

INTEGRATING ESSBASE AND OBIEE+: THE FUTURE OF ORACLE EPM & BI

Edward Roske, interRel Consulting

OVERVIEW

If your company has a data warehouse or is thinking of building one, this whitepaper will help you learn how Essbase and OBIEE, Fusion Edition can further drive value from your warehouse solution. We'll discuss the benefits of Essbase and OBIEE+ and how they fit into your overall information delivery platform. Essbase and the Oracle technologies combine to provide seamless integration between your applications, relational databases, and multi-dimensional databases. The OBIEE+ Reporting and Analysis components provide a single interface into any type of data source, combining data from your data warehouse and data marts into in nicely formatted reports, user friendly dashboards, and easy-to-use analytic screens.

WHAT ARE THE QUESTIONS

Each different level of an organization is concerned with questions pertaining to their area of specialty. Since any well run organization strives for synergy between departments if one department makes a decision to undertake a project it is imperative for the other departments to be able to accommodate and support their goals. Perhaps the marketing organization suggests an initiative to increase sales via a mid-year sales promotion. This suggestion has implications for the entire company. The manufacturing organization needs to access whether or not they can produce the additional product. The distribution team needs to understand if they have the capacity to handle the additional product shipments. Finance needs input into whether or not they can fund this initiative. And ultimately, the entire organization needs to understand if the tactics they have employed have helped them achieve their stated objectives.

Now in order for organizations to have the ability answer these questions (correctly) and come up with relevant solutions they need to be able to aggregate and analyze relevant data. The problem is that data is *everywhere* and a number of problems arise through the striving for common data access. One of the issues is that there may be no single strategy for information delivery. Ad hoc & data mining against transaction systems, resulting in performance degradation to transaction systems, developer time required to create and maintain results and limited capabilities in ad hoc queries & data mining. There are silos of information with a multitude of tools and technologies for accessing that information, limited integration with the existing databases and data marts as well as limited visibility to metrics for operations management and budget evaluation and finally limited interoperability with proprietary databases.

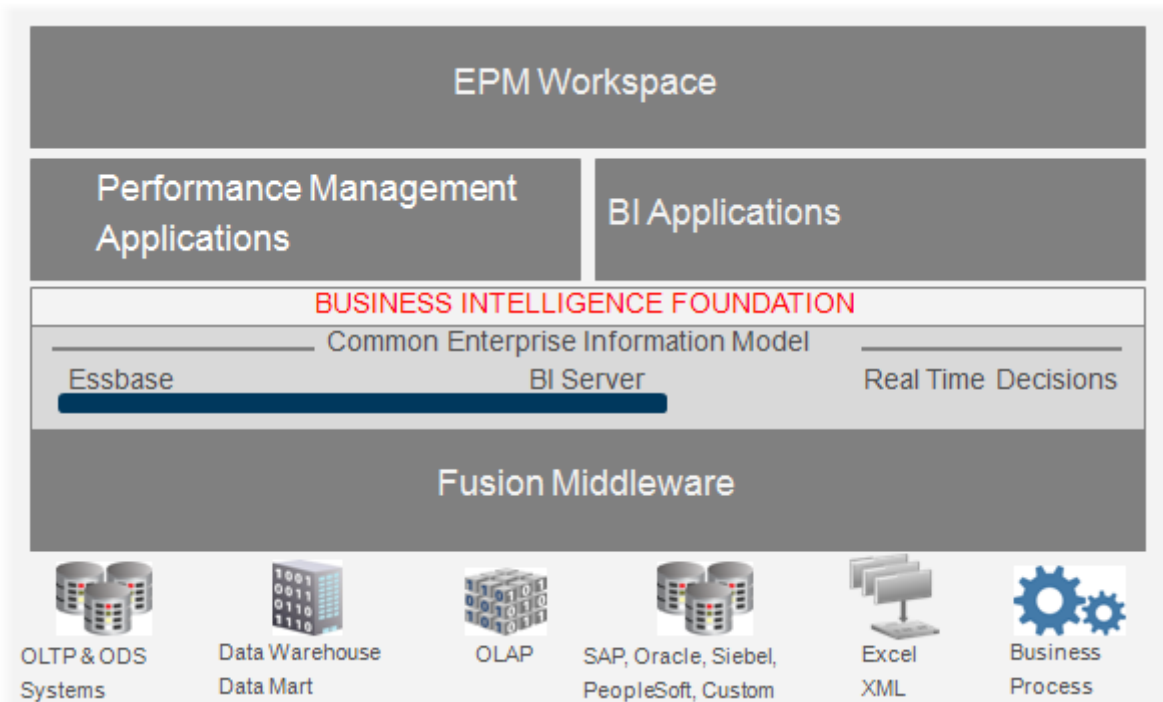
Different user communities have different interfaces and delivery needs: IT staff vs. data experts vs. analysts vs. executives and management. All the different user interfaces and delivery needs creates a knowledge gap which only gets worse over time. As time goes by the amount of information that Information Systems are able to deliver does not tend to increase at the same high right as user demand for information, thereby creating a growing knowledge gap.

DATA WAREHOUSE SOLUTION?

That's why I have a data warehouse, right? In part yes, but data warehouse solutions only get you part of the way. In real life what often occurs are multiple data stores and marts, multiple end user tools; essentially no consistent place to get common enterprise information. Common data warehouse issues include time consuming changes and updates. New sources require modeling, ETL, integration to overall data warehouse design, mart design, reporting and analysis, and possible updates to hardware / underlying architecture. The update frequency is not fast enough for the end users. Performance for reporting can be slow requiring aggregate structures for reporting and analysis. And don't forget the politics and constant debate of what should be stored, when, and how.

THE SOLUTION

Essbase and OBIEE+ are the enabling technologies to help close the knowledge gap that occurs between source systems and business critical “information”. Essbase and OBIEE+ address and resolve a number of the common data warehouse issues, delivering event driven business intelligence, pervasive use, and real time, predictive data. Both Essbase and OBIEE+ are part of Oracle’s Business Intelligence layer within the Oracle EPM System.




OBIEE+ is the product suite containing Oracle BI Server and supporting end user tools like Answers and Dashboards. Oracle BI Server (BI Server) is the main engine of the OBIEE+ suite.

OBIEE +	Product	Overview
	BI Server	Create a virtual data warehouse of many sources Data resides in the transaction systems (could be a burden on transaction systems)
	Answers	Analysis over the web for OBI sources
	Dashboards	Easy-to-create and use dashboards for OBI Sources
	Delivers	Delivery to emails, dashboards, web services Alerting capabilities
	Mobile Analytics	Content access via mobile devices
	BI Publisher	Formatted reports, pixel perfect reports Not a tool for analysis
	Financial Reporting	Printed, formatted reports for Oracle EPM System (Essbase, FM, Planning)
	Interactive Reporting	Analysis, reporting and dashboards for relational and Essbase sources
	Web Analysis	Analysis and dashboards for Oracle EPM System (Essbase, FM, Planning, HPCM)
SQR Production Reporting	Pixel perfect, large volume reporting	
Essbase	<i>The multidimensional database engine</i>	

Now to fully understand how these tools can extend your data warehouse, let's deep dive into the fundamentals of Essbase and the Oracle BI Server.

INTRODUCTION TO ESSBASE

Essbase is the biggest and most powerful OLAP engine ever to compete in the BI/EPM space, providing fast query capabilities, powerful calculation engine, full featured write back capabilities, and more. Originally Essbase was created by Arbor in 1991, acquired by Hyperion in 1998 and finally Oracle in 2007.



To put it simply, all databases prior to Essbase were built for the purpose of storing transactions. The goal for these systems was to get individual records into the database as quickly as possible and to get those same records back out again as quickly as possible.

Essbase, an online analytical processing solution (OLAP), was one of the first databases designed purely to support analysis. Essbase is a multi-dimensional database. Dimensions are grouping of data elements into one or more hierarchies for reporting and analysis. Dimensions are created based on how the user needs to answer questions related to her business; a business view of the information. Dimensions can contain thousands or even millions of data elements (a.k.a members). Example dimensions include Time, Accounts, Product, Customer, Employee, and Asset.

Multi-dimensional simply means that any of the dimensions set up in a database could be put in the rows or columns or pages of a report whereas relational databases are two dimensional. While any relational database can be set up to give the appearance of having multiple dimensions, it takes a lot of up front work by developers. Essbase and other OLAP databases have dimensionality built-in.

- Measures
 - Profit (+)
 - Margin (+)
 - Sales (+)
 - COGS (-)
 - Total Expenses (-)
 - Marketing (+)
 - Payroll (+)
 - Misc (+)
 - Inventory (~)
 - Opening Inventory (+)
 - Additions (~)
 - Ending Inventory (~)
 - Ratios (~)
 - Margin % (+)
 - Profit % (~)
 - Profit per Ounce (~)

- Year (Alias: Period)
 - Qtr1 (+) (Alias: Q1)
 - Jan (+)
 - Feb (+)
 - Mar (+)
 - Qtr2 (+) (Alias: Q2)
 - Apr (+)
 - May (+)
 - Jun (+)
 - Qtr3 (+) (Alias: Q3)
 - Jul (+)
 - Aug (+)
 - Sep (+)
 - Qtr4 (+) (Alias: Q4)
 - Oct (+)
 - Nov (+)
 - Dec (+)

Essbase has repeatedly had no comment to the allegations that it is like a spreadsheet on “steroids”.

Essbase databases are also optimized for retrieval at any level of the hierarchy, even the very topmost number that might represent every dollar the company has ever made in its history. OLTP databases (relational databases) were nicely optimized for retrieval of detailed records but definitely not hierarchical information. By pre-saving summarized information, Essbase allows analysis to happen from the top down with no decrease in performance.

Because hierarchy is inherent to OLAP databases, drill-down (sometimes known as “slicing and dicing” but never known as “making julienne data”) is inherent as well. Essbase is great at doing adhoc analysis because it knows that when a user double-clicks on Qtr1, she wants to see Jan, Feb, and Mar. This is because the roll-up of months to quarters is pre-defined back on the Essbase server. Drill down capabilities are available for all of the dimensions in the Essbase database, allowing you to zoom up and down level by level or skipping levels even. Unfortunately, this article can’t do this feature justice; it is one of those things you have to see for yourself.

The Essbase database structure is defined through an object called the outline. The Essbase outline contains *all* of the hierarchical information. Unlike a relational database where the hierarchy is applied to data already stored in tables, in Essbase, the outline (and as such, the hierarchy) directly controls *how* data is stored and indexed. When you modify your outline, you restructure your database. Building dimensions, ordering them, adding hierarchies, adding members with member attributes, and creating member formulas in the outline all drive the performance of the Essbase database.

Essbase provides two types of databases: Block Storage Databases and Aggregate Storage Databases. Block storage option (BSO) databases were *the* original Essbase databases or cubes. Data is stored in a way that business users intuitively understand the hierarchies and could perform analysis really, really fast – just as we discussed earlier. Users of BSO databases can write data back to any point in a hierarchy with the right access, allowing for information collection, budgeting, forecasting and “what-if”ing for any subject area. BSO databases also come with a powerful calculation engine for centrally housed consolidations of the database and complex business rules (calculations are defined once for all users across all interfaces and reports). With over 250 built in functions, financial intelligence, time intelligence and more, enterprise performance management is easier than ever with Essbase as you can perform profitability and cost allocations, what if scenarios, driver based plans, and more .

Aggregate storage option (ASO) databases were introduced in Essbase 7.1 as a new alternative to BSO databases. They were created specifically to deal with requirements for very large sparse data sets with a high number of dimensions and potentially millions of members. ASO utilizes a new kind of storage mechanism that allows improved calculation times from ten to one hundred times faster than BSO databases – the calculations just aren’t as complex. ASO can also store up to 2⁵² dimensional combinations. If you aren’t that great at math, just know this is a really, really big number. New types of Essbase databases are now possible like customer analysis on potentially millions of customers, logistics analysis where we can analyze near real-time updates of product shipments, and market basket analysis where we can analyze what products are purchased along with other products for our online businesses.

INTRODUCTION TO THE ORACLE BI SERVER

The Oracle Business Intelligence Server is a scalable, efficient query and analysis server that federates data from any source including relational, multi-dimensional, pre-packaged applications, and unstructured files. The BI Server is the central component in the OBIEE+ product suite which also includes presentation tools like Answers, Dashboards, Publisher and Delivers.

The primary purpose of the BI Server is to federate data across the enterprise through virtual schemas. Instead of physically moving the data into a data warehouse, the BI Server moves the data “virtually”, creating a single logical view of information. In a sense you are building a data warehouse without building a data warehouse. Because BI Server is pulling data dynamically, real time information access is possible for end users. Think of the value placed into the users’ hands; real time data access side-by-side with historical information for reporting and analysis. The reporting and analysis available are relational in nature, more of a rows and columns type of reporting environment.

Imagine you need to combine sales data from the GL and customer data from the CRM system and you don’t have the resources or budget for a full scale data warehouse implementation. BI Server can easily connect to and unify this information into one virtual schema for reporting and analysis. (Do we recommend sourcing directly against transaction

systems? In most cases no, but we want to illustrate the possibilities). Behind the scenes, efficient logical SQL is executed against the sources and presented in a user friendly format.

To this end, BI Server provides an abstraction layer, taking complicated technical structures and terms and putting them into a logical, business view of the information with business terminology. E.g. how many financial analysts know that the column CUSTDSC is Customer Name? Well, maybe most analysts because this was an easy one but most of our source systems don't provide easily understood naming conventions. We also want something a bit more formal to display on our reports. So within the BI Server physical layer, the CUSTDSC is queried but in the presentation layer, the designer has abstracted the field and renamed it Customer Name. This is just a very simple example. Now think about more complex requirements. BI Server can combine tables across sources, adding the customer contact column from CRM to the main customer table in the GL into one logical table. Or BI server can combine data in tables across sources, for example, joining sales data from multiple systems in a distributed environment into one main fact table. Let's review an example that illustrates another benefit to this feature: in a quest for world juggling domination, Juggling Wolverine Company is quickly acquiring juggling independents across the globe. With their acquisitions, they need to quickly integrate GL information from all of the new systems for reporting. If they owned the BI Server, they could simply pull in the physical table from the new source system and map it to existing logical GL tables, with the new content available in reports. If needed, the BI Server can perform transformations across from the new sources, mapping data elements to an "apples-to-apples" view in the logical table.

BI Server delivers all these power punches via three layers in the administrative console: Physical, Business Model and Mapping, and Presentation layer. The Physical layer is where you define the data sources, either creating manually or importing the metadata into the BI Server. Join relationships are defined as necessary for all of the sources in the Physical layer.

After you have created your sources, you then define the logical or business view of the data in the Business Model and Mapping layer. This is the first step in taking more complex IT structures and multiple sources and grouping them into models or content focused areas using business terminology. At this stage, you are defining your fact tables (tables that contain the measures or "numbers") and dimension tables (remember dimension from earlier: Time, Product, Customer). You create logical tables, columns, dimension hierarchies and joins to create the business view of the data.

BI server can group logical columns from a single dimension table into hierarchy levels, allowing users to "drill in" to dimensions to get more detailed data. The catch is that users have no real visibility into the drill paths; the dimension hierarchies are not visible from the presentation layer (whereas in Essbase, they are and users can use a number of ways to pick and choose how they want to navigate the hierarchy).

You can also create calculated columns and measures in the Business Model and Mapping layer, defining the business logic once for use across any query or report. Although be careful; these can be expensive from an end user performance standpoint. Calculated measures could include variances for actual and budget, ratios and more.

The Presentation layer is the final layer that further adds abstraction over the Business Model and Mapping layer. This is what the end user sees. Here you organize the columns into logical catalogs and folders, customizing naming for specific user groups as needed, and generally make the information as easily and logically accessible as possible. Catalogs show different views of a business model to different users via presentation tables which look like folders to the end user perspective. The folders contain the grouping of columns that make sense to the end user. Users can easily navigate through the catalogs and folders, picking and choosing the desired data elements and pull them into a detailed report, pivot report, chart or graph, apply filters or add in report side calculations.

For the most part BI Server does not store data but efficiently queries the original data sources. In some cases for large and/or complex queries or queries requiring aggregation, performance could be slow. A number of different query performance enhancing alternatives are available (don't worry, they are totally legal). One such option is the use of aggregate tables in the BI Server to store pre-computed results of measures for upper levels of dimension hierarchies. Because the values are already calculated and stored (also results in fewer rows), reporting and analysis is faster. The BI Server intelligently understands when to use aggregate tables or the data source. For those folks out there with Oracle databases, materialized views are a great way to speed up reporting similar to aggregate tables except that you are creating objects outside of the BI Server. These refreshable objects are created by the DBA, contain the aggregated information as needed for analysis, and most importantly use the super power of your RDBMS to summarize data. BI server also offers caching to save query results in memory "intelligently".

In summary BI Server provides us a common enterprise model for information delivery across an enterprise from almost any source with efficient querying capabilities. BI Server provides the benefits federating data sources and detailed reporting. Combine this with the powerful analytic capabilities of Essbase and you are meeting the reporting and analysis needs of users.

THE FRONT END

The EPM Workspace present contents content from both the BI Server and Essbase using tools like Answers, Dashboards, Web Analysis, Financial Reporting, and Interactive Reporting. You can also publish other types of external documents like pdfs, spreadsheets and PowerPoint presentations. Most of the end users don't really have to know or care which tool created a document; they simply run the reports or use the dashboards. Today, the Workspace is the best way to incorporate content across the Oracle BI platform and present it together in a common web end user interface.

WHY INTEGRATE OBIEE+ AND ESSBASE WITH YOUR DATA WAREHOUSE?

Common Data Access Issues	The Solution
No single strategy for information delivery	Essbase and OBIEE+ along with your data warehouse delivers a complete information delivery system
IT vs. end user queries	Common workspace with user friendly tools puts reporting and analysis in the hands of end users
Silos of information	Oracle BI server delivers a common enterprise model to eliminate silos of information across databases and systems
Limited visibility	Essbase and BI Server provide better visibility to key metrics and factors that drive business performance
Limited correlations	Essbase and BI Server provide a way to correlate relevant information across database and systems
Different end user needs	A common workspace provides information access for both advanced and basic end users BI Server and Essbase allow customization to different end user groups sourced from the common enterprise model

Common DW Issues	The Solution
Time consuming data warehouse changes	Incorporate changes and new sources to the data warehouse in BI Server (virtually) vs. in your data warehouse (physically) Immediate reporting access for new sources as they are included in logical schemas
Politics of what data should be stored	Gain more flexibility in this debate as no data is stored in BI Server and/or efficiently stored in Essbase
Data latency	Provide real time / near real time access with BI Server and Essbase incremental loads
Costly aggregate structures	Essbase's main purpose in life is to store aggregated data for fast reporting and analysis; Remove aggregations from the DW and use Essbase
Ever changing Business Definitions	Both BI Server and Essbase provide a model for easy updates to business definitions
Common dimensions and common definitions	Both BI Server and Essbase provide a common enterprise model for information; at the same, they allow customized presentations and views of end user specific terms and dimensions

Businesses today are taking advantage of proven integration between ERPs, data warehouses and Essbase / OBIEE+ (Oracle's BI Foundation Layer). Essbase and OBIEE+ fill in the knowledge gap between your source systems and end users, addressing some of the constraints and limitations in your data warehouse and other systems. Essbase and OBIEE+ provide reporting, dashboards, and advanced Analytics for *all* users regardless of experience level through a central interface. Essbase and OBIEE+ provide more dynamic ways to present information with reduced maintenance and deployment costs. Essbase and OBIEE+ is the solution to integrate with your data warehouse to deliver an "information platform" across the enterprise.

