

The Role of the OLAP Server in a Data Warehousing Solution



Contents

Introduction	1
Core Components of a Data Warehouse Solution	1
Data Warehouse Access	3
OLAP Requirements	3
OLAP Applications	12
Best-practice Data Warehousing/ OLAP Architecture	13
Summary	14

Introduction

Data warehousing is not a product but a best-in-class approach for leveraging corporate information. Data warehousing accommodates the need to consolidate and store data in information systems that provide end users with timely access to critical information so they can make decisions that improve their organization's performance. Within a data warehousing strategy, online analytical processing (OLAP) supports strategic, high return on investment (ROI) applications. Specifically optimized for OLAP, Hyperion® Essbase® OLAP Server from Hyperion Solutions is an essential component of an enterprise data warehouse or data mart solution.

This paper describes how OLAP fits into data warehousing and data mart strategies and how Hyperion Essbase delivers a unique, compelling solution for the reporting, analysis, modeling and planning component of an enterprise decision-support system.

Core Components of a Data Warehouse Solution

Sears: Data Warehouse Institute OLAP Award Winner

Sears' data warehouse includes the Informix RDBMS and Hyperion Essbase on IBM servers. Its award-winning OLAP system, the Enterprise Planning Information Center (EPIC), uses Hyperion Essbase to deliver budgeting, planning and analysis, forecasting and reporting applications that leverage warehouse data. EPIC encompasses seven OLAP applications and has more than 1,000 users. Sears deployed its first OLAP applications in only two months and has identified tens of millions in cost savings from the EPIC system.




Core Components of a Data Warehouse Solution

The role of the OLAP server is best understood in the context of the six key components of a data warehousing or data mart strategy:

- **Policy** — Rules governing business issues such as goals and scope of the data warehouse, operational issues such as loading and maintenance frequency and organizational issues such as information security are as critical to successful data warehouse and data mart implementations as are product and technology choices.
- **Transformation** — Before raw data can be stored in a data warehouse, it must be transformed and cleansed so it has consistent meaning. This transformation may involve restructuring, redefining, filtering, combining, recalculating and summarizing data fields from disparate transaction systems.
- **Metadata** — Metadata is reusable information about data structures, objects, applications and business rules in the data warehouse that is defined once and is centrally stored. This allows all tools and applications accessing the data warehouse to use a common set of conventions which greatly reduces maintenance and assists both users and information technology professionals in finding the information they need and understanding what it means.

- **Storage** — As data is moved from various transaction systems into the warehouse, it must be stored in a way that maximizes system flexibility, manageability and overall accessibility. Because the information stored in the warehouse is read-only, historic in nature and includes detailed transaction data, the best data warehousing technology is the relational database.
- **Analysis** — This component provides a means to understand what is driving business performance. The analytic component must support sophisticated, autonomous end-user queries, rapid calculation of key business metrics, planning and forecasting functions and what-if analysis on large data volumes. Unlike the storage component of a data warehouse, the analytic solution must encompass historical, derived and projected information. While the data warehouse storage component is read-only, the analytic solution must support multiple users simultaneously updating and recalculating information to enable “what-if” and “what-next” modeling and planning applications. The OLAP server is the best technology for the analytic engine in the data warehouse.
- **Presentation and Access** — User tools, such as spreadsheets, query tools, Web browsers, statistical packages, GIS systems, visualization tools and report writers, provide a way to select, view, manipulate, visualize, analyze and navigate through data.



OLAP Server: Analysis component for the data warehouse.

Data Warehouse Access

Both data warehouses (e.g., comprehensive, enterprise-wide, etc.) and data marts (e.g., subject- or application-specific) must be accessible to a wide variety of users to satisfy their information needs. In general, the user tools for accessing data warehouses and data marts have historically been placed in one of two main categories:

- **Personal-productivity tools** (e.g., spreadsheets, word processors, statistical packages and graphics tools) allow users to manipulate and present data on individual PCs. Developed for stand-alone environments, personal productivity tools address applications requiring desktop manipulation of small volumes of warehouse data.
- **Data query and reporting tools** deliver warehouse-wide data access through simple interfaces that hide the SQL language from end users. These tools are designed for list-oriented queries, basic drill-down analysis and report generation. Though they provide a view of simple historical data, these tools do not address the need for multidimensional reporting, analysis, modeling and planning. They are also not capable of computing sophisticated metrics that users need to understand what is driving the business.

The Role of the OLAP Server in the Data Warehouse

OLAP servers deliver warehouse applications such as performance reporting, sales forecasting, product line and customer profitability, sales analysis, marketing analysis, what-if analysis and manufacturing mix analysis — applications that require historical, projected and derived data. With OLAP servers' robust calculation engines, historical data is made vastly more useful by transforming it into derived and projected data. Users gain broader insights by combining standard access tools with a powerful analytic engine.

OLAP Requirements

OLAP applications share a set of user and functional requirements that cannot be met by query or personal-productivity tools that work directly against the historical data maintained in the warehouse relational database. An OLAP server provides functionality and performance that leverages the data warehouse for reporting, analysis, modeling and planning requirements. These processes mandate that the organization looks not only at past performance, but more importantly, at the future performance of the business. It is essential to create operational scenarios that are shaped by the past yet also include planned and potential changes that will impact tomorrow's corporate performance. The combined analysis of historical data with future projections is critical to the success of today's corporation.

Who is building data warehouses?

An August 1996–1997 survey of Hyperion Solutions' customers indicates:

- 50 percent already have production data warehouses
- 80 percent are building data warehouses

A META Group data warehousing survey found that:

- 90 percent of Fortune 500 companies are currently building a data warehouse

Requirements for the OLAP component of a data warehouse or data mart strategy include:

- The ability to scale to large volumes of data and large numbers of concurrent users
- Consistent, fast query response times that allow for iterative speed-of-thought analysis
- Integrated metadata that seamlessly links the OLAP server and the data warehouse relational database
- The ability to automatically drill from summary and calculated data, which is managed by the OLAP server, to detail data stored in the data warehouse relational database
- A calculation engine that includes robust mathematical functions for computing derived data (aggregations, matrix calculations, cross-dimensional calculations, OLAP-aware formulas and procedural calculations)
- Seamless integration of historical, projected and derived data
- A multi-user read/write environment to support users what-if analysis, modeling and planning requirements
- The ability to be deployed quickly, adopted easily and maintained cost-effectively
- Robust data-access security and user management
- Availability of a wide variety of viewing and analysis tools to support different user communities

Neither personal-productivity tools nor query and reporting tools are designed to meet these requirements by themselves. ROLAP (relational OLAP) tools meet only one or two of these requirements. To deliver reporting, analysis, modeling and planning applications within a data warehousing strategy, an OLAP server must be implemented. The Hyperion® Essbase® OLAP Server is optimized for data warehousing applications and meets all of the above requirements.

Bell Canada: Web-enabled OLAP and Data Warehousing Success

Bell Canada's data warehouse includes the IBM DB2 RDBMS running on MVS integrated with Hyperion Essbase on several Windows NT servers. Bell Canada replaced 16,000 mainframe reports and retired a VAX-based reporting system with its Hyperion Essbase OLAP applications. It has improved performance and reduced maintenance — slashing the IT reporting backlog. Bell Canada's users access Hyperion Essbase using Excel spreadsheets, Cognos Powerplay and Netscape Navigator via Hyperion Web Gateway.

OLTP vs. Data Warehouse RDBMS vs. OLAP Server

System	OLTP	Data Warehouse RDBMS	OLAP Server
Purpose			
System charter	Operational	Historical and detail data	Analytic
Access			
Access type	Read/write	Read-only	Read/write
Access mode	Atomic, singular, simple update	Singular, list-oriented queries and reports	Iterative, comparative analytic investigation
Access process	IT-supported queries	IT-assisted or preplanned queries and reports	IT-independent, ad hoc navigation and investigation drill-down
Response characteristics	Fast update, varied query response	Varied, potentially very slow query response	Fast, consistent query response
Data Storage			
Content scope	Application-specific Actual/vertical Limited historical	Warehouse: cross-subject data Data mart: single subject area Historical data	Many cubes. Each cube is a single subject area: historical, calculated, projected, what-if, derived data
Data detail level	Transaction detail	Cleansed and lightly summarized	Summarized, aggregated and calculated using sophisticated analytics
Data structure	Normalized	Normalized or denormalized	Dimensional, hierarchical
Data structure design goal	Update	List-oriented query	Analysis
Data volumes	Gigabytes	Gigabytes/terabytes	Gigabytes
Implementation			
Deployability	Slow (multi-month/year)	Slow (multi-month/year)	Fast (days/weeks)
Adaptability	Limited—requires significant resource	Low—requires significant resource	High—easily modified
Computer hardware investment required	Moderate to expensive	Moderate to extremely	Minimal to moderate expensive

OLTP servers handle mission-critical production data accessed through simple queries, while OLAP servers handle management-critical data accessed through an iterative analytic investigation. This chart shows the evolution of data as it moves over three server platforms, how the servers differ in purpose, the means of access, the data they store and how they are implemented.

The following sections summarize the capabilities of Hyperion Essbase in each of these critical areas:

1. The ability to scale to large volumes of data and large numbers of concurrent users

Hyperion Essbase has the capacity to handle large amounts of data. It supports an unlimited number of dimensions (e.g., product, time, market and channel) and an unlimited number of members (e.g., year, quarter, month and week) within a dimension. Hyperion Essbase's scalability enables users to perform detailed analysis without worrying about the memory and storage capacity of their individual PCs.

Hyperion Essbase supports parallel loading and calculation to take advantage of scalable computing hardware. Today, major corporations with thousands of users have deployed Hyperion Essbase in production with over 150 gigabytes of information in a single OLAP cube and more than 500 gigabytes in a single application made up of multiple OLAP cubes. Each Hyperion Essbase server supports multiple applications, so the effective capacity of a single server is even larger.



The OLAP server manages multiple subject-area cubes of information within the data warehouse or data mart, each targeted for a particular user community or analytic need.

Hyperion Essbase features a thin-client architecture that enables corporations to deploy OLAP applications to large numbers of users over wide area networks. Hyperion Essbase is fully multi-threaded and leverages Symmetrical Multi-Processing (SMP) hardware to take advantage of available processing power on modern server computers. Multi-threading and SMP support enable Hyperion Essbase to scale and support thousands of concurrent users. In audited OLAP benchmark results, Hyperion Essbase proved its ability to support a simulated user community of nearly 7,000 concurrent users on a single four-processor server. Hyperion Essbase's efficient storage and indexing deliver outstanding concurrency, as only minimal resources are required to process user queries.

2. Consistent, fast query response times that allow for iterative speed-of-thought analysis

Hyperion Essbase consistently delivers fast query response times that make an iterative environment for analytic queries possible. OLAP users' queries are neither predictable nor fixed, and the results of one query often frame the requirements of the next. In this environment, responses must be forthcoming in seconds— not minutes or hours — or analysts will cut short the investigative process to meet management deadlines. To be effective, an analysis session must be interactive and keep pace with the analyst's speed of thought.

With Hyperion Essbase, most users receive answers to their queries in a fraction of a second. Even the most complex queries take only a few seconds. In audited OLAP benchmark results, Hyperion Essbase processed more than 6,800 complex queries per minute on a four-processor server — an average response time of just 0.00876 seconds per query.

While query tools are becoming increasingly sophisticated, their performance is still limited by the response time of the data source. Hyperion Essbase consistently provides fast response by allowing designers to optimize performance based upon an application's unique requirements for query performance, calculation complexity, calculation window (the amount of time available to load and calculate the application), user concurrency and disk utilization. Hyperion Essbase achieves this flexibility through three calculation options: precalculate, calc on the fly and calc on the fly and store. Together, these three calculation strategies let designers maximize flexibility, capacity and performance.

Precaculate calculates the data required to answer user queries in advance. This strategy minimizes query response time and maximizes user concurrency.

Calc on the fly calculates the data that is required to respond to user queries at query time, rather than in advance as with precalculation. This strategy minimizes batch calculation time and disk utilization and eliminates the phenomenon of database explosion associated with traditional OLAP systems and with aggregate table explosion in ROLAP implementations.

Calc on the fly and store calculates data when the first user requests it and then stores the calculated data on disk. This strategy provides virtually immediate availability of data and delivers optimal performance for very large user communities that are querying similar information.

Designers can combine all three calculation strategies in any combination to balance query performance, batch calculation and resource utilization to optimize overall system performance. The low-level access methods used by Hyperion Essbase are also highly efficient. Instead of scanning through thousands of rows of tabular data, Hyperion Essbase directly retrieves the requested data using bit mapped indexing and offset addressing. Hyperion Essbase minimizes input/output operations and consistently delivers fast response time regardless of data size.

3. Integrated metadata that seamlessly links the OLAP server and the data warehouse relational database

Hyperion® Integration Server uses metadata to integrate Hyperion Essbase with the data warehouse relational database. Hyperion Integration Server dramatically reduces the time and expense to create, deploy and manage OLAP applications from data warehouses and data marts using Hyperion Essbase OLAP Server. Hyperion Integration Server allows OLAP applications to share common metadata by creating a centralized catalog of enterprise OLAP metadata. This catalog stores reusable, documented definitions of data sources, mappings of relational data OLAP data, OLAP structures such as dimensions, hierarchies, calculations and predefined OLAP application templates. The centralized metadata catalog ensures that all OLAP applications across the enterprise use a common and consistent set of definitions for dimensions, hierarchies and calculations, and that all of these definitions are always synchronized with the data warehouse relational database.

The data warehouse contains massive volumes of information that change every day. It is imperative that OLAP applications have the ability to adapt automatically to this changing information. For example, whenever a data warehouse table containing product information changes, all of the OLAP applications that access product information must be synchronized with the changed data. Centralized OLAP definitions greatly simplify OLAP application administration, promote reusability and make it easy to deliver data-driven applications that are tightly focused on particular analytic requirements. Once the Hyperion Integration Server metadata is defined, new analytic applications can be automatically generated on-demand by combining and customizing objects stored in the shared metadata catalog. Hyperion Integration Server automatically generates optimized SQL, including multi-pass SQL, to transfer both metadata and data to Hyperion Essbase for analysis.

4. The ability to automatically drill from summary and calculated data, which is managed by the OLAP server, to detail data stored in the data warehouse relational database

With Hyperion Integration Server, end users can extend their analyses by drilling from summarized, calculated and derived data, which is managed by the OLAP server, into detail data stored in the data warehouse relational database. Driven by the OLAP metadata catalog, Hyperion Integration Server Drill-through automatically generates optimized SQL statements to retrieve desired information. Users can drill from the “bottom” of the data managed by Hyperion Essbase into massive volumes of detail data stored in the data warehouse relational database. This dramatically increases the volume of information users can access within a single OLAP application.

Hyperion Integration Server Drill-through reports are context-sensitive so users who want to access detail information in the data warehouse relational database can be presented with the specific detailed reports that are relevant just to the data they are analyzing in Hyperion Essbase. For example, a district manager who is analyzing the profitability of the stores in his region may want to know when the stores first opened and who manages each store. The store managers, who are analyzing profitability for their particular store, may want to retrieve recent invoices for particularly profitable or unprofitable products. Hyperion Integration Server allows users to run default reports or use a graphical wizard to customize their requests. Hyperion Integration Server automatically generates all SQL statements required to retrieve the desired information.

5. A calculation engine that includes robust mathematical functions for computing derived data (aggregations, matrix calculations, cross-dimensional calculations, OLAP-aware formulas and procedural calculations)

The need for robust OLAP calculation capabilities is often overlooked by data management professionals, yet sophisticated OLAP calculations deliver the highest return on investment in data warehousing by turning massive volumes of raw data into actionable business information. OLAP applications require the ability to calculate summarized, derived and projected information from the atomic or source data in the data warehouse. Neither relational databases nor the SQL language has significant calculation capabilities; therefore, an external OLAP calculation engine is required. Major relational database vendors including Oracle, IBM and Microsoft have acknowledged this by endorsing the need for external OLAP calculation engines.

Without a calculation engine, deploying and maintaining an OLAP application requires extensive, ongoing manual intervention. This means creating hundreds or thousands of summary tables, as well as developing calculation-specific C or COBOL programs, massive index management, query optimizer manipulation, performance tuning and more. The overhead involved in creating sophisticated OLAP calculations directly in the warehouse RDBMS without an OLAP server is very severe. Companies that have chosen ROLAP engines as the sole analytic component within their data warehousing strategies have experienced this pain directly. Many of these companies have switched from ROLAP engines to OLAP servers after failing to realize the benefits they had hoped for with their ROLAP implementations. In fact, most organizations that try to deploy OLAP solely using relational databases alone find themselves unable to deliver the highest value and highest ROI because, by their nature, these applications require the most sophisticated OLAP calculations.

There are five main categories of OLAP calculations: aggregations, matrix calculations, cross-dimensional calculations, OLAP-aware formulas and procedural calculations.

OLAP Calculation Capabilities

OLAP Server	RDBMS/SQL
✓Aggregations	✓Aggregations
✓Matrix Calculations	✓Simple Matrix Calculations
✓Cross-dimensional Calculations	
✓OLAP-aware Formulas	
✓Procedural Calculations	

Aggregations are the simplest type of OLAP calculations and the only type possible to perform natively in a relational database using the SQL language. Aggregations are simple additions of detail information into higher-level summaries. For example, daily sales figures are added together (aggregated) into weekly summaries and weekly summaries are aggregated into monthly summaries and so on.

Matrix Calculations are similar to the familiar calculations used in spreadsheets such as Excel or 1-2-3. Matrix calculations are used to define calculated relationships both vertically and horizontally across a grid of numbers. Ratios are the most common matrix calculations. For example, creating the percentage of different products' sales in comparison to each other requires a simple matrix calculation. A period-to-period comparison (e.g., the percentage of sales of product A versus product B for this month versus last month) is slightly more complex matrix calculation. Relational databases and the SQL language are incapable of performing most matrix calculations because they do not support inter-row (only intra-row) calculations. Only trivial matrix calculations can be accomplished using standard SQL, and even these require highly complex, slow running correlated sub-queries and self-joins.

Cross-dimensional Calculations are similar to calculations between multiple linked spreadsheets or calculations across an N-dimensional spreadsheet. Cross-dimensional calculations define calculated relationships across different dimensions and different levels of aggregations within a single dimension. Cross-dimensional calculations require a built-in understanding of hierarchical relationships such as parents, children, ancestors and descendants. For example, market share for a particular product is calculated by computing the sales of that product in comparison to the total sales of all products. Relational databases and the SQL language are incapable of performing cross-dimensional calculations for the same reasons they cannot perform matrix calculations.

OLAP-aware Formulas are similar to spreadsheet formulas — such as statistics, forecasting functions, financial functions and time calculations — that have been extended to support OLAP. OLAP-aware formulas make it easy to deliver the rich calculated information that business analysts need. Just like leading spreadsheets, Hyperion Essbase provides a library of several hundred built-in formulas that have been extended to be OLAP-aware. With Hyperion Essbase, a single calculation can easily span 100 gigabytes of information. For example, a single OLAP-aware formula can calculate standard deviation, variance, interest payments or Net Present Value across all combinations of all dimensions. Relational databases and the SQL language do not provide formulas for developing and managing business calculations.

Procedural Calculations are used to transform source and historical data based on complex business rules. Hyperion Essbase includes procedural calculation languages used to develop very sophisticated OLAP calculations based on the needs of business analysts. There are a vast number of OLAP applications that cannot be implemented without procedural calculations. Some common examples include:

- All types of profitability analysis (profitability must be calculated before it can be analyzed)
- Financial consolidation
- Activity-based costing
- Budgeting and forecasting
- Merchandise planning

Since relational databases do not have OLAP-aware procedural languages or calculation engines, they are incapable of creating the calculated information that enables the most sophisticated and highest ROI OLAP applications.

6. Seamless integration of historical, projected and derived data

Hyperion Essbase rapidly incorporates and calculates historical, projected and derived data. The data warehouse contains just read-only historical data, or in the words of Bill Inmon “subject-oriented, consistent, time-variant, nonvolatile information.” Historical data warehouse data is most often sourced from legacy systems, production OLTP applications or external data providers. Projected information usually comes from spreadsheets, external data feeds, or direct user inputs. Derived data is computed by the OLAP server’s calculation engine.

By integrating all three types of data, Hyperion Essbase users are able to perform analysis that looks forward as opposed to only report on what has happened in the past. Support for historical data alone limits a system to what-happened questions. Support for projected and derived data enables a system to additionally support what-if and what-next analysis. The ability of an OLAP server to deliver historical, projected and derived data side by side enables forward-looking applications, which are otherwise impossible to deploy, that provide significant value to end users.

7. A multi-user read/write environment to support what-if analysis, modeling and planning

While the data warehouse relational database is a read-only historical storage environment, important OLAP applications such as planning, modeling and forecasting require that the end user has the ability to change business assumptions and immediately view the impact of those changes. For example, in a what-if planning application a user might need to analyze the effects on both sales and profitability of raising advertising expenses by 10 percent for a particular product line. These what-if and what-next applications combine traditional analysis and operational decision making.

Systems being used for analysis, modeling and planning applications must allow users to update information. Without being able to change assumptions there is no way to see the effects of those changes on the business model. Another way to say this is that OLAP servers must provide multi-user read/write access to support planning, modeling and forecasting applications. Hyperion Essbase enables users to update projections and forecasts and run iterative what-if scenarios in a secure, shared environment.

Hyperion Essbase provides a server-based read/write environment and robust security system that ensures data consistency and integrity. Most organizations agree that read/write OLAP applications should not be performed directly against the historical archive in the data warehouse relational database. With a data warehousing strategy, Hyperion Essbase serves as a read/write analytic data mart to enable reporting, analysis, modeling and planning applications. Historical information is provided by the data warehouse; plans and forecasts are captured, modeled and updated iteratively in Hyperion Essbase. At the end of the planning cycle, the completed plan can be exported out of Hyperion Essbase back into the warehouse relational database for historical storage. In this fashion, the warehouse RDBMS can operate in a read-only mode. Users can reap the benefits of read/write planning, modeling and forecasting applications, and IT can maintain control of the consistency of the information in the data warehouse.

8. The ability to be deployed quickly, adopted easily and maintained cost-effectively

Hyperion Essbase is easily deployed and highly adaptable. Prototype OLAP applications can be deployed in days, and the average time to deploy production OLAP applications is measured in weeks. Hyperion Essbase provides graphical tools that enable the developer to visually design and manage OLAP applications and automatically build and maintain all summaries and indices.

Once deployed, end users can operate independently, thus dramatically reducing the support burden on IT. IT no longer needs to write reports for users, and there is no ongoing requirement to build new summary tables or indices to support unplanned user queries.

Benefits of the OLAP Server in the Data Warehouse

User Benefits	IT Benefits
Insulates users from SQL language	Eases system management
Insulates users from relational model	Automates maintenance of indices and summaries
Improves query performance and user scalability	Reduces load on data warehouse RDBMS
Vastly enhances business calculation capability	Relieves IT from generating reports for users
Enables read/write operational OLAP applications	Enables central control over analytic data
Supports wide range of client tools	Deploys quickly at low risk and expense

User independence is provided because Hyperion Essbase does not require that users learn new tools or query languages. Users are able to access Hyperion Essbase without using a query language from a wide array of desktop tools including spreadsheets, query tools, Web browsers, report writers, statistical packages and more.

9. Robust data-access security and user management

As a data warehouse, by definition, stores sensitive business information in a centralized location, robust data security and user management capabilities are critical. Hyperion Essbase provides highly granular security that controls whether a user has no access, read access or read/write access to information all the way down to the individual data cell level. Hyperion Essbase also features robust user-management tools that provide a graphical interface to manage user passwords, access rights, user groups and security filters. In fact, Hyperion Essbase provides significantly more security than is possible natively in the data warehouse relational database. Relational databases have only table or column-level security, which makes it impossible to control access on a row-by-row basis within a single table. While column-level security is sufficient for OLTP applications, it is inadequate for OLAP applications that require much more granular control. Robust cell-level data security must be a requirement for all organizations implementing data warehousing strategies.

10. Availability of a wide variety of viewing and analysis tools to support different user communities

A key premise of data warehousing is to make a large amount of information available to a large community of end users. The more users who are able to access data warehousing applications, the more value data warehouses will provide to organizations. This implies that to maximize success and ROI, the OLAP server in the data warehouse must also be accessible from a wide variety of end-user tools.

Hyperion Essbase is unique in its integration with more than 30 different Essbase-Ready™ front-end tools. These products, from a wide variety of leading vendors, are natively integrated with the Hyperion Essbase API (application programming interface) and seamlessly access the data managed by Hyperion Essbase. These tools may be grouped into several distinct categories, each of which serves different end-user communities. Products ranging from spreadsheets, query tools, application development environments, Web browsers, enterprise reporting systems, statistics packages, data mining tools, executive information systems, visualization tools and packaged applications are all available for use with Hyperion Essbase.

OLAP servers with limited third-party front-end tool choices (or even worse, offer only a single proprietary front-end) will, by definition, only be appropriate for a small user community. Thus, these servers provide less value than an open OLAP server such as Hyperion Essbase.


OLAP Applications

Spreadsheets	
• Lotus 1-2-3	• Microsoft Excel
Query Tools	
• Brio Technology Brio Enterprise	• Internetivity db Probe 3.0
• Business Objects	• IQ Software IQ/Vision
• Cognos PowerPlay	• Lighten Advance
• Comshare Commander Decision and DecisionWeb	• Lingo Computer Design FISCAL
• CorVu Integrated Business Intelligence Suite	• Management Science Associates Business Web
• Hummingbird BI/Analyze	• Seagate Crystal Info
• Hyperion® Analyzer	• Sterling Vision Dimensions
• InfoSpace SpaceOLAP	
Web Browsers	
• Microsoft Explorer	• Netscape Navigator
Application Development	
• Alphablox Enlighten	• John Galt Development ForecastX
• C, C++, Java	• Microsoft Visual Basic
• Hyperion® Objects	• Painted Word Mocha Blend
• Hyperion® Web Gateway	• Sybase Powerbuilder
• Inprise Delphi	
Statistics and Data Mining	
• Data Mind MarketOne	• SPSS Netview and SPSS Base 8.0 for Windows
• IBM Intelligent Miner	
EIS	
• arcplan Information Services inSight	• TEMTEC Executive Viewer
• Pilot Lightship (Mobile Software)	• Track Objects Inc. Track Objects 32
Visualization	
• AVS Express	• Visible Decisions in 3D
• ESRI MapObjects	
Enterprise Reporting	
• Hyperion Reporting®	• Seagate Crystal Reports and Seagate Crystal Info
• Hyperion® Enterprise Reporting	
Server Management and Data Integration	
• Acta ActaLink	• Noetix Views
• Blue Isle Software InTouch	• Painted Word Inc. Telltale and Clockwork
• Constellar Warehouse Builder	• Prism Solutions Warehouse Executive
• Decisionism Aclue Decision Supportware	• Sagent Datamart Solution
• Hyperion Integration Server	• Sybase PowerDesigner
• Informatica PowerMart	• Vision Enterprises EssInfo
• Leonard's Logic Genio	
Packaged Applications	
• Activity-based Costing	• Links to Packaged Applications
• Budgeting	• Planning
• Consolidation	• Risk Management
• Demand Planning	• Many more!
• Financial Planning	

A wide range of Hyperion® Essbase-Ready solutions are available to leverage Hyperion Essbase.

Best-practice Data Warehousing/ OLAP Architecture

Hyperion Essbase integrates easily into both data warehousing and data mart strategies and provides a robust solution for reporting, analysis, modeling and planning applications. Hyperion Solutions recommends a multi-tier or “hub-and-spoke” data warehousing architecture to help organizations get the most out of their data warehousing investment. This architecture, endorsed by industry analysts and data warehousing gurus such as Gartner Group, Meta Group, the Data Warehouse Institute, Bill Inmon, Ralph Kimball and Doug Hackney, results in the highest ROI and lowest risk for companies implementing data warehouses.



The multitier or “hub-and-spoke” data warehousing features a general purpose relational data staging area as the “hub,” coupled with OLAP-based application-specific data marts as “spokes” to deliver information efficiently.

Relational databases are the preferred platform for the data warehouse staging area because they provide excellent data capacity, manageability and extensibility. However, relational staging areas do not provide the fast query performance, ease of deployment, business calculation capabilities or robust analytics required by users to gain business benefits from data warehouse information. OLAP Servers provide exceptional query performance, easy deployment, sophisticated business calculation capabilities and robust analytics but do not match relational databases in terms of data capacity, manageability and extensibility. The best practice strategy, therefore, is to couple an application-neutral relational data warehouse staging area (the hub) with application-specific OLAP data marts (the spokes). This architecture has been successfully implemented by a large number of Hyperion Solutions customers.

The application-neutral data staging area should be built using a relational database and include transaction- and historical-level detail information. The data warehouse staging area should contain minimal summarization and must be designed to be extensible so it can be built incrementally. The

Hyperion Headquarters

Hyperion Solutions Corporation
1344 Crossman Avenue
Sunnyvale, CA 94089

tel: 408 744 9500
fax: 408 744 0400

info@hyperion.com
www.hyperion.com

data warehouse staging area should be architected to include information for cross-functional, cross-departmental purposes that address a broad set of an organization's needs. The data warehouse staging area should be built incrementally over time. The best practice is to start by populating the staging area with the data required by a single application and gradually expand to the data required by additional applications and subject areas. An architected approach allows initial applications to be deployed quickly, which is critical for data warehousing success.

The "spokes" of the data warehouse architecture should be built using OLAP data marts that deliver specific business applications. Application-specific OLAP data marts deliver fast query response, ease of deployment, powerful business calculation capabilities and robust analytics that users require. The OLAP data marts should be designed to solve particular business problems, as opposed to collections of departmental data. For example, the goal should be to deliver sales analysis and customer profitability applications, not a marketing data mart. The application-specific OLAP data marts source consistent, cleansed historical data from the data warehouse staging area and then summarize, calculate and derive additional information for users to access and analyze.

Combining application-neutral data warehouses with application-specific OLAP data marts meets the combined needs of both IT professionals and end users throughout the enterprise. This best practice data warehousing architecture reduces the initial risk of a data warehousing strategy and ensures low ongoing maintenance costs. For these reasons, leading industry analysts, data warehousing gurus and Hyperion Solutions have endorsed the multitier, hub-and-spoke architecture as the best-practice architecture for data warehousing.

Summary

Data warehousing is a best-in-class approach to leveraging corporate information. Data warehousing addresses a broad range of decision-support requirements that are served by a variety of personal-productivity, query and reporting tools, as well as OLAP servers. OLAP servers fulfill a very wide set of user-driven functional requirements for reporting, analysis, modeling and planning applications.

An organization's inherent requirement for reporting, analysis, modeling and planning applications is to use information from the data warehouse to help drive improved business performance. To successfully meet this requirement an organization must first understand past performance and second, have the ability to prepare for and manage the future.

Hyperion Essbase meets all key requirements of OLAP while integrating smoothly into a data warehousing strategy. With Hyperion Essbase, companies can extend their decision-support systems by moving beyond a historical focus to proactively chart the future direction of the business.