



SimbaProvider for OLAP Business White Paper

www.simba.com

Contents

Overview	1
The Data Analysis Environment ..	2
Online Analytical Processing	2
OLE DB for OLAP, XML for Analysis, and MDX	4
Building a Customized OLAP Data Provider	4
SimbaProvider for OLAP- Strategic Data Connectivity	5
SimbaProvider Advantages	8
Criteria for Integrating SimbaProvider for OLAP	9
Conclusion	9

This white paper introduces business reasons for and benefits you will receive, if you connect your OLAP or multi-dimensional data to the possibilities of OLE DB for OLAP (ODBO), XML for Analysis (XMLA), and Multi-Dimensional Expressions (MDX) using SimbaProvider for OLAP.

Overview

Independent software vendors concerned with building end-to-end data analysis solutions understand that connecting to third-party products is a necessity, if they are to provide a complete range of analysis capabilities for their customers. The OLE DB for OLAP (ODBO) and XML for Analysis (XMLA) open standards make it possible for OLAP servers and data stores from different vendors to communicate consistently with 32-bit/64-bit ODBO and XMLA compliant analysis tools.

The greater challenge for most vendors, however, is connecting their proprietary applications to the ODBO or XMLA interfaces. Adapting a proprietary system to communicate with the ODBO or XMLA interfaces is difficult and most vendors do not have the expertise to develop an OLAP data provider themselves. SimbaProvider for OLAP is a software development kit that enables developers to build a customized OLAP data provider faster and more easily than starting from scratch. SimbaProvider SDK enables the development of robust ODBO Providers and both .NET-based and Java-based XMLA Providers, ensuring the utmost in interoperability and cross-platform enabled solutions—all you need to do is choose your development language and proceed to develop your data provider, as explained in SimbaProvider's comprehensive documentation.

The MDX Language is core to ODBO and XMLA, and understanding and implementing a MDX Parser and MDX Execution Engine is a

Simba Technologies Inc.
938 West 8th Avenue
Vancouver, BC Canada
V5Z 1E5

Tel. +1 604 633 0008
Fax. +1 604 633 0004
solutions@simba.com



major effort even for the most experienced OLAP server teams. Additionally, the MDX Language continues to evolve. Keeping up with the changes adds an additional maintenance burden that would otherwise take your OLAP team away from the core work of enhancing the OLAP Server. In this regard, SimbaProvider SDK contains Simba's MDX 2005 compliant MDX parser and MDX engine, for which Simba assumes the task of ensuring is up to date and offers the highest level of interoperability and compliance with standards. SimbaProvider SDK removes the MDX development burden by providing a full-fledged MDX engine, enabling your team to remain focused on core work for your product.

The Data Analysis Environment

Through technology, the nature of information has changed the way business is conducted today. Modern networked businesses are dynamic; they are characterized more by the flow of information than by the amount they store. With employees and diverse teams spread across an organization and outside of it, the separation between what is inside and what is outside becomes blurred—the notion of the enterprise is changing.

All businesses face the dilemma of having to use technology to compete, while having to compete with technology at the same time. This puts pressure on those software vendors that supply businesses with products designed to collect, analyze and distribute data. These vendors understand that providing a range of advanced analysis capabilities for their customers is important. They also understand that connecting to third-party products is a necessity, if they are to make this possible.

The current market for data analysis products is characterized by many vendors that produce an array of task-specific "point" products. Specialized data analysis tools represent a range that spans many categories, such as statistical analysis, forecasting, data mining, and data visualization.

For business-specific applications that can't be addressed by stand-alone products out of the box, application development environments are used extensively.

User Types and Analysis Requirements

Individuals in organizations that use data analysis tools are often differentiated by their requirements—the type, size, style and urgency of their analyses vary dramatically. Users also analyze corporate data in a variety of ways:

- Some of them analyze data frequently and predictably, according to formal business rules. Their process is characterized by small queries focused on summarized data.
- Some of them restructure detailed historical data extensively. Subsequent analysis steps are based on the results of preceding analyses. Their process is characterized by very large queries on huge volumes of data.
- Some of them analyze subsets of data stored in the enterprise data warehouse, uncovering patterns and trends. Their process is characterized by large queries on disparate data stores.
- Some of them access the Web extensively to initiate wider analyses within an organization. Their process is characterized by small, random, wide-ranging queries targeted at unstructured data.

Analysis is the common activity that promotes effective decision making throughout a modern business organization. The data passed between analytical applications must be easy to import and export, it must conform to a consistent structure and format, and it must be secure. An open end-to-end solution that includes an array of data analysis capabilities creates a forum in which people can focus on discussing ideas and making well-informed decisions. The most important consideration for organizations selecting front-end analysis tools is that they be flexible enough to share data and the data analysis results.



The multi-dimensional data model makes it simple for users to formulate complex queries, arrange data on a report, switch from summary to detail data, and filter or slice data into subsets.

Online Analytical Processing

In the past, the bulk of investment in corporate computing was in business departments that generated or captured data, such as accounting, order processing, manufacturing and customer relations. Data warehousing developed out of a need to add value to this data by collecting, cleansing and sifting data from a variety of operational systems, and making the resulting information available to a wide audience of end users for analysis and reporting.

Today, organizations are investing in applications and technologies that deliver additional value from this collected data. Most information in a relational data warehouse is too arcane or unwieldy for the average business user. Advanced query tools hide the database complexity to some degree, but for a large class of applications where the end user is viewing multi-dimensional data, Online Analytical Processing (OLAP) technology provides a better solution.

The OLAP Data Model

In an OLAP data model, information is viewed conceptually as cubes, which consist of quantitative values (measures) and descriptive categories (dimensions).

For example, typical measures could include dollar sales, unit sales and inventory; typical dimensions could include time,

geography and product. Within each dimension, data can be organized into a hierarchy that represents levels of detail in the data. A time dimension, for example, could include levels for years, months and days. The multi-dimensional data model makes it simple for users to formulate complex queries, arrange data on a report, switch from summary to detail data, and filter or slice data into subsets.

All organizations have some form of multi-dimensional data, and the complexity of this data is not necessarily a function of company size. Even the smallest company would like to track sales by product, salesperson, geography, customer and time. OLAP applications also deliver information more quickly to end users by preparing some computed values in advance, rather than at execution time. The combination of easy navigation and fast performance lets end users view and analyze data more quickly and efficiently than is possible with relational database technology alone.

OLAP Analysis Capabilities

In the past, OLAP tools have been used for specialized financial applications, such as budgeting, forecasting and reporting. Today, OLAP technology is used to accelerate analysis and decision making for business critical applications throughout the enterprise, such as Enterprise Resource Planning (ERP), manufacturing, Customer Relationship Management (CRM), and supply chain management.

OLAP is not a new concept. In 1993, Dr. E.F. Codd, the database researcher and inventor of the relational database model, coined the term in a white paper, which outlined 12 rules that defined the characteristics of OLAP applications.



Advanced OLAP capabilities include data mining and data visualization. Advanced visualization is useful for comprehending complex data by displaying it from several dimensions at once using color, shapes, maps, charts, perspective and animation. Visual queries make it easier and faster to navigate and select complex multi-dimensional data in a purely visual way.

In addition to advanced analysis capabilities, pre-built analytical applications for budgeting, sales analysis and performance management are widely available. These applications are quicker to install and are usually more functional than applications developed in-house.

The cost and availability of OLAP technology is no longer a barrier to deployment, but selecting the appropriate OLAP technology requires a good understanding of particular business needs, performance requirements, data volumes, user skills and system architecture.

OLE DB for OLAP, XML for Analysis, and MDX

Historically, OLAP server technology has been tightly linked to proprietary client technology, meaning that customers have had little choice in their selection of mixed, best-of-breed products. This restriction has led to high implementation costs and often inadequate choices for applications that require both client/server and web-based OLAP analysis capabilities.

In 1998, Microsoft established the OLE DB for OLAP (ODBO) open standard, which makes it possible for OLAP data stores from different vendors to communicate consistently with a growing number of front-end analysis tools and OLAP server products. Customers have the opportunity to combine best-of-breed products in an integrated solution.

ODBO is a set of Component Object Model (COM) interfaces that extend OLE DB to access multi-dimensional data. OLAP analysis clients can access multi-dimensional data in a third-party OLAP server through the ODBO interface.

For expressing queries to this data, ODBO employs Multi-dimensional Expression (MDX) language, which is the multi-dimensional equivalent of SQL. MDX defines multi-dimensional data selections and calculations in the ODBO API.

The ODBO standard expands the availability of OLAP data analysis capabilities in the business world by lowering the costs of acquisition, implementation and maintenance. It also redefines scalability to serve both large systems and individual users.

In 2001, Microsoft and Hyperion established the XML for Analysis (XMLA) open standard, which builds upon the success of ODBO and MDX. XMLA is a standard that allows client applications to talk to multi-dimensional or OLAP data sources. The communication of messages back and forth is done using web standards – HTTP, SOAP and XML.

As the Web has grown in importance to business users, ODBO and XMLA have also helped facilitate the evolution of web-based interfaces for accessing OLAP applications. Today, applications that adhere to the ODBO and XMLA standards can connect to a variety of desktop and web-based analysis tools to access catalogs of reports and other types of data. All major vendors, including Microsoft, Oracle, SAP and IBM now support ODBO, XMLA and MDX.

Building a Customized OLAP Data Provider

Many vendors build analysis products that take advantage of OLAP technology—through an optimized or proprietary architecture—to provide specialized capabilities for a specific industry or business application. The challenge for these vendors is one of providing direct access to all available front-end analysis tools. ODBO and XMLA offer an ideal solution; they provide them with a standardized way to connect to third-party analysis applications without having to change their system architecture to do it.



Once it becomes clear that ODBO and XMLA are the way to provide third-party analysis capabilities, the challenge then becomes one of connecting the proprietary application to the ODBO and XMLA interfaces. However, adapting the syntax of a proprietary system to communicate through an ODBO or XMLA interface via MDX is not a trivial task—queries from third-party tools must be able to communicate with the proprietary system, and when possible, write back changes to the underlying data store.

The Role of an OLAP Data Provider

An OLAP data provider maps a data query to the proprietary or optimized architecture of an underlying data source through the standards-based ODBO or XMLA interface via MDX. Any front-end application that connects to the ODBO or XMLA interface will then be able to query the proprietary data store. Without the requisite expertise, building an OLAP data provider from scratch is very difficult and time consuming. Vendors who try will likely fail on their first attempt, and many of them will fail outright after several attempts.

SimbaProvider for OLAP—Strategic Data Connectivity

While creating a data provider from scratch is very difficult and time consuming, it is far from impossible. But most vendors don't have the time or the expertise to develop an OLAP data provider as quickly as they would like. Nor do they want this expertise—while connectivity is core to their product, connectivity skills are not core to their business.

SimbaProvider for OLAP is an SDK that enables proprietary OLAP-based products to access third-party analysis tools—without modifications to the proprietary architecture. It shields the complexities of ODBO, XMLA and MDX from developers and lets them build a customized OLAP data provider for any proprietary system.

The following is a list of SimbaProvider's key features.

MDX Engine

SimbaProvider includes our 32-bit and 64-bit MDX Engines—a first in the industry. Flexible, extensible and easy to understand, the MDX Engine and sample components make it easy to add advanced MDX support, such as calculated members to your provider. In fact, we estimate you'll save over 80% in development effort compared to writing the MDX capabilities yourself. MDX, or multi-dimensional expression language, is the multi-dimensional equivalent of SQL. It is used to define multi-dimensional data selections and calculations in the ODBO and XMLA APIs.

MDX query resolution and evaluation capabilities are essential so that your users can analyze trends and perform time-based analysis. The MDX Engine controls the following processes during a multi-dimensional query (see Figures 1 and 2):

- Creates a parsed representation of the MDX query
- Interprets the meaning of the MDX query (aka Resolution)
- Coordinates metadata data requests required to solve the MDX query (aka Evaluation)
- Formats the results of the MDX query and returns them to the data provider

Simba has invested many person years in understanding the MDX grammar. The effort required just to understand the MDX language and develop an EBNF version of the MDX grammar alone is probably greater than the effort required to integrate SimbaProvider for OLAP with your proprietary OLAP data source.

Complete Object-Oriented Development Environment

SimbaProvider for OLAP consists of a set of easy to understand, object-oriented interfaces designed and documented to follow a step-by-step process for developing a prototype



OLAP data provider. These interfaces shield the complexities of the ODBO and XMLA model from your developers, which means you can save your most senior resources for other projects.

With SimbaProvider for OLAP SDK, you can offer a XMLA (implemented in 32-bit or 64-bit .NET and Java) interface, as well as 32-bit and 64-bit ODBO interfaces (see Figures 1 and 2). In this way, the OLAP Provider that you build will support both the ODBO and the XMLA open standard APIs for the same cost and development effort.

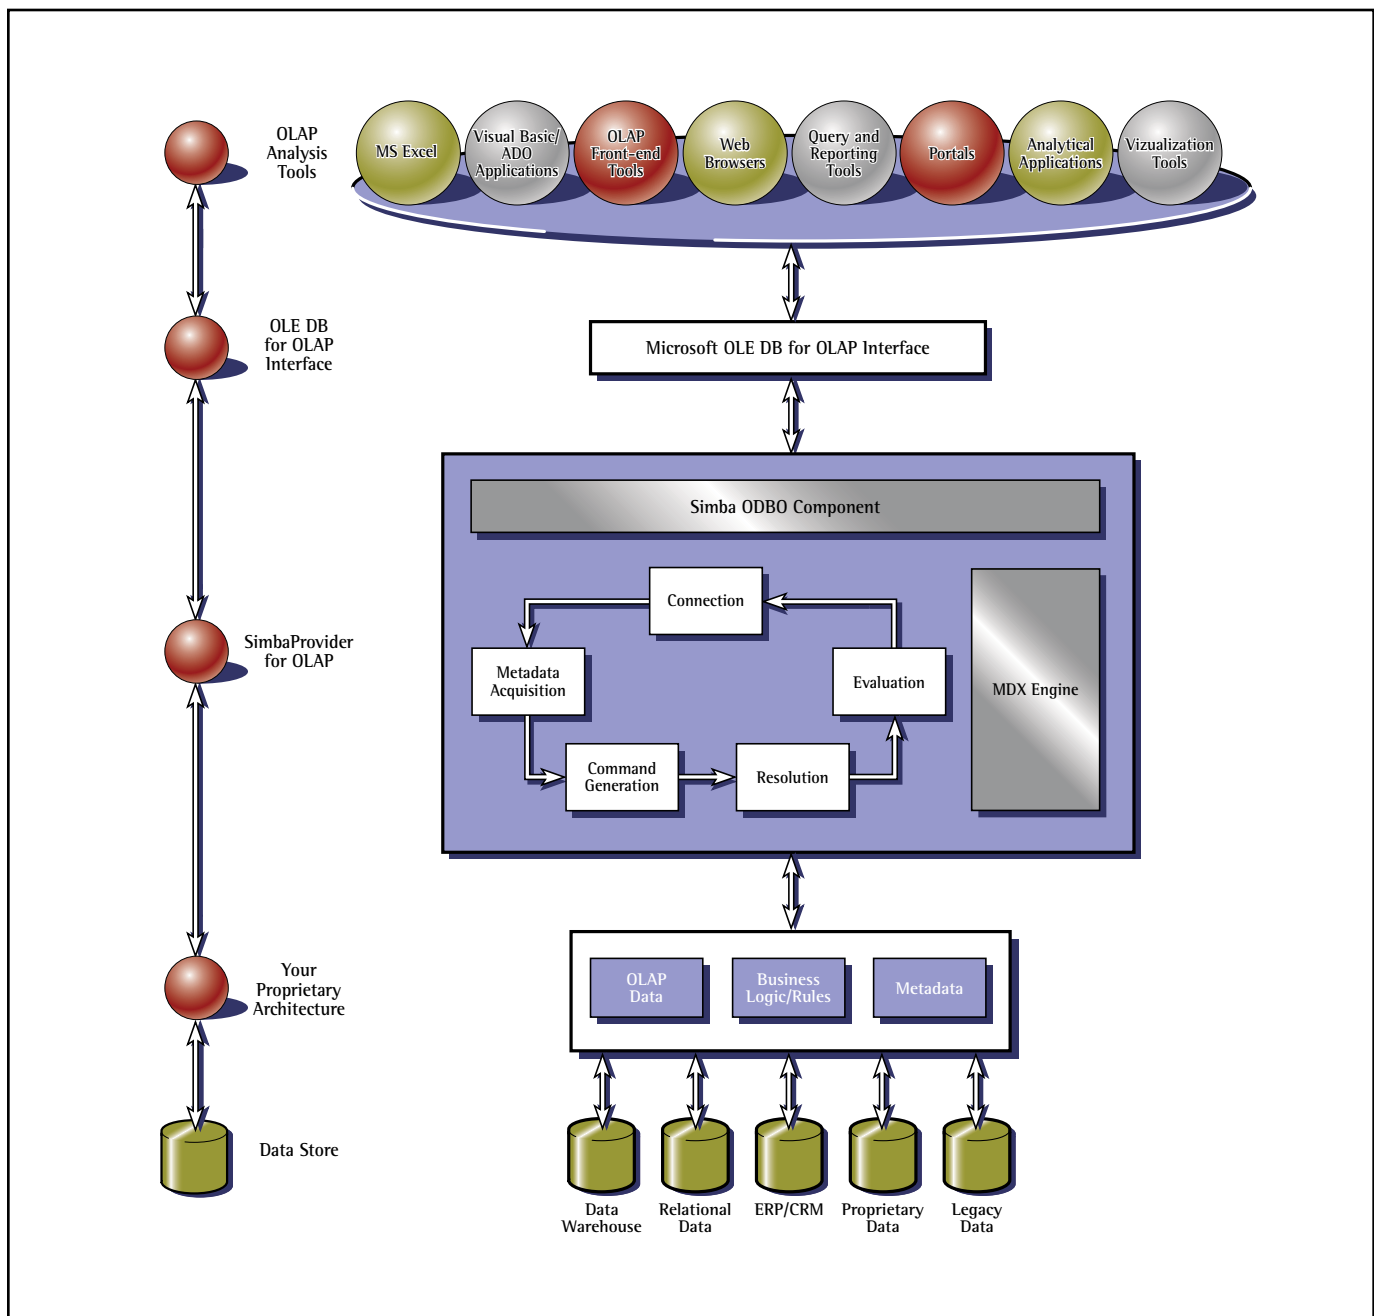


Figure 1: Protecting your proprietary architecture with SimbaProvider for OLAP between your data and the third-party analytic tools that follow the OLE DB for OLAP (ODBO) standard.

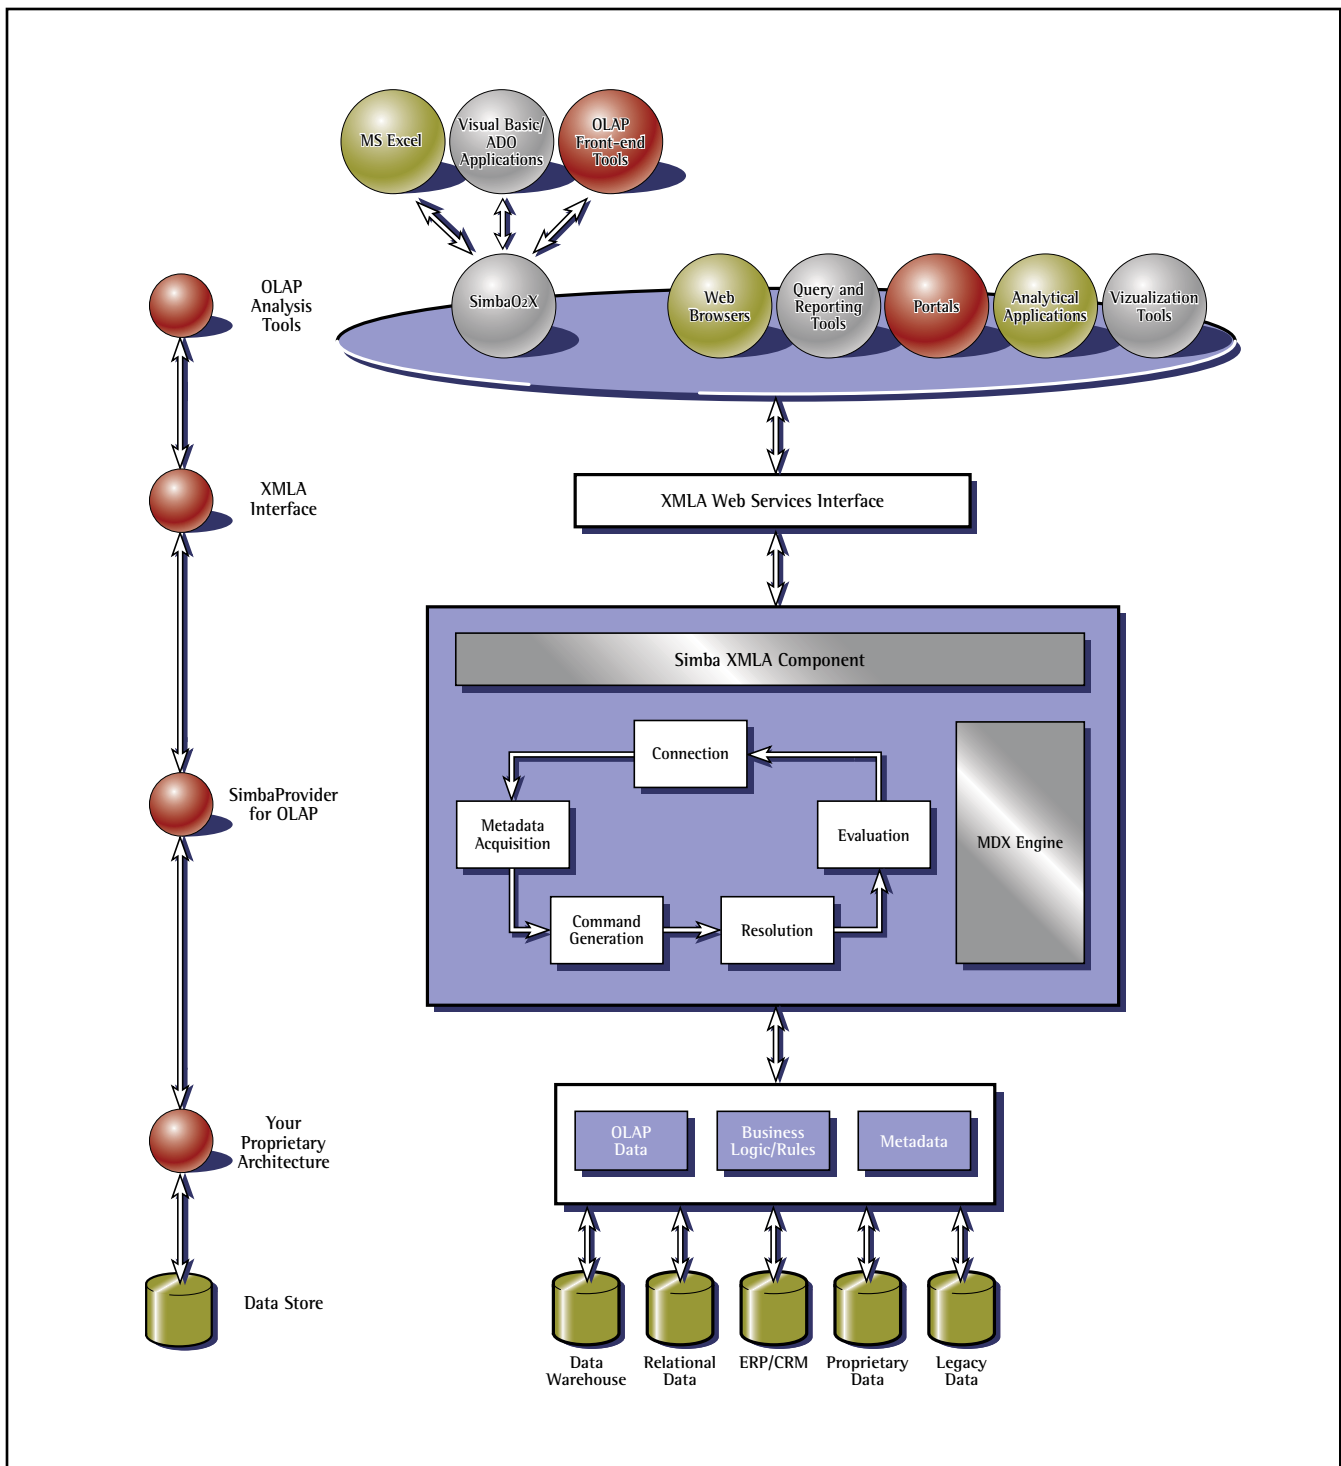


Figure 2: Protecting your proprietary architecture with SimbaProvider for OLAP between your data and the third-party analytic tools that follow the XML for Analysis (XMLA) standard.



Transaction Support

Transactions ensure data integrity by detecting exceptions in advance. A transaction is a sequence of operations within a session that can be performed as a single unit. SimbaProvider for OLAP supports all transaction capabilities on multi-dimensional data, such as two-phase commit and roll-back. Advanced exception detection is useful in large applications to uncover and display data anomalies at detailed hierarchy levels.

SimbaProvider Advantages

Although the benefits of a customized data provider depend on the way it is applied, many vendors share similar concerns with respect to building end-to-end data analysis solutions. With SimbaProvider, vendors can accomplish the following:

Add ODBO or XMLA Interfaces on Proprietary Architectures

Vendors that want to connect to third-party analysis tools do not have to alter the proprietary or optimized architectures on which they depend for competitive advantage. Using a standards-based interface also means that they do not have to develop native connectivity for every kind of third-party analysis tool that they want to support.

Reduce Development Time, Cost and Resource Impact

The SimbaProvider for OLAP SDK follows a packaged "framework" development process. It includes a sample

"SimbaProvider for OLAP is a software development kit that enables developers to build a customized OLAP data provider faster and more easily than starting from scratch."

OLAP data provider, comprehensive technical documentation, prewritten MDX parser, evaluator and resolver, and OLAP development assistance from Simba Technologies, if required. Your team could save more than 80% in development time and be comprised of less senior resources than if you developed your provider from scratch. Developers are shielded from having to learn the ODBO, XMLA and MDX specifications in depth. And, as the ODBO, XMLA and MDX specifications evolve, periodic updates to the SDK by Simba Technologies reduce development life cycle costs.

The Simba advantage is highlighted by our customers' experiences with the introduction of the XMLA standard. Every SimbaProvider for OLAP customer was able to add the XMLA interface to their OLAP Provider with virtually no extra development and minimal testing. This experience stood in stark contrast with those companies who had attempted a go it alone approach.

Interoperate with ODBO and XMLA-Compliant Analysis Products

SimbaProvider for OLAP enables proprietary server solutions to interoperate with all other analytical client applications and web-based analysis products that support the ODBO and XMLA standards. Periodic toolkit updates ensure that proprietary systems aligned with front-end analysis products evolve as the ODBO and XMLA specifications change. Vendors can also extend an end-to-end data analysis solution as required by adding clients and technologies that support ODBO and XMLA.

Distribute OLAP Capabilities via the Web

The Web provides an effective way to distribute data and analysis capabilities throughout the enterprise. And, many leading front-end analytical application vendors already offer web-enabled versions of their products. Use your OLAP data provider built with SimbaProvider for OLAP to enable analytical applications to connect to proprietary architectures via ODBO and XMLA. These applications can then be distributed throughout the enterprise via the Web,



which reduces software support requirements and makes it possible to establish enterprise-wide OLAP computing standards. With web-enabled OLAP capabilities, users can analyze data with a consolidated set of tools, consistency of analysis results is improved, and training costs decrease.

Criteria for Integrating SimbaProvider for OLAP

SimbaProvider for OLAP is designed specifically for independent software vendors and system integrators who already have some sort of multi-dimensional data or OLAP server to exploit. In some cases, the multi-dimensional data in their products may not yet be contained in a true OLAP data structure. For these types of vendors, opening up a system completely to another OLAP server architecture is not an alternative. Generally, the vendors who have the most to gain from using SimbaProvider for OLAP fit the following criteria:

- They are looking for ways to connect their multi-dimensional data to third-party analysis tools.
- Their products use multi-dimensional data or an OLAP calculation engine.
- They want to connect their products to ODBO, XMLA or MDX.
- They want to retain the proprietary or optimized architectures of their products.

Connecting to third-party data analysis products is critical for these vendors to remain competitive and to satisfy their customers' demands for advanced data analysis capabilities.

Conclusion

Analysis is the common activity that promotes effective decision making throughout a modern organization. Today, OLAP technology is used to accelerate analysis and decision making for business critical applications throughout the enterprise—the cost and availability of OLAP technology is no longer a barrier to deployment.

Vendors whose products incorporate some form of proprietary OLAP-based architecture understand that connecting to third-party products is necessary, if they are to provide a range of advanced analysis capabilities for their customers. ODBO and XMLA provide them with a standardized way of connecting to third-party analysis tools without having to change their system architecture to do it. However, adapting the syntax of a proprietary system to communicate through an ODBO or a XMLA interface via MDX is difficult.

SimbaProvider for OLAP enables proprietary OLAP-based products to access third-party analysis tools—without modifications to the proprietary architecture. It shields the complexities of ODBO, XMLA and MDX from developers and lets them build a customized OLAP data provider for any proprietary system.

SimbaProvider SDK contains:

- A MDX parser and execution engine compliant with the latest MDX 2005 standard (available on Windows and Linux)
- A 32- and 64-bit ODBO Provider on Windows
- A Java-based XMLA Provider on 32- and 64-bit Windows and Linux
- A .NET-based XMLA Provider on 3- and 64-bit Windows
- Proven compatibility with products like Microsoft Excel 2010 and 2007, SAP Business Objects Voyager, IBM Cognos and many others.

"SimbaProvider includes our MDX Engine—a first in the industry... you'll save over 80% in development effort compared to writing the MDX capabilities yourself."



About Simba Technologies

Simba Technologies Inc. is the recognized world leader in standards-based data access and analytics solutions. Simba works with the world's leading software companies to deliver first class data connectivity solutions.

Simba is a pioneer in ODBC, MDX, ODBO and XMLA. Since 1991, Simba has developed advanced data access solutions for thousands of end users. Today, more than half of all MDX providers have been built with Simba technology, and through a partnership with Microsoft, Simba's SQL technology has been installed on more than 30 million desktops worldwide. Simba's firm commitment to delivering the highest customer value through innovative solutions and expert support has gained the company a reputation as the industry leader for data connectivity solutions.

For more information, visit www.simba.com, call us at 604.633.0008, or email us at solutions@simba.com.

©2011 Simba Technologies Inc. All Rights Reserved. Simba and the Simba logo are trademarks of Simba Technologies Inc. All other trademarks or service marks are the property of their respective owners. Printed in Canada.