

A case study of new operations/interfaces/services

JeffreyGHarrison Tuesday, May 22, 2001 10:55 AM

If we want to add a new GetView Operator to the Web Map Server Interface, do we need new elements in the Capabilities documents. Is the change just adding GetView? What are the "gotchas"??

John Davidson Wed, 30 May 2001 10:02:46 -0400

The Military Pilot Project Phase 1 (MPP1) has identified two interesting service design issues that I think deserve further and broader discussion. I hope that by discussing these cases we achieve two objectives: a) a better understanding and a possibly a refinement of the OGC Service Model and b) a way forward for MPP1 that matches the OGC Service Model scheme. (Disclaimer: these are proposed new operations that are emerging out of the IP process. Their origins can be traced to the MPP1 sponsors' requirements for Web accessible services.)

Case 1: GetView operation (<http://alpha.skylinesoft.com/services/ogc/wMS/test.html>)

The GetView operation generates a Terrain View... a picture of a place that represents the normal viewpoint of a person (i.e., a three-dimensional, perspective view of the world). For MPP1 this operation is currently implemented as an extension of the WMS interface. This seems easy and, as a practical matter, allows the implementers to reuse all the interfaces of their existing WMS implementation. Extending the WMS interface for a GetView operation, however, requires modification of the WMS GetCapabilities response schema.

Alternatively, a GetView operation could form the basis of a new OGC Web Service interface. In this case, an entirely different interface from that defined by WMS would be specified (although many of the WMS elements would presumably be "reused" including the GetCapabilities operation and the basic types and structure of its response schema. These "reusable" elements would be defined as part of BSM).

Trying my best to use the B' Service Model definitions (<http://www.intl-interfaces.net/servicemodel/vote-definitions.html>), the question is this: Should GetView a) extend the existing WMS interface with an additional (but optional) operation, b) form the basis of an entirely new interface that can be used to extend the WMS Web Service and presumably other services or c) define an entirely new kind of OGC Web Service? It seems that sound arguments can be made for any of these approaches. Does the OGC Service Model inform us of the preferred approach for extending existing interfaces, services or defining new ones?

Case 2: Annotation Operations: AddText, AddPolygon, AddPolyline
(<http://alpha.skylinesoft.com/services/ogc/wMS/test.html>)

These proposed operations are used to place text or simple geometric objects into a view via an HTTP GET request. Pretty much the same questions as above. Assuming these operations are useful and have been accepted, should they a) extend existing WMS for new portrayal operations, b) form the basis of a new interface that packages specialized operations for portrayal of annotation (e.g., a "simple annotation" interface) or c) define and otherwise implement a new kind of OGC Web Service (e.g., a simple annotation portrayal service)?

For these cases, would you: a) extend the existing WMS service interface with a new set of operations (e.g., GetView, AddText, etc)?

b) define new interfaces (e.g., called "GetView" and "Annotate") that could be used by more than one kind of Web Service, including WMS?

c) define new kinds of OGC Web Services whose role are to 1) support "terrain view" portrayal (e.g., a new Web View Service(WVS)) and

2) support image/map annotation portrayal (e.g., a new Web Annotation Service (WAS))?

d) other?

Jeff de La Beaujardiere Wed, 30 May 2001 10:58:13 -0400 (EDT)

> Case 1: GetView operation

Very interesting. At first glance, I would suggest that this be provided by a new OGC Web Service (Web Viewpoint Service or something) rather than overloading WMS.

The WMS/WCS/WFS paradigm always looks "down" and thereby makes it simple to overlay content from multiple sources without worrying about viewshed obstructions.

> Case 2: Annotation Operations: AddText, AddPolygon, AddPolyline

Doesn't XIMA address some of this (at least the annotation format, if not the Add* operations)? Seems like another distinct OWS as well, for general use by WMS/WVS/etc.

Nuke Goldstein Wed, 30 May 2001 11:21:40 -0400

In the last Terrain View/XIMA TEM we arrived to the conclusion that the XIMA specs are a good base to this, moreover I would even take it a step further and add to the 'GET' method request something similar to the FILTER which provide annotations.

something like:

`http://.../WFS?request=GETMAP&...&addfeatures="<..XIMA-type-data..>"`

off course there would be some work to do on the XIMA specs to achieve that but it seems like a good solution to give the client some more power.

Jeff Lansing Wed, 30 May 2001 09:43:24 -0700

By looking "down" (from an "infinite distance"), the WMS/WCS/WFS paradigm provides a simple reversible transformation of ground coordinates into image coordinates. This allows a client to work in image coordinates directly, in order to do traditional client side tasks such as overlaying, positioning and alignment, annotation, etc.

Apparently there is no equivalent of this transformation for terrain views. So besides being a good reason for defining a new interface, this lack is also a good reason to have server side analogs for the traditional client side tasks.

Jeff Lansing Wed, 30 May 2001 10:01:40 -0700

Once we start placing text and drawing polygons and polylines via these kind of operations, we are also going to need a way to style that text and those drawings. Controlling color is crucial (black text on a black background is useless). Controlling the width of polylines, the filling of polygons, and text size are also very important. Other operations such as line style (dashed, etc) would be useful.

If we are going to provide some of these styling operations (and we have to provide some), wouldn't it be even better to provide a complete, versatile interface?

Craig Bruce Wed, 30 May 2001 13:29:26 -0400

If there is nothing to distinguish this requirement from rendering specific graphics, why not just define an optional `SVG={url}` parameter. Or build an annotation WMS that will render annotations into an image layer that viewer clients can place on top of maps coming from other sources?

Jeff Lansing Wed, 30 May 2001 10:45:01 -0700

Isn't this annotation WMS a lot like the user layers of a SLD enabled WMS? I think the only difference would be that in the case of the annotation WMS the client supplies the geometry, rather than saying which server to get it from. In the case of the Terrain View server, would it be better to have an annotation WMS that can do terrain view transformations, or would it be better to have a Terrain View server that supports operations for annotations? (I think this is related to John D's original questions.)

John Davidson Thu, 31 May 2001 01:30:15 -0400

I think there are probably several valid use-cases for annotating images/maps that the two complimentary approaches (XIMA document and HTTP GET requests) being discussed address. Now Craig's comments have reminded me that there's yet another (as we discussed in the last SLD TEM)... an SLD-based approach to annotation. To get clarity,

perhaps the use cases (i.e., the requirements) for annotation need to be more clearly described.

Just an observation.

But discussing the merits or technical details or even the requirements of these approaches was not the intent of my original message. The nature of my original message was more abstract or philosophical. Jeff has re-stated the question this way: In the case of the Terrain View server, would it be better to have an annotation WMS that can do terrain view transformations, or would it be better to have a Terrain View server that supports operations for annotations?

So far I heard Jeff DLB say he prefers not to overload WMS with terrain view and annotation operations (and therefore implement them as separate service(s)).

I "heard" Craig Bruce say that he'd prefer to implement annotation operations as an SLD-enabled WMS. Didn't hear anything about GetView.

I heard Jeff Lansing say that they should be separate service(s).

Have I characterized your positions accurately?

Ron Lake Thu, 31 May 2001 09:49:10 -0700

It seems to me that there are two aspects to annotation and they should not be confused.

One view of annotation as content is really a kind of metadata description. For example in RDF (and in some of the GML: 3.0 proposals) one can make comments about other data (or persons or services or anything.) The proposed GML 3.0 syntax for this is:

```
<gml:Description of = "uri">
  <gml:property1> .. </gml:property1>
  ...
  <gml:propertyn> .. </gml:propertyn>
</gml:Description>
```

Here we have a description of some feature or coverage or piece of same. These properties could include identification of type or more or less any "comment". This could include geometric properties as well and the uri could point to an image or to an element (area of interest) that refers to an area within an image.

Another view of annotation as content is to denote a set of feature hypotheses (e.g. I think this is a nuclear fuel processing plant). In this case we can represent these hypothesized features in GML (this is the XIMA view) and on each feature provide a link to the data source (image, voice record, text document) from which it is derived.

Both these content views are consistent with one another.

The other view of annotation is as a portrayal issue. This must be dealt with VERY carefully. This comes down to a portrayal of the annotation content in the senses expressed above. For example - I have some "statements" that I want to make about an area of interest in an image - this can clearly be stated as a Description as above without regard to its appearance. Now I want to say that I would like these comments to appear as text with an arrow pointing to the area of interest, AND that I want the area of interest to be denoted by a red circle.

To capture this I need a call-out style that I apply to the Description content (for an example of call-out styles look at the Call-Outs in Microsoft PowerPoint under Autoshapes | Callouts.

I think these issues need to clear in our minds before we proceed with looking at annotation as part of this spec or another.

Jeff Lansing Thu, 31 May 2001 10:22:31 -0700

The content of an annotation is not necessarily a statement in some language. For example, in the following:

{Image}

There are two kinds of annotations, a scale marker, and some scaled and numbered pointers to other images which are more detailed cross sections of the thing that this is an image of. Whatever the content of these annotations is, I think it is not either metadata description or hypothesis.

Raj Singh Thu, 31 May 2001 13:35:23 -0400

We have kept GetView closely linked to GetMap because GetMap is really a special case of GetView. It is becoming apparent that aside from any technical considerations, working with terrain views--from both the client and server perspective--is quite different. I think as we go forward there will be very little reason to maintain a link between plan and perspective views, other than the level of commonality that the Basic Services Model provides.

John Davidson Thu, 31 May 2001 15:26:58 -0400

So you're saying that the GetView operation should be separated from the "WMS interface" (in the B' sense of "interface"), correct? What operations would a "Terrain View interface" require in addition to GetView and GetCapabilities? What BSM things do they share?

Joshua Lieberman Thu, 31 May 2001 23:35:41 -0400

could be specialized and b) services (implementations) could support more than one interface. That was the idea at least. This work was unfinished pending further discussion of the BSM and research into mechanisms for specifying Web-based interfaces for services.

Raj Singh Fri, 01 Jun 2001 10:58:39 -0400

Regarding BSM things that GetView would use, most of the service request parameters apply, such as version, format, exceptions, srs, and vendor-specific parameters. We currently have HTTP GET request support, and foresee an HTTP POST request in the near future to fully support annotations, either through SLD and/or XIMA.

Ron Lake Fri, 01 Jun 2001 08:58:50 -0700

Development of a service type hierarchy is extremely important and I think one of the key challenges of the OGC going forward.

The BSM should NOT define, however, any geo-specific specific services but rather provide the language for service description, and perhaps the root of the service tree and the rules for the definition of a service hierarchy. Am I missing something?

John Davidson Fri, 1 Jun 2001 12:08:56 -0400

Nope. Except that I think you'll find that the BSM document (OGC project document 01-022r1) pretty much already does as you suggest.

Lou Reich Fri, 1 Jun 2001 12:13:18 -0400

I think B' had the most votes but there was no consensus and certainly not enough votes to think we are developing a new BSM (I have talked to Allan Doyle and he doesn't think we have come to any conclusions. The issues you are talking about and the generalization and specialization hierarchies for interfaces are key to getting an intelligent view of OGC services. In a classic sense the WMS could be viewed as a specialization of the WVS and anyone searching a catalog for appropriate WMS instances should also receive appropriate WVS instances.

I think the discussions of MPP1 services are a good real world instance of the abstract service discussion and we need to carefully look at the implications of each option before documenting results in the Basic Service Model.

PS: I still vote for A. I think it would better handle this case.

John Davidson Fri, 1 Jun 2001 12:39:11 -0400

Thanks for your comments. As I said when I opened this thread, the purpose was to bring these real-world issues up and out for broader discussion by the TC. I hope this

discussion will continue and culminate with some concrete conclusions (or at least next-steps) at the TC next week!

One more thought. There's a two-way interaction that has to occur between the TC and the IP initiatives. As you know, the IP process is designed to percolate (implement and test) new/emerging interfaces and technologies from which may come proposals (discussion papers, IPRs and DIPRs, RFCs, etc) for consideration by the TC. On the other hand, I believe the IP needs specific guidance from the TC on the abstract models to use (and specifically now the Service Model/BSM). It is imperative that we (the OGC) come to a consensus understanding of the Service Model so that the current and upcoming IP initiatives can be as productive and fruitful to the OGC as possible.