

Open Geospatial Consortium, Inc.

Date: **2007-12-29**

Reference number of this document: **OGC 08-000**

Version: **0.0.9**

Category: **OGC Discussion Paper**

Editor: **Raj Singh**

OGC® Canadian Geospatial Data Infrastructure Summary Report

Copyright © 2008 Open Geospatial Consortium, Inc. All Rights Reserved.
To obtain additional rights of use, visit <http://www.opengeospatial.org/legal/>.

Warning

This document is not an OGC Standard. This document is an OGC Discussion Paper and is therefore not an official position of the OGC membership. It is distributed for review and comment. It is subject to change without notice and may not be referred to as an OGC Standard. Further, an OGC Discussion Paper should not be referenced as required or mandatory technology in procurements.

| | |
|--------------------|---------------------------|
| Document type: | OpenGIS® Discussion Paper |
| Document subtype: | NA |
| Document stage: | Approved Discussion Paper |
| Document language: | English |

Preface

This report summarizes the work performed under the Canadian Geospatial Data Infrastructure Pilot. The purpose of this pilot was to test the utility of certain OGC standards, in particular the Geography Markup Language (GML) and Web Feature Service (WFS), in the implementation of a spatial data infrastructure. OGC documents 08-001 and 08-002 are more technical companions to this document.

Suggested additions, changes, and comments on this draft report are welcome and encouraged. Such suggestions may be submitted by email message or by making suggested changes in an edited copy of this document.

The changes made in this document version, relative to the previous version, are tracked by Microsoft Word, and can be viewed if desired. If you choose to submit suggested changes by editing this document, please first accept all the current changes, and then make your suggested changes with change tracking on.

| Contents | Page |
|--|-------------|
| 1 Introduction..... | 5 |
| 1.1 Scope..... | 5 |
| 1.2 Document contributor contact points..... | 5 |
| 1.3 Revision history | 5 |
| 1.4 Future work..... | 6 |
| 2 References..... | 6 |
| 3 Terms and definitions | 6 |
| 4 Conventions | 7 |
| 4.1 Abbreviated terms..... | 7 |
| 5 Overview..... | 7 |
| 6 Enterprise View | 9 |
| 6.1 Stakeholders & Participants..... | 10 |
| 6.2 Requirements & Use Cases..... | 10 |
| 6.2.1 System Roles..... | 11 |
| 6.2.2 Use Cases..... | 12 |
| 6.3 Policy Concerns | 16 |
| 6.3.1 Schema versioning..... | 16 |
| 6.3.2 Feature versioning..... | 16 |
| 7 Information View..... | 17 |
| 7.1 Feature Schema, Data Sets & Their Mapping | 17 |
| 7.1.1 Local schema | 17 |
| 7.1.2 National schema..... | 17 |
| 7.1.3 Mapping..... | 17 |
| 7.2 Transactions | 18 |
| 7.3 Events & Notifications..... | 18 |
| 7.4 Application State or Map View Context Documents | 18 |
| 8 Computational View | 18 |
| 8.1 OpenGIS® Web Feature Service..... | 18 |
| 8.2 OpenGIS® Web Feature Service with Transactions (WFS-T)..... | 19 |
| 8.3 OpenGIS® Catalog Service for the Web (CS/W) | 19 |
| 8.4 OpenGIS® Web Mapping Service (WMS) | 19 |
| 8.5 Update and Feedback Services via Atom Publishing Protocol (APP)..... | 19 |
| 9 Engineering View | 19 |
| 9.1 Components | 20 |
| 9.1.1 Provider Server | 20 |
| 9.1.2 Geobase Server | 20 |
| 9.1.3 Geobase Application Portal | 20 |
| 9.1.4 Desktop GIS..... | 20 |
| 9.2 Protocols | 20 |

- 9.2.1 HTTP Transport20
- 9.3 Workflows.....20
 - 9.3.1 Dataset Publish – Subscribe – Harvest20
 - 9.3.2 Correction submittal.....21
 - 9.3.3 Event Notification.....22
 - 9.3.4 Context publication.....23
- 10 Technology View.....23
 - 10.1 Single point of access to closest-to-source data.....28
 - 10.2 Geographic data updates — Update Feed Service.....29
 - 10.3 Anonymous end-user feedback — Feedback Feed Service.....29
 - 10.4 Emergency event notification — Emergency Event Feed Service.....31
- 11 Findings.....32
 - 11.1 Shared Schemas are A Key Requirement32
 - 11.2 Synchronization Standards Are Needed32
 - 11.3 Service Performance Issue33
- 12 Demonstrations33

| Figures | Page |
|--|-------------|
| Table 1 — Use Case #1: Data Dissemination Portal | 13 |
| Table 2 — Use Case #2: Data Update | 14 |
| Table 3 — Use Case #3: COP application for analysis using vector data | 16 |
| Table 4 — Participants and Software Deployments | 24 |
| Figure 5 — GeoBase Discovery/Browse/Download Client | 28 |
| Figure 6 — Feedback Acceptance Client | 30 |
| Figure 7 — User Feedback Creation Client | 30 |
| Figure 8 — Emergency Event Client | 31 |

| Tables | Page |
|--|-------------|
| Table 1 — Use Case #1: Data Dissemination Portal | 13 |
| Table 2 — Use Case #2: Data Update | 14 |
| Table 3 — Use Case #3: COP application for analysis using vector data | 16 |
| Table 4 — Participants and Software Deployments | 24 |

Canadian Geospatial Data Infrastructure Pilot Discussion Paper (GCDI Pilot Engineering Report)

1 Introduction

1.1 Scope

This OGC™ Discussion Paper provides a summary of the work performed under the Canadian Geospatial Data Infrastructure Pilot. It describes the deployment of a federated network of Web Feature Services used for closest-to-source spatial data sharing and update.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. The Open Geospatial Consortium Inc. shall not be held responsible for identifying any or all such patent rights.

Recipients of this document are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of the standard set forth in this document, and to provide supporting documentation.

1.2 Document contributor contact points

All questions regarding this document should be directed to the editor or the contributors:

| Name | Organization |
|-----------------|--------------------------|
| Raj Singh | OGC |
| Peter Vretanos | Cubewerx Inc. |
| Josh Lieberman | Traverse Technologies |
| Peter Rushforth | Natural Resources Canada |

1.3 Revision history

| Date | Release | Editor | Primary clauses modified | Description |
|-----------|---------|-----------|--------------------------|--------------------------------|
| 28 Jan 08 | 1 | Raj Singh | | 1 st public release |
| | | | | |
| | | | | |

1.4 Future work

As a report on work performed in a Pilot that has ended, no future work on this document will occur. However, this Pilot suggested that future work on the Web Feature Service standard and related standards may be advisable.

2 References

The following documents are referenced in this document. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. For undated references, the latest edition of the normative document referred to applies.

The Atom Syndication Format, <http://www.ietf.org/rfc/rfc4287.txt>

The Atom Publishing Protocol, <http://www.ietf.org/rfc/rfc5023.txt>

GeoRSS: Geographically Encoded Objects for RSS feeds, <http://georss.org>

OGC 06-121r3, *OpenGIS® Web Services Common Specification*

OGC 04-094, *OpenGIS® Web Feature Service Implementation Specification*

OGC 03-105r1, *OpenGIS® Geography Markup Language (GML) Encoding Specification*

OGC 06-049r1, *OpenGIS® GML 3.1.1 simple features profile (1.0.0)*

3 Terms and definitions

For the purposes of this report, the definitions specified in Clause 4 of the OWS Common Implementation Specification [OGC 06-121r3] shall apply. In addition, the following terms and definitions apply.

3.1

Atom

Atom Syndication Format

Atom is an XML-based document format that describes lists of related information known as "feeds". Feeds are composed of a number of items, known as "entries", each with an extensible set of attached metadata. For example, each entry has a title. Atom is defined by the Internet Engineering Task Force at <<http://www.ietf.org/rfc/rfc4287.txt>>.

3.2

GeoRSS

GeoRSS describes a number of ways to encode location in RSS and Atom feeds. There are two encodings of GeoRSS: GeoRSS Simple and GeoRSS GML. Both encodings support the encoding of basic geometries (point, line, polygon, and box). GeoRSS

Simple geometries may only support one coordinate reference system—WGS-84 latitude/longitude in decimal degrees. GeoRSS GML is a formal GML Application Profile, and allows the coordinate reference system to be specified as with any GML feature. Both formats are designed for use with Atom 1.0, RSS 2.0 and RSS 1.0, although it can be used just as easily in non-RSS XML encodings. The definition of GeoRSS is at <<http://georss.org>>.

3.3

Atom+GeoRSS

Atom+GeoRSS refers to an XML encoding using Atom version 1.0 with the GeoRSS extension.

3.4

GML Application Schema

an XML schema which imports some or all definitions from the GML Schemas in order to model spatial phenomena in GML. Can be designed by the database designer or adopted from a vertical environment. Examples of vertical standard GML Application Schema are GeoSciML and CityGML.

4 Conventions

4.1 Abbreviated terms

| | |
|------|--|
| APP | Atom Publishing Protocol |
| CGDI | Geoconnections Canada Canadian Geospatial Infrastructure Pilot |
| GML | Geography Markup Language |
| IT | Information Technology |
| SDI | Spatial Data Infrastructure |
| SOA | Service Oriented Architecture |
| WFS | Web Feature Service |
| XML | eXtensible Markup Language |

5 Overview

In today's environment where users and organizations are collaborating over the creation and on-going maintenance of geospatial data sets, it is of critical importance for all Spatial Data Infrastructure (SDI) partnerships to be able to share geographic information in real time, and make the costs associated with sharing—in labor and IT resources—negligible. CGDI tested the ability to share data using the OpenGIS® GML and WFS standards. The project also highlighted the need for an interoperable means of *updating* these shared geographic information resources, so a large part of the project became the development of open interfaces for requesting and receiving updates.

In this report the activity of enabling cross-jurisdiction updates to geographic data is often termed *federated geo-synchronization*. Some have approached this requirement from the database perspective, expecting collaborating organizations to use database replication technology for their geographic data sharing needs. This strategy cannot work in most cases because, most notably, every organization needs to work with a different information model. For example, a local jurisdiction often needs to maintain a large set of properties on their roads data, such as maintenance-related items like pavement quality and type. At the national level, however, that information is unnecessary. Another limitation of replication technology is that all parties must be using the same software. This is rarely the case in a multi-organization setting. Finally, replication involves reaching deep into another organization's IT infrastructure, requiring a level of trust between the parties which isn't always there.

The CGDI Pilot sought to minimize these barriers to participation by developing an interoperable solution to the geographic data synchronization problem that has none of these drawbacks. We also believe that this solution to the federated geo-synchronization issue is applicable to ongoing SDI initiatives across the globe.

This report is written using the Reference Model of Open Distributed Processing, ISO/IEC 10746-1 to ISO/IEC 10746-4 (RM-ODP). It was felt appropriate to take this approach as the CGDI is an excellent example of an open, distributed system.

There are five viewpoints defined in RM-ODP:

- **Enterprise viewpoint:** Software systems must support the enterprise or business. Whether they are initially driven by an organizational need or they arise from a technology opportunity, software systems exist to support the business of that organization. The Enterprise Viewpoint, therefore, defines the business problems and associated processes in clear and consistent terms that can be realized as system requirements and articulated the same way throughout the other viewpoints.
- **Information viewpoint:** A viewpoint on the system and its environment that focuses on the semantics of the information and information processing that is to be performed to address the identified problems.
- **Computational viewpoint:** A viewpoint on the system and its environment that presents the system as a functional decomposition of services that interact at interfaces. This viewpoint captures the details of these components and interfaces without regard to distribution.
- **Engineering viewpoint:** A viewpoint on the system and its environment that focuses on the mechanisms and functions required to support distributed interaction between objects in the system.
- **Technology viewpoint:** A viewpoint on the system and its environment that focuses on the choice of technology in that system.

The Enterprise, Information, and Computational viewpoints were developed in preparation for, and refined during the project. The Engineering and Technology viewpoints represent the execution of the project in software systems residing in organizations—the work performed by software vendor and government partner participants to realize the CGDI.

6 Enterprise View

The enterprise viewpoint describes the purpose, scope and policies for the Pilot's communities of interest. Effective software systems must support a community. Whether the system is initially driven by organizational needs or arises from a technology opportunity, a software system persists because it meets the needs of an enterprise. The Enterprise View, therefore, defines the business objectives supported by the system detailed in the other views.

The goal of this project is to reduce the cost, effort, and complexity associated with shared spatial data management and delivery to stakeholders. There are two general cases under consideration. The first is one where data is needed on a seamless, nationwide level. The second is where data is needed at a larger scale—a region within a province, or a single municipality—and richer information is required than what is provided in the aggregated, nationwide product. In all cases, the goal is to make these data products represent the most current information available, and have them available digitally and immediately upon request.

The Canadian Geospatial Data Infrastructure (CGDI) is the underlying foundation needed to share geographic information (e.g. maps, satellite images) over the Internet. It consists of four key components:

- National framework data
- Common data policies
- Technical standards
- Enabling technologies

The CGDI supports the shared decision making necessary for resolving many horizontal and inter-jurisdictional challenges. With its four key components, the CGDI serves as a common foundation for key government information systems, as well as for third-party service delivery. The goals of the CGDI Interoperability Pilot are closely aligned with the goals of the CGDI as a whole.

The CGDI “enterprise” consists of: data providers at the federal, provincial, and municipal levels; various other public agencies, commercial firms, and the public. Stakeholders in this Pilot include the data providers and users above, plus software vendors and systems integrators who will aid in building out this spatial data infrastructure.

6.1 Stakeholders & Participants

Finance and Administration Division
Resource Information Management Branch
Alberta Sustainable Resource Development
Province of Alberta

Base Mapping and Geomatic Services Branch,
Integrated Land Management Bureau,
Ministry of Agriculture and Lands,
Province of British Columbia

Surveys and Mapping Division,
Province of Newfoundland and Labrador

Direction générale de l'information géographique,
Ministère des Ressources naturelles et de la Faune,
Gouvernement du Québec

GeoNOVA Program,
Service Nova Scotia and Municipal Relations
Province of Nova Scotia

Land Information Ontario,
Ontario Ministry of Natural Resources
Province of Ontario

Emergency Management Ontario
Ministry of Community Safety and Correctional Services
Province of Ontario

Office of Geomatics Coordination,
Information Technology Office,
Province of Saskatchewan

6.2 Requirements & Use Cases

The geographic data sharing requirements of the CGDI Pilot were as follows:

1. Establish mechanism for supporting feature synchronization and data currency in a distributed and collaborative SDI environment.
2. Support and promote closest-to-source ownership of geospatial data for all data providers.
3. Ensure real-time access to current geographic features in Canada.

4. Establish a permanent collaborative environment for geospatial data as a key component for creating a sustainable and current spatial data infrastructure in Canada.
5. Consolidate the CGDI as the permanent spatial information infrastructure for accessing current geographic data content across government departments in Canada.

The CGDI Pilot project aims at establishing distributed geospatial data resources from data providers at regional, provincial and national levels in Canada and maintaining common and integrated data content at the national level using shared data models and open interfaces for supporting feature synchronization.

Use cases will capture the expected way users will interact with the system and are split into scenarios describing the steps taken to accomplish a required task, using the system as a tool. There will be 3 use cases, each demonstrating WFS capabilities in the following areas:

1. GeoBase portal-type data dissemination
2. Access and update using desktop GIS.
3. Public-safety scenario for data download and analysis

6.2.1 System Roles

There are a number of actors involved in the CGDI IP use cases. These include:

Data Provider

A data provider maintains a locally or regionally bounded vector dataset for their own use and wishes as well to contribute to regional or national dataset aggregation and/or access.

Portal Provider

A portal provider supports user access to nationwide feature datasets.

Application End User

End users wish to discover, view, and obtain current feature datasets which may cover any part of Canada but which are customized to the user's area of interest.

Correspondent

A Correspondent is an end user or other non-provider with access to feature data who provides feedback regarding data correctness, currency, or quality.

EM Analyst

An Emergency Management analyst accesses nationwide feature data to both create a local operational view of an emergency situation, and alert others to appropriate response actions, such as mobilization and evacuation.

6.2.2 Use Cases

Title: Data Dissemination Portal

Summary

Use case describing an end user going to a web portal to discover/locate/select relevant data for download to the end user computer system.

Preconditions

End user has access to web browser and internet

Portal is accessible

Distributed WFS network is in place with most current data.

Triggers

End user wishes to obtain some data as GML

Basic course of events

1. User navigates to the portal website in his web browser
2. User accepts the unrestricted-use licence agreement
3. User executes a search function by entering a keyword to provide a list of feature types of stored in a registry
4. User selects the dataset that he is interested in
5. User clicks on a link to examine the metadata
6. User clicks on a link to bring up a map viewer which displays the dataset and checks the data for suitability to the task
7. User specifies styling parameters to display the dataset in a different colour - **optional**
8. User navigates to his area of interest by (one or more of these options)

- *draw a bounding box on the map*
 - *query admin boundary dataset (via WFS) to return an envelope*
 - *query NTS lookup service to return bounds of NTDB map sheet index*
 - *query placename dataset (via WFS) to return a point (set zoom level appropriately)*
9. User specifies filter parameters so select a subset of the dataset (e.g. road “class”) - **optional**
 10. User selects a subset of attributes to be returned with the dataset (i.e. *element sets*) - **optional**
 11. User clicks on an “obtain data” link and a WFS query is generated to the data dissemination service
 12. A GML response is returned to the user which he saves on the disk

Alternative paths

Saving: user saves the map view in an OWS context document to share the document with another user

Post conditions

User has the data he is interested in in-hand

Table 1 — Use Case #1: Data Dissemination Portal

Title: Data update

Summary

A use case to demonstrate the potential of WFS technology to meet data currency requirements by using closest-to-source data architecture.

An authorized user makes additions/deletions/modifications to the data being served via WFS. Changes may be the result of feedback from an end user, or from updates initiated by a source agency.

Preconditions

An operator trained in the use of the application.

A pre-established trust relationship between the systems in the network.

Technology in place to enforce the trust relationship and allow authenticated access.

Triggers

End user notices an error in the location and attribution of a feature

A municipality provides geometry and attributes describing a new subdivision

One municipality merges with another and municipal boundaries change

Basic course of events

1. Source agency receives notification of updates including geometries and attribution
2. Operator accesses an application with read/write access to the source database.
3. Operator validates the new information according to framework data standard
4. Operator updates the affected feature objects with the new information
5. Operator saves the changes
6. Updates are propagated to other systems using WFS-T
7. End-users have near real-time access to the data.

Alternative paths

Variations in how the operator is notified of changes: directly from end-users, or from source databases

Trusted end-users make the corrections directly in a secured application

Post conditions

Corrected data is available in near real-time after the update is saved.

Notes

The use case applies to all geometry types under test.

Table 2 — Use Case #2: Data Update

#3 COP application for analysis using vector data

Iteration

Summary

A Community of Practice decision support application accesses data via WFS and performs some analysis on the resulting vector data.

The scenario is a toxic plume emission that requires the evacuation of communities in its path over time. The requirement is to identify communities that must be evacuated, the evacuation routes and evacuation deadlines.

Preconditions

An operator trained in the use of the application.

Triggers

A sour gas well blowout in Southern Saskatchewan, near Estevan, emits a toxic plume.

Basic course of events

1. Provincial Emergency Operations Centre (EoC) receives notification of the event
2. Operator accesses the application and navigates an interactive map to the geographic area for the event.
3. Operator generates a plume model for the next 48 hours from meteorological data (canned model).
4. Operator loads the plume model into the mapping application
5. Operator performs an intersection of the plume with administrative boundaries and generates a list of municipal EoC telephone numbers.
6. Operator performs an intersection of the plume with roads and styles those roads differently in the mapping application. (roads closed)
7. Operator inverts the plume/road intersection and styles the highest capacity roads differently in the mapping application. (evacuation routes)
8. Operator saves the state of the map.
9. Operator performs an intersection of the plume with point place names and generates a list of municipalities to be evacuated, at one hour intervals.

10. Application updates a GeoRSS feed containing the products generated above (to be consumed by regional police headquarters and municipal EoCs)

Alternative paths

The municipal EoC in Estavan, SK, triggers the event by updating its own GeoRSS feed with a description of the event.

Add an application that consumes the GeoRSS feed.

Notification of events through alerts (e.g. CAP alerts)

Population and dwelling count data are also used to generate an additional information product (optional)

Post conditions

A map context is available and the EoC RSS feed is updated.

Notes

The above interaction should occur in less than 2 minutes.

Table 3 — Use Case #3: COP application for analysis using vector data

6.3 Policy Concerns

6.3.1 Schema versioning

For framework geodata sets in general and the three types of feature data in the CGDI IP in particular, there are differing local and national feature schemas. Schemas both vary between organizations and evolve over time. A concern of this project was to respect local differences in feature schemas which address local needs, but work toward national feature schemas which are strict subsets of local ones, i.e. minus the attributes which are of more local concern.

6.3.2 Feature versioning

A challenge for the CGDI IP was to balance accessibility with currency and accountability. This required explicit rules and agreements for frequency and workflow of feature updates, as well as an update process that is both robust and transparent.

7 Information View

The Information Viewpoint is concerned with the roles, structures, and semantics of information and information processing. It defines schemas for geospatial or geospatially related information and methods.

7.1 Feature Schema, Data Sets & Their Mapping

The CGDI Pilot information view was mainly concerned with the GML application schemas developed to model three classes of geographic features:

- Geographic Names
- Roads
- Administrative Boundaries

These each had an information model designed for expression as seamless, nationwide data sets, which is generally a subset of the locally available information model. Data providers may have chosen to design a richer information model for local use if the national model did not satisfy all of their specific needs (e.g. feature-level management metadata), but that fact was generally transparent to the Pilot.

In order to mediate between local and national views of geodata, stable feature schemas for each and mappings between were required. The information view took into account the existing information models/schema of GeoBase layers, namely GeoNames, administrative boundaries and National Roads Network (version 1).

These schema issues are more fully described in the companion Engineering Report, Canadian Geospatial Data Infrastructure WFS and GML Best Practices (OGC document #08-002).

7.1.1 Local schema

A local schema is defined as one that is maintained by a local data provider (province, locality, specialized community) that represents the same geographic feature types across multiple organizations, but includes geometries, attributes, and/or constraints that are particular to that locality.

7.1.2 National schema

A national schema represents a more or less strict common subset of all the local schemas that are to be aggregated and/or viewed in a national context.

7.1.3 Mapping

A mapping describes either the element subset or more involved rules for translating geodata from a local schema to the corresponding national schema.

7.2 Transactions

Update and propagation of feature data was accomplished using WFS-T transactions. The transaction requests themselves served very well as the information structure to represent such feature dataset transactions. Additional information was required, however, to represent both the new data produced by a transaction, the status / process of the transaction, and the prior state of the dataset which has been altered (older feature versions and/or deleted features). Such additional information served both to document dataset updates and make earlier versions of features and feature datasets retrievable through reverse application of the transaction record.

The work on information modeling for transactions is fully described in the companion Engineering Report, Loosely Coupled Synchronization of Geographic Databases in the Canadian Geospatial Data Infrastructure Pilot (OGC document #08-001).

7.3 Events & Notifications

A number of asynchronous interactions occurred in the operation of the CGDI IP, e.g. “there is an update”, “an error has been discovered”, “an update has been applied”, “a view of an emergency situation is available”. These require the exchange of additional information beyond the usual synchronous interactions of OGC Web Services requests and responses. Both publish-subscribe interactions and notification-action interactions require persistent event information as provided by such information structures as Web feeds (e.g. RSS or Atom). In addition, alert information such as CAP alerts were important for serving additional notification channels (e.g. email, SMS).

The work on information modeling for event notification is also described in OGC document #08-001).

7.4 Application State or Map View Context Documents

OWS Context documents represented the view state of applications and dynamic maps (such as the layers and their data sources, ordering, styling, annotation, etc) in the Pilot. They functioned both as a way of persisting (e.g. storing, discovering, recovering) such views, and a way of collaborating by exchanging the context of a feature or features being communicated from one user to another.

8 Computational View

The computational viewpoint is concerned with the functional decomposition of the system into a set of services that interact at interfaces. This viewpoint captures the details of these components and interfaces without regard to distribution.

8.1 OpenGIS® Web Feature Service

The Pilot was designed to use WFS as the service for spatial data exchange. This service was heavily exercised by all participants, and successfully played that role.

8.2 OpenGIS® Web Feature Service with Transactions (WFS-T)

WFS-T allowed clients to send messages relating to making changes to a geospatial database. This service interface was used by many participants to apply changes to their spatial databases. It was also used to describe changes that a user would like to have applied to a database (i.e. geographic feedback).

8.3 OpenGIS® Catalog Service for the Web (CS/W)

CS/W defines a common interface that enables diverse but conformant applications to perform discovery, browse and query operations against distributed heterogeneous catalog servers.

This service was used to discover all the other services available in the Pilot.

8.4 OpenGIS® Web Mapping Service (WMS)

WMS enables the creation and display of registered and superimposed map-like views of information that come simultaneously from multiple remote and heterogeneous sources.

Many participants provided a WMS implementation mainly as an alternative to WFS, when the client application needed a map view of the data instead of the data itself.

8.5 Update and Feedback Services via Atom Publishing Protocol (APP)

APP is an application level protocol for publishing and editing web resources. The protocol at its core is the HTTP transport of Atom-formatted Web feeds.

During the course of the Pilot, it was discovered that no standard interface existed for negotiating the changes in status and content of spatial databases. This need was filled by designing new services for spatial data updates, as well as end-user suggestions for updates (feedback).

These services are more fully described in the companion Engineering Report, Loosely Coupled Synchronization of Geographic Databases in the Canadian Geospatial Data Infrastructure Pilot (OGC document #08-001).

9 Engineering View

The engineering viewpoint relates the Enterprise, Information, and Computation viewpoints to specific components linked by a communications network. This viewpoint is concerned primarily with the interaction between distinct computational objects: its chief concerns are communication, computing systems, software processes and the clustering of computational functions at physical nodes of a communications network. The engineering viewpoint also provides terms for assessing the “transparency” of a system of networked components – that is, how well each piece works without detailed knowledge of the computational infrastructure.

9.1 Components

The characteristics of classes of software component are defined here.

9.1.1 Provider Server

A provider server is the platform for generating and maintaining feature datasets and supporting local users, as well as for engaging in update notification (publication-subscription), propagation, feature roll-up, and metadata harvesting with larger-scale servers such as the GeoBase or intermediate-level servers. It consisted of a WFS-T as well as client and middleware functions for carrying out update and notification workflow.

9.1.2 Geobase Server

The GeoBase server acted as the central aggregating and/or cascading platform for national feature datasets. It also engaged in notification, harvest, and update exchanges with provider servers and cascaded some larger-scale feature requests to more local WFSes. Its other function was to support demonstration scenarios with services such as WFS-T, WMS, and CS/W.

9.1.3 Geobase Application Portal

The Application Portal provided end-users with applications for discovering, browsing, and downloading feature datasets, in support of the use cases described earlier.

9.1.4 Desktop GIS

Desktop GIS applications were the principal means of creating and updating feature datasets. It is also expected that EM analysts will largely use desktop GIS and/or image-processing applications to generate both alerts/notifications and new derived operating view datasets.

9.2 Protocols

9.2.1 HTTP Transport

Interactions between distributed components in the Pilot used the HTTP protocol, using the HEAD, GET, and POST methods. This included browser access to the GeoBase Portal, Web service requests, and exchange of event notification feeds.

9.3 Workflows

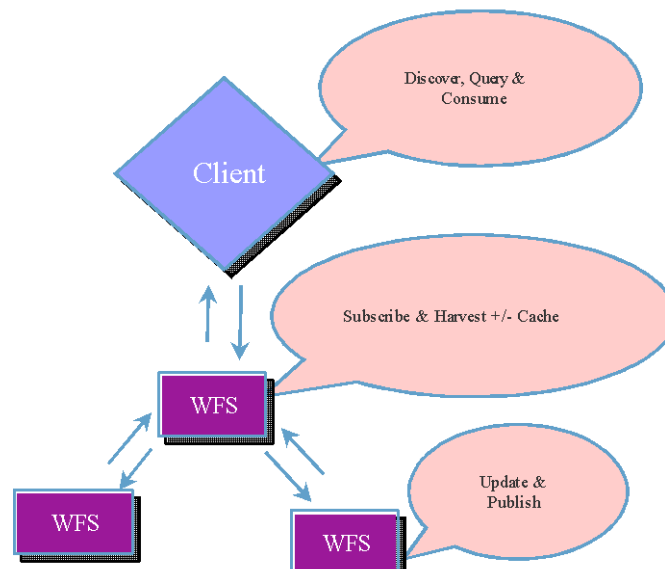
Workflows are often the “solution” matching information transformations with user actions, particularly across distributed processing components.

9.3.1 Dataset Publish – Subscribe – Harvest

Description:

1. Partner updates a feature set and publishes its availability to the network
2. GeoBase server has subscribed to the partner update notification and harvests metadata about the updated features.

Figure 2 – Dataset Publish, Subscribe, Harvest

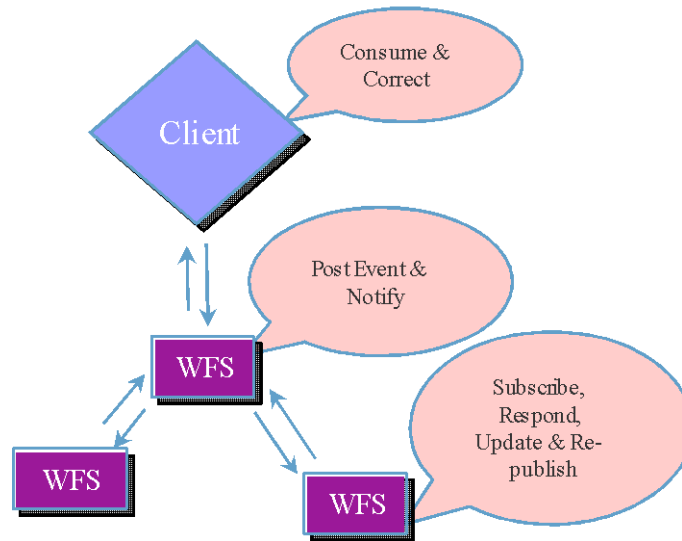


3. User discovers, queries, and then downloads updated features from the partner
4. Optionally, GeoBase server aggregates the partner data locally in the national level feature schema for central access

9.3.2 Correction submittal

Description:

1. User provides feedback and correction to a feature
2. GeoBase server posts a correction event and proposed feature update
3. Partner server has subscribed to corrections, fetches correction, considers it, then responds, updates its feature holdings, and re-publishes to the network

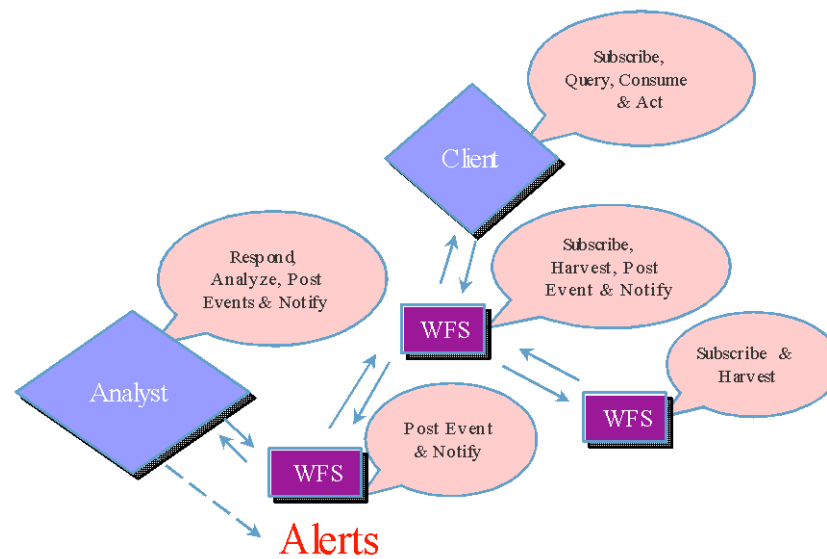
Figure 3 — Correction Submittal

9.3.3 Event Notification

Description:

1. Analyst produces, posts new features, posts event, optionally issues alerts
2. Partner server publishes event and new features
3. GeoBase server harvests features and event
4. Users receive event and query new features
5. Optional flow of features and events from other partners

Figure 4 — Event Notification



9.3.4 Context publication

An OWS Context is an XML document. Publication of such a document can be accomplished in many ways. In the Pilot, these documents were made accessible via a URL or advertised via notification feed.

10 Technology View

A technology viewpoint describes the realization of a system in software. It describes the implementation and deployment environment using technologies, standards and products available at a point in time.

Table 1 shows all software services and clients deployed for the CGDI. The following sections of the Technology view reference these deployments and describe how they were integrated to meet the business requirements and use cases expressed in the Enterprise view.

Table 4 — Participants and Software Deployments

WFS=Web Feature Service; WMS=Web Mapping Service; UFS=Update Feed Service; CSW=Catalog Service for Web

| Organizational Deployment | Services Deployed | Software | Vendor | Service Location | Endpoint |
|---------------------------|-----------------------------------|-----------------|--------------------|------------------|---|
| - | WFS, WMS, and CSW clients | Gaia, CarbonArc | The Carbon Project | Federal | (desktop client) |
| - | Feedback, Emergency event clients | CarbonArc | The Carbon Project | Federal | (desktop client) |
| GeoConnections | WFS, WMS, and CSW clients | Mapbuilder | LisaSoft | Federal | http://wolf.cits.rncan.gc.ca:8080/mapbuilder/demo/ |
| GeoConnections | Feedback creation client, feed | Mapbuilder | LisaSoft | Federal | http://wolf.cits.rncan.gc.ca:8080/mapbuilder/demo/feedback_client/index.html |
| GeoConnections | Update Feed client | Mapbuilder | LisaSoft | Federal | http://wolf.cits.rncan.gc.ca:8080/mapbuilder/demo/notification/index.html |
| GeoConnections | Emergency Event Feed client | Mapbuilder | LisaSoft | Federal | http://wolf.cits.rncan.gc.ca:8080/mapbuilder/demo/event/index.html |
| GeoConnections | Cascading WFS, Aggregating WFS | Cubeserv | Cubewerx | Vendor | http://www.pvretano.com/cgdi/geoserver/cwwfs.cgi? |

| Organizational Deployment | Services Deployed | Software | Vendor | Service Location | Endpoint |
|----------------------------------|--------------------------|-----------------|---------------|-------------------------|--|
| GeoConnections | CSW | Cubeserv | Cubewerx | Vendor | http://demo.cubewerx.com/cgdi/cwwrs.cgi |
| GeoConnections | Aggregating WFS | Geoserver | LisaSoft | Federal | http://wolf.cits.rncan.gc.ca:8080/geoserver/wfs? |
| GeoConnections | WMS | Geoserver | LisaSoft | Federal | http://wolf.cits.rncan.gc.ca:8080/geoserver/wms? |
| Alberta | WFS, UFS | Cartalinea | Galdos | Province | WFS with GML3: https://securexnet.env.gov.ab.ca/cgdi/http? WFS with GML2: https://securexnet.env.gov.ab.ca/cgdi2/mapserv? |
| British Columbia | WFS | Geoserver | Refractions | Vendor | http://office.refractions.net:21880/geoserver/wfs? |
| Manitoba | WFS, WMS | Cubeserv | Cubewerx | Vendor | WFS: http://demo.cubewerx.com/cgdi/cubeserv.cgi?config=mb&datastore=Manitoba WMS: http://www.cubewerx.com/cgdi/cubeserv.cgi?CONFIG=mb |
| Newfoundland and Labrador | WFS, UFS | Geoserver | Sunertek | Vendor | WFS: http://nfl.sunertek.com:8090/wfs UFS: http://nfl.sunertek.com:8090/updatefeed.xml |

| Organizational Deployment | Services Deployed | Software | Vendor | Service Location | Endpoint |
|---------------------------|-------------------|-----------|----------|------------------|--|
| Northwest Territories | WFS, WMS | Cubeserv | Cubewerx | Vendor | WFS: http://demo.cubewerx.com/cgdi/cubeserv.cgi?config=nt&atastore=NWT WMS: http://www.cubewerx.com/cgdi/cubeserv.cgi?CONFIG=nt |
| Nova Scotia | WFS | Mapserver | - | Province | http://sack.irmapping.com/CGDIXSL/Default.aspx?map=C:\Kraig\Work\CGDI\NS_WFS.map |
| Nunavut | WFS, UFS | Geoserver | Sunertek | Vendor | WFS: http://nunavut.sunertek.com:8080/wfs UFS: http://nunavut.sunertek.com:8080/updatefeed.xml |
| Ontario | WFS, UFS | Geoserver | Sunertek | Vendor | WFS: http://ont.sunertek.com:8090/wfs UFS: http://ont.sunertek.com:8090/updatefeed.xml |
| Prince Edward Island | WFS, WMS | Cubeserv | Cubewerx | Vendor | WFS: http://demo.cubewerx.com/cgdi/cubeserv.cgi?config=pe&atastore=PEI WMS: http://www.cubewerx.com/cgdi/cubeserv.cgi?CONFIG=pe |

| Organizational Deployment | Services Deployed | Software | Vendor | Service Location | Endpoint |
|---------------------------|-------------------|-----------|----------|------------------|--|
| Québec | WFS, WMS | Cubeserv | Cubewerx | Vendor | WFS: http://demo.cubewerx.com/cgdi/cubeserv.cgi?config=pq&datastore=Quebec WMS: http://www.cubewerx.com/cgdi/cubeserv.cgi?CONFIG=pq |
| Saskatchewan | WFS | ArcIMS | - | Province | http://www8.swa.ca/wfsconnector/com.esri.wfs.Esrimap/WFS_Roads? |
| Yukon | WFS, WMS, UFS | Geoserver | Sunertek | Vendor | WFS: http://yukon.sunertek.com:8080/wfs UFS: http://yukon.sunertek.com:8080/updatefeed.xml |

10.1 Single point of access to closest-to-source data

The primary goal of the project was to provide a single point of access to closest to source data. The Web-based client from LisaSoft can access WFS, WMS and CSW services, and provided that single point of access to all data services in the federation. **Error! Reference source not found.** shows this client in use. In the Keywords box, a user has searched the catalog by typing ‘place’ in the ‘Keywords’ field, and selecting ‘WFS’ in the ‘Service Type’ drop-down menu. In the bottom right corner of the figure are the search results. The user has chosen to add ‘Alberta Place Names’ to the map.

Note that through this single interface the user has been able to discover data and map services located on any of the servers participating in the federation. The data residing in WFS services can also be downloaded in GML format (XML) from within this interface.

The screenshot displays the CGDI-IP Geobase Discovery/Browse/Download Client interface. The main map shows a geographical view of Canada with various cities labeled, including Whitehorse, Yellowknife, Iqaluit, Edmonton, Regina, Winnipeg, Québec, Toronto, Ottawa, and Halifax. The interface includes a search bar with the keyword 'place' and a service type dropdown set to 'WFS'. Below the search bar, there are fields for loading external OWS context (URL: context.xml) and adding GML layers (URL: plumes.gml). A search results panel on the right lists several services, with 'Alberta Place Names' selected.

| Layer Control | Service Type | Layer Name |
|-------------------|--------------|--------------------------|
| 1 [checked] [5] | WFS | Alberta Place Names |
| 2 [unchecked] [2] | WFS | CW - Administrative Area |
| 3 [unchecked] [2] | WFS | CW - Road Segment |
| 4 [unchecked] [2] | WFS | CW - Place Name |
| 5 [unchecked] [2] | WFS | LS - Administrative Area |
| 6 [unchecked] [2] | WFS | LS - Road Segment |
| 7 [unchecked] [2] | WFS | LS - Place Name |
| 8 [checked] | WMS | CGDI NRCan WMS |

| Search Results | Service Type | Layer Name |
|----------------|--------------|------------------------------|
| [checked] | WFS | PlaceName_Type |
| [checked] | WFS | Alberta Place Names |
| [checked] | WFS | Place Names (Manitoba) |
| [checked] | WMS | PlaceName_Type |
| [checked] | WMS | PlaceName_Type |
| [checked] | WFS | PlaceName_Type |
| [checked] | WFS | Place Names (Geoserver - CW) |
| [checked] | WMS | PlaceName_Type |

Figure 5 — GeoBase Discovery/Browse/Download Client

10.2 Geographic data updates — Update Feed Service

In the course of the project, it was discovered that in order to have an information federation function efficiently, there needed to be a way to know what information has changed since last accessed. The WFS service was found lacking in this area because that interface only allows a client to retrieve the current data, putting the burden of figuring out what has changed on each individual client. It was deemed important to remove the uncertainty from the update process by designing a service that was much more explicit about data updates.

A number of provinces implemented a prototype Update Feed Service (UFS) for this purpose. A federation partner (aggregator) that wants to keep up-to-date with another partner's data service should start with a bulk data load. After the bulk load, the aggregator can simply query the partner's UFS to get the recurring updates. The UFS will either return complete geographic features in GML, or a WFS query URL that returns GML. This is described in more detail in OGC document 08-001.

Update Feed Service deployments are listed in Table 1 (with a UFS notation). The Cubewerx aggregating WFS was a UFS client that kept the GeoBase WFS in synchronization with provincial partners. Cubewerx also implemented a UFS at the federal level, effectively providing a single point of access to updates from all federation partners.

10.3 Anonymous end-user feedback — Feedback Feed Service

Another requirement of the project was to provide a way for CGDI users to submit update requests, or correct errors in data sets they were using. We found that the update feed service architecture was suited for this purpose as well. The only difference is that instead of the updates being about a well known, authoritative, existing data source, the updates are coming from relatively anonymous, untrusted parties.

To ensure that this type of geographic data update was relatively structured, LisaSoft developed a 'feedback client' shown in Figure 7 that required the user (the feedback submitter) to start from an existing feature schema—in our case either roads, place names, or administrative boundaries. The user could then modify an existing feature, or create a completely new one. This client is accessed at the federal level, but the feedback is then reported back to the affected province for action. The Carbon Project modified CarbonArc, their desktop client shown in Figure 7 (an ArcGIS extension), to serve as a feedback review, quality assurance, and acceptance tool. The Feedback Feed Service is described in more detail in OGC document 08-001.

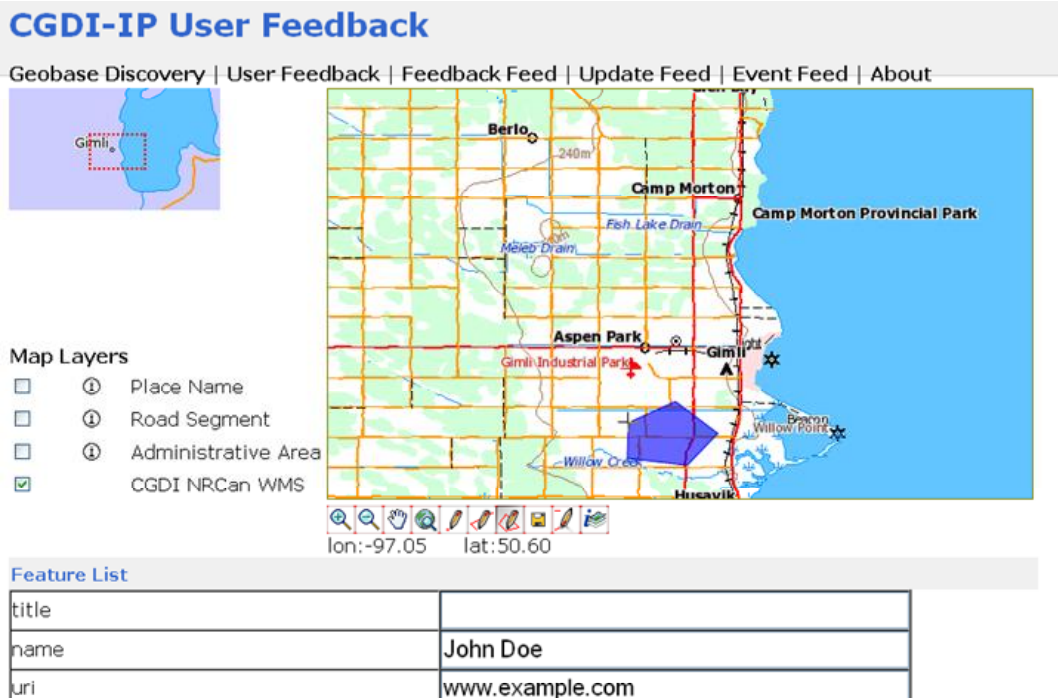
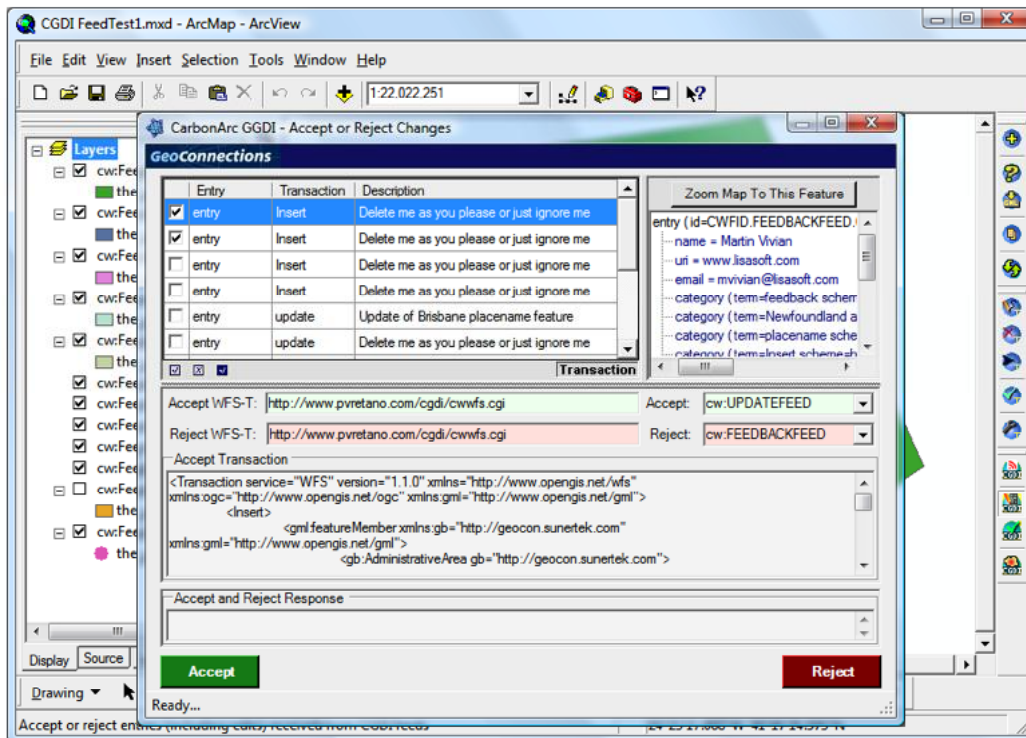


Figure 7 — User Feedback Creation Client



Copyright © 2007 Carbon Project Inc.

Figure 6 — Feedback Acceptance Client

10.4 Emergency event notification — Emergency Event Feed Service

The final requirement of the CGDI was to use the architecture to support emergency response, and through the use of the geographic resources provided by the CGDI, help ensure that affected jurisdictions could be quickly and accurately identified and notified of important events. Once again, the feedback architecture was used to satisfy this use case. The only difference being that the geographic features being reported represent new information in a new schema — for example a toxic gas plume. These features would most likely not be saved into an existing geographic data store, but used as a clipping boundary to identify municipalities that require notification, or identify roads that should be closed.

The Carbon Project modified CarbonArc to provide emergency event visualization and spatial intersection functionality. LisaSoft provided a Web-based emergency event visualization client. The Emergency Event Service is described in more detail in OGC document 08-001.

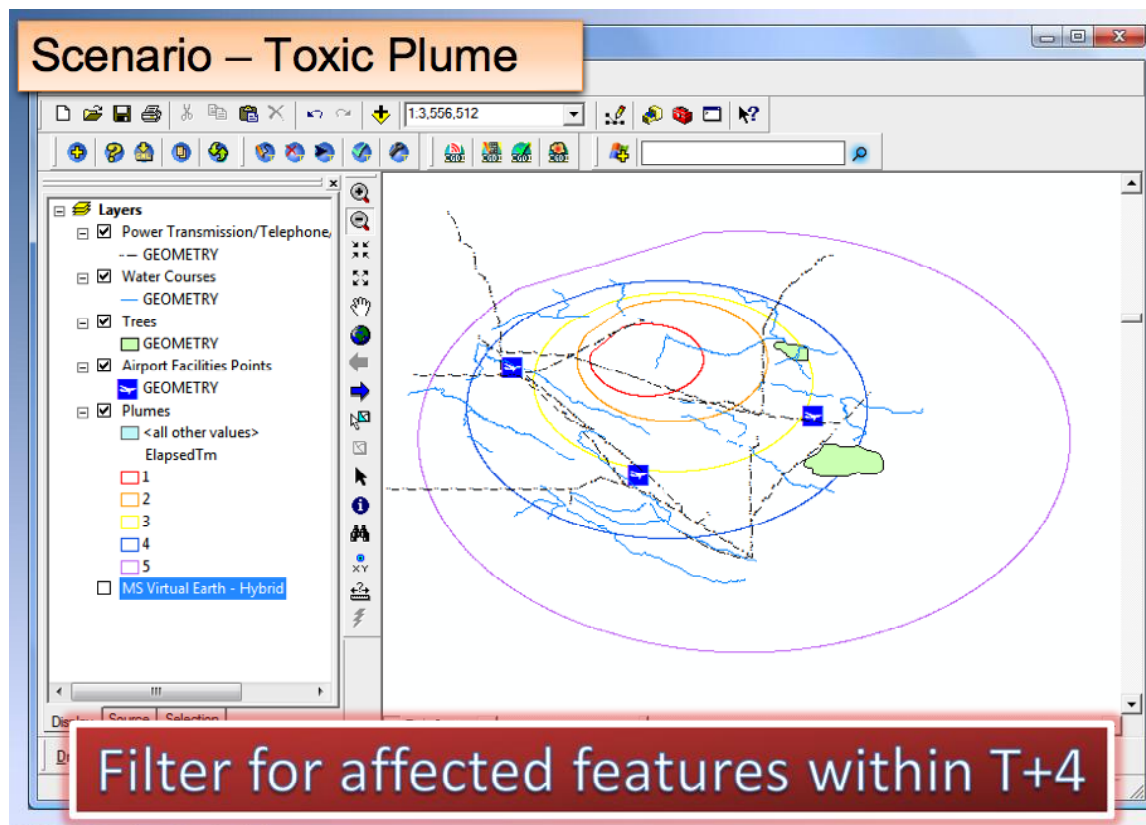


Figure 8 — Emergency Event Client

11 Findings

11.1 Shared Schemas are A Key Requirement

The use case scenarios designed for this pilot were identified based on their applicability to situations which could occur in any jurisdiction represented by one of the WFS network nodes. To achieve interoperability of client applications across jurisdictions, it was necessary to design GML application schema for all uses cases based on a shared or “community” information model.

The challenges in agreeing on content semantics were largely addressed by two considerations:

- i) The GeoBase community had to a great extent done much of the hard work involved in reaching agreements on semantic interoperability for feature definitions.
- ii) The nature of a pilot project allows simplifications which would not be possible in a production scenario. Where GeoBase had not already devised a community schema, it was possible to adapt de facto application schema standards, such as for AdministrativeArea.

As a result, a community GML Application Schema was readily agreed upon by pilot participants, to underpin the use case scenarios, and is presented in **OGC 08-002 Appendix C – CGDI IP Community GML Application Schema**.

A key challenge that the pilot faced (and overcame) was to implement the community schema within the various WFS deployed within the network. This forms the basis of a recommendation to WFS implementers, as described in **OGC 08-002 Section 6.1 Difficulty of implementing community application schema**.

11.2 Synchronization Standards Are Needed

The use cases implemented by the pilot all reflected, or placed an emphasis on especially three of the fundamental principles of the CGDI, namely that CGDI data content will be: 1) up to date, 2) maintained closest to source and 3) secure. To achieve this alignment, updates were to be performed in a secure, authenticated and authorized environment by provincial data custodians, based on integrated feedback received from external users via GeoRSS GML notification feeds. External users would transmit feedback and receive automated notification of acceptance / rejection / some combination of the two of their feedback via a dedicated GeoRSS service.

The CGDI Interoperability Pilot demonstrated the efficacy of the loosely-coupled mechanism to synchronize databases, and the next step is to agree on standards for this activity pattern. See OGC 08-001 for further details.

11.3 Service Performance Issue

The CGDI interoperability pilot did not set specific performance requirements that had to be met under given load conditions. The pilot goals sought to validate the use of WFS, WMS and related technologies for the problem domain, while demonstrating best practises for optimizing interactive application performance. These goals were attained, and some considerations regarding performance issues in use of WFS are described in OGC 08-002 Sections 6.3, 6.4 and 6.5. Recommendations from pilot participants identified the future goal of quantifying performance under predetermined load conditions in representative data processing environments.

12 Demonstrations

The project results were demonstrated live and in face-to-face demonstrations to all the participating organizations, in Ottawa, November 6-8, 2007. The same demonstration was performed live over the internet on November 30, 2007, to an audience estimated at over 500, from across North America and around the world. These demonstrations were recorded and are available for download from OGC Network at <http://www.ogcnetwork.net/cgdi> or GeoConnections at <http://geogratias.cgdi.gc.ca/download/documentation/>.