zapthink
zapnote

Doc. ID: ZTZN-0151-1
Released: April 18, 2002

ZAPTHINK ZAPNOTE™

IPEDO
AN IN-MEMORY NATIVE XML DATA STORAGE SYSTEM

Briefing Date: September 18, 2001

Analyst: Ronald Schmelzer

Abstract

Ipedo has developed a proprietary, hierarchical, in-memory database engine aimed at storing XML documents "natively". The Ipedo XML Database uses memory techniques to get substantial performance gains versus some other NXD approaches. Built as an all-Java server, the Ipedo system is meant to be easily integrated with typical application server environments. Ipedo has extended XPath for search across multiple XML documents and has also included XSLT transformations within the data store itself.

Native XML Data Storage

Emerging from roots in SGML, HTML, text, object-oriented, and hierarchical storage roots comes a new category of XML storage solution: the "Native" XML Data Store (NXD). Rather than being forced to map XML documents to an alternate data representation format or insert an entire XML document into a single field, the NXD can accept XML documents for storage and retrieval without any modifications, mapping, or transformations. The term NXD does not imply storage architecture, but rather implies that XML data can be stored without any explicit transformation, mapping, or manipulation. Basically, to the user it should seem as if they are simply inserting an XML document. Any mapping, transformation, or other manipulations must happen completely behind the scenes, invisible to the user.

Since an NXD doesn't specify any particular data storage architecture, quantitative or qualitative measurements such as system performance, footprint, and other such metrics cannot be construed as valid benefits of NXD systems. NXD systems can use hierarchical, object-oriented, file system, relational, and binary storage formats that exhibit highly variable performance metrics. As a result, the primary value of NXDs is not so much their performance, but rather their inherent understanding of XML capabilities.

NXDs' most valuable attribute is their ability to store arbitrary and highly variant XML documents. XML-enabled RDBMS systems require explicit mappings to XML documents, and by their very nature are unable to deal with XML documents that have a highly variable structure and take advantage of XML's extensibility capabilities. Ipedo presents a Native XML Database simply called the Ipedo XML Database that serves to address the requirements of NXD systems using an in-memory model.

Ipedo XML Database

Ipedo has developed a proprietary, hierarchical, in-memory database engine aimed at storing XML documents "natively". The Ipedo XML Database uses memory techniques to get substantial performance gains versus some other NXD approaches. Built as an all-Java server, the Ipedo system is meant to be easily integrated with typical application server environments. All of the components that make up the run-time engine of the Ipedo XML Database are accessible through standard Java programming interfaces. The Ipedo data store can be deployed in either a client/server model or in an embedded model depending on the end implementation.

The system is built as an in-memory database technology that stores XML document indexes with content in memory. 128 Mbytes of RAM is required to use the system, but users don't have to use the main memory option, but users can flexibly choose how much main memory

to dedicate to the database. Of course, more memory implies better overall system performance. Ipedo is capable of indexing XML documents down to the individual element level across a collection of XML documents through several different kinds of indexes: t-Tree, b-tree, hash-map, and others. Ipedo's technique for indexing is called "hot indexing," which allows you to define indexes on elements and attributes within documents and to manage those indexes in-memory for quick navigation and access. Ipedo uses a GUI or APIs to do the import, organization, and indexing of documents in DB. Advanced memory management technology keeps most often used information in memory, and the XML caching system is user configurable allowing users to move any collection of XML nodes into or out of memory dynamically to tweak performance.

Ipedo has extended XPath for search across multiple XML documents and has also included XSLT transformations within the data store itself. The result is that the system can operate on results in XML directly and output results in whatever format desired. Ipedo has a built-in Web Services interface for Query and update over SOAP. The system can be accessed at the API level (look like a DOM), XQuery, or SOAP-based using XPath. Ipedo also leverages JAXP for Java-based document parsing. The system uses XML Schema to organize documents, can also input schema-less and DTD documents as well. The schema manager performs incremental validation when information is updated to maintain conformance. Ipedo XML Database allows updates, inserts and deletes at the attribute, individual element or node level. Modifications can also be performed at the collection or document level. XML elements to be updated are located using XPath expressions.

Ipedo's centralized XML management allows users to group XML documents together based on common attributes. The security manager centralizes access to system resources by providing username and password authentication. Ipedo also provides the capability for definition of roles, allowing for quick assignment of access control rules to given users. All user, authentication, and access control information is accessible through standard Java Naming and Directory Interface (JNDI) APIs. The Ipedo Integration Manager provides a mechanism for integrating Ipedo server with external data sources based on a proprietary HTTP-based XML messaging protocol. This is combined with versioning capabilities for auditing and document-level control and the locking mechanism makes concurrent updates to the database possible. The end result is an enterprise-class system capable of some transactional capabilities including backup/restore, journaling, and bulk loading. The current release of the Ipedo XML Database supports single-level transactions with plans to support nested transactions in a future release.

Customers & Release History

Ipedo launched its first version late in the spring of 2001 and subsequently went to production version in summer of 2001. They announced their second version at the XML Edge conference in October of 2001. In future releases, Ipedo XML Database will add support for both the XQuery standard and intelligent free-form text search. At this time, they have around 15 customers, $2M in annual revenue, 30 employees and are growing at a 300% year to year growth rate. Most of their sales are focused in the US but have some interest from UK and Asia companies. The company sells direct and is inking agreements for OEM and ISV distribution.

Competition & Alternatives

Competition is heating up for XML Data Storage technologies (see ZapThink XML Data Storage Technologies and Trends Report for more detail). The space roughly is split between XML-enabled RDBMS vendors such as **IBM, Microsoft,** and **Oracle**, and the NXD vendors including **Software AG (Tamino), Excelon, IXIASoft, XML Global,** and many others. Ipedo competes in this latter category but differentiates itself based on its performance capabilities, in-application XSLT transformation, and simplified data management. Ipedo claims that

adding memory (on the order of Gigabytes) would dramatically improve performance, which is not necessarily the case with systems such as Tamino or Excelon.

Key Conclusions & Recommendations

- Companies seeking a Native XML Data store with performance and transformation as key decision factors should consider Ipedo's XML Database solution.
- Ipedo will face increased competition from XML-enabled RDBMS and other Native XML Data store vendors as the market for XML storage and retrieval continues to heat up. Ipedo can continue to differentiate itself by focusing on performance, value-added data manipulation, and support for advanced XML query and storage specifications.

Profile: Ipedo (April 2002)
Date Founded: 2000
Funding: Privately-held, Venture-backed
Draper Fisher Jurvetson and Silicon Valley angel investors
CEO / President: Tim Mathews
Employees: 30
Products:
- Ipedo XML Database

Address:
1001 Marshall St
Redwood City, CA 94063
URL: www.ipedo.com
Main Phone: (650) 306-4000
Contacts:
David Finklestein david@cormancom.com
Tim Mathews tim@ipedo.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)
- *Excelon* ZapNote (ZTZN-0205)
- *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)

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