
interpolationSet(C,r)

c. subsetExpr

The subsetExpr element specifies spatial and temporal domain subsetting.
It encompasses spatial and temporal trimming (i.e., constraining the result coverage domain to a subinterval), slicing (i.e., cutting out a hyperplane from a coverage), extending, and scaling of a coverage expression. All of the subsetExpr elements allow to make use of coordinate reference systems other than a coverage's image CRS.

d. trimExpr

The trimExpr element extracts a subset from a given coverage expression along the dimension indicated, specified by a lower and upper bound for each dimension affected. Interval limits can be expressed in the coverage's image CRS or any CRS which the coverage supports.

e. extendExpr

The extendExpr element extends a coverage to the bounding box indicated.

f. sliceExpr

The sliceExpr element extracts a spatial slice (i.e., a hyperplane) from a given coverage expression along one of its dimensions, specified by one or more slicing dimensions and a slicing position thereon."

9. Conclusions

The OGC standards summarized here appear to correspond to the IVOA standards for the working groups dealing with Semantics (*UCD Controlled Vocabulary*), Registry (*Resource Metadata; Registry Interfaces)*, Applications (*VOTable Format)* and Data Access Layer (*Simple Search)*. However, they are not exact equivalents, for obvious reasons, and coverage of *VOTable Format* for tabular information deserves further examination. Semantics (*Vocabularies*) are less visible in these standards.