

ORACLE BUSINESS INTELLIGENCE FOUNDATION SUITE

Technical Overview

January 2011



TABLE OF CONTENTS

INTRODUCTION.....	5
Enabling Enterprise Business Intelligence	5
Product Overview	6
Server Components	6
End-User Delivery Components.....	7
Systems Management Components	7
Oracle BI Product Strategy	8
FOUNDATION SUITE SERVER TECHNOLOGY	9
Oracle BI Server	9
Common Enterprise Information Model.....	9
The Oracle BI Server Provides the Following Key Capabilities	10
Query Parsing and Compilation	10
Code Generation.....	11
Parallel Execution Engine	11
Information Reliability	12
Oracle BI Database Gateways	12
Mission Critical Performance, Scalability, and Reliability	12
Accessing Oracle BI Server Information	14
Multiple Layers of Security.....	15
Physical Data Storage Independence.....	15
Complex Business Measures	16
Integrated Segmentation Engine	16
Oracle Business Intelligence Administration.....	17
Multi-User Development Environment.....	18
Oracle Essbase.....	19
Component Overview and Deployment Architecture.....	19
Optimized Multi-dimensional Storage.....	20
Performance, Scalability and Availability	21
Flexible Business Model Development	22
Multi-Dimensional Calculation Engine	23
Essbase Application Administration.....	24
Oracle BI Server and Essbase Integration.....	26
Oracle Essbase Analytics Link for Oracle HFM.....	26

Oracle EAL Architecture Overview	27
Benefits of EAL.....	28
FOUNDATION SUITE END-USER CAPABILITES	28
Oracle BI Presentation Services.....	28
Features of Oracle BI Presentation Services	29
Homepage	30
Integrated BI Search	30
BI Analysis and Reporting	30
Using BI Analysis and Reporting	31
Features of BI Analysis and Reporting.....	31
Blended Relational and Multidimensional Analysis.....	32
BI Visualizations	33
Interactive Dashboards.....	34
Features of BI Interactive Dashboard	34
Using BI Interactive Dashboards	35
Oracle Scorecard and Strategy Management.....	36
Features of Oracle Scorecard and Strategy Management	37
Suite-wide Integration	38
Oracle BI Publisher.....	38
Highly Formatted and Interactive Output.....	39
Data from Anywhere.....	39
End User Reporting with Choice	40
Deliver to Anywhere	40
Communicate Globally	41
Unsurpassed Performance.....	41
Open Standards. Easy Integration.....	41
Pre-Delivered Content	41
ACTIONABLE INTELLIGENCE.....	41
Oracle BI Delivers.....	41
Using Oracle BI Delivers.....	42
Features of Oracle BI Delivers	42
Guided Analytics with BI Interactive Dashboards	43
BI Action Framework.....	43
BI ON-THE-GO.....	44

BI Briefing Books.....	45
Oracle BI Mobile.....	45
BI Office Integration	46
BI Office Plug-In.....	46
Oracle Smart View for Office.....	47
CUSTOM DEVELOPMENT AND INTEGRATION	48
Application Development Framework Views.....	48
Collaboration, WebCenter and Portals	49
Web Services	50
Sample Application (SampleApp).....	50
SYSTEMS MANAGEMENT	51
BI Systems Management	51
SUMMARY	53
Integrated Foundation.....	53
Oracle BI Foundation Suite Key Differentiators	54
Conclusion.....	56

INTRODUCTION

ENABLING ENTERPRISE BUSINESS INTELLIGENCE

Many organizations today use a collection of Business Intelligence (BI) tools and applications to allow experts to gather information from a variety of sources, analyze it, and share it with managers and staff. However, ever-increasing business dynamics and increased competition means businesses now require a much higher level of value from their BI investments. BI must now help drive profitable growth, change, and many other operational and financial performance goals. Not only does BI need to deliver significant Return on Investment (ROI), but it also needs to be deployed in a manner that minimizes Total Cost of Ownership (TCO).

Enterprise Business Intelligence must give managers and staff much more than tools that access information. It must provide a broad set of capabilities, from self-service monitoring of performance and processes to driving action based on insights. Enterprise Business Intelligence requires not just a comprehensive BI tool set, but pervasive BI that provides insight to all employees within the context of their workflows. It needs to unify the many fragmented systems into a coherent enterprise view, while aligning forward-looking information to real time and historical data. It must be integrated ensuring accuracy and integrity of information across all delivery channels and resulting in lower cost of ownership. It must be open, meaning it will plug into the company's existing middleware architectures and data infrastructure. It needs to be fully secure to protect all enterprises information assets. It needs to support BI applications that scale from single-node departmental to multi-node enterprise-scale solutions regardless of user population or whether on-premise or on the cloud is the desired deployment model.

To achieve this vision the Oracle Business Intelligence (BI) Foundation Suite delivers the most complete, open, and integrated business intelligence tools and technologies on the market today. The Oracle BI Foundation Suite provides comprehensive and complete capabilities for business intelligence, including enterprise reporting, dashboards, ad hoc analysis, multi-dimensional OLAP, scorecards, and predictive analytics on an integrated platform. The Oracle BI Foundation Suite enables access to information through multiple channels such as web-based user interfaces, industry standard portals, mobile devices, and the Microsoft Office Suite of applications. A powerful enterprise information model unifies disparate data systems within an organization and provides a platform for BI tool integration. The Oracle BI Foundation Suite is completely open: (1) supporting both Oracle and non-Oracle data sources ranging from file-based data, to all popular relational database management systems, and to leading multi-dimensional sources; (2) supporting prevalent middleware solutions including application servers and security systems; and (3) providing open-APIs for integration with a range of enterprise systems. A strong and flexible security model ensures that information is accessed and delivered by those with the appropriate privileges. The Oracle BI Foundation Suite simplifies systems deployment and management through integrated systems management tools that offer single-click scale out capabilities that can support a range of deployments with proven capabilities for applications that reach tens of thousands of users accessing multi-terabytes of data. Finally, the Oracle BI Foundation suite offers best-in-class capabilities for managing the development lifecycle for BI applications with proven support for hundreds of geographically disperse developers.

In summary, traditional BI tools are not designed to enable the insight-driven enterprise. A fundamentally different infrastructure and business intelligence solution set is required to meet this need. The Oracle Business Intelligence Foundation with its complete, open and integrated modern

architecture and broad range of analytical capabilities is the only business intelligence solution designed to meet the needs of today's insight-driven organizations.

PRODUCT OVERVIEW

The Oracle BI Foundation Suite provides powerful capabilities that offer significant value for BI applications across the enterprise. The Oracle BI Foundation Suite consists of Oracle Business Intelligence Enterprise Edition 11g, Oracle BI Publisher, Oracle Essbase, Oracle Scorecard and Strategy Management, and Oracle Essbase Analytics Link (EAL). Following is an overview of the key components and features of the Foundation Suite.

Server Components

- *Common Enterprise Information Model:* The semantic model of OBIEE. It is accessed via an open API, making it available to any Oracle or non-Oracle delivery channel, thus providing a common version of the truth for all Business Intelligence users and applications.
- *Oracle BI Server:* A highly scalable, highly efficient query and analysis server that integrates data via sophisticated query federation capabilities from multiple relational, unstructured, OLAP, and pre-packaged application sources, whether Oracle or non-Oracle.
- *Oracle Essbase:* The industry-leading multi-dimensional online analytical processing (OLAP) server, providing a rich environment for effectively developing custom analytic and enterprise performance management applications.
- *Oracle Essbase Analytics Link:* Enables the delivery of effective management and financial analytic reporting to a broad user community by facilitating the real-time or on-demand transfer of financial information from Oracle Hyperion Financial Management to Oracle Essbase.



Figure 1 - Oracle BI Foundation Suite Overview Architecture

End-User Delivery Components

- *Enterprise Reporting:* Oracle Business Intelligence (BI) Publisher (formerly XML Publisher) is an enterprise reporting solution for authoring, managing, and delivering highly formatted documents, such as operational reports PDF forms, shipping labels, checks, sales and marketing letters, and much more. Built on open standards, reports can be designed using a feature-rich online layout editor or through familiar desktop products and viewed online or scheduled for delivery to a wide range of destinations. While Oracle BI Publisher is fully integrated with OBIEE 11g, it can also be deployed separately.
- *Ad hoc Query and Reporting:* A powerful ad-hoc query and analysis environment that works against a logical view of information from multiple data sources in a pure Web environment. This single interface is designed to seamlessly handle both relational and OLAP style analysis.
- *Interactive Dashboards:* Rich, interactive pure Web dashboards that display personalized information to help guide users in effective decision making.
- *Scorecard and Strategy Management:* Extends the Oracle BI Enterprise Edition (OBIEE) with capabilities that enable strategic goals to be communicated across the organization and monitoring progress over time. Oracle Scorecard and Strategy Management includes visualizations that graphically communicate strategy & strategic dynamics using Strategy maps, Cause and Effect diagrams, and Custom views. Scorecard metadata objects and visualizations are treated just like any other OBIEE 11g metadata object and can be easily embedded in dashboards, ad-hoc query and analysis views and can be monitored as alerts.
- *Actionable Intelligence:* Consists of an Action Framework that provides the ability to invoke a workflow, web services, web content, additional BI content, java method, and other custom procedures from any delivery channel and an alerting engine that captures and distributes notifications via multiple channels in response to pre-defined business events and/or data exceptions to speed exception based decision making.
- *Integrated Search:* Ability to search existing content based on full indexing of Dashboards, Analyses, Views, Prompts, KPIs, Scorecards, Publisher Reports, Agents, Actions, Catalog, and Folders. Ability to drill into BI with context; Index metadata & prompts. Search results can be secured via SSO integration.
- *BI on the go:* Consists of capabilities to provide Business Intelligence content when the user is not directly connected to the enterprise network. Includes Briefing Books — reports that capture a series of snapshots of an Oracle BI Dashboard or report allowing the information to be viewed offline in presentation style; rich integration with Microsoft Office allowing for interaction with BI content and access to pre-built analysis and mobile from Office products.

Systems Management Components

- *Oracle Enterprise Manager Integration:* Providing centralized, comprehensive web based management of small to enterprise level systems. This enables an Oracle BI system administrator to manage a multi server enterprise system from a single interface.

ORACLE BI PRODUCT STRATEGY

The Oracle BI Foundation Suite is designed to meet the requirements for a new class of enterprise business intelligence solutions. It consists of a broad set of capabilities including ad-hoc query and analysis, interactive dashboards, scorecards, reporting, proactive intelligence and alerts, mobile analytics, and more. The Oracle BI Foundation Suite is designed around the following principles:

- *Unified Enterprise View of Information:* Virtually every organization has information fragmented in multiple repositories and enterprise applications. The Oracle BI Foundation Suite enables organizations to define a single, consistent, and logical view of enterprise information across these heterogeneous systems such as data warehouses, multidimensional sources, and operational transaction systems. It provides the business with a unified, enterprise view of their information.
- *Unified Semantic View of Information:* The Oracle BI Foundation Suite allows an organization to model the complex information sources of their business as a simple, semantically unified, logical business model. It provides facilities to map complex physical data structures including tables, derived measures, and OLAP cubes into business terms - abstracting how a business user expresses calculations. It translates familiar, easy-to-understand business concepts into the technical details required to access the information. The Oracle BI Foundation Suite is unique in the market because it defines an enterprise semantic layer that spans across the unified enterprise view of information.
- *End User Self Service:* The Oracle BI Foundation Suite provides business users with the ability to access the information they need without for the need for assistance from professional analysts. Because end-users work with the unified, semantic view of the information, they are provided with self-service access to analyses across multiple sources via multiple delivery channels while maintaining a consistent definition of the information. Business users only need to understand a single, business-oriented view of their information.
- *Real-time Information Access:* With technologies like trickle feed ETL, Business Activity Monitoring, Business Event Management and federated data access directly from transaction processing systems, the Oracle BI Foundation Suite allows users to combine historical and real-time information to get an up-to-the-minute view of their business. In addition, Oracle BI can combine data from real-time systems with data in the Data Warehouse to give unparalleled insight into the business.
- *Insight-driven Action:* The proactive intelligence facilities of Oracle BI Delivers and the Guided Analytics facilities of the Interactive Dashboards are designed to help business users navigate information quickly and to effectively troubleshoot problems and take action proactively in response to business events.
- *Unified Platform:* The Oracle BI Foundation Suite is an integrated suite sharing a service-oriented architecture; integrated data access services; integrated analytic and calculation infrastructure; integrated metadata management services; a common semantic business model; an integrated security model and user preferences; and integrated administration tools which improve access to information and lower operational costs.

FOUNDATION SUITE SERVER TECHNOLOGY

The Oracle BI Foundation Suite features the industry’s best-in-class server technologies for relational and multidimensional analysis. This section describes the rich capabilities of the Oracle BI Server and Oracle Essbase as well as the associated tools to develop and maintain applications and metadata.

ORACLE BI SERVER

Oracle BI Server is a highly scalable, highly efficient query, reporting and analysis server that provides services that enable the other components of the Business Intelligence Suite such as Analysis & Interactive Reporting, Dashboards, Data Mining and Analytic Applications.

The Oracle BI Server exposes its services through standard ODBC and JDBC-compliant interfaces. Clients of the Oracle BI Server see a logical schema view independent of the source physical database schemas. Oracle BI Server clients submit “Logical” SQL, which ultimately gets translated by the server to native, source-specific data source query languages like SQL and MDX. Intermediate processing to calculate complex business metrics and integrate multiple data sources occurs within the Oracle BI Server Execution Engine. The Oracle BI Server infrastructure includes facilities such as session and query management, cancellation, statistics logging, monitoring, and other server administration functions.

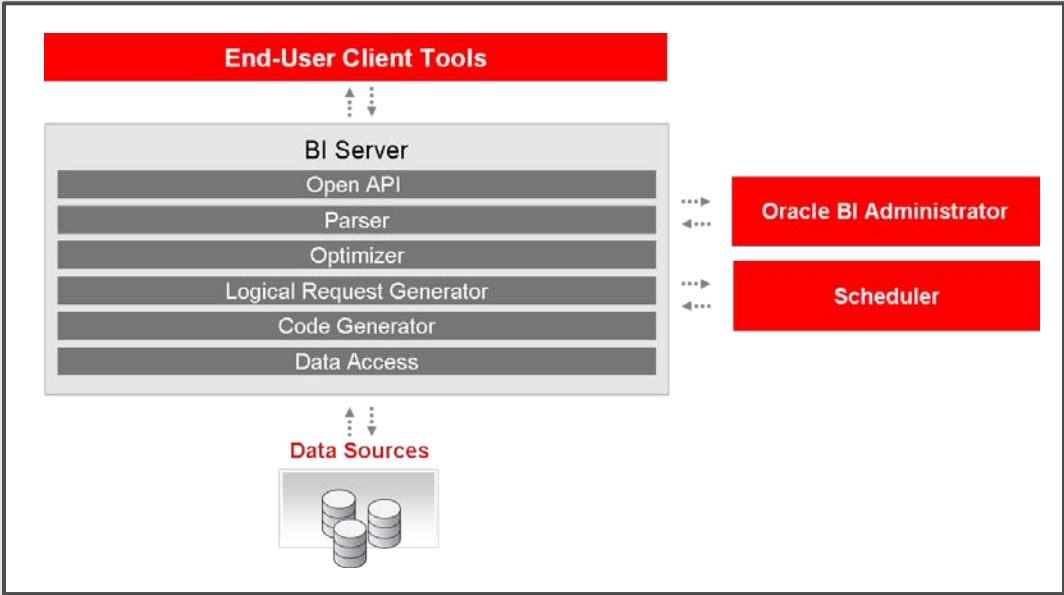


Figure 2 - BI Server Functional Components

COMMON ENTERPRISE INFORMATION MODEL

The advanced semantic layer structure used in the Oracle BI Server is called the Common Enterprise Information Model. This model provides a single version of truth for all BI system users and applications. It takes advantage of all the features of the BI Server. This model is layered to provide flexibility and maintainability:

- *Physical Layer:* models each physical source's connection parameters and schema. In the case of a relational source, the schema would include tables, columns, joins, and security parameters. Metadata rich multi-dimensional sources like Oracle Essbase, Oracle OLAP Microsoft SQL Server Analysis Services or SAP BW, the full metadata models are imported including measures, dimensions and hierarchies. This is the only layer that is aware of the physical nature of the source, such as whether it is relational 3rd normal form, star, snowflake, multidimensional cubes or XML. If the source is a database, this layer is the only one that is aware of what brand and release the database is, and what functions it does or doesn't support.
- *Business Model and Mapping Layer:* models the way the business elements function: conformed dimensions and hierarchies, measures (including aggregation rules, complex business calculations, dimensionality and time series), data security rules, and human-readable attributes and dictionary definitions. The mappings from the semantic objects back to the physical objects define the federation and aggregate navigation across multiple sources. Because of this layering and mapping, the physical source can migrate to a different brand of database, or even add an aggregate, without impacting the business model, presentation layer or reports.
- *Presentation Layer:* organizes the semantic objects, or "logical columns," into "logical tables" that can be exposed to users by role. Presentation tables and columns are completely localizable allowing a single implementation to consistently support users across languages around the globe. These are normally the only objects in the semantic layer that are exposed via the ODBC and JDBC interfaces, whether the client is Oracle BI, a custom program or a 3rd party BI tool. This allows the administrator to provide subject organization to make objects easy for users of Oracle BI Foundation clients or other third party client tools to find, as well as to apply role-specific security.

THE ORACLE BI SERVER PROVIDES THE FOLLOWING KEY CAPABILITIES

Query Parsing and Compilation

At a simplified level, the internal layers of Oracle BI Server have two primary functions: (A) compile incoming query requests into executable code, and (B) execute the code. Query compilation is composed of the following five phases: (1) parsing, (2) logical request generation, (3) navigation, (4) rewrites, and (5) code generation. The final output of the query compiler is executable code. This code is passed to the execution engine that is then responsible for executing the code in parallel. The Oracle BI Server has ground breaking innovation in query parsing and compilation techniques; content aware data federation; parallel execution; connectivity adapters; custom memory management and latch contention.

- *Parsing:* In the first compilation phase, the multi-threaded parser accepts the full ANSI SQL compliant syntax (including sub-queries, derived tables, set operations, etc...) and generates a parse tree as its output. Subsequently, the logical request generation component is responsible for instantiating the inferred aggregation in the simplified SQL supported by the Oracle BI Server.
- *Logical Request Generation:* The navigation and rewrite phases do the bulk of the work in compiling a query. The output of these two major phases is an execution plan that is then fed into the code generation phase. The navigator is responsible for the "content-

aware” data federation capabilities, such as partitioning by source, partitioning by value range, and aggregate navigation; the input to the navigator is a logical request tree describing the precise semantics of the requested data while its output is the initial physical execution plan. The navigator exploits knowledge of content to eliminate the predominant majority of traditional multi-database joins. It also has built-in support for common business analytics such as time series comparisons, shares, and dimension-specific aggregation rules.

- *Rewrite/Optimizations*: Once the navigator generates the initial physical execution plan, the rewrite phase of the compiler is responsible for distributed relational query optimization and optimal native SQL or MDX generation. This phase covers (i) Multi-database join plan generation; (ii) Function shipping; (iii) Functional compensation analysis; and (iv) Optimized SQL or MDX generation. (Functional Compensation means the BI Server executes the query function when the physical source database does not support the function requested by the logical query, such as RANK() against a Microsoft Access database.) The Oracle BI Server’s join engine is seamlessly invoked when necessary, as determined by the following: physical location of tables, SQL functionality supported by the source database(s), and analytical complexity of the original logical query. Join plans are constructed to maximize collective function shipping down to the source databases. Two types of internal join strategies are currently supported: (1) sort/merge and (2) parameterized nested loop joins (PNLJ). (PNLJ optimizes cross-database joins by fetching a small result set from database A and joining it to a large table in database B using a parameterized query, thus avoiding the slow operation of fetching the large result to the BI Server to join it there.) The BI Server further supports federated lookup tables allowing disparate systems to lookup values across one another and merge result on the fly. Optimal function shipping reduces loads on the source database and the network. The most important query processing elements to function ship include GROUP BY and aggregation; Filters; and Multi-pass SQL operations.
- *Equivalence Preserving*: aggregate and filter rewrites may push aggregates and filters through the tree (past operators such as joins, UNION ALLs, etc.) down to the database, thus reducing database load and network traffic. Both WHERE and HAVING filters may also be pushed to the database, depending on the GROUP BY clause.

Code Generation

Code generation is responsible for producing the native query processing language directives to communicate with heterogeneous, remote databases (i.e. physical SQL or MDX generation). It is also responsible for generating the code for any remaining query processing that has not been function shipped to remote databases. This includes the insertion of parallel execution directives for the Analytics execution engine.

Parallel Execution Engine

- The Oracle BI Server execution engine is a state-of-the-art; parallel SQL execution engine extended with analytical execution operators. It leverages the sophisticated technology and architectural concepts developed over more than 20 years in the database research community. Some of its key features:

- **Function-Shipping:** The Oracle BI Server ships directives for native SQL or MDX query strings; directives to execute one or more aggregation passes; and directives for various types of filters to the source database.
- **Parallel Query Execution:** The Oracle BI Server allows multiple queries to be submitted and executed in parallel, perhaps on different machines. Any cancellations would also be done in parallel.
- **Sort Optimizations:** If sorts required for the FULL OUTER JOIN cannot be pushed to the databases, the Oracle BI Server has facilities to allow sorts to be done in parallel. It ensures that no rows are lost between the two queries.
- **Merge:** The Oracle BI Server has sophisticated join facilities to merge two or more result sets from several parallel queries.
- **Ranking and Filtering:** The Oracle BI Server can rank and filter rows efficiently.

Information Reliability

Oracle BI Server defines and stores all the elements of analytic calculations as metadata in a central repository. This provides a centralized, consistent definition of measures for all users. Should the definition of a measure need to change, it needs only be changed within the central repository and all analyses and existing reports automatically use the new definition. This eases the maintenance burden and lowers cost of ownership.

Oracle BI Database Gateways

Oracle BI Server has an extensible and open connectivity layer with a set of adapters that are responsible for communicating with source data servers. An Oracle BI Gateway is a dynamically loaded library that can be configured to run within the Oracle BI Server process itself or in an external process. Individual adapters have been built to communicate with for the following systems:

- Relational Database System including Oracle Database, Oracle Exadata Database Machine, Oracle TimesTen In Memory Database, DB2, DB2, , Microsoft SQL Server, Teradata, Netezza, Informix, Sybase and other ODBC compliant data sources
- OLAP Sources including Oracle Essbase, Hyperion Financial Management, Oracle Database OLAP Services, Oracle RPAS, Microsoft Analysis Services Cubes, and SAP BW Infocubes.
- XML Data Sources including access to other types of data servers (e.g., other non-relational servers), Microsoft Excel spreadsheets, and Web Services.

Mission Critical Performance, Scalability, and Reliability

Oracle BI Server has a number of performance, scalability, and reliability optimizations to provide optimal performance and scalability whether users are constructing new analyses; changing the visualization of an existing analysis; or refreshing several analyses embedded on a single dashboard. The most important performance and scalability features are described below.

Highly Efficient Oracle BI Server Design

The Oracle BI Server offers several performance and scalability optimizations including custom heap memory management to avoid memory contention issues; hashing to avoid central locking; specialized synchronization mechanisms such as spin latches; parallel query and computation execution engines; and high-throughput connectivity adapters. When performance requirements exceed the capability of a single server, Oracle BI Servers can be clustered together with session replication and automatic fail-over. Oracle BI is architected to leverage the capabilities and scalability of modern 64-bit operating systems.

Highly Efficient Data Sourcing and Aggregation

Oracle BI Server minimizes data retrieval time by selecting the most efficient data sources to satisfy user queries. It is aware of and automatically selects "aggregate tables" in relational databases or cubes in multidimensional sources like Oracle Essbase or Oracle OLAP. Pre-aggregating and storing additive information is the standard practice for improving the query performance of relational databases. When users request information at a high "grain" of aggregation, the Oracle BI Server can use the pre-aggregated sources instead of requiring the database to add up the detail at report time. Oracle BI Server can select appropriate summary tables in lieu of the detail table based on where the requested columns are located in their respective hierarchies.

Oracle BI Server further builds its own summary aggregates through its data mart automation feature. This feature builds, refreshes and queries summary data stored in standard relational databases or in memory databases like Oracle TimesTen Database.

Aggregate navigation or transparent query rewrite across federated relational, in memory and multidimensional data sources ensure that the entire available data architecture is fully leveraged.

Exploiting Database Facilities

Oracle BI Server also optimizes performance and minimizes network traffic by exploiting the native capabilities of the available relational and multidimensional database platforms. When generating SQL (or other query languages such as MDX), the Oracle BI Server is aware of the functions and language constructs that the database supports and generates highly optimized target-specific queries. The Oracle BI Server "function-ships" this optimized SQL or MDX to the database conducting as much processing as possible in the database itself. Examples of such differences between databases include string processing, statistical and mathematical functions; logical if-then-else statements; expression maps in HAVING clause; and others. Conversely, if the database platform does not support a function or a SQL feature, the Oracle BI Server will itself compensate for the missing functionality using its own computation and data processing engine. By doing so, it exploits the advances in query optimization, indexing, data partitioning and other technologies in relational databases. Note that the Oracle BI Server can perform a superset of the data manipulation and calculation capabilities of SQL compatible database products. This ability to customize the query language to the platform and to compensate for missing functionality is unique to the Oracle BI Server.

Oracle BI natively (or via ODBC) supports virtually all major relational and multidimensional data sources but has unparalleled optimizations for the industry leading Oracle Database and Oracle Exadata Database machine. Oracle BI's Oracle Call Interface (OCI) integration, query gateway, extensive use of Oracle SQL grammar and integration with other defining features like Oracle Virtual Private Database, Oracle Spatial and Locator, Oracle OLAP Option and Oracle Data Mining make Oracle BI the industry's standard for Oracle Database. No 3rd party BI vendor has the capability or knowledge to integrate better with Oracle Database than Oracle BI.

Connection Pooling

The Oracle BI Server can be configured with one or multiple connection pools for each database. The administrator can specify a maximum number of database connections to keep open until they are unused for a specified period. As the query load increases, the number of open connections increases in the connection pool. When the maximum number is reached, the server will queue new connection requests. This prevents database servers from being overloaded. With more than one connection pool configured per database, specific users or groups of users can be assigned to specific connection pools. This allows an administrator to give certain groups higher priority.

Query Reuse and Caching

When multiple users access the Oracle BI Server, many queries will have similar content allowing the Oracle BI Server to intelligently re-use previous query results, a capability called "query caching". These are the caching methods available:

- *Web Server:* Oracle Analytics' Web Server caches queries and query results. When a user submits a query, the web server examines the logical SQL to see if it matches an existing cached query. If it does, then the Web Server uses the results without re-submitting logical SQL to the Oracle BI Server. As a user generates new data views, manipulates a pivot table, or returns to a recently viewed dashboard page, the Web Server uses cached results. The user can explicitly "refresh" the query if needed.
- *Oracle BI Server:* Query caching is a highly differentiating feature that also occurs inside the Oracle BI Server. The Oracle BI Server saves each component of a logical query, the text of the logical SQL component, the time and date of the query, the list of physical tables used in the SQL (or other query language), and the results of the query. The Oracle BI Server will analyze each new query it receives and determine whether it can answer it using cache. Oracle BI Server will refresh reports leveraging a mix of on the fly data source queries and cache, when available, to provide the fastest possible end user query experience.
- *Database Server:* The Oracle BI Server also allows queries that require extensive database processing to be pre-scheduled to run so that results are already available when users open their dashboards.

A frequently experienced benefit of caching is improved dimensional browsing performance. Since it has been estimated that 80% of user queries to a data warehouse are pure dimensional browses, this results in a significant reduction in database activity and improves the responsiveness of the system.

Scalability and Availability

The Oracle BI Server supports clustering for high availability and scalability. The clients themselves may be active-active clustered, as in the case of the Oracle BI Presentation Services, BI Publisher or Delivers.

The BI Server cache is cluster-aware in order to maximize the performance benefit of cache seeding. Also, on-line metadata changes can be made against the BI Server designated as the Master, and then automatically synchronized with the other BI Servers in the configuration to maintain information reliability. Oracle BI publishes benchmarks on industry leading hardware and operating systems demonstrating linear scalability for 10s of thousands of concurrent users.

Accessing Oracle BI Server Information

Oracle BI Server presents itself to other applications as ODBC or JDBC data source or also as web services. This means that virtually any ODBC or JDBC-capable report writer or web service enabled query tool can use the Oracle BI Server as if it were a relational database. When it does, the

query/reporting tool: (i) does not need connectivity to underlying data sources; (ii) is completely insulated from changes in source tables and database platforms; (iii) benefits from BI Server caching, aggregate awareness, and other performance accelerators; (iv) automatically takes advantage of the built-in security and connection pooling of the Oracle BI server, and (v) can use all the tables and columns of the Presentation Layer subject area of the Common Enterprise Information Model as if they were stored in a single simple database schema. This enables reporting tools to leverage all the derived measures contained in the logical data model the same as any other column. Users of these tools are insulated against returning erroneous results as a result of incorrect table joins or missing data – SQL traps sometimes known as chasm traps, fan traps, or missing data traps.

Multiple Layers of Security

Oracle BI Server enforces multiple layers of security across objects and data: Data access at row-level (implemented either in the repository or in the database), object permissions and query limits (governors). Oracle BI Server will leverage session or user level variable to dynamically apply security rules to each incoming query. This enables fully personalized environment and secured data access for each end user.

Physical Data Storage Independence

The Oracle BI Server and its Common Enterprise Information Model eliminates the need for business users to understand physical data storage and enables them to combine data from multiple enterprise information sources quickly and easily. Some of the key features of Oracle BI Server in this area are:

- *Combining Structured Data from Multiple Sources:* Oracle BI Server allows users to combine data from multiple applications or databases in a single calculation. For instance, to compare sales forecasts, quotas, and actual revenue to accurately predict revenue growth, a business user may need to combine data from three sources – the forecasting system, the sales system, and the general ledger. Within the Common Enterprise Information Model and Oracle BI the three sources appear as one logical source to the business user.
- *Combining Relational and OLAP Data Sources:* Oracle's Common Enterprise Information Model allows users to combine data from a relational system and an OLAP source in a single calculation. For example, a user can compare sales forecasts from an Oracle CRM System with budget data from an Oracle Essbase planning application. To clients of the Common Enterprise Information Model, the forecast and budget data appear to be from the same logical source.
- *Combining multiple Relational databases, in memory databases, or Relational and OLAP, for Aggregate Navigation:* Pre-building measure aggregations during the load window is the most important data warehousing practice for achieving good query performance on large datasets, usually producing three or more orders of magnitude improvement overall. The Oracle BI Server has sophisticated “aggregate navigation” features to take advantage of all available aggregates transparently - users see the performance improvement without being aware of the extra tables. The BI Server uniquely allows the aggregate tables to exist in a different database than the detail tables, or for the aggregates to use a multidimensional source such as Essbase for lower TCO.

- *Combining Relational and Spreadsheet Data Sources:* Oracle BI allows users to combine data from relational databases with non-relational data from Excel spreadsheets, for example, in a single calculation to compute a complex metric.
- *Combining Transactional Data with Data Warehouse Information:* Finally, the Common Enterprise Information Model allows users to combine data from a data warehouse with information from transaction processing systems in a single calculation to get the most up-to-date value of a metric.

Unlike the Common Enterprise Information Model, many business intelligence tools restrict users by allowing access to only information from a single data source for a specific calculation or analysis. In addition, some of these tools even restrict a user's access to a single data source during an entire session. The Common Enterprise Information Model uniquely enables pervasive access to information to answer a business question that may require data from one or more sources.

Complex Business Measures

Oracle BI's Common Enterprise Information Model allows users to define complex business measures — such as market share changes versus a year ago or sales percentage changes versus a year ago — in calculations. Some of the key features of Oracle BI business measures are:

- *Complex Business Measures:* are a challenge to compute in SQL or in most commonly used reporting products because they either: (i) involve "row to row" comparisons, something SQL was not designed to do, or (ii) involve queries that combine multiple levels of aggregation. The Oracle BI Server allows complex business measures to be calculated at query execution time without having to pre-calculate and store data.
- *Eliminates Time-based Reporting Tables:* The Oracle BI Server eliminates the need to create and store complex time-based reporting tables. For instance, most organizations have tables structured with N*M columns representing the last N periods of data for M measures plus N*M more showing the variance from last year and so on. Oracle BI Server makes these measures available by simply defining them in Common Enterprise Information Model metadata, thereby eliminating the need to build and physically maintain such tables.
- *Derived Measures:* Oracle BI Server simplifies the use of derived measures, i.e. measures that are computed on a query result set, such as ranks, Ntiles, standard deviations, running totals, moving averages, and moving medians. These derived measures are difficult to compute in SQL but are very useful — moving average and moving median are valuable functions for smoothing data and discerning trends. Oracle's Common Enterprise Information Model allows users to define new formulas using existing measures.

Integrated Segmentation Engine

Oracle BI Server includes a flexible segment and list designer engine that leverage the Common Enterprise Information Model to build highly targeted and sophisticated lists or segments, Oracle BI Server optimizes the complex SQL generation and dynamically persists relevant information. The result is an optimal query design for retrieving of lists, samplings and record counts of complex filtering operation made on extremely large datasets.

Oracle Business Intelligence Administration

The BI administration tool is used by administration-role users to create, manage and maintain the Common Enterprise Information Model described above. The administration tool has been designed with wizards, utilities, and interface design elements to help the administrator work efficiently with real-world, large-scale enterprise metadata.

- *Calculation Wizard:* helps administrators write formulas (e.g. percent share) and assures their correctness
- *Metadata Import Wizard:* connects to each type of data source and populates the physical catalog metadata for that source. In the case of multidimensional sources such as Essbase and Hyperion Financial Management, even business model semantics such as dimensions, hierarchies and aggregations are imported and populated.
- *Open BI Server XML Metadata API:* Oracle BI provides an XML based API for Common Enterprise Information Model metadata exchange. This API enables conversion of the entire OBIEE 11g repository to XML and back. This open API enables extraction, reuse and manipulation (add/update/delete) of OBIEE 11g metadata.
- *Aggregate Persistence Wizard:* enables the administrator to use the Common Enterprise Information Model metadata to design and automate the deployment and loading of aggregate tables, and to automatically create their mappings in the metadata. This significantly lowers the TCO of the very important performance technique of pre-aggregation.
- *Global Change Utilities:* A rename wizard makes it easy to change the tech-oriented names of multiple physical data objects to more human-friendly names at once, substituting text, changing case, and adding prefixes or suffixes. Similarly, the administrator can set the aggregation rule for dozens of measures all at once, rather than one column at a time.
- *Dependency and Impact Analysis within the Common Enterprise Information Model:* A query utility allows the administrator to find metadata objects by type, while filtering on properties and relationships to other objects. For example, an administrator could find all logical columns that are dependent on specific physical table or column to determine which subject area columns will be affected if a certain physical column is deleted in the database.
- *Dictionary:* The administration tool provides facilities to export Common Enterprise Information Model metadata, such as formulas and human-readable object description fields, to create dictionary functionality for end users. Answers users will see dictionary information in roll-overs of catalog objects, with links to the HTML dictionary page associated with the object. Links on that page lead to definitions of other objects it is derived from.
- *Session Management:* The administration tool offers a way to view (and terminate) current user sessions; see the variables being used in each session; list the available cache entries by subject area, user, or physical table; and report on the recent history of cache usage. Usage logs written by the Oracle BI Server(s) can provide a basis for understanding usage patterns, response times, and load variations. This information is useful for diagnosing and tuning systems. Security rules enforced in the source databases can be used together with security rules enforced in Oracle BI.




Figure 3 - BI Administration Tool

Multi-User Development Environment

Two distinctive features of the Common Enterprise Information Model are its enterprise scalability, and its support for portable BI applications. The Administrator environment provides the ability to develop and manage applications of this scale and portability.

- *Three-way Merge:* One reason why the Oracle BI Server is the only BI platform with a successful set of BI applications is its powerful three-way merge. This enables customers to update a configured application to the next release without losing their changes. The rule-based algorithm automatically resolves conflicts when possible, and presents the developer with a simple decision list to resolve the remainder.
- *Branching:* As in code development, organizations can use branches to manage parallel projects on different schedules. Project check-out ensures each branch or sub-branch is self-consistent and unit-testable. A streamlined form of three-way merge is used to check branches back in so they can be integration-tested with other projects and migrated to production.
- *Development Sandboxes:* Individual developers can check out smaller projects to do their development and unit testing on their own, private BI stack.
- *BI Server XML API:* The Common Enterprise Information Model has an XML schema and utilities to enable export, import and altering of individual objects.
- *Patch Creation and Application:* Developers can create XML patches to incrementally migrate content from one model to another.
- *Bug Fixing:* When a production bug must be fixed without impacting large projects under development for the future, the developer can either use an XML patch to apply a few

individual object changes, or check out a new branch from production to make bigger changes.

- *Migration:* Enterprise Manager migrates new repository versions from development to production, and enables a zero-downtime rolling restart. The XML API provides the ability to automate any parameter changes required.

ORACLE ESSBASE

Oracle Essbase is the market-leading multi-dimensional OLAP server that enables the development of advanced forward-looking analytic applications that enable speed-of-thought analytics. By leveraging its self-managed, rapid application development capabilities, business users can quickly model complex business scenarios. For example, line-of-business personnel can simply and rapidly develop and manage analytic applications that can forecast likely business performance levels and deliver "what-if" analyses for varying conditions. Oracle Essbase supports extremely fast query response times for vast numbers of users, large data sets, and complex business models.

Component Overview and Deployment Architecture

Essbase incorporates powerful architectural features to handle a wide range of analytic applications across large multiuser environments. The following illustration provides a high-level view of the information flow between the three tiers of the Essbase architecture. The client tier (on the left) includes Essbase Server clients, such as the Oracle BI Server, Oracle Hyperion Smart View for Office, and administration interfaces. The middle tier (in the center) includes services, such as Oracle Hyperion Provider Services, Oracle Essbase Administration Services, and Oracle Essbase Studio Services. The database tier (on the right) is made up of the Essbase Agent and Essbase Database. Communication between the client and middle tiers, and the middle and database tiers, is through HTTP. Communication between the client and database tiers is through TCP/IP or HTTP. Clients access is through an open API interface that includes support for the MDX language. Communication between data sources and the metadata catalog with the middle and database tiers is through ODBC and JDBC drivers.

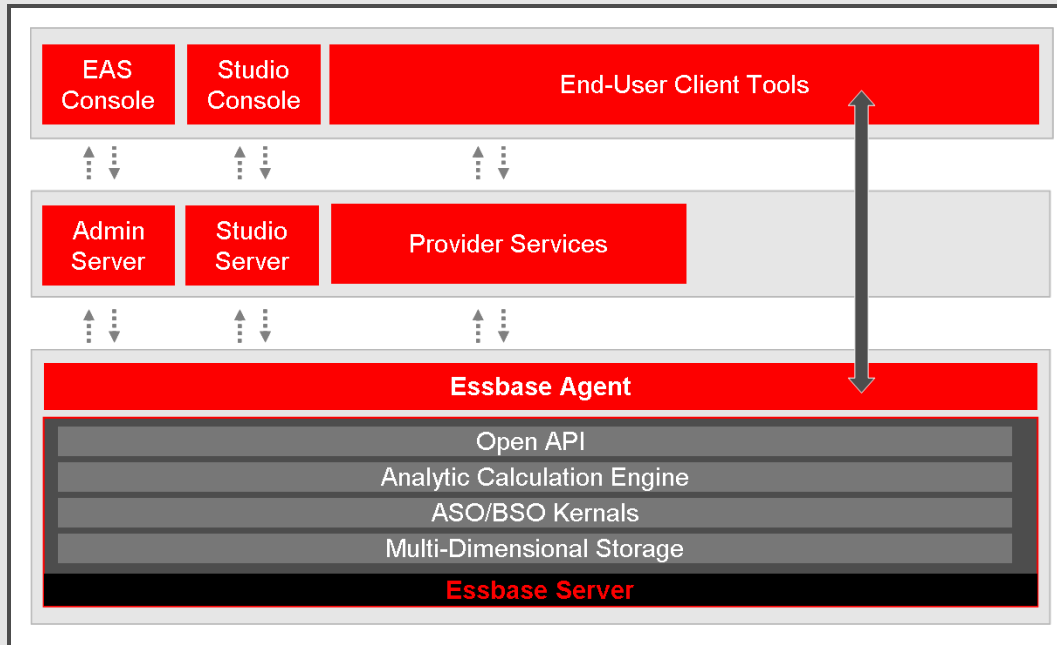


Figure 4 - Essbase Functional Architecture

All Essbase application components, including database outlines and calculation scripts, application control, and multidimensional database information, reside on a server. With Essbase, you can configure server disk storage to span multiple disk drives, enabling you to store large databases. Essbase requires a server to run a multi-threaded operating system so a server can efficiently manage simultaneous requests. A server also runs a server agent process that acts as a traffic coordinator for all user requests to applications managing communications and security. The Essbase Server leverages Oracle Process Management and Notification Server (OPMN) to manage starting and stopping agent processes. OPMN also enables Essbase high-availability services.

Optimized Multi-dimensional Storage

The Essbase server provides advanced multi-user read and write capabilities, including data update and multi-user recalculation. Business users with front-end tools can write data back to a server and recalculate the data on a server using calculation scripts—key functionality to support sophisticated modeling and planning applications.

The Essbase database is a multi-threaded OLAP database that takes advantage of symmetric multiprocessing hardware platforms. The server acts as a shared resource, handling all data storage, caching, calculations, and data security. The Essbase Server client needs only to retrieve and view data that resides on a server.

The Essbase database provides multiple storage options with unique and complementary capabilities:

- Block storage option (BSO) arranges dimensional members into dense and sparse dimensions and stores data in dense hyper cubes that are indexed by sparse dimension members. BSO is optimized for write-back and procedural calculation operations.

- Aggregate storage option (ASO) is designed to handle high-dimensionality sparse data sets and supports rapid aggregation of the data. ASO can calculate aggregate values dynamically or as needed administrators can materialize aggregate views for frequently accessed dimensional levels.
- In addition to ASO and BSO, Essbase also provides a hybrid storage model called XOLAP. With XOLAP metadata is stored in an ASO outline and the data itself resides in relational. The benefit of XOLAP is reduced data redundancy. The trade-offs include some feature limitations, including lack of write-back capabilities and performance is dependent on RDBMS abilities.

ASO databases complements BSO databases and enable dramatic increases in database dimensionality. Using aggregate storage, Essbase serves a wide range of analytic needs—financial analysis, planning, budgeting, sales analysis, marketing analysis, supply-chain analysis, and profitability analytics—all from a single analytic infrastructure. Essbase partitioning capabilities allow ASO databases to be combined with BSO databases to create a single application view allowing end-users to take advantage of the benefits of both storage models.

Essbase provides for several compression options that optimize the use of physical storage including, bitmap compression, run length encoding (REL), zlib compression, and Index Value pair compression.

Performance, Scalability and Availability

Essbase is a true enterprise-class multi-dimensional OLAP server offering unparalleled user and data scalability on a high performance infrastructure. Essbase applications have been successfully deployed in departmental BI solutions and have also demonstrated scalability to tens of thousands of users and billion cell databases. Essbase supports both 32- and 64-bit computing across both Windows and Unix operating environments. Essbase includes a number of features that support high availability and scalability.

Efficient Cache-Architecture

Essbase offers a number of memory caches to improve performance for query, load, and calculation operations. Essbase provides default size settings for each cache; however cache settings can be adjusted as necessary to optimize performance needs based upon available memory, database size, service level commitments, and batch windows.

High-Availability Options

Essbase provides both active-passive and active-active clustering capabilities. Active-passive Essbase clusters support failover with write-back to databases. Essbase failover clusters use the service failover functionality of the Oracle Process Manager and Notification (OPMN) server. A single Essbase installation is run in an active-passive deployment, and one host runs the Essbase agent and two servers. OPMN stops, starts, and monitors the agent process. Active-active Essbase clusters support high availability and load balancing. An active-active Essbase cluster supports read-only operations on the databases and requires the use of Provider Services.

Trickle-Feeds

Essbase offers a number of techniques to load incremental data into databases while minimizing maintenance windows. Database slices in ASO allow data to be trickle-fed into a database while online. With database slices, data can be stored in multiple slices. For example, a real-time slice can be loaded without impacting the historic slices, allowing users to remain active in the database and immediately see data as it is updated. In addition, by loading into a specific slice of data, incremental loads performance is optimized.

Flexible Business Model Development

Essbase offers many key advantages to help business users develop effective multidimensional applications. Business analysts can quickly develop forward-looking applications and quickly model complex business scenarios. For example, line-of-business personnel can simply and rapidly develop and manage analytic applications that can forecast likely business performance levels and deliver "what-if" analyses for varying conditions.

The basis of an Essbase analytic application is a business outline which defines the dimensions, dimension members, hierarchical relationships, member attributes, calculations, and business rules for an application.

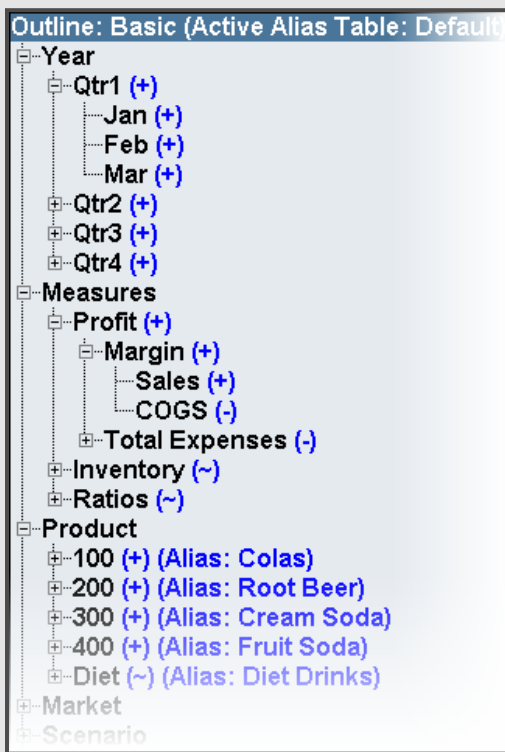


Figure 5 - Essbase Outline

Using out-of-the box tools that are delivered with Essbase, business users can:

- Use graphical interfaces to define and manage a business outline

