J11.4 PROTOTYPING INSIGHTS: *A WEB COVERAGE SERVICE DESIGNED FOR A CONSOLIDATED NEXRAD AND TDWR SERVICE FOR NEXTGEN WEATHER*

Keith Bourke*, Russ Sinclair, Mike Haertjens Harris Corporation, Melbourne, Florida

1. INTRODUCTION

As the volume of weather data is increasing, the question "What weather products should we provide?" is becoming a question of "How should we provide access to weather data"? The NextGen vision describes how new technologies including Web Services and Service Oriented Architecture (SOA) will help transform the weather domain and provide benefit to users beyond the traditional NAS. This paper describes how Harris Corporation established a NextGen-like, net-centric environment in our Weather Research Lab and used it to implement and evaluate OGC standards-based Web Coverage Services (WCS) request/response and SOA technologies publish/subscribe messaging. WCS design and lessons learned are discussed in detail.

In order to put the description of the Weather Lab NextGen net centric environment and the discussion of WCS request/response in context, this paper first describes a notional NextGen net-centric concept of operations (CONOPS) for the Weather Domain and then shows how the Harris Weather Research Lab and the WCS implementation fits into it.

2. NOTIONAL NEXTGEN NET-CENTRIC WEATHER DOMAIN CONCEPT OF OPERATIONS

In the NextGen vision, net-centricity within the FAA Enterprise is accomplished by establishing a SOA relationship between Data Providers and Data Consumers. In this SOA concept, Data Providers are Services that supply data to other Services within the Enterprise that use the data in some way, e.g., to make decisions or combine data from multiple sources to add value. Services that use data are referred to as Data Consumers. A System Wide Information Management System (SWIM) will be used as the foundation of the SOA in the FAA Enterprise. The SWIM will facilitate data sharing between weather Data Providers and Data Consumers by coordinating access to the weather Registry, Repository, Data Catalog and Messaging Services provided by the NextGen Network Enabled Weather (NNEW) system. The weather Data Catalog is a portal that presents all weather data (pub/sub products and weather services; e.g. WCS, WMS, etc.) available to Data Consumers. The Data Catalog accesses and presents information contained in the Registry and Repository to Data Consumers so they can get a comprehensive view of each weather data element. Data Consumers will use the weather Data Catalog to discover and access weather data available within the FAA Enterprise. An FAA Enterprise Governance organization will manage SWIM and control the sharing of data by the various domains across the enterprise, e.g., the Weather Domain, the Surveillance Domain and Trajectory Based Operations (TBO) Domain. The Weather Domain is made up of all the systems, applications, governance, etc., that are specifically needed to carry out the weather mission within the NAS. Key Weather Domain components include Weather Data Governance and Weather Data Providers. [Note: Weather Data Providers may also be Weather Data Consumers. That occurs if they add value to data that they first consume and subsequently provide it to consumers.] The Weather Data Governance (WDG) component of the Weather Domain also manages weather data collection, production and dissemination needs and coordinates them with the FAA Enterprise and other Domain Governance organizations.

----

* *Author address:* Keith Bourke, Russ Sinclair, Mike Haertjens Harris Government Communications Systems Division, P.O. Box 9800 Melbourne, FL 32902-9800;
e-mail: keith.bourke@harris.com

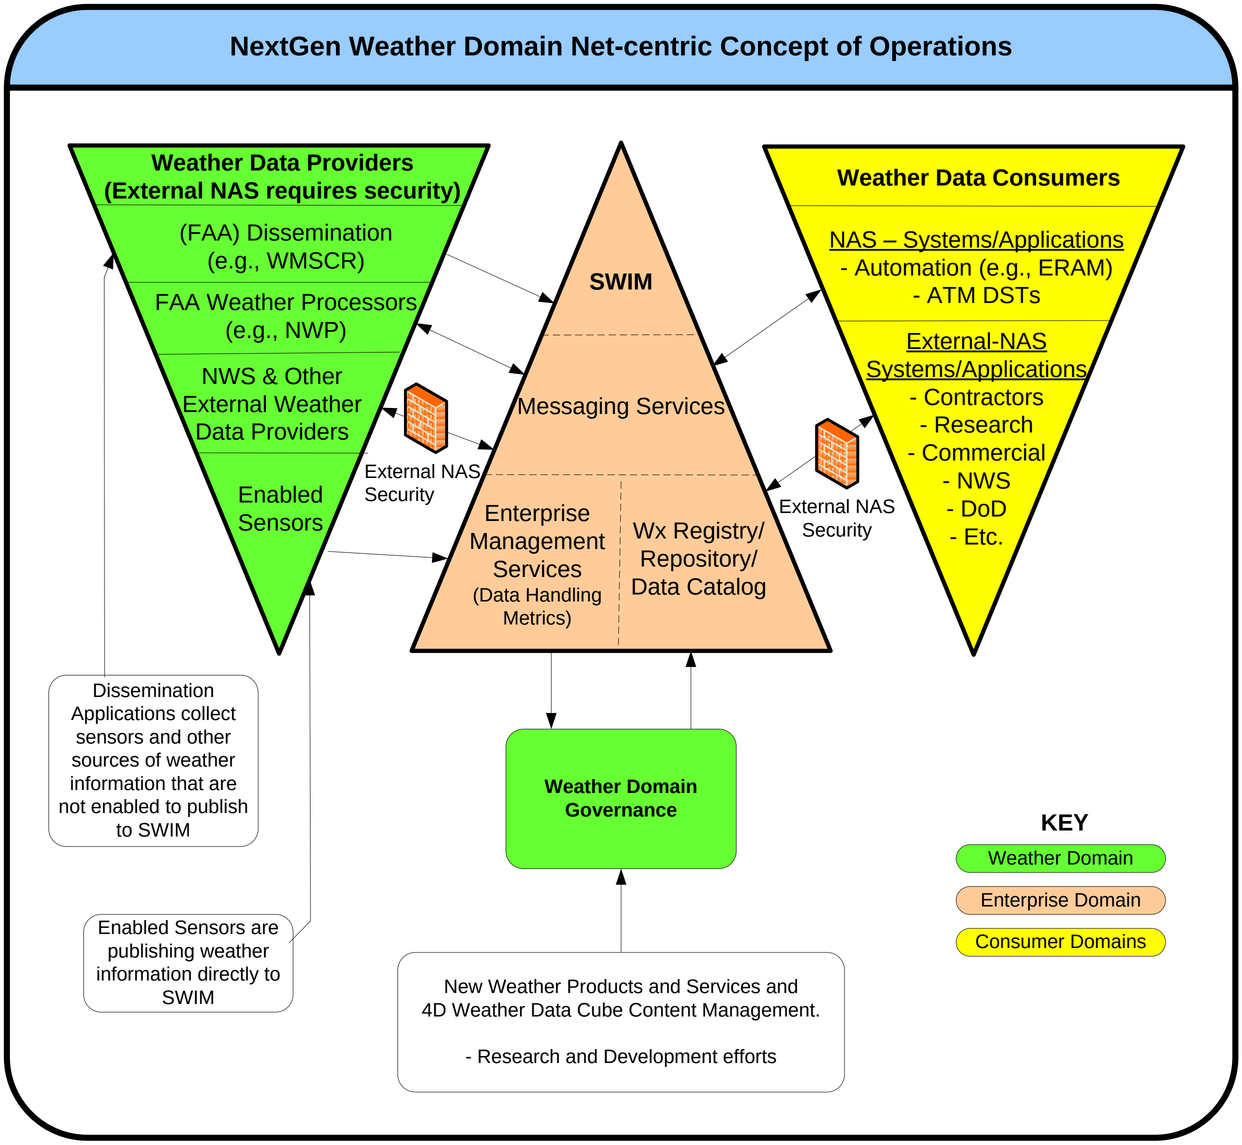

Figure 1 - Notional NextGen Weather Domain Net-Centric ConOps

Figure 1 graphically portrays a notional net-centric CONOPS for weather data sharing in the NextGen era. At the highest level, the green, beige and yellow triangles in Figure 1 respectively represent Weather Data Providers, SWIM, and Weather Data Consumers. Examples of applications representing notional Weather Data Providers, typical SWIM Core Services, and notional Weather Data Consumer applications are identified within the subsections of each triangle. An example of a Weather Data Provider application is the NextGen Weather Processor (NWP). A typical SWIM Core Service is Messaging. An example of a Weather Data Consumer application is ERAM. If a Weather Data Provider, e.g., NWP, also consumes weather data, it is depicted only once as a provider; it is not also depicted as a consumer. The arrows between the triangles in Figure 1 represent provider and consumer interactions with SWIM, as well as data flow from providers to consumers via SWIM.

Weather Data Providers advertise and facilitate access to the data they produce or collect through the weather Registry and Data Catalog services. Weather Data Consumers discover services and data offered by Weather Data Providers by interacting with the weather Registry/Repository and Data Catalog. For example, consumers can discover the “location” of provider services, access methods offered, and the weather products available. It is likely that some of the consumer interaction will happen at consumer application “design time” rather than at runtime. Data from Weather Data Providers is accessed by Weather Data Consumers through publish/subscribe and request/response messaging services. The publish/subscribe service is a very efficient way to share data because it requires that a provider only publish once to have its data consumed many times. WCS is an example of a NextGen era request/response service that would be "discovered" via the weather Data Catalog, be invoked via SWIM and return data to a Weather Data Consumer via SWIM, except for when very large files requested would likely be returned out of band, e.g., via FTP.

The green box in the center lower half of Figure 1 represents Weather Domain Governance (WDG), which controls all aspects of the Weather Domain, including authorization of Weather Data Providers to offer and advertize data and data services for consumers.
The WDG would authorize all changes made to the weather Registry/Repository and Data Catalog. The WDG would coordinate and reconcile service advertisement and data delivery with the NAS Enterprise Governance organization. The Weather Domain Governance could monitor weather data handling at the NAS enterprise level by retrieving metrics collected by the Enterprise Management Service in SWIM; metrics on processing and data handling parochial to Weather Data Providers could be published to SWIM, where they could be consumed by the appropriate WDG applications. The WDG could use these metrics to evaluate if products/services need to be expanded, improved or decommissioned. The metrics could also help identify where infrastructure improvements are needed. Additionally, the WDG would have the insight needed to identify and recommend new weather capabilities for research and new development efforts.

The Registry/Repository and Data Catalog resources could be provided by the Enterprise Governance and administrative access and privileges could be granted to the Weather Domain Governance. This approach provides ownership and control of Weather-specific information to the WDG, while minimizing costs related to software licenses and hardware resources. Many aspects of the Registry/Repository and Data Catalog would likely need to be administered by the Enterprise Governance, where in-depth SWIM knowledge is required.

3. HARRIS WEATHER RESEARCH LAB NEXTGEN ENVIRONMENT

Harris constructed a NextGen Net-centric environment that is graphically depicted in Figure 2 in the context of the notional architecture described in the previous section. The green triangle represents Weather Data Providers; the beige triangle represents SWIM; and the yellow triangle represents Weather Data Consumers. Key elements of the lab environment are based on several Harris prototypes that are aligned with the FAA Weather Infrastructure Roadmap. The Harris Advanced Radar Processor (HARP) is the prototype of the NextGen Weather Processor (NWP) real-time mosaic application and the Data EXchange (DEX) is the Harris prototype for SWIM (Segment 2). The HARP is hosted within the High Performance Weather and Climate (HPWC) environment, which is a Harris NWP prototype. The HPWC is based on the architecture being implemented by Harris in the development of the NOAA/NASA GOES-R Ground Segment program; the architecture evolved through numerous years of Harris Independent Research and Development (IRAD) investment and has been modified for use in NextGen with the addition of external interface capabilities. The HPWC external interface enables the exchange of data via the DEX.
The DEX Messaging, Registry/Repository, and Data Catalog Services provide the SOA technology-based data sharing services in the lab environment. The HARP, hosted in the HPWC infrastructure, acts as a Weather Data Service Provider offering radar mosaic data services to Weather Data Consumers via DEX.

The Enabled Sensors Simulator is a second Weather Data Provider in the lab environment. The Enabled Sensors Simulator acts as NEXRADs and TDWRs that are enabled to publish their data to SWIM. The Enabled Sensors Simulator captures CONUS-wide Level III NEXRAD and TDWR base radar reflectivity products from NOAAPORT and publishes them to DEX, which delivers them to the HARP for real-time mosaic generation. A Consumer Simulator, based on a custom display application, acts as a NextGen Weather Data Consumer in the lab environment. The Consumer Simulator was developed to visualize the results of data sharing activities between Weather Data Providers and potential Weather Data Consumers. In the NextGen era Weather Data Consumers primarily would be applications that process weather data and integrate it with non-weather data rather than display it in its original form.

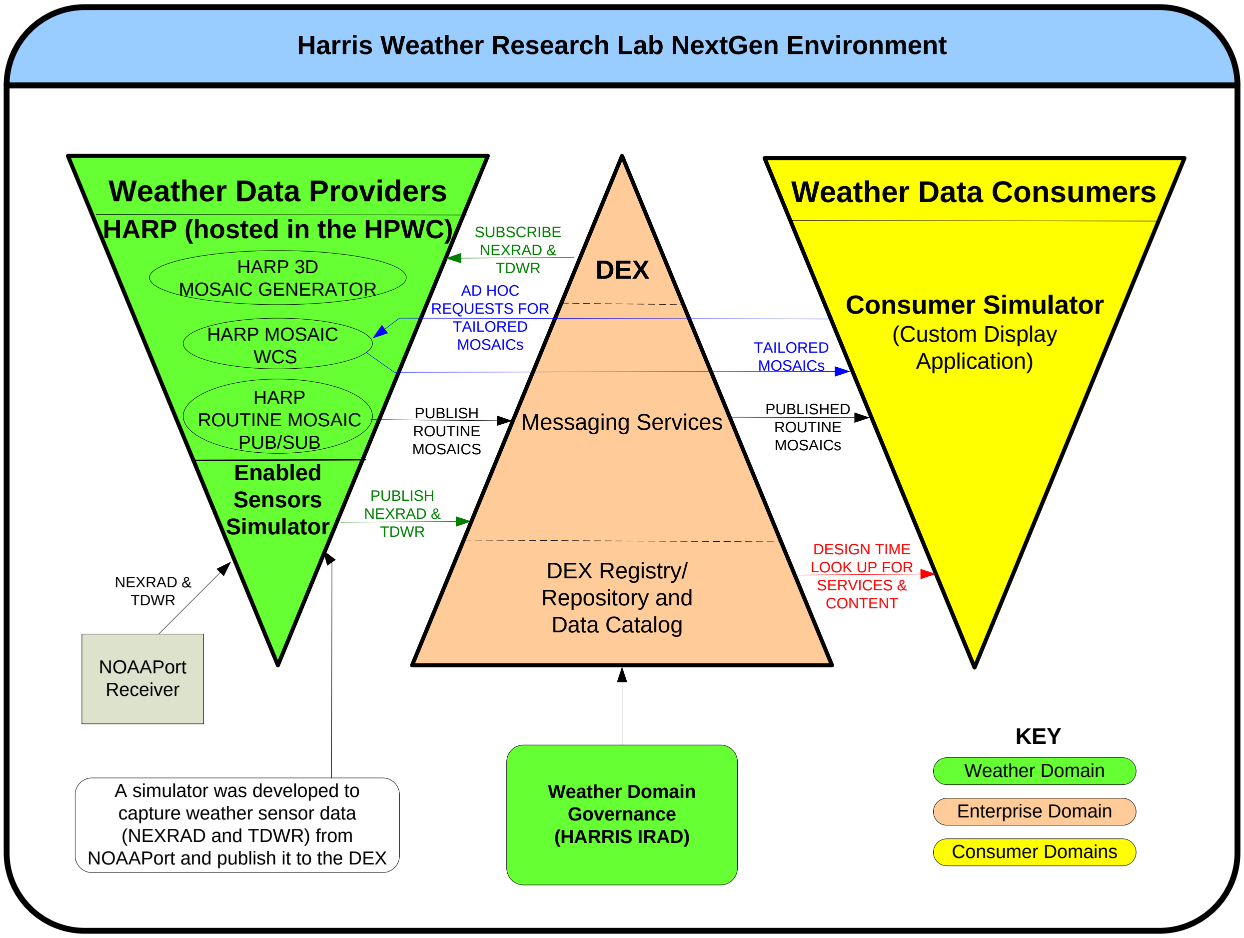

Figure 2 – Harris Weather Lab NextGen Environment

As shown in Figure 2, there are three key components of the HARP. The first component, the 3D Mosaic Generator, acts as a data source for the other two HARP components: the HARP Mosaic WCS and the HARP Routine Mosaic Pub/Sub. As the HARP receives each NEXRAD or TDWR base reflectivity product, the 3D Mosaic Generator component incorporates it into a high resolution 3-D representation of the CONUS airspace in memory. The HARP Mosaic WCS and the HARP Routine Mosaic Pub/Sub leverage this 3D, continuously updated representation, to produce and expose products to the Consumer Simulator. The Consumer Simulator can subscribe to receive routine mosaic products generated by the HARP or make ad hoc requests for uniquely tailored radar mosaics products (e.g., a product with atypical geographical coverage).

The HARP Routine Mosaic Pub/Sub component produces mosaic snapshots similar to those available in the NAS today, but with improved accuracy and spatial resolution. Routine products are generated by “taking a snapshot” of the 3D representation of the base reflectivity in memory and applying the appropriate algorithm to produce the selected product, e.g., Echo Tops. Routine products are produced at pre-determined, configurable time intervals. Each product can be defined in terms of area of coverage, resolution and map projection (e.g., polar stereographic). As routine snapshots are produced they are published to DEX and DEX then pushes them to subscribers – in this case, the Consumer Simulator. Utilization of the routine snapshot publish/subscribe service makes available a virtual stream of radar mosaics for consumers that can improve common situational awareness.

Examples of routine radar mosaic products that can be produced by HARP are a National 0-60,000 ft Composite Base Reflectivity Mosaic Product and a Northeastern US Regional Echo Tops Product.

The HARP WCS component supports ad hoc requests for radar mosaics covering specific geospatial volumes of interest through an OGC standards-based WCS interface. This interface allows weather consumers to submit custom requests for radar mosaics in terms of a specific atmospheric volume, product type, and preferred response format. In our current Lab environment, the Consumer Simulator can make WCS requests for composite layer products and other radar mosaic products having 2D geographic coverage, e.g., Echo Tops. While the WCS protocol can handle requests that cover 3D volumetric coverage, the current HARP WCS component is limited to handling 2 horizontal dimensions – a task to expand it to 3 D is planned for the coming year.

To give some perspective on data flow in the lab environment, four major data flows are depicted in Figure 2. The design time flow of information from the DEX Registry/Repository and Data Catalog to the Consumer Simulator developers is represented by the red arrowed line between the DEX and the Consumer Simulator. This line represents the flow of the metadata describing the available services (e.g., HARP WCS and HARP Routine Mosaics Pub/Sub) and data (e.g., Northeastern US Regional Echo Tops Product) to the Consumer Simulator developers, who use it to design the application interfaces to access the desired data. The green arrowed lines in Figure 2 represent the flow of real-time base reflectivity data from the Enabled Sensors Simulator to the HARP. The Enabled Sensor Simulator publishes real-time NEXRAD and TDWR products to the DEX, and the HPWC/HARP, which has a previously established subscription for them, receives them from the DEX as soon as they are published. The administrative work that established this pub/sub delivery was performed at design time by the HARP developer and the DEX administrator. The flow of Routine Radar Mosaics from the HARP to the Consumer Simulator is represented by black arrowed lines. As the HARP Routine Mosaic Pub/Sub service generates mosaic snapshots, it publishes them to the DEX, which immediately pushes them to subscribers – in this case the Consumer Simulator. The final flow is depicted by the two blue arrowed lines in Figure 2. It covers the request/response messaging with a WCS standards-based request for tailored radar mosaics by the Consumer Simulator. The blue line with the arrow head touching the HARP Mosaic WCS represents ad hoc requests for mosaics that have been tailored by the HARP in accordance with specific user needs. The blue line with the arrowhead touching the Consumer Simulator represents the WCS data delivery, which in the lab implantation is actually a web service pull. The blue lines are shown crossing over the DEX triangle in Figure 2 because in the current implementation these messages do not rely on DEX at all. In the NextGen era, the request and WCS response would likely be proxied by DEX, which would enforce security policies and track consumer data access, but as in the currently lab implementation, the large data files would be returned outside of DEX (out-of-band).

3.1 WCS Implementation Details and Lessons Learned

As discussed previously, the HARP WCS provides an OGC standards-based WCS interface that allows custom requests to be submitted to the HARP in an ad hoc fashion.

The HARP WCS consists of two major components: a fairly generic front-end WCS component and a HARP-specific back-end WCS component. A Java Message Service (JMS) message queue carries the WCS request and response messages between the two front- and back-end WCS components and a generic HTTP server (e.g. Apache) serves the radar data mosaic coverages generated by WCS GetCoverage requests. The JMS message queue and the HTTP server are technically parts of the HPWC architecture in which the HARP Mosaic WCS and HARP 3D Mosaic Generator execute.

The front-end WCS component provides the client-facing interface to the HARP Mosaic WCS. It accepts WCS requests in one of three formats, performs (partial) validation on each request, and transforms each request to a common format. The supported input formats are:

- HTTP GET - in which the WCS request parameters are encoded within the HTTP URL itself. These are also known as KVP or (Key Value Pair) requests.
- HTTP POST - in which the WCS request is encoded as an XML document.
- SOAP – in which the WCS request is encoded as an XML document and wrapped in a SOAP envelope XML document.

The validation performed by the front-end WCS component is “partial” in that (a) it is primarily syntactic in nature and (b) semantic validation does not enforce any HARP-specific constraints. For example, the front-end WCS component will ensure that a 2-dimensional bounding box is specified according to OGC standards, with corner coordinates within +/-90 degrees latitude and +/-

180 degrees longitude. The HARP-specific constraint that 2-dimensional bounding box corner coordinates are limited to a smaller CONUS bounding box are not validated by the front-end WCS component.

In order to simplify implementation and messaging between the two HARP WCS service components, the HARP-specific back-end WCS component supports WCS XML document messages only. This requires the front-end WCS component to transform HTTP GET and SOAP requests to the equivalent XML WCS request document before sending the request to the back-end WCS component via a JMS message queue. Conversely, responses from the back-end WCS component are transformed as necessary to a format compatible with the original request.

Limiting the responsibilities of the front-end WCS component to simple protocol transformations and data validation operations independent of HARP-specific constraints makes it more straightforward to implement, less likely to change in the future, and a better candidate for re-use.

When a WCS request is received by the back-end WCS component, additional HARP-specific validation is performed on the request. Valid requests are submitted to HARP 3D Mosaic Generator, which in turn generates a radar data coverage corresponding to the WCS request parameters.

It should be noted at this point that unlike most WCS implementations, which typically interact with some form of persistent storage (e.g. a relational database) to generate/extract requested coverage(s), the HARP back-end WCS component interacts directly with the HARP 3D Mosaic Generator, which maintains an up-to-date 3D base reflectivity mosaic in memory. So while the HARP 3D Mosaic Generator is not a typical coverage data source, the fact that it provides only near-real-time data (i.e. no historical data) removes the need for the back-end WCS component to deal with any temporal constraints when processing a WCS GetCoverage request. Temporal constraints are in fact considered invalid.

Regardless of the unique non-temporal nature of radar mosaic coverages provided by the back-end WCS component, the radar mosaic coverages can still be very large (several megabytes) in size. For this reason (and others to be discussed below), all mosaic radar coverages provided by the HARP Mosaic WCS are returned to the client in an out-of-band fashion. Specifically, generated radar mosaic coverages (which takes the form of a standard GeoTIFF image file or a standard NetCDF4 file) are posted to a generic web server that is a part of the HPWC architecture and the URL of posted files are returned to clients in the GetCoverage response. This approach has several benefits:

- GetCoverage responses are significantly smaller in size. This allows the thread processing an individual GetCoverage request to run to completion more quickly.
- The overhead of ending very large GetCoverage response messages across the JMS queue is eliminated.
- The issue of attaching opaque data to an XML document is neatly side-stepped. This is especially beneficial when SOAP messaging is involved, as there are at least five different approaches to attaching binary data to SOAP messages.
- The mechanism for retrieving the coverage data (an HTTP GET for some URL sent to a web server) is simple, well understood, and can be easily implemented by service consumers.
- Neither HARP WCS implementation (front-end or back-end) is responsible for “cleaning up” coverage data after it has been published to an HTTP server running in the HPWC architecture. Rather, the HPWC architecture itself becomes responsible for the periodic removal of all HTTP server data marked as having a finite lifetime, and can do so without regard to the source of the published data.

It should be noted that the separation of the HARP WCS into a generic front-end component and a HARP-specific back-end component communication via a JMS message queue lays the groundwork for a scalable, fault-tolerant system. Multiple HARP 3D Mosaic Generators, each associated with a HARP-specific back-end WCS component instance, could be deployed within and managed by the HPWC architecture. All back-end WCS component instances would listen for and receive WCS requests on the same JMS queue. Because they all share a common request queue, load balancing across a cluster of HARP 3D Mosaic Generators and back-end WCS component instance pairs can be achieved.

Furthermore, the temporary loss of a HARP 3D Mosaic Generator or its associated back-end WCS component instance would result in degraded system performance rather than total failure.
The HARP Mosaic WCS is implemented using the FUSE releases of Apache CXF, Apache Servicemix, and Apache Active MQ.

- CXF provides the necessary HTTP/JAXB/ JAXWS/SOAP capabilities.
- Servicemix provides a flexible ESB with OSGI container for hosting the front-end WCS component
- Active MQ provides the JMS message functionality.

FUSE was selected because of its stable release and support model. In addition, the FAA and NWS have stated a preference to use Open Source products such as FUSE in NextGen systems.

SUMMARY

Harris used a Notional NextGen CONOPS to guide the implementation of a NextGen-like Net-centric Environment in our Weather Research Lab. Existing prototypes, the HARP/HCWP and the DEX, were leveraged to build a NextGen-like environment where SOA and OGC Web standards-based technology relevant to weather data sharing were studied. Key lessons-learned are:

- Separation of the WCS into a generic front-end component and a Weather Data Provider specific back-end component facilitates scalability and fault-tolerance.
- Limiting the responsibilities of the front-end WCS component to simple protocol transformations and data validation makes it more straightforward to implement, less likely to change in the future, and a better candidate for re-use.
- When requested data coverages tend to result in the return of large volume data to consumers, there are significant benefits to returning the data in an out-of-band fashion, e.g., via a web server (URL reference):
    - Quick response, because they are small in size;
    - JMS queue overhead eliminated;
    - Eliminates attaching opaque data to an XML file (e.g., SOAP messages);
    - Simple interface to consumers;
    - WCS is no longer responsible for data cleanup; becomes a simple Web Server task.

FUTURE PLANS

- Enhancing the WCS to support request for 3D NetCDF4 data volumes is planned for the next software release;

- Addition of Canadian radar into the HARP; this will be an easy addition once a data source is identified that will add value to both the snapshot mosaic products and the WCS;

- HARP updates to provide higher resolution outputs where possible; currently, the maximum resolution of all HARP products and services is 1km but in the areas where TDWR coverage exists, higher resolution is possible;

- Evaluate and prototype HARP-DSR/ERAM interface solutions using SOA technologies. Understanding the latencies involved will help mitigate risk associated with a consolidated radar processing approach for the NWP.