

Illustra

Illustra's Web DataBlade Module

An Illustra Technical White Paper by John Gaffney

Illustra Information Technologies, Inc.
1111 Broadway, Suite 2000,
Oakland, CA 94607

Tel: (510) 652-8000
Fax: (510) 869-6388
Email: info@illustra.com
URL: <http://www.illustra.com>

Illustra(TM) and the Illustra logo are trademarks of Illustra Information Technologies, Inc. All other products are trademarks or registered trademarks of their respective companies.

Abstract

Illustra's Web DataBlade module is a comprehensive toolset for creating Web-enabled database applications that dynamically retrieve and update Illustra database content. You can construct simple query front ends in a matter of minutes and powerful Web applications in just a few hours with the Web DataBlade module. The Illustra Web DataBlade makes it easy for you to take full advantage of the Illustra server's many important features, including extensible data types, an underlying rules system, and Time Travel capabilities, all of which make Illustra the database of choice for managing all types of content on the World Wide Web.

The Evolution of Content Management in Cyberspace

Over the past few years, with the introduction of the HTTP protocol and HTML technology, a paradigm shift has begun to revolutionize the way information is conveyed. Initially developed for ad hoc information access to the military and universities via the Internet, the Web has evolved from a collection of simple text files to a wide variety of interactive multimedia applications that support dynamic content management. This shift offers a number of benefits, from reduced costs for publishing information to the ability to collect customer feedback through intelligent Web applications.

Most first generation Web applications were simply a collection of HTML documents connected by hyperlinks. But as the requirement to access new and varied types of information grew, so did the complexity of HTML documents. Developers soon recognized the need to dynamically generate some portion of the HTML content. The Common Gateway Interface (CGI) was adopted for this purpose. With CGI, programmers build applications that generate static HTML and then access a data source to retrieve content to build the dynamic portion of the Application Page. This second generation approach requires that programmers develop proprietary CGI applications, usually in C or perl.

Second generation applications have a number of inherent problems:

- These applications do not handle rich data types well, or readily accept the introduction of new data types, such as HTML documents, PDF, VRML, JPEG, and GIF, or audio. Typically, rich types are managed as BLOBs (binary large objects), which need to be loaded into proprietary client applications to be processed. This is an inefficient process that exacts a high overhead.
- These applications do not provide any built-in mechanism for sophisticated analysis of the Web server log. This means, for example, that you cannot easily find out what types of information visitors to your Web site are retrieving, or how often visitors are logging in to the site. Instead, you are required to write your own analysis tools, or buy a third party product.
- Integrating the database server and Web server poses a formidable challenge. Creating proprietary CGI applications is time-consuming. Typically, such home-grown applications do not take advantage of the latest developments in Web server technology, or manage the database-Web server interface well.

The Solution - the Illustra ORDBMS and the Web DataBlade Module

Illustra solves all of these problems with the core Illustra ORDBMS and the Web DataBlade module. You can easily construct interactive, intelligent Web sites and other Web-enabled applications with Illustra's ORDBMS and Web DataBlade module. For example, through a catalog application customers choose from product categories - CDs, videos, electronic merchandise, and so forth. Each product listed in a category has a detailed description page that includes still images, video, and sound. Users can view categories of products, search for specific products, or insert new products into the database, tailoring catalogs to a specific vertical interest. At the same time, with the Illustra ORDBMS, a Web site administrator can analyze usage patterns and derive other important information about an application.

Illustra's Extensible Data Types for all Your Web Content

Today, Web applications need to manage and support rich data types such as images, graphics, formatted text, and audio samples. The Illustra object-relational DBMS, is the database of choice for managing the sophisticated content utilized by today's multimedia Web-enabled database applications. With the Illustra ORDBMS you are not restricted to the handful of data types found in ordinary database systems. Instead, you can define your own data types, or use additional data types and access methods provided through other Illustra DataBlade modules.

Creating Intelligent Home Pages

Instead of merely acting as a passive data repository, the Illustra ORDBMS is also an active, intelligent system capable of notification and analysis. This is accomplished through Illustra's rules system, programmed in easy to maintain industry-standard SQL.

Illustra's rules system makes click stream analysis possible so that you can understand exactly who your customers are and how they view your application. Click stream analysis means that you can track where and how often a user clicks on a particular application area. With the feedback gathered through click stream analysis, you can build a business model around the usage of an application. You can, for example, give popular products more prominence in your application, or redesign an infrequently visited portion of your application.

Rules also enable specified events to trigger notification messages. These events can include inventory running low or a new order being placed. Notification provides useful feedback in many application areas.

Managing Your Web Site Over Time

Illustra supports the creation of multiple versions of an object and its components through a unique feature called Time Travel. Time Travel can be used to build novel applications (such as the Graffiti Wall on Illustra's Web site), without the overhead of writing sophisticated application code.

In Illustra, with archiving enabled, changes to an object do not destroy the old version of the object. Instead, new information is written and old information is archived. This old information is available online until the database administrator explicitly removes it. While the archive information is online, an application can ask for a query to be answered "as of" a certain time. This makes it possible for corporate home page administrators to maintain older versions of Application Pages and content without having to program in any special application logic. For application developers, this feature provides the benefits of version control by enabling developers to easily retrieve past versions of Application Pages. For users, Time Travel can provide quick access to information such as last week's prices in an inventory list, or last month's front page story from a news service.

Illustra Web DataBlade Module Overview

By reducing development and maintenance time, the Illustra Web DataBlade module frees you to focus on attracting new customers to your business. In standard Web database applications, most of the logic is in the CGI application, typically written in perl, Tcl, or C. Generally, this CGI application connects to a database, builds and executes SQL statements, and formats results. By providing complete integration between the Web server and the database, the Web DataBlade module eliminates the need to develop a proprietary CGI application to dynamically access database content.

Application Pages Simplify Development Effort

With the Web DataBlade module, instead of programming complex CGI code, you create Application Pages by using Web DataBlade tags. This enables you to write your own SQL statements directly in the HTML document, thus avoiding CGI scripting.

Application Pages are stored in the database. When they are retrieved, the SQL statements are executed and the results are dynamically formatted. Because they are stored in the Illustra database, you can manage Application Pages over time, retrieving past versions, if desired. This method of Web-based application development provides an industry-compliant approach, that embeds the action language

directly in the HTML document. Thus, you can extend application pages by programming with other action language tags, for example, by calling a Java applet.

Web DataBlade tags provide a rich framework with which to manage SQL statements, format specifications for data retrieval, manage variables, implement error handling, and call system commands.

Architecture of the Web DataBlade Module

The Web DataBlade module consists of two core pieces: the WebExplode function, and Webdriver, Illustra's CGI driver:

- WebExplode parses the Application Page. The Application Page is stored in the database. When WebExplode finds Web DataBlade tags, it issues SQL statements and formats the results based on your formatting instructions.
- Webdriver handles all aspects of database interface and enables you to customize Web applications based on information obtained from a configuration file, the CGI environment, URLs, and HTML forms. Webdriver manages the database connection and retrieval of the Application Page from the database. Webdriver's Application Page Builder, Figure 1, with its convenient graphical interface, is a Web DataBlade application that you can use to create, edit, delete, and execute HTML pages in Web applications.

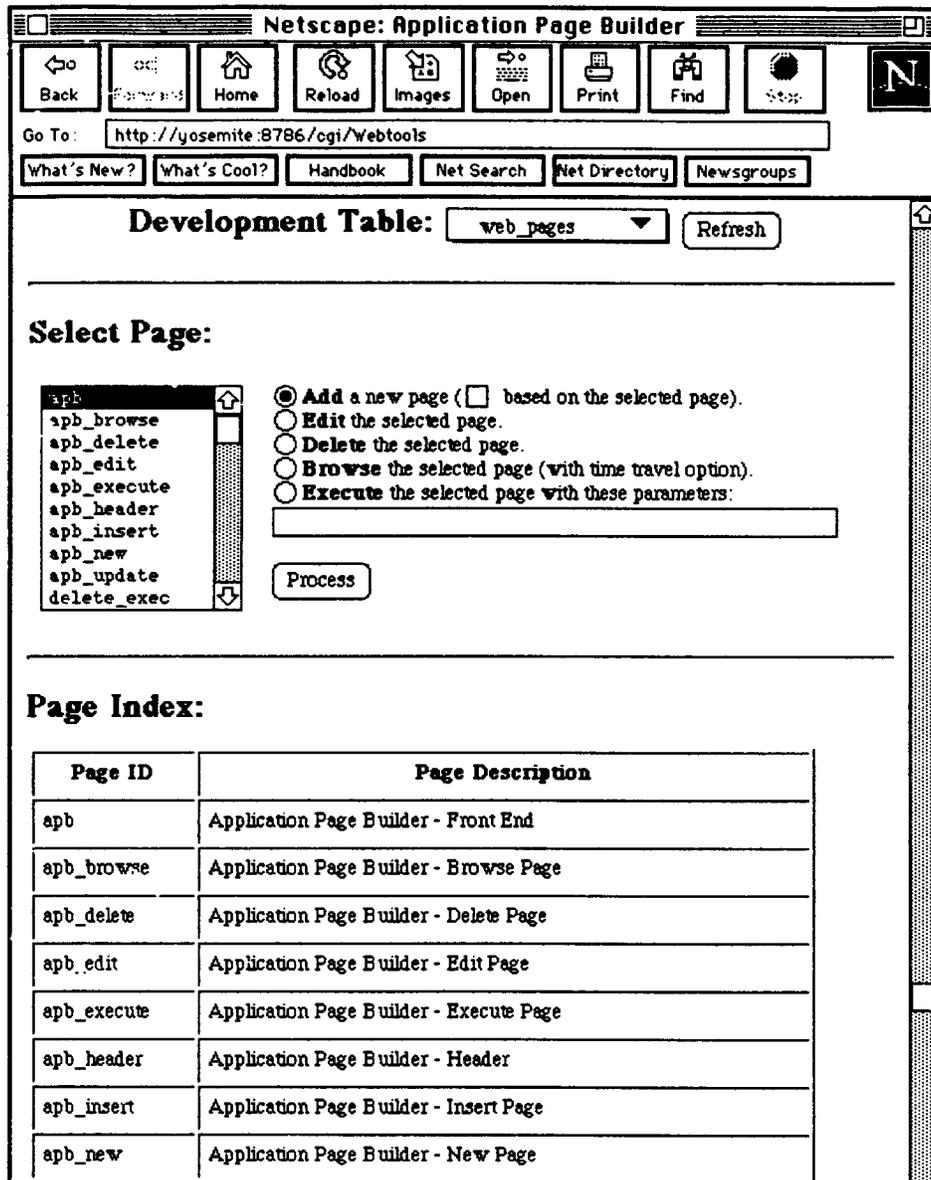


Figure 1. The Application Page Builder Interface

Performance, Load, and Scalability

The Web DataBlade module also provides Webclient, a tool for managing high-load web sites. Webclient, which can be used as an alternative to Webdriver, is a CGI-based application that can be integrated into any proprietary API that supports CGI.

Webclient optimizes connection management between the Web server and the Illustra database. Web site administrators can configure a predefined number of database connections, so that no matter how many CGI processes are invoked, the requests to the database are all managed and processed through these already-established connections. This feature means that administrators can customize their system configuration to optimize the underlying CPU architecture and efficiently handle peak loads.

How the Web DataBlade Works

The core functionality of the Web DataBlade module is the WebExplode function. Most applications designers will call this function through the Webdriver interface. Webdriver extracts an Application Page from the database by issuing an SQL select statement. WebExplode parses the results of the select, executes the embedded SQL statements and generates an Application Page, which is then rendered to the browser. Unlike parsing functions for other commercial products, WebExplode executes inside the Illustra server for high performance.

The diagram below illustrates how the Web DataBlade module executes queries and processes results.

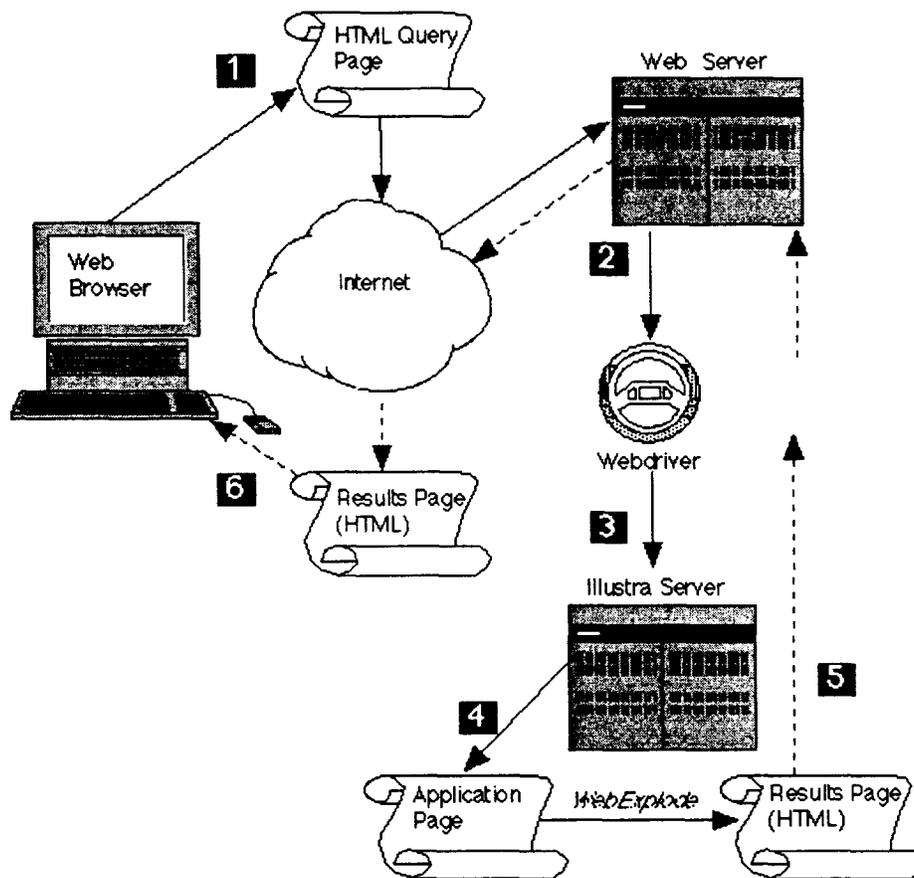


Figure 2. Web DataBlade Architecture

1. The Web browser queries the Web server with a form submission request or other query.
2. The Web server launches the client database application (Webdriver or Webclient).
3. Based on a defined set of parameters, the driver connects to the specified database as a specified user, and then selects a specified Application Page, stored as a row in a table.
4. WebExplode parses the extracted pages, executes the embedded queries, and then formats the results. This is accomplished within the database server.
5. Based on the results of WebExplode, an Application Page is placed in the database server memory space and returned to Webdriver (or Webclient).
6. Webdriver (or Webclient) renders the Application Page through the Web server to the client

browser program.

With the Web DataBlade module, a Web-enabled database application becomes a collection of HTML pages containing embedded SQL statements stored in the database. A small (but powerful) set of database-oriented markup tags embed the SQL in the HTML. These Web DataBlade tags adhere to all the properties of SGML processing tags, and at the same time make it possible for you to embed any SQL statement within the Application Page. The following table lists the tags and their function.

Tag	Purpose
<?MISQL>	Specifies the SQL, formatting instructions. Provides an optional conditional flag.
<?MIVAR>	Manages variables within the Application Page.
<?MIERROR>	Manages error processing within the Application Page.
<?MIEXEC>	Enables the application designer to call system services from within the HTML page.
<?MIBLOCK>	Enables you to delimit logical blocks of HTML based on a conditional statement.

Core Features of the Web DataBlade Module

- Embeds industry-standard SQL in HTML
- Is not server-dependent; runs on any server
- Implements error handling
- Manages variable processing
- Supports conditional statements
- Provides a set of template examples
- Provides an application development tool for building Application Pages
- Stores and manages applications in database
- Works with all industry-standard Web servers.

Advanced Features of the Web DataBlade Module

- Scalable driver manages and handles high loads
- Set of server functions enables sophisticated application development
- Provides comprehensive error reporting and logging
- Maintains transactional integrity within an Application Page

Benefits of the Web DataBlade Module

Saves Time.

The Web DataBlade module enables rapid application development and maintenance. Application designers are freed from low-level coding work and can focus on application flow and design. Applications are structured modularly in HTML pages and are therefore easy to maintain.

Reduces Development Costs.

The Web DataBlade module eliminates costly hours of CGI interface programming in perl, C, or Tcl. With the Web DataBlade module, application designers with HTML and SQL knowledge can quickly build Web-enabled database applications. The Application Page Builder development tool provides an

easy to use interface for assembling Application Pages.

Offers an Open Development Environment.

The Web DataBlade module works with all standard HTTP and standard Web servers. It runs on most UNIX platforms, and Microsoft Windows NT.

Benefits of Using Illustra for Web-Enabled Content Management

The Illustra ORDBMS offers the following unique benefits:

- Supports rich data types such as HTML, VRML, PDF, JPEG, GIF and audio, as well as traditional data types such as integer and text.
- Has the unique built-in capability to provide content search on all data types.
- Illustra's rule system supports click-stream analysis in Web applications, and provides for events being triggered on select statements.
- With Time Travel you can now easily manage your application and content over time.

Test Fly These Examples

Several demonstration applications are available in the Test Flight area of Illustra's home page. These applications illustrate many of the concepts discussed in this paper.

- **The Illustra Graffiti Wall Demo** uses Illustra's no-overwrite storage to Time Travel over archived data in the database.
- **The Illustra/IUMA (Internet Underground Music Archive) Demo** shows the use of complex multimedia data, in this case audio and video data types, in an Illustra database.
- **The Visual Intelligence (VIR) Demo** stores over 1200 GIF images. You can search these images by actual content: shape, color, pattern, and so forth.

Conclusion

Illustra's Web DataBlade module provides application developers with a toolset for quickly creating attractive, media-rich, database-driven Web sites. The Web DataBlade module solves many of the problems inherent in previous generations of Web application development products, including difficulties surrounding the management of complex data types and the need to write proprietary CGI access and management code.

References

Colton, Malcolm. *Multimedia Asset Management*. Copyright, *Illustra Information Technologies*, 1995.

Introduction to Illustra, Illustra Information Technologies. Part # ILL0795-01III *Illustra Web DataBlade User's Guide*, Release 2.1 Beta. Sept. 1995. Part # WEB-00-12-UG