Application Development Trends

DEVELOPMENT SOLUTIONS FOR CORPORATE SOFTWARE MANAGERS
WWW.ADTMAG.COM

Application Engineering/Development Tools

More than a framework for Web-based transactions, CS2K has evolved into a foundation for integrated e-business.

Microsoft Commerce Server: The next generation

MICROSOFT'S FOURTH-GENERATION ELECtronic commerce product, Commerce Server 2000 (CS2K), is not only a framework for Web-based transactions; it is a platform for doing business online. Arguably, the most improved aspects of CS2K are the tightly integrated business intelligence analytics and customer knowledge acquisition, management and data mining.

Like its predecessors, Commerce Server 2000 should not be considered a "product solution." Rather, CS2K is a platform solution upon which end-user developers, third parties or a combination of the two can develop finished Web solutions. Unlike its predecessors, CS2K has been totally reengineered for XML, integrated from top to bottom and made more accessible to developers and third-party add-ons. The process was possible both through naming conventions that flow throughout the platform and by refining the architecture of the platform so that developers can add functionality without having to rip-and-replace major portions of the out-of-the-box product.

Finally, and perhaps most importantly for companies that are choosing a commerce platform that will either make or break their Internet strategies, Microsoft has both integrated the product with the other members of the .NET server family and tightened up the third-party performance standards through independent testing labs and certifications. The integration's purpose was to ensure that ill-conceived or poorly executed add-ons do not create performance bottlenecks or introduce memory leaks or component instabilities that contribute to less than rock-solid reliability.

The entire product, platform and related implementation and management infrastructure is beyond the scope of a single article. This article will focus on the underlying architecture for analytics, the components of the platform that address analytical issues and the tools available to tap the customer knowledge accumulated by the platform.

Overview

Before discussing the analytical architecture, understanding a little about the overall architecture of

BY RICHARD L. WARREN

Figure 1: Commerce Server 2000 architecture

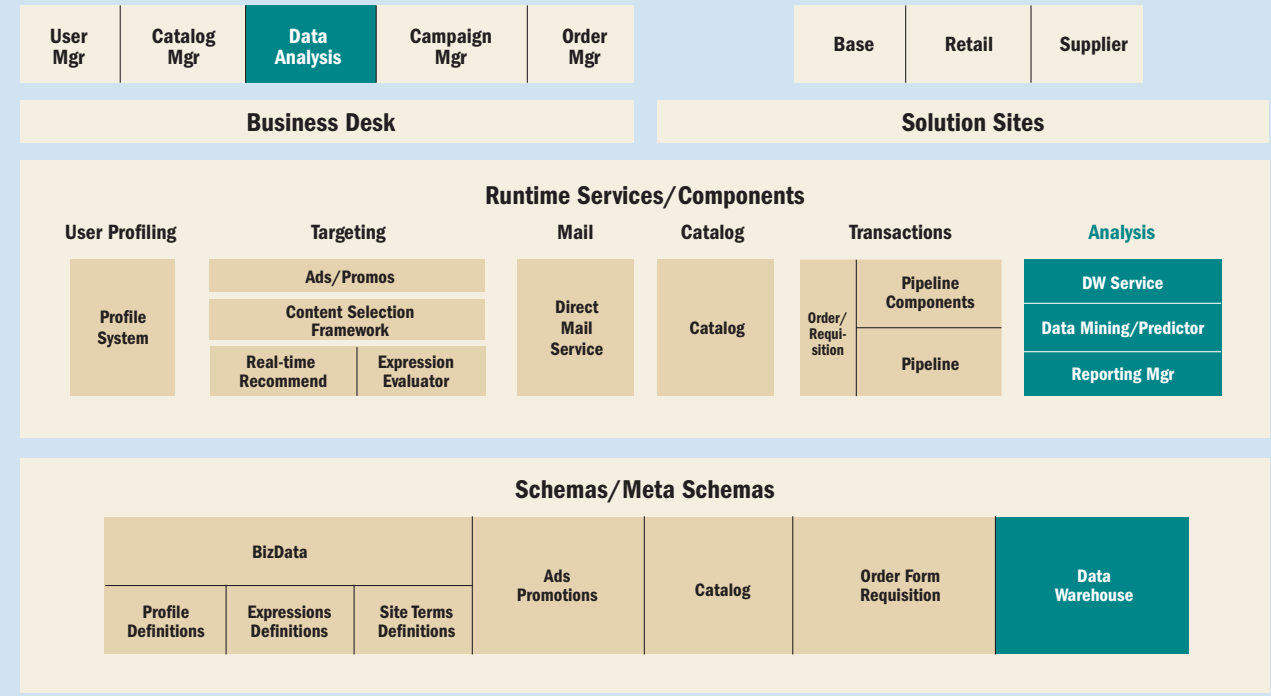

Source: Richard L. Warren

the Commerce Server 2000 (shown in Fig. 1) would be beneficial.

At the foundation of Commerce Server is a robust data warehouse hosted on SQL Server 2000. (Commerce Server 2000 can, in fact, operate in SQL Server 7.0 in what Microsoft refers to as an "alternative configuration.") Running above the storage services are the Runtime Services and Components. This is where all the platform intelligence of Commerce Server sits. It includes

- user profiling that can be integrated into full eCRM suites;
- targeting that deploys the user profiling to selectively present not only ads and promotions but an entire content selection framework that can be used to target site content in conjunction with Interwoven's TeamSite, a "lite" version that ships with Commerce Server;
- direct Mail Services that can be used for both E-mail and as a merge source for "snail mail";
- a strong, yet totally replaceable, Catalog Service that, in conjunction with the profiling system, can generate 1:1 pricing and product presentation/availability;
- a much more robust Transaction Service than was present in the previous version, aided and abetted by a certification program that will ensure third-party adherence to Microsoft's architectural and performance criteria before it can bear the Commerce Server 2000 logo; and, finally,
- Analysis Services encompassing not only data warehouse services, but data mining and prediction services as well as the more rudimentary, yet critically important reporting services.

CS2K provides the architecture for the integration of not only e-commerce event data, but related CRM, ERP and SCM business intelligence as well.

Sitting atop these services are two sites: one managerial and the other actual, running the commerce site itself.

Having taken two years to evolve from its Site Server Commerce Edition 3.0 predecessor, Commerce Server 2000 is too large a topic for a single article to even introduce, much less drill down on any particular facet of the platform. The highlighted aspects of Figure 1 are essential to understanding the Analytical Architecture. Without dwelling too much on their integration with the other services and components at the middle tier of the platform (which are extensive), the remainder of this article will focus on the foundation, services components and managerial interface that make it remarkable not just for transactional e-commerce sites, but as the enabling platform for run-of-the-mill corporate Web presence sites as well.

Analytical architecture

At the highest possible level, Commerce Server 2000 provides the architecture for the integration of not only e-commerce event data, but related CRM, ERP and SCM business intelligence as well. As shown in Figure 2, many sources other than transactions are made available through the Commerce Server Business Data Warehouse and the OLAP Store, but all information

flows through those stores to the Business Intelligence components on the right side of the flow diagram. The goal of the architecture is to provide a "360° view" of the business, whether the transactions take place through traditional means or directly through the Commerce Server's own transaction services, while at the same time providing closed-loop feedback to itself for segmentation and targeting purposes.

Generally, the analytics system processes two kinds of information: data from events that occur on the site such as transactions, page views and other "clickstream-related" data, and data from more static sources such as catalogs, user lists or profiles and other databases. As shown in Figure 3, both event-related and non-event-related information flow into the data warehouse through SQL Server's Data Transformation Services (DTS). DTS provides developers with a set of tools to extract, transform and consolidate data from any number of sources by graphically creating a business-function-related group of tasks, including their sequencing and configuration, and then scheduling the execution of the task group (called a "package") using SQL Server's scheduler. This, too, is a major change from the previous version's reliance on custom log processing routines and utilities. Using DTS, developers are working with the same interfaces and processes they use for all other SQL Server processing. Coupled with Internet Information Server's ability to create log files every hour, business managers can be kept up-to-date on the status of their online business with very little to no manual intervention by their IT support staff.

The data warehouse itself is tightly integrated with Commerce Server and is fully compliant with the W3C extensible Web log file format. Beyond the out-of-the-box schema, the warehouse is completely extensible by database developers who can both import and extend the initial schemas and access data and OLAP services through standard development interfaces. The net effect is offering a far greater range of processing flow and analytics through far fewer interfaces — in a word, simplicity.

Figure 2: Overall information flow

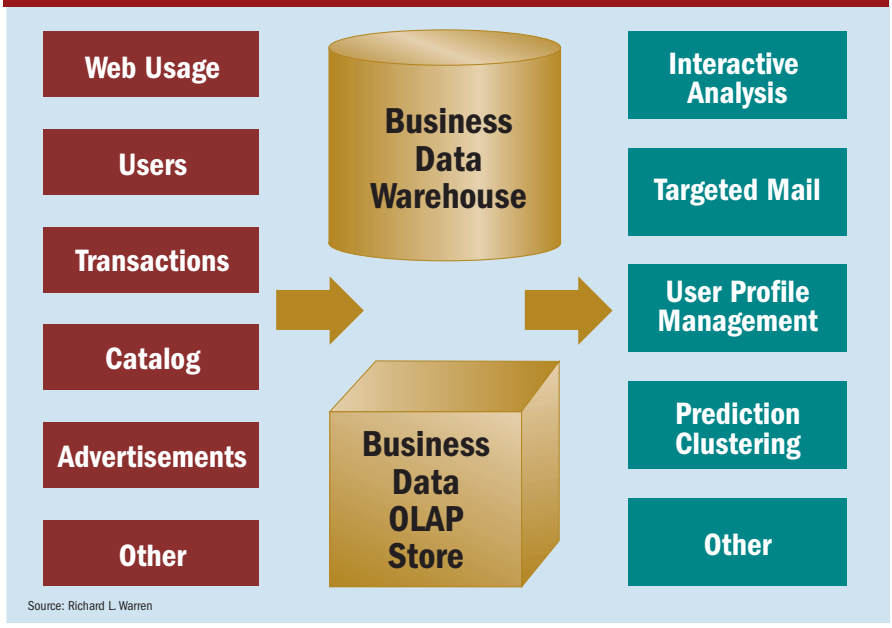

Source: Richard L. Warren

Analytical interface: The Business Desk

The end point of the process flow for analytical information within the Commerce Server platform is the Business Desk, shown on the extreme right-hand side of Figure 3

Figure 3: Overall process flow

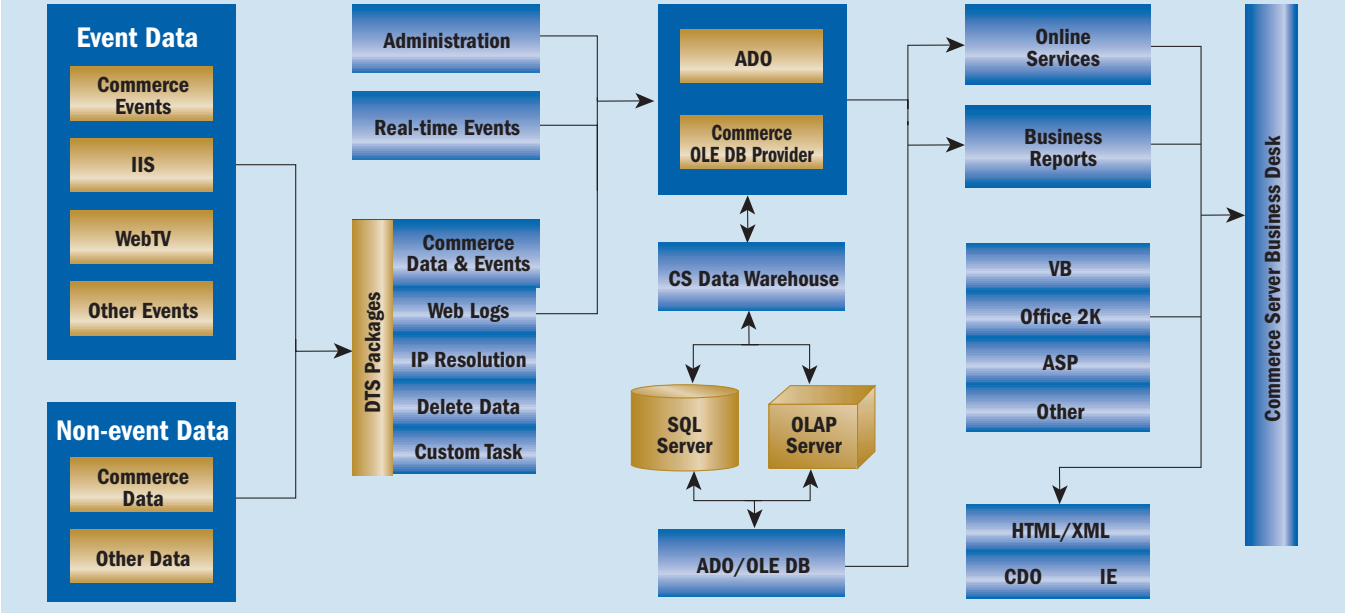

Source: Richard L. Warren

Figure 4: The Business Desk

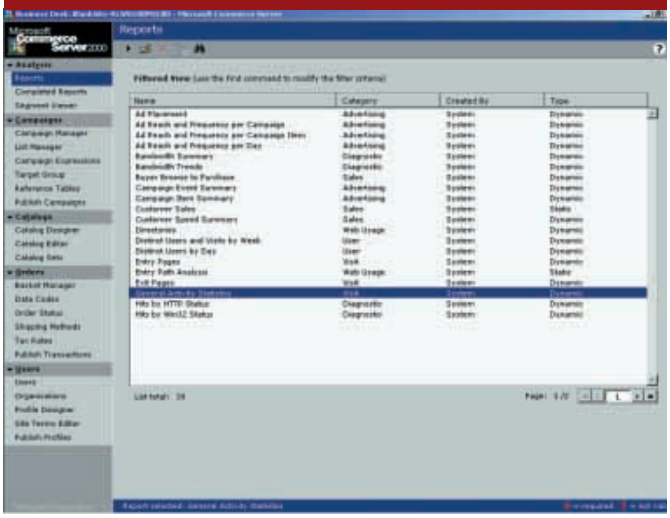

The Business Desk displays both the major components of the console and the built-in standard analytical reports.

and in the Figure 4 screen shot. This is another major departure from previous versions of the platform: the separation of a commerce site's technical functionality from its business functionality.

A developer or technical manager can manage a Commerce Server site through a Microsoft Management Console (MMC) "snap-in," similar to all the other Microsoft server MMCs. The Business Desk, on the other hand, is an interface that allows the IT team to segment business functionality into a console, allowing business managers to run their own online enterprises.

In earlier versions of the product, there was no distinction between the business and technical functionality of the platform. The administrative console that came with the "starter stores" was more focused on the structure of the site and not the business flow of information within the site's components. What was worse, the out-of-the-box consoles were not modular, the components were hard (if not impossible) to reuse and new functionality was difficult to add. In short, the administrative consoles were not "production quality."

In Figure 4 the right pane lists about half of the standard analytical reports that come built into the "blank" solution site. Along the left-hand pane are the other major components of the console including support for advertising campaigns, multiple catalogs and order and user management.

From a developer's perspective, the three most important aspects of the Business Desk are how easy it is to customize and extend its functionality, its modularity with an underlying foundation of built-in and common functions and its "production quality" out-of-the-box so users can have both an example of best practices and a solid starting point.

The Business Desk is built as an HTML Application (HTA) with the categories contained in an XML file called "bizdesk.xml" located in the Config directory of each Business Desk. The modules within the categories are also driven by an XML file and security for the modules can be controlled by using Access Control Lists (ACLs) within the regular NT security model. If an individual business manager does not belong to the group the administrator has designated as having access to a particular function, that function does not show up under the category. This integrated security model lets administrators develop the complete Business Desk functionality, then make it appear different for different groups of business managers. To change the way a manager sees the desk, the system administrator only has to change which groups the users are members of, not the Business Desk itself.

The "analysis" category for the simplest solution site, the "blank site," is driven from the analysis.xml file. Instead of using a separate file for each module within the category, this file defines all of the modules in just one file. This approach should be taken when all users will either have or not have access to all modules within a category. By setting the ACL on the analysis.xml file, access to all analytical functions can be controlled.

Applying the module configuration file analysis.xml to the architecture of the Business Desk (shown in Fig. 5), generates the three modules shown in Figure 4 — Reports, Completed Reports and Segment Viewer — along with the task buttons and menus needed by the Business Desk.

The "standard" reports that come with Commerce Server 2000 and the OLAP and SQL cube and table definitions are found in a SQL file called csreports.sql that is processed when a site and its associated Business Desk are created. There are four different kinds of reports: static SQL-based reports, static OLAP-based reports, dynamic SQL-based reports and dynamic OLAP-based reports. The Software Developers Kit (SDK) contains a sample of each script. The key differences are the data source and when the report is run.

Static reports are run as soon as they are requested and then stored with the report data in the data warehouse. The typical use of a static report occurs when the user needs to be able to display a report on a Web site or export the report for display within the Business Desk itself. Good examples of static reports include usage reports for periods that are already concluded (for example, "Last Month's Site Usage," "Last Week's Site Usage" and so on). Dynamic reports, on the other hand, contain only the report definition and run against the data warehouse every time the report is requested. The Reports Module of the Business Desk can be used to run these reports as well, but a key differentiator is that dynamic reports can be produced in both PivotTable and PivotChart formats, allowing business managers to play "what if" scenarios with Excel spreadsheets instead of with the analysis server.

The goal of the architecture is to provide a 360° view of the business, while at the same time providing closed-loop feedback to itself for segmentation and targeting purposes.

The SDK and forthcoming Resource Kit (which is not expected until Spring 2001) also contain the documentation that development teams need to start to tear into the object model underlying the Commerce Server platform.

Third-party tools

Having created this rich development platform for analytics, it is no surprise that several of Microsoft's key ISVs have leapt in to offer pre-built extended functionality for Commerce Server 2000.

Harmony House (www.harmony.com) will be revving their Harmony 360° analytical engine and releasing it in early 2001. The only information publicly available on the product at press time was a 21 Mb multimedia demo that, with a good connection and some patience, provides the gist of the eventual product's capabilities. The company intends to make the version both Windows 2000- and NT-compatible.

Knosys also has only very minimal information available at their site (www.knosysinc.com), but their ProClarity 3.0 OLAP Client software is available now and will work against the OLAP cube in the Commerce Server 2000 data warehouse. A version tailored to Commerce Server is also expected in Spring 2001.

Finally, Visual Insights (www.visualinsights.com) has a beta of their eBizinsights CS for Commerce Server 2000 available online as well as product specifications and expected product availability.

All of the Analysis ISV partners have had a considerable amount of time to bring their products along given the length of time Commerce Server has been in development. But, with the RTM version of the product only now reaching resellers, expect some delay in the availability of add-ons, especially if they undergo Microsoft certification.

CS2K will provide business managers the kind of instrumentation they need to drive the site looking through the windshield instead of the rear-view mirror.

Summary

Microsoft's Commerce Server 2000 has all the advantages of a fourth-generation platform in terms of functionality, performance, extensibility and maximum leverage of various acquisitions during the past few years.

Unlike its predecessor, Commerce Server 2000 has been totally redeveloped and reengineered to be as completely XML-based as possible. Even many of the proprietary COM/DCOM component-to-component links were removed and reworked in XML to provide a more open, loosely coupled extensibility and integration platform than would otherwise have been the case.

The Business Intelligence driven by the underlying Analytical Architecture is impressive, and the ability to extend the baseline capabilities will provide business managers the kind of instrumentation they will need to drive the site looking through the windshield instead of the rear-view mirror.

The key is to not think of it purely as a transactional, e-commerce development platform but, rather, as a foundation for integrated e-business whether or not money actually changes hands online. The analytical architecture, tools, third-party offerings and "hooks" for eCRM integration make a very strong case for, at the very least, evaluating it as a company's primary site architecture and platform. ●

Richard Warren is chief strategy officer for Internosis, an Arlington, Va.-based Microsoft-centric consultancy. He can be reached via E-mail at rwarren@internosis.com.

Figure 5: Business Desk architecture

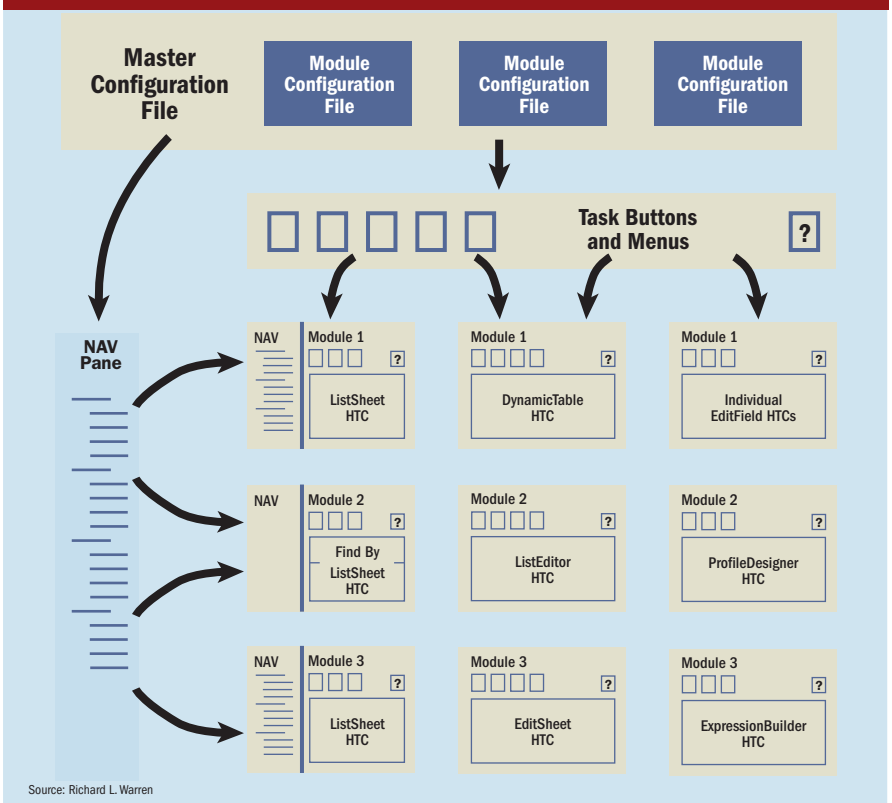

Source: Richard L. Warren