

ZAPTHINK ZAPNOTE™

EXCELON: XIS *ROBUST NATIVE XML DATA STORAGE AND MANAGEMENT*

Briefing Date: December 6, 2001

Analyst: Ronald Schmelzer

Abstract

There are two major categories of XML data store: extensions to relational database systems (RDBMS) and a new category of "Native" XML data stores (NXDs). While extensions to RDBMS systems simply enable RDBMS databases to map XML documents to relational tables, NXDs allow users to insert XML documents directly into the system without need for mapping or interacting with anything besides the XML document. eXcelon's eXtensible Information Server (XIS) is an "XML data management system (XDBMS)" that is aimed squarely at the problem of storing arbitrarily structured XML documents. Among other features, XIS provides node level management of XML data, dealing with XML document information at the element level, rather than the document level.

All Contents Copyright © 2001-2002 ZapThink, LLC. All rights reserved. Reproduction of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. ZapThink disclaims all warranties as to the accuracy, completeness or adequacy of such information. ZapThink shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice. All trademarks, service marks, and trade names are trademarked by their respective owners and ZapThink makes no claims to these names.



Native XML Data Stores

As XML proliferates in the enterprise, the need to store and retrieve it is increasing at a rapid pace. Until recently, the only means for persisting XML documents was by using the file system. This approach, while simple, is tremendously inefficient at storing XML and ineffective at querying XML documents. As a result, a variety of new solutions are available for storing XML, each having different advantages and trade-offs. There are two major categories of XML data store: extensions to relational database systems (RDBMS) and a new category of "Native" XML data stores (NXDs). While extensions to RDBMS systems simply enable RDBMS databases to map XML documents to relational tables, NXDs allow users to insert XML documents directly into the system without need for mapping or interacting with anything besides the XML document.

NXDs themselves are based on a variety of different architectural platforms – some even sitting on top of RDBMS systems. What differentiates them from the "extension" market is that they present an XML-only interface to users, rather than an RDBMS or OODBMS specific interface. This difference allows users to abstract the actual storage mechanism from the document insertion process. Basically, to the user it seems that they are simply inserting an XML document, rather than having to go through any steps to prepare the database for XML insertion or retrieval. This provides clear benefits for users who have a large number of XML documents that may vary in structure. However, XML extensions to RDBMS systems may make more sense when data is originally stored in relational tables and they simply need to be exposed as XML.

The eXtensible Information Server (XIS)

One of the early developers of NXD systems is eXcelon Corporation, formerly known as Object Design Inc. (ODI). The company has leveraged their technology for storage of object-oriented components as a potent solution for storage, retrieval, and management of XML documents. The solution, known as the eXtensible Information Server (XIS) is an "XML data management system (XDBMS)" that is aimed squarely at the problem of storing arbitrarily structured XML documents.

The system provides node level management of XML data, dealing with XML document information at the element level, rather than the document level. XIS allows users to create, read, update, and delete element or XML document fragment data, thus avoiding the need to fetch and parse entire XML documents. By managing XML at the node level, XIS performs all access, search, and update functions in the engine, without having to lock entire XML documents, which can be megabytes in size, and load them into memory. Another side effect of its element-level granularity is the system's ability to accept arbitrary XML documents

TAKE CREDIT FOR READING ZAPTHINK RESEARCH!



Thank you for reading ZapThink research! ZapThink is an IT market intelligence firm that provides trusted advice and critical insight into XML, Web Services, and Service Orientation. We provide our target audience of IT vendors, service providers and end-users a clear roadmap for standards-based, loosely coupled distributed computing – a vision of IT meeting the needs of the agile business.

Earn rewards for reading ZapThink research! Visit www.zapthink.com/credit and enter the code XISROB. We'll reward you with ZapCredits that you can use to obtain free research, ZapGear, and more!

For more information about ZapThink products and services, please call us at +1-781-207-0203, or drop us an email at info@zapthink.com.

without requiring a DTD or XML Schema definition. The system can then manage this data at low-level without storing it as a monolithic data entity or BLOB. The system provides a sophisticated indexing and a patented caching architecture for query support and real-time XML processing. XIS provides built-in support of XPath 1.0 and the W3C XQuery XML document query standards, and supports XSLT transformation out of the box.

The XIS database engine, at its base core, is based on the Object Store technology, but users are not even made aware of its presence. The lowest level API that's exposed to users is the DOM API, which while based on Object Store, exposes only an XML-based interface. In this manner, XIS is a true NXD. To protect the XML document from improper access, only stored procedures and triggers can walk the persistent DOM trees. The technology provides all the capabilities of an RDBMS including ACID transactions, a scalable query architecture, incremental update, caching, locking, and database management tools. The system uses a modified version of the Xerces SAX parser internally, which parses and generates tree of XML nodes that are persisted internally. In this manner, the system can support any document regardless of document structure or encoding. XIS provides three types of indexing at the node level: structural indexes that allow users to query particular branches of documents, value indexes to allow numeric or text searches on certain nodes for comparison, and word indexes that allow free-form word searches within text nodes.

XIS supplies back-end connectivity functionality through "XConnects." XConnects transforms data and content from legacy sources into standard XML. The system includes native connectors for over 100 different structured and unstructured sources as well as a flexible toolkit for developing custom adaptors. Through XConnects and iWay's EDA, XIS supports RDBMS systems, enterprise and legacy data sources such as SAP and Siebel, file-based formats such as Paradox and dBase, mainframe systems such as CICS, Lotus Notes, delimited ASCII, Btrieve, and XML-based formats including HL7 and DIALOG. The system provides a visual tool to define mappings from these data sources into and out of XML format as well as mappings between XML formats. The tool can also generate custom mappings for unstructured data sources such as mainframe screens, word processing documents, and HTML files. All of this can be controlled through Java APIs and a scripting language, based on Microsoft Visual Basic.

At the core of XIS is the Dynamic XML Engine (DXE) that provides the actual node-level management of XML data. DXE presents a file system-like interface for managing XML data, which can be partitioned across multiple systems and disk storage arrays. In addition to a high-performance storage engine, DXE provides an XSLT transformation as well as support for DOM Level 1 and DOM Level 2 (Core) API standards. The system also supports a number of document linking mechanisms including the notion of "Binders" that can bind related XML documents together into a single virtual document view. Binders can include different document types and allow for aggregated document query and retrieval. DXE allows users to index individual XML elements or groups of elements, and supports both value indexes for element and attribute comparison searching and word indexing for text searches within document nodes. Queries can span single documents or multiple documents aggregated in a binder.

DXE provides support for "live" updating of XML documents, using XML-based update expression for updating individual elements or groups of nodes. DXE server extensions or DXE servlets run inside the server's process using the embedded Java 1.3 virtual machine and provide functionality similar to stored procedures and triggers. Triggers can be attached to any node. When the node gets its content updated, the trigger can be fired. In addition, users can extend the product using Java or COM extensions and a JCA resource adapter that enables XIS to participate in container-managed transactions. Finally, the DXE Manager provides centralized management for distributed processes, routing policies, and load distribution.

XIS also includes a suite of visual tools for simplifying and enhancing the XML experience. The Stylus Studio is a visual environment for creating XML, XSLT stylesheets, and DTDs with a built-in Java editor and debugger that can handle all XML documents and XSLT extension functions that are written in Java. Other tools include a visual, explorer-like browser for interactively browsing and managing DXE repositories, a visual environment for creating and executing query and update expressions, a means for managing distributed cache processes, defining routing policies between applications, and creating caches for load balancing through the DXE Manager. The tools provided also include a way to manage online backups and restoration of DXE repositories, offline replication, and bulk loading of XML data from other systems.

Customers & Release History

Since the company is over 12 years old, and had a very large percentage of the object database market (with their ODI product), the company has been able to successfully leverage its leadership by migrating its existing customers and attracting new ones to its XIS platform. The recent merger with C-Bridge has also given it greater capability to tailor its products to specific vertical industries.

The XIS system was first released in 1999, which was very early for NXDs. As a result, they have seen hundreds of customers for their product line, making it a strong contender for leading the market in sales. Historically, they have sold very well in the Asia-Pacific and European markets, with over 50% of their revenue being derived from non-US markets. As a result, the company expects continued rapid growth in both the US as well as international markets.

The product is priced at around \$25,000 for a developer license and \$40,000 per CPU for an in-production system. For accurate pricing, please contact the company for more details.

Competition & Alternatives

The NXD market has been heating up over the past two years. As a result, eXcelon is not alone in providing capabilities for storing XML data. While they were the first to ship with an in-production product in March 1999, many competitors of different sizes have emerged since then. One of the largest of these is Software AG with their Tamino product line. Tamino differs from XIS in many respects, foremost of which is its need for a DTD or XML Schema in the document insertion process. EXcelon also claims to be months ahead of Tamino in development and license sales. So far, it seems that Software AG is leveraging their existing ADABAS customers for new license sales of their Tamino product line, so it's possible the two companies aren't directly competing for new license sales at the moment.

XIS also competes with products from emerging NXD vendors including Ipedo, IXIASoft, Neocore, Coherity (PyBiz), XYZFind, X-Hive, B-Bop, and others. These products span very different architectural, footprint, performance, and pricing models. As such, general comparisons are not helpful. Users looking to adopt an NXD database should make specific feature or performance comparisons between XIS and those vendors as are appropriate. The only major difference is that eXcelon has a much longer track record than the majority of those companies, many of which have been formed since 2000 and have fewer than 100 license sales.

Key Conclusions & Recommendations

- Users considering a mechanism for storage of XML documents and data should utilize Native XML Databases (NXD) for persistence and querying of XML documents.
- eXcelon's XIS provides a robust, time-tested, and deeply functional NXD system that leverages its years as leader in the object-oriented database market.

- Users should be aware of the rapidly changing NXD market and be aware of new vendors, standards, and technologies that will assist in NXD vendor selection.

Profile: eXcelon	(December 2001)
Date Founded: 1989	
Funding: Publicly-held (NASDAQ: EXLN) [Merged with C-Bridge in September, 2001]	
CEO: Joe Bellini	
Employees: 500	
Products:	
<ul style="list-style-type: none">• eXtensible Information Server (XIS)• Stylus Studio	
Address:	
25 Mall Road 5th Floor Burlington, MA 01803	
URL: www.exceloncorp.com	
Main Phone: 781-674-5000	
Contacts:	
John Bachman jbachman@exceloncorp.com	
Coco Jaenicke coco@exceloncorp.com	
Jeremy Nazarian jnazarian@exceloncorp.com	

Related Research

- *XML Data Storage Technologies and Trends* Report (ZTR-ST101)
- *XML Data Storage Multi-Client Study* (ZTR-ST102)
- *Web Services Technologies and Trends* Report (ZT-WEBSRV)
- *B-Bop* ZapNote (ZTZN-0204)
- *Coherity* ZapNote (ZTZN-0144)
- *Ipedo* ZapNote (ZTZN-0151)
- *NeoCore* ZapNote (ZTZN-0146)
- *Software AG Tamino* ZapNote (ZTZN-0116)
- *X-Hive* ZapNote (ZTZN-0200)
- *XAware* ZapNote (ZTZN-0154)
- *Xyleme* ZapNote (ZTZN-0326)
- *XYZFind* ZapNote (ZTZN-0117)

-- continued --

About ZapThink, LLC

ZapThink is an IT market intelligence firm that provides trusted advice and critical insight into XML, Web Services, and Service Orientation. We provide our target audience of IT vendors, service providers and end-users a clear roadmap for standards-based, loosely coupled distributed computing – a vision of IT meeting the needs of the agile business.

ZapThink's role is to help companies understand these IT products and services in the context of SOAs and the vision of Service Orientation. ZapThink provides market intelligence to IT vendors who offer XML and Web Services-based products to help them understand their competitive landscape and how to communicate their value proposition to their customers within the context of Service Orientation, and lay out their product roadmaps for the coming wave of Service Orientation. ZapThink also provides implementation intelligence to IT users who are seeking guidance and clarity into how to assemble the available products and services into a coherent roadmap to Service Orientation. Finally, ZapThink provides demand intelligence to IT vendors and service providers who must understand the needs of IT users as they follow the roadmap to Service Orientation.

ZapThink's senior analysts are widely regarded as the "go to analysts" for XML, Web Services, and SOAs by vendors, end-users, and the press. They are in great demand as speakers, and have presented at conferences and industry events around the world. They are among the most quoted industry analysts in the IT industry.

ZapThink was founded in October 2000 and is headquartered in Waltham, Massachusetts. Its customers include Global 1000 firms, public sector organizations around the world, and many emerging businesses. ZapThink Analysts have years of experience in IT as well as research and analysis. Its analysts have previously been with such firms as IDC and ChannelWave, and have sat on the working group committees for standards bodies such as RosettaNet, UDDI, CPExchange, ebXML, EIDX, and CompTIA.

Call, email, or visit the ZapThink Web site to learn more about how ZapThink can help you to better understand how XML and Web Services impact your business or organization.

ZAPTHINK CONTACT:

ZapThink, LLC
11 Willow Street
Suite 200
Waltham, MA 02453
Phone: +1 (781) 207 0203
Fax: +1 (786) 524 3186
info@zapthink.com