3010 - Keys To Success For Data-Centric Business Applications for the Web

John Kirby
Dynamic Information Systems Corporation (DISC)
5733 Central Ave Boulder CO 80301
303.444.4000

Abstract

As the World Wide Web revolutionizes how people obtain information and transact business, it places new demands on today's information systems. Once satisfied with simple documents stored in static HTML files, many organizations are now developing more comprehensive Web sites that integrate data from 'live' databases. This data comes from databases such as Oracle, Sybase, Informix, and Image/SQL.

The HP3000 platform has traditionally supported transaction-intensive OLTP applications. Now, organizations are developing Web applications that integrate data from their HP3000's, turning the HP3000 into a Web database server.

These "data-centric" Web applications are low cost, easily deployable equivalents to their Client/Server cousins. They are especially useful for company Intranets and electronic commerce applications, however, they often have performance problems. As Information Technology departments endeavor to deliver information from corporate databases to internal employees or external customers via the Web, they face common obstacles:

- Search capability is inadequate - users want Web search flexibility by the data content
- Connectivity to 'live' databases is costly to establish and maintain
- "Stateless" connections are inefficient for transactional environments

OMNIDEX for the Web provides solutions to these issues. Users obtain intuitive access to live business data immediately from their Web browser. Flexible searches are made possible through the **OMNIDEX Database Search Engine**, which allows users to instantaneously execute multi-dimensional searches and perform iterative drill downs into information in a database. Additionally, the **OMNIDEX State Manager** maintains client session information and physical connections between the database and the Web server for maximum performance.

OMNIDEX for the Web can transform your HP3000 or HP9000 into a powerful Web database server. It supports Image/SQL, as well as the major relational databases, with a single, uniform approach.

OMNIDEX for the Web - Overview

Web search engines like YAHOO! or Excite allow searches by any combination of words located anywhere on a Web page (an HTML document). Internet users now expect the same capabilities they use in those "document-centric" retrievals to be available for "data-centric" access. Most database applications do not support this powerful keyword search capability, or are too slow to meet the ever increasing demands of impatient users who want immediate responses to their queries.

Dynamic Information Systems Corporation (DISC) delivers the keyword search capabilities and speed required for highly demanding Web applications involving dynamic business data. DISC's OMNIDEX software enhances the Image/SQL database, as well as popular relational databases, with a high-speed keyword search engine, providing unlimited multi-dimensional analysis and dynamic data aggregations.

OMNIDEX allows access to information the way people think -- by the content of the data. It provides the flexibility of selecting by keywords embedded within text fields of any length, and by any combination of fields or columns using AND, OR, NOT. For example, when a user enters the keywords "Internet OR Web", OMNIDEX instantly finds matching data, even in phrases such as "OMNIDEX for the Web". This powerful keyword search technology far surpasses conventional database indexing.

OMNIDEX for the Web allows data from multiple databases and platforms such as online applications and data warehouses, as well as documents, **all** to be available from a company's Web site. With OMNIDEX on the Web, users immediately obtain the information they need, with one consistent method to access data across the enterprise.

OMNIDEX for the Web is a custom consulting solution developed by DISC consultants to meet your Web needs. OMNIDEX solutions have been implemented in over 4000 organizations for Client/Server, Data Warehouse, Documents, and OLTP applications, and now the Web.

OMNIDEX for the Web - Architecture and Integration Options

OMNIDEX Access using Standard CGI

The current methodology for writing data-centric Web applications is through the Common Gateway Interface (CGI). CGI programs are usually written in C/C++ or Perl. A typical CGI program provides the embedded database logic and generates dynamic HTML pages.

OmniAccess is the core OMNIDEX API that allows developers to write C/C++ applications to access databases on the various platforms OMNIDEX supports. The flexibility and portability of this API makes it well suited to developing data-centric applications for the World Wide Web (WWW) as well. OmniAccess supports the CGI environment and is available in C/C++ or in a Perl version (OAPerl).

Slow "Stateless" Performance

In developing data-centric CGI applications, DISC has found that the performance is not up to the standards that DISC and our customers expect. The performance problem often lies with the "stateless" nature of the Web. Each time a Web client (e.g. a Web browser) accesses a CGI program, the CGI program must establish a "state". This is very expensive in time and computer overhead -- the initialization overhead to re-establish the database connection and set up cursors for retrievals is repeatedly incurred. Add the uncontrollable latency of the Internet, and you can see why some database CGI applications have the well-deserved reputation of the "World Wide Wait"!

Maintaining State with Cookies

Some Web developers use cookies to maintain state. Cookies are files written to the client's local disk to store state information. There are several problems with cookies:

- Physical connections to databases usually need to be re-established
- Security - A client may not want a Web application to have write access to their local disk
- Lack of session management (the ability to control, maintain, and monitor individual client connections to the database)
-

Maintaining State using Open Market's FastCGI

Another technique is to store state information on the server side. Open Market, a leading Web server vendor, has developed an extension to CGI called FastCGI. FastCGI is an open API (freeware) that caches information on the Web server to maintain state. FastCGI applications are persistent so they can maintain database connections as well as other important state information such as cursor handles. FastCGI can greatly improve database performance.

The disadvantages of FastCGI are that it does not provide session management, and its functionality is not consistent across all supported Web servers (e.g., session affinity).

DISC's Answer - OMNIDEX State Manager

OmniAccess supports all CGI platforms, as stated earlier, and also FastCGI on the Open Market Web Server.

What's more, DISC developed a state manager to solve the problems of statelessness on the Web. OMNIDEX State Manager is a DISC product written in C/C++ that allows state to be maintained and monitored between OmniAccess-supported databases, the Web Server, and the Web client. Specifically, State Manager is a UNIX daemon process that sits between the database and the Web server, providing session management and persistent database connectivity through pre-allocated or dynamic UNIX named pipes (FIFO).

OMNIDEX State Manager is accessed through CGI programs using the State Manager version of our OmniAccess API (either C/C++ or Perl). State Manager uses our Client/Server architecture to allow data access to remote databases anywhere on a TCP/IP network. The databases can be located on any of our supported platforms. For example, with our State Manager architecture you can have your Web server and State Manager located in Denver, query a database in New York, and fetch data from Seattle from a client's browser located in Paris. This is truly distributed computing!

One of the most powerful features of State Manager is its session management capabilities. The State Manager API provides configurable session time-out capabilities to allow automatic disconnection of inactive client sessions. State Manager also includes an administration tool (ODXSMgr) to allow your Web administrator to monitor Web sessions and configure State Manager.

OMNIDEX for the Web - Architecture

Future Directions

DISC software will continue to evolve to take advantage of new technology as the nature of the Web changes, and as OMNIDEX further develops. Some of the new technology being proposed for the Web is to replace the stateless Web protocol (HTTP) with the new persistent-message based protocols of IIOP (Internet Inter-ORB) and DCOM (Distributed Component Object Model) from Microsoft. IIOP- and ORB-enabled Web browsers are included in release 4.0 of Netscape Navigator and MSIE. (An ORB is an Object Request Broker - an application framework that provides interoperability between objects, built in (possibly) different languages, running on (possibly) different machines in heterogeneous distributed environments.)

For more information, please contact your DISC Account Representative at (303) 444-4000 ext. 2400, or via e-mail at *info@disc.com,* or visit our Web site at *www.disc.com*.