

---

# **d-cyphaMap – an Internet delivered spatial decision support system**

Matthew Lythe

GIS Technical Architect

*d-cypha Limited*

## **ABSTRACT**

The provision of dynamic mapping and spatial analysis over the Internet poses challenges both in terms of technology but also bandwidth and performance. d-cypha has met this challenge with the implementation of a three-tier technology solution called d-cyphaMap currently being offered as an ASP service in New Zealand. Underpinned by Oracle Server, the d-cyphaMap solution also incorporates software from GIS vendor ESRI and MacroMedia Cold Fusion.

d-cyphaMap differs from most GIS implementations in that it employs an Oracle database as a form of inter-process communication, adding information over and above that transferred through the GIS software. In contrast to protocols like XML and Open GIS, d-cyphaMap uses data stored in the Oracle database to generate a full-featured Internet mapping application.

## **INTRODUCTION**

Until only recently Geographic Information Systems (GIS) applications were considered to be overly data and processor intensive to be delivered in a fully distributed manner over the Internet. However a steady improvement in the supply of digital geographic data, the development of new web based technologies from leading GIS vendors and the increased accessibility to broadband communications have collectively enabled the provision of spatial data and services over the Internet to become reality.

Through the integration of leading technology solutions and a comprehensive digital landbase, d-cypha have developed a subscription based GIS service now commercially available in New Zealand. Launched in November 2002, **d-cyphaMap** is a full-featured mapping and spatial analysis application that provides spatial content and services securely to a wide range of businesses from corporate enterprises to small business. **d-cyphaMap** is an extension to d-cypha's established strategy to offer high quality data management services to targeted industry sectors.

Incorporating several foremost software components in specialised roles **d-cyphaMap** uses an Oracle database as a form of inter-process communication to deliver a complex mapping and spatial application. This paper presents the spatial ASP deployment model, the development and design of the **d-cyphaMap** service with particular focus on the integrating function Oracle server plays in the delivery of a robust and scalable application.

## APPLICATION SERVICE PROVISION

The spatial ASP model can range from very simple to extremely complex. There are many examples of successful ASP deployment models providing 'read-only' spatial services e.g. Multimap.com, MapQuest.com, Whereis.com. These deployments require a fairly streamlined architecture able to deliver relatively simple spatial content e.g. map, route, aerial photograph or location of nearest service reliably on request to a mass user base.

A conventional GIS transaction on the other hand, requiring 'read-write' capability, data creation or editing and spatial analysis is complex and data intensive. In a traditional LAN client/server environment it is not uncommon for such operations to take several seconds to complete. The outputs associated with GIS operations typically include a feature-rich map together with some attribute data.

To successfully bring a complex 'read-write' GIS application to market using a spatial ASP model d-cypha had to meet the following design criteria;

- The ability to deliver rich content and functionality to light-weight clients
- High quality and up-to-date spatial content
- High performance in a web initiated medium-high level transactional environment
- Scalability and reliability
- Strong availability profile (frequently referred to as 'up-time')
- Effective and robust security

Through advances in data, software, database architecture and chip speed many of these targets are becoming achievable and in future the bottleneck to delivering complex spatial solutions will become primarily bandwidth capacity.

Like many places in the world New Zealand has seen a significant increase in the quantity and quality of fibre cable especially in urban areas. This capacity provides a robust network through which a spatial ASP model can be delivered. Also, as networking and optical technology continues to advance, the cost and performance of delivering spatial services through an ASP model will continue to improve.

Of course implementing a spatial ASP model must have a persuasive economic reality. One of the immediate attractions is the reduced capital expenditure for GIS software and spatial data. The outsourced model also eliminates the overhead costs associated with internal spatial support services. Having a customisable and flexible application available at users desktops (via browser) with a technical helpdesk and the latest content are other compelling economic points.

The following sections describe the **d-cyphaMap** technology and data model used to integrate the software components.

## APPLICATION OVERVIEW

**d-cyphaMap** is a three tier technology solution comprising six core components each of which were chosen after substantial evaluation of available software (Table 1). Divided into presentation, application/logic and data tiers we consider these components to be the best available technologies in their class.

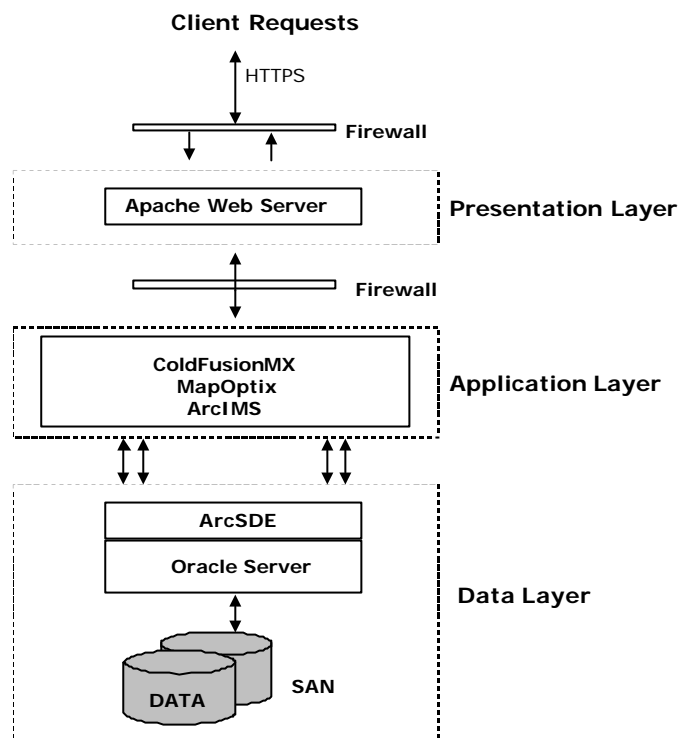
TIER	SOFTWARE COMPONENT
Presentation	<i>Apache</i> HTTP server – HTTP server
Application	<i>Macromedia ColdFusion</i> – web application server <i>GeoNorth MapOptix</i> – extensible administrative extension to ESRI ArcIMS <i>ESRI ArcIMS</i> – internet GIS server application software
Data	<i>ESRI ArcSDE</i> – GIS gateway to relational database <i>Oracle Server</i> – relational database management system

**Table 1. Core software components in d-cyphaMap.**

While at first inspection it may appear that **d-cyphaMap** is complex and unwieldy, collectively these components form a powerful, scalable and extensible architecture. The three tier configuration is presented schematically in Figure 1.

The core GIS software is ESRI ArcIMS. Deployed in a conventional manner ArcIMS provides an environment suitable for the deployment of industrial strength Internet GIS services. It offers several viewers including a thin HTML interface and two thicker Java interfaces that have extensive GIS capabilities. Indeed ArcIMS has been used for spatial ASP deployments in many international sites. However, while ArcIMS provides a robust environment for developing and serving spatial applications over the Internet it is limited in that it is totally dependent on the ArcXML protocol for communication between its core components and a web server. Also there are other limitations to ArcIMS in terms of authoring and interface development. ArcIMS does however provide powerful servers for the generation of maps (image server), spatial queries (query server) and metadata (metadata server).

In order to provide a more feature rich, dynamic and flexible interface d-cypha integrated a software component called MapOptix from United States based software provider *GeoNorth*. MapOptix is an extension to ArcIMS that provides numerous value-added features such as tabular data integration, content control, remote authoring, metadata viewing, image catalog, optimised image output and more. The MapOptix software and protocol allows a developer to deliver customised content and query capabilities to specific users through secure user profiles. Built on ESRI's ArcIMS technology and Macromedia's ColdFusion, MapOptix has a fully configurable user interface, works with any ODBC compliant database and has a scalable and extensible architecture. Using MapOptix and ColdFusion MX we have developed a powerful data model to store map properties and other configuration settings. This data model controls communication between client and application adding information on top of the specific spatial information generated by ArcIMS.



**Figure 1. d-cyphaMap architectural components**

A principal **d-cyphaMap** design specification was thin client access (i.e. client does not require a client plug-in or applet). The MapOptix software complies fully in this regard without compromising the richness of the application. It achieves this through tight coupling with the ColdFusion MX application server functions. With this architecture **d-cyphaMap** can generate specific functionality such as an image, query result, coordinates, buffer result all based on a single request.

The data tier is important in two crucial areas. Firstly, as briefly discussed above it provides configuration information and control to client requests using JDBC. This will be discussed in more detail in the next section. Secondly through ESRI ArcSDE it acts as an application server to deliver spatial data to ArcIMS. Essentially ArcSDE is a gateway that facilitates the management of spatial data in a RDBMS. To ensure application security the presentation and application layer components are located within separate DMZ's. All database transactions are exchanged from the ColdFusion Application Server and the database server located within the secure network.

## **d-cyphaMap INTERFACE**

The interface is designed for simplicity of use. It is comprised of four parts;

- Toolbar - providing navigation, query, selection and system functions,
- Layers/Legend – providing layer list, legend, search tools and buffer function

- Map frame – the principle map frame
- Results frame – for displaying selection sets

Key functions include map request (nationwide scale through to 1:500), aspatial database query, spatial query (intersect, overlay, geosearch, buffer), graphical query (select by rectangle, polygon, line) and report. In contrast to other conventional GIS applications the majority of database transactions in **d-cyphaMap** are direct JDBC database select or insert statements. **d-cypha** in particular make considerable use of Oracle InterMedia and ColdFusion Verity to match patterns in database fields. As a consequence of this and the repeated requests of core landbase data across the country the buffer hit ratio of **d-cyphaMap** is around 98% while practically all sorts are within memory. Considerable time is spent to ensure the database instance is optimally configured. This ensures the required performance targets are met both during periods of heavy load and in medium bandwidth environments.

The application also provides powerful selection set management tools whereby query results used in a particular session may be combined to solve problems and generate information products. **d-cyphaMap** also has a flexible reporting capability. For example an A4 print ready asset report including map, chart and attribute data may be generated based on a geographic query. The application is designed to integrate company data with the comprehensive landbase embedded in the application. Through integration of company data with this landbase meaningful information may be extracted. Figure 2 shows several examples of the interface.

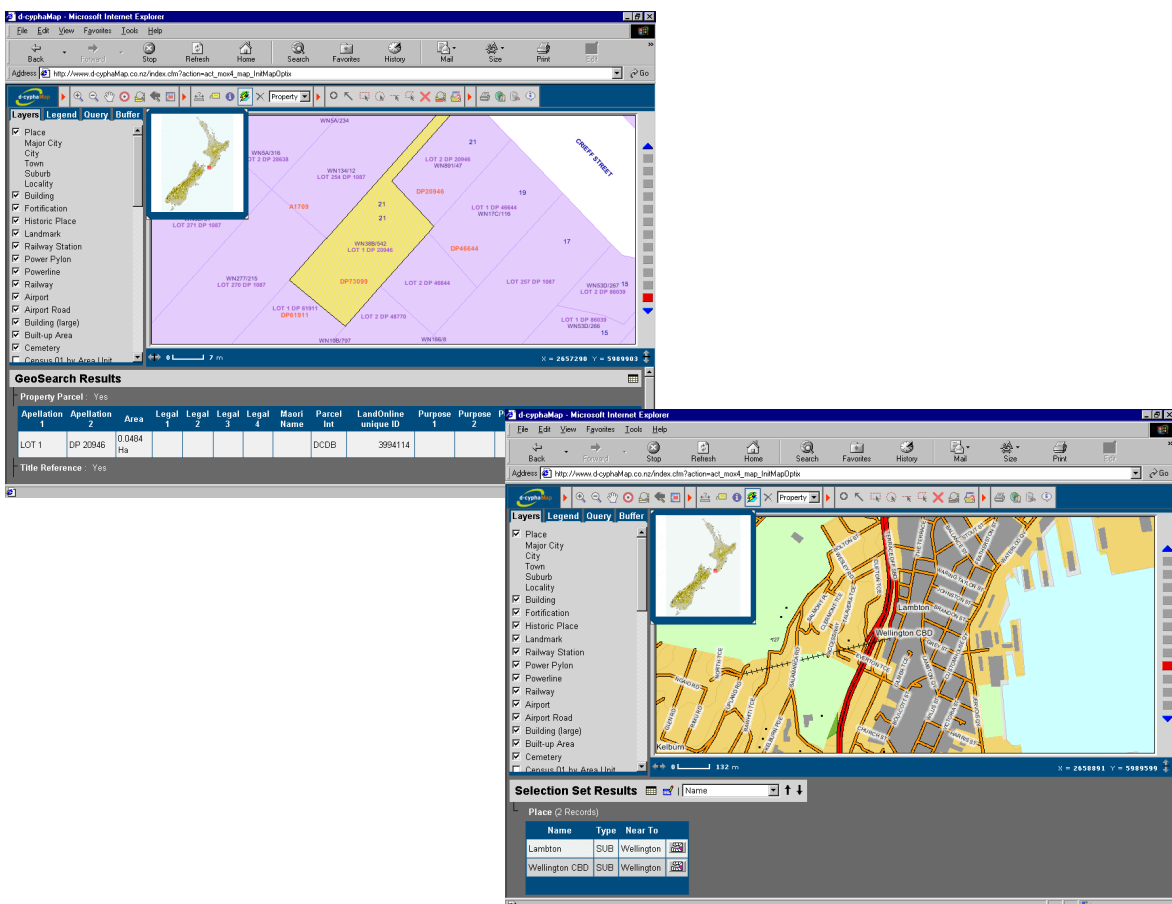


Figure 2. d-cyphaMap interface

## d-cyphaMap DATA MODEL

A fundamental point of difference between **d-cyphaMap** and other similar spatial ASP deployments is the role of the database. In most GIS applications proprietary cartographic representation documents or functional scripting languages are used to store map properties and configuration settings. In the web sphere numerous web protocols are also now available to configure spatial applications e.g. ESRI (ArcXML), MapServer, Open GIS and MapGuide. While these protocols enable the development of rich applications they require considerable development and often incorporate substantial redundancy.

d-cypha chose to use a carefully structured data model stored within Oracle server as a form of inter-process communication for the following reasons;

1. Non proprietary protocol placed less emphasis on any one software component
2. Primary commands are few, all information is referenced through unique database id's.
3. Configuration information is stored within a database for fast and easy access/update
4. ColdFusion can read configuration information before generating and processing a request to the map server

The **d-cyphaMap** data model is best described through the anatomy of a transaction (Table 2). The steps controlled by the data model are italicised.

STEP	ACTION
1	<b>Client authenticates to application</b>
2	<i>ColdFusion application reads configuration details based on user id and presents interface</i>
3	<i>Client issues request, Cold Fusion application validates ArcIMS parameters before request is passed to the ArcIMS application</i>
4	ArcIMS requests data from ArcSDE and Oracle.
5	ArcSDE searches and retrieves data and passes feature set back to ArcIMS to construct map image or attribute packet
6	<i>ColdFusion updates selection set tables and image reference and configures result page</i>
7	<i>ColdFusion presents result templates back to web server and client</i>

**Table 2. Anatomy of a d-cyphaMap client transaction**

The data model approach provides several useful advantages. It enables d-cypha to closely track sessions and transactions. For example selections sets returned from queries are stored within the database for future use. This means that a user may return to a session and continue with existing selection sets. The database also allows us to pinpoint faults quickly helping us provide a rapid solution to any application problems.

While d-cypha don't believe the data model approach is a replacement for other web protocols it does provide a robust and extensible method for developing fully featured spatial services. Using this approach, changes may be made to interface and display parameters, map render style and spatial functions quickly and seamlessly without interruption to service. In addition, through the use of an application server like ColdFusion other third party applications may be integrated within a single application in a consistent manner. In short, the data model protocol allows us to utilise specific software applications like ArcIMS and ColdFusion in an optimum manner.

## WHY ORACLE?

As the database server plays such a crucial role both in terms of spatial data search and retrieval and inter-process communication it was imperative that a robust and scalable technology was chosen. d-cypha have chosen to deploy an Oracle 8.1.7.4 database license running on Tru64 V5.0 Unix. Oracle was chosen for the following reasons;

- considered to be the optimum platform for ArcSDE
- provides extensive tuning flexibility to optimise spatial data repository
- provides powerful text searching tools to facilitate query capability
- suitability for online transaction processing (OLTP)

The **d-cyphaMap** service generates a mixed workload including spatial data search and retrieval, aspatial queries, spatial queries (e.g. intersect, overlay, buffer) and selection set management. In order to service this workload with multiple client connections d-cypha have carefully tuned the Oracle server and maintain 24x7 monitoring which generates alerts of any events.

A key feature of **d-cyphaMap** is the ability to search and retrieve information about location, property and land title. At a national level these data sets comprise several million records incorporating large varchar fields containing information such as land owners, road-name and legal description. In order to access these long varchar fields d-cypha has developed a number of queries which access catalog indexes built using Oracle 8.1.7 intermedia text. The catalog index provides text indexing that is free from maintenance overhead and text searching that has structured query support and faster response time than the intermedia context index. The operator catsearch is used to query the CTXCAT index. This operator is much simpler than the contains operator, with the intent that web-style queries can be passed through to catsearch without extra parsing.

Advantages with this procedure are;

- Multiple words are treated as an AND.
- Vertical bar is used for OR.
- Double quotes delimit phrases

Using CTXCAT queries d-cypha provide full and powerful text searching of large unstructured property information at a national level.

## THE GEODATABASE

Included within the **d-cyphaMap** landbase are over 60 separate spatial data sets and related aspatial join tables. These core data include the cadastre (property boundaries, land titles and address points), transportation links (road centreline, railway, airport, ferry links, tunnels), topography, census data (1996,2001), climate, infrastructural (powerlines, buildings), land classification and imagery.

The principal reason GIS is a data intensive application is because a typical map request requires approximately 1 MB of data to render a map. Data search, access and retrieval are therefore vitally important in the provision of spatial data using an ASP model. ArcIMS and ArcSDE work together as an integrated back office solution for fast Internet access to vector, raster, and survey data stored in a relational database. ArcSDE provides a powerful set of configuration options to optimise the storage of complex spatial data types. Through careful optimisation of the Oracle instance and ArcSDE d-cypha have been able to restrict access and retrieval times of ArcIMS spatial data requests to less than 1 second in 90% of all cases. ArcSDE also provides a native API and a C and Java API for the development of powerful spatial queries.

## CONCLUSION

The **d-cyphaMap** application provides a low risk strategic alternative to conventional software licensing for consumers interested in investing in GIS services. Through the integration of a comprehensive landbase, leading software applications and a non proprietary data model complex spatial applications may be supported on a subscription basis. Oracle server and the **d-cyphaMap** data model plays a central role in the application providing inter-process communication between the components and ensuring transaction times are as fast as possible. Downstream, as software vendors improve web-based technologies, such applications are likely to become more streamlined and capable of delivering even more rich functionality.